Washington DC ServiceNow AI Platform Administration

Last updated: 12/16/2025

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

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

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

Table of Contents

Administer the ServiceNow AI Platform

As a platform administrator, you have the power of the ServiceNow AI Platform at your fingertips. The ServiceNow AI Platform is an application platform as a service that automates business processes across the enterprise.

Get started

Choose one of these tiles to get started.

Start here to learn about the ServiceNow AI Platform

Learn how the ServiceNow AI Platform works.

Configure ServiceNow AI Platform core features

Configure settings that affect global ServiceNow AI Platform behavior and behavior for ServiceNow AI Platform applications.

Maintain and monitor the ServiceNow AI Platform

Identify health issues and settings that you need to change. Optimize ServiceNow AI Platform performance using performance monitoring and diagnostics.

Upgrade to the latest version

Plan and manage your ServiceNow AI Platform upgrades using Upgrade Center. Monitor upgrade progress and view a summary of upgrade results with Upgrade Monitor.

Additional Resources

Upgrade your platform administrator experience to the next level

Start here to learn about the ServiceNow AI Platform

Learn about the ServiceNow AI Platform and these important applications:

Admin Center

A central hub for platform owners and admins to access platform capabilities, discover new applications, and get intelligent, actionable insights. Start with Adoption blueprints, which tell you which plugins to install to meet business objectives.

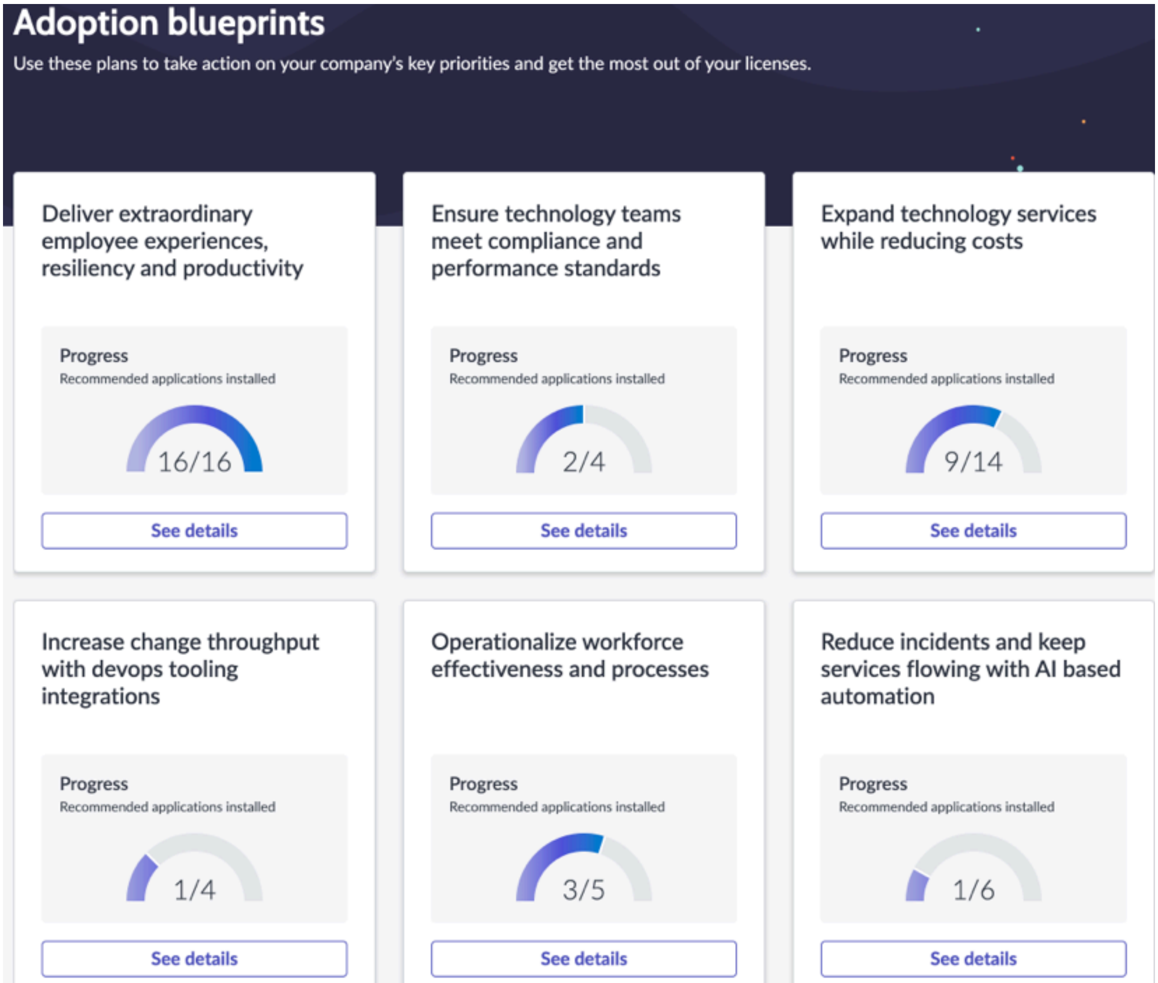

Subscription Management

A centralized management experience to help you understand your subscriptions and entitlements.

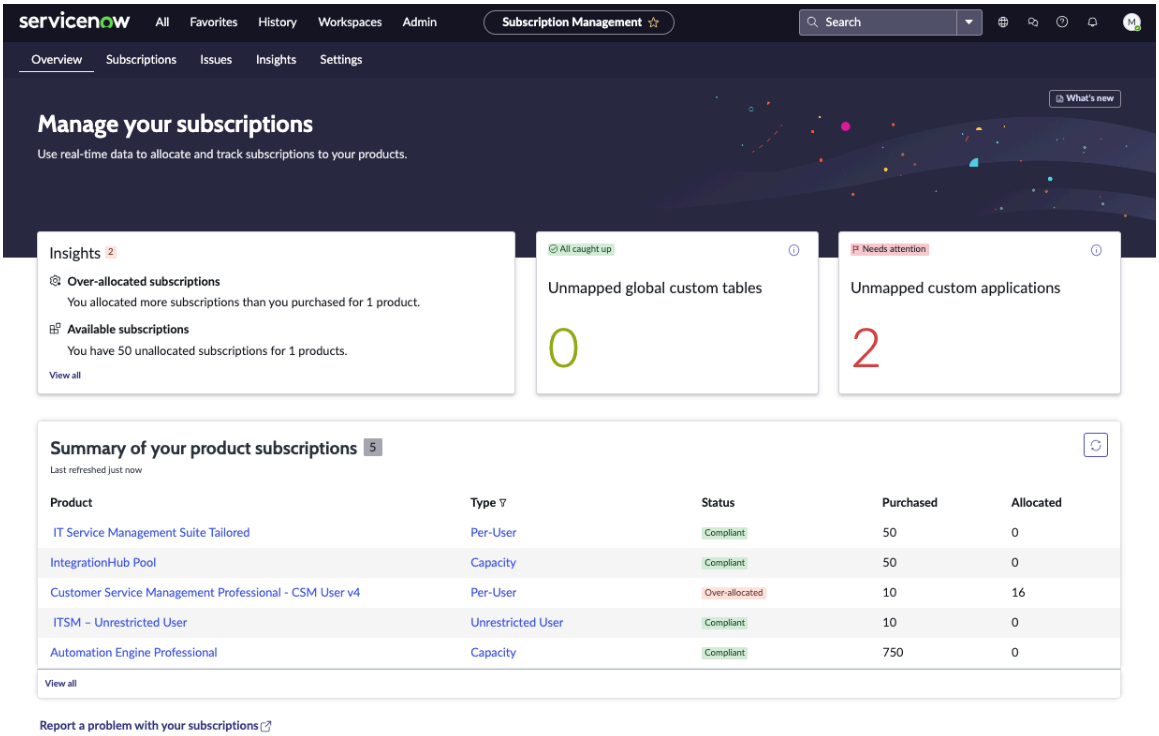

Configure ServiceNow AI Platform core features

Explore core features that affect global ServiceNow AI Platform behavior as well as behavior in some ServiceNow AI Platform applications.

User administration

Configure users, groups, and roles to provide access to ServiceNow AI Platform applications and settings.

Now Platform tables and data

Learn how to work with ServiceNow AI Platform records in the classic environment.

Search administration

Configure search engines and interfaces for the ServiceNow AI Platform. Select AI Search for a consumer-grade search experience, or use the legacy Zing text search engine. Use Contextual Search to let service agents view relevant results based on an Incident or Problem record's query context and user intent.

Time configuration and currency configuration

Configure data and time fields, time zones for your users, and calendars.

Translation and localization

Provide service around the globe with language translation and localization for specified regions.

Integrations

Learn about integrating third-party applications into the ServiceNow AI Platform.

Notifications

Use Notifications to manage system email, create system notifications, and configure how your system responds to inbound email.

Service Level Management

Configure the Service Level Management (SLM) application to gather your service requirements, monitor your SLAs, and report on your SLA compliance.

Maintain and monitor the ServiceNow AI Platform

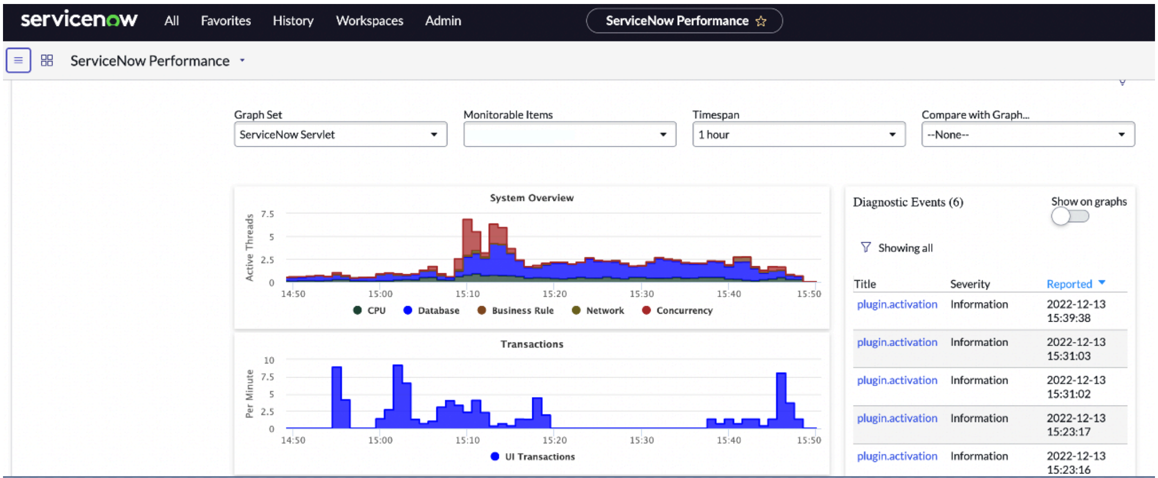

Instance Scan

Identify ServiceNow AI Platform health issues and incorrectly configured settings.

Events

Monitor system events and determine which the system should process them.

Platform performance

Ensure system optimization through performance monitoring and diagnostics.

Plan and manage your ServiceNow AI Platform upgrade

Upgrade Center and Upgrade Monitor

Plan and manage your ServiceNow AI Platform upgrades using Upgrade Center. Use Upgrade Monitor to monitor the progress of your upgrade. After the upgrade completes, you can view a summary of upgrade results.

Applications and features

- Core configuration
- Admin Center
- Subscription Management
- User administration
- Table administration
- Data management
- Field administration
- Form administration
- List administration
- Search administration

- Currency administration
- Time configuration
- Dynamic Translation
- System Localization
- State Management
- Integrate with third-party applications and data sources
- Events
- Metrics
- Platform performance
- Upgrading the Now Platform

Getting started on the ServiceNow AI Platform

ServiceNow AI Platform is the exclusive ServiceNow platform. It provides a range of options to improve and automate your business processes.

How the ServiceNow AI Platform works

The ServiceNow AI Platform is the foundation of ServiceNow. It enables data to flow seamlessly across applications and existing systems.

Basic system configuration

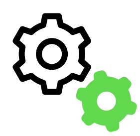

Customize UI and forms, create a mobile experience, and adjust other settings in applications across the ServiceNow AI Platform.

Admin Center

Admin Center provides a central hub for you to access platform capabilities, discover new applications, and get intelligent, actionable insights.

Subscription Management

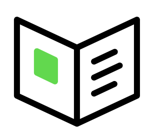

With the ServiceNow® Subscription Management application, you can manage how your ServiceNow subscriptions are used on your instances.

How the ServiceNow AI Platform works

Learn how the ServiceNow AI Platform can provide the solutions that you need to improve your business and manage risk effectively. Powerful automation, AI, and embedded analytics optimize your workflow processes, with a wide variety of resources made to help you respond to business changes quickly and easily.

Basic system configuration

Modify your UI to suit your business operations, with various core configuration settings at your disposal. Use Form Administration to customize the appearance of forms and lists. Provide tailored UI through User Interface Configuration, and improve the navigation and search experience for users with Navigation and UI. Enhance your instance and applications with Localization and Time settings, and stay on the go with Mobile Device Support. A wide range of plugins expand the platform even further, with each customizable to your needs.

Admin Center

Discover the capability of available applications and how to implement them to work toward achieving your business goal. Your administrators will be able to get started quickly with setting up apps, loading initial data, and establishing roles, groups, and users. The admin center gives your team the power to customize user interfaces, update versions, applications, and modules configurations, handle user access, monitor system performance and UI policies, and much more.

Subscription Management

Manage your ServiceNow subscriptions and how they're used on your instances with Subscription Management. This app helps you allocate and monitor subscriptions, and display current subscriptions in your instance. Active management of user-based subscriptions and monitoring of capacity subscriptions lets you track usage and adjust allocation efficiently. Enabled by default, this management app means you can plan for subscription renewals with no added costs.

How the ServiceNow AI Platform works

ServiceNow® offers a flexible, secure platform with a vast set of options, all following a single data model. Begin here to learn more about the platform.

https://player.vimeo.com/video/980759261?badge=0&autopause=0&player_id=0&app_id=58479

Watch the video above for an introduction to ServiceNow AI Platform architecture.

ServiceNow Mobile Apps Overview

Watch the video above for an introduction to the Now Mobile app.

Overview

The ServiceNow AI Platform is a highly versatile, scalable cloud solution enterprise for IT, operations, and business management, and process automation.

"The single source of truth" in your business is your data, and ServiceNow holds all your data on a single, secure platform, using a singular data model for all operations. Enterprise applications are delivered through this system, integrating seamlessly with each other and with external services. The ServiceNow AI Platform helps you address your business needs, boosting efficiency and productivity from the individual user through your entire enterprise, saving time and money.

Architecture

The ServiceNow AI Platform offers easy customization, live reports and monitoring, and strong data privacy to save your business money and boost productivity. Services on the ServiceNow

AI Platform are provided on a multi-instance basis, instead of running one instance with multiple users. Administrators can segregate their data between business entities in a multi-instance architecture, offering different experiences per instance while they all remain centrally managed through the platform. Deep customizability is available in all your instances, with easy software installation, data importing, and creation of workflows, processes, applications, or UIs.

Administration

The ServiceNow AI Platform Admin Center gives business administrators easy subscription management, a full-service catalog for services and products, and total control over workflow process optimization. Localization is available in more than a dozen languages and currencies. In-house and third-party translation services handle communications, virtual agent chat, user input translation, and more. Analytics tools provide detailed feedback on user activity, security, and other important aspects of your digital enterprise.

Integration

The Automation Engine enables connections between instances and external services on the ServiceNow AI Platform. Both unidirectional and bidirectional data integrations are available, along with email-based connection, and various other methods and technologies, all through the platform's MID Server system. Connection and automation occur in one place, with full instance and data replication, importing and exporting, and pre-built processes to address various tasks and third-party services.

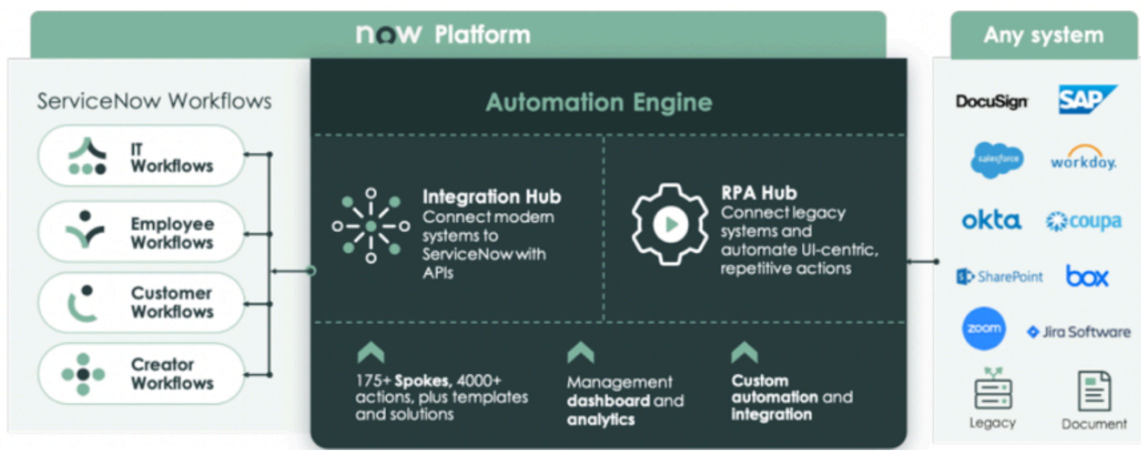

Tools

ServiceNow has a full series of options available to ServiceNow AI Platform subscribers. Ready-made tools and templates all mean easy and fast setup of your instance from the start, along with quick creation of interfaces and applications.

App Engine Studio

Speed up—and scale up—low-code development for your instance.

Flow Designer

Recreate processes as digital workflows.

UI Builder

Create user interfaces and improve user experiences.

Integration Hub

Natively integrate ServiceNow with any external service or product.

Predictive Intelligence

Perform smart automation of workflows.

Virtual Agent

Create intelligent chatbots to address users' problems and services.

Performance Analytics

Analyze and maximize workflow performance.

AI Search

Offer a consumer-grade search experience for users.

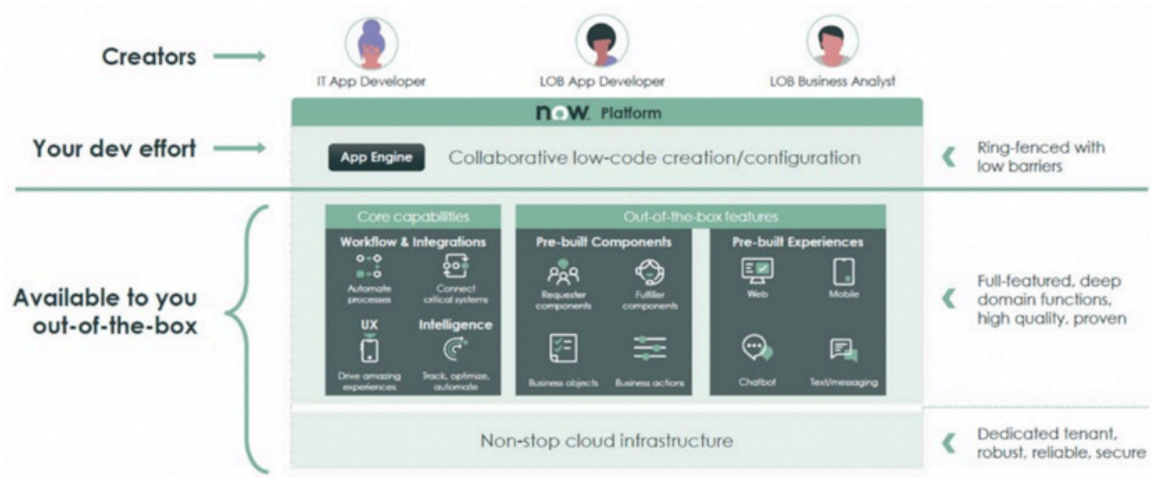

Platform Analytics

Distribute and consume Platform Analytics through data visualizations and dashboards with optional filters. Explore KPIs and receive insights into significant events in the data. Use dashboards to create a visual story with Platform Analytics data that you can share with multiple users. Use data visualizations based on table, indicator, or other data in dashboards across the Analytics Center. Administrators can leverage Usage Insights to get information on user engagement with both mobile and web applications, including user population, retention, and behavior. Process Mining helps you discover ways to streamline your business, find bottlenecks and deviations in workflows, filter and compare processes.

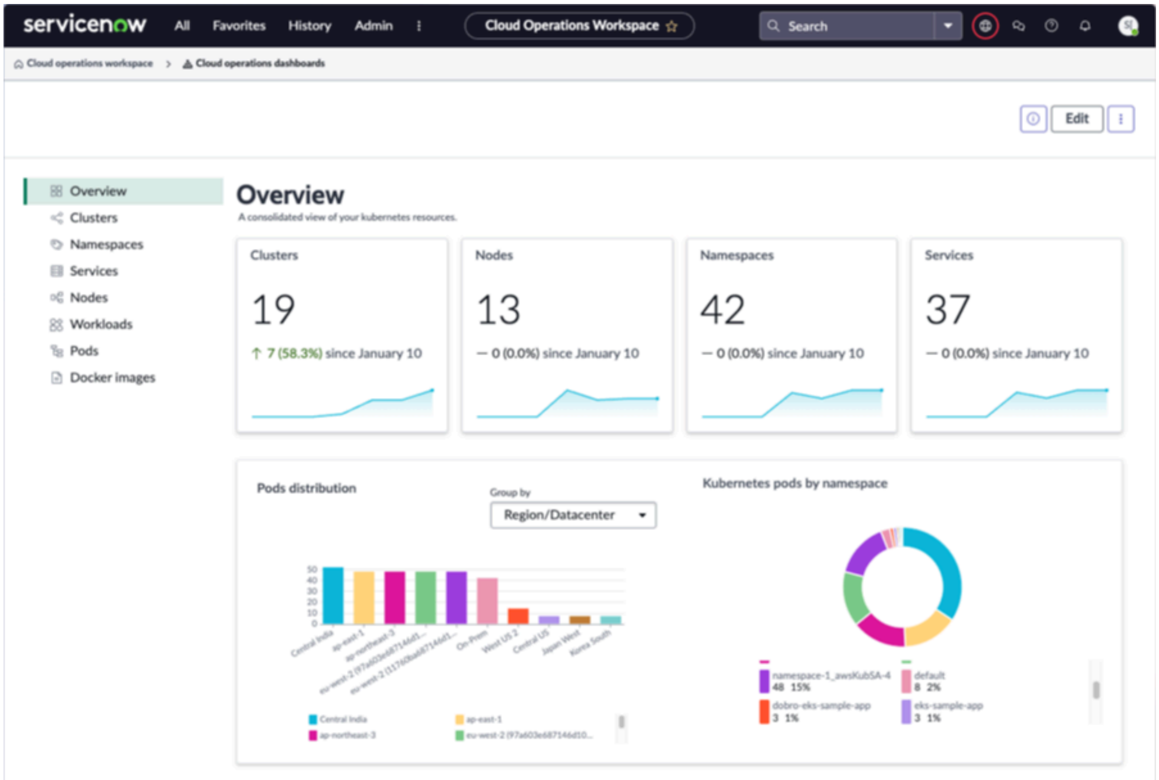

Mobile

Desktop access to the ServiceNow AI Platform is matched by our mobile applications. The Now Mobile app offers a consumer-grade experience for your customers and users to access available services, including help articles and live or virtual agents. ServiceNow mobile, geared toward service providers in your enterprise, provides tools for task scheduling and handling, communications, GPS and mapping, multimedia uploading, and offline services. Easy no-code configuration and in-depth developer tools, including reusable components, views, and a robust mobile publishing system, give you the tools to build mobile apps yourself.

Knowledge

ServiceNow provides not only the tools you need for your business to thrive, but the know-how to make use of all the options on the platform. Detailed knowledge base articles enable users and agents to get the information they need, to handle specialized tasks or processes, troubleshoot, or find other help. The Knowledge Management system enables you to create, store, and share information easily within your organization, monitor usage, and collect feedback (including full analytics) to improve your knowledge base.

For more information, see Knowledge Management.

Artificial Intelligence

AI is standard across the ServiceNow AI Platform with Intelligent Experiences. This purpose-built AI interprets user requests and messages, assists agents with task management and service owners with oversight and knowledge management. You can tailor the AI to meet instance demands for Natural Language Understanding or database querying, data extraction from documents, and prediction of future data trends. AI Search is part of the standard ServiceNow AI Platform package, for easy filtering of information according to user needs. Customized Virtual Agent help is available 24/7. Full AI-based analytics give you detailed data on the lifecycle of workflows and processes in your business.

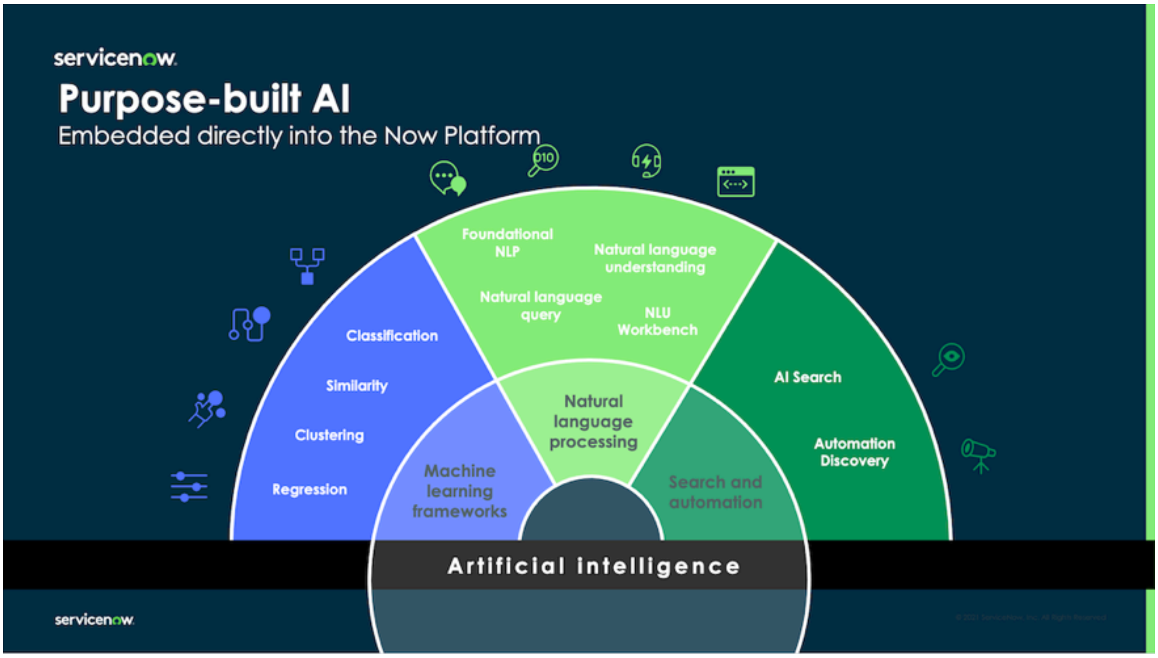

Basic system configuration

Basic system configuration encompasses changes made to the platform as well as supporting applications. These changes can affect global settings as well as settings for particular applications.

Learn about these ServiceNow AI Platform basic system components and settings:

- ServiceNow plugins
- Find components installed with an application
- Available system properties
- Query join and complexity size limits
- Web proxy

ServiceNow plugins

Plugins are software components that provide features and functionalities within a ServiceNow instance.

Plugins and applications video, approximately two minutes long.

Differences between plugins and applications

Plugins and applications are separate components on the ServiceNow AI Platform.

- A plugin adds features to installed applications and extends functionality on the ServiceNow AI Platform. For additional details about plugins, see KB0716626.
- An application is a standalone piece of code that implements a collection of features on the ServiceNow AI Platform. You can install applications from the ServiceNow Store or create your own applications.

Activating plugins

All plugins are provided by ServiceNow, but some plugins are activated differently.

- A set of plugins are activated on your instance by default. You can activate additional plugins on your instance depending on your business needs. For details, see Activate a plugin.
- Other plugins require activation by ServiceNow personnel. To request activation, follow the steps in Request a plugin.

There are some plugins that you can't activate or upgrade on your instance and the button is greyed out. This situation can happen in the following cases:

- The plugin requires a subscription that you haven't purchased. To purchase a subscription, contact your ServiceNow account manager.
  - Some features are not included in the ServiceNow AI Platform and require a separate subscription. To purchase a subscription or to find plugin costs, contact the account manager in your company.
  - When you purchase a subscription, certain plugins are activated automatically. If a paid plugin isn't activated automatically, you can manually activate it from the All Applications list in your instance. If the plugin doesn't appear in the instance, submit a request via the Now Support Service Catalog.
  - You can't activate a plugin in your production instance. However, you can test a plugin by activating it in a non-production instance. This includes plugins that are associated with a for-fee subscription. When you're ready to activate the plugin in your production instance, request the plugin from Now Support.
- The plugin depends on other plugins being activated.
  - When you activate a plugin with dependent plugins, the system notifies you of the dependencies before it completes the activation. If you opt to continue with activation, the system automatically activates all the plugins required.
  - In some cases, the plugin works, but certain features that don't apply to your system aren't installed. When you ask to activate a plugin in this category, the system notifies you of the dependencies before it completes the activation. If you opt to continue, the system activates the requested plugin but doesn't include the features that depend on other plugins.
- The plugin is already activated on your instance.
  - To determine if a plugin is already activated, see Plugin Activation Status [KB0678767].
  - After a plugin is activated, you can't disable or deactivate it. If needed, you can hide the functionality. You should thoroughly test, in a non-production instance, the features that are installed when you activate the plugin before using them in a production instance.

After a plugin is activated, you cannot disable or deactivate it. You should thoroughly test, in a non-production instance, the specific features and functionalities that are installed when you activate the plugin before using them in a production instance. See Uninstall a ServiceNow plugin [KB0716414].

- If necessary, you can hide or disable the functionality by restricting access through the use of roles and ACLs.
- Some plugins support rollback contexts. For details on using rollback contexts, see .

A warning message appears if the alter table or create table operation fails when you activate or upgrade a plugin. Contact Customer Service and Support for help with processing any skipped tables.

Personal developer instances

You can test a plugin on your personal developer instance (PDI) before launching it in a production instance. If a plugin isn't active by default on a PDI, you can activate the plugin without requesting or purchasing it.

To activate a plugin on a PDI, use the All Applications page on your instance. If the plugin isn't available on the All Applications page, activate it from the ServiceNow® *Developer Site*. For more information, see Activate a plugin on a personal developer instance.

Note:
Some plugins aren't available for activation on PDIs.

Related topics

List of plugins (Washington DC)

Activate a plugin

You can activate most plugins to expand features and functionalities in the ServiceNow AI Platform.

Before you begin
Role required: admin.

Note:
Before requesting or activating a plugin, check whether the plugin has already been activated on your instance. For details on how to check a plugin activation status, see the Plugin Activation Status [KB0678767] article in the Now Support knowledge base.

Procedure

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

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

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

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

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

Related topics

List of plugins (Washington DC)

Activate a plugin on a personal developer instance

You can test a plugin on your personal developer instance (PDI) before launching it in a production instance.

Before you begin

Request a PDI on the ServiceNow® *Developer Site*.

About this task

If a plugin isn't active by default on a PDI, you can activate the plugin without requesting or purchasing it.

To activate a plugin on a PDI, use the All Applications page on your instance. If the plugin isn't available on the All Applications page, activate it from the ServiceNow® *Developer Site*.

Note:
Some plugins aren't available for activation on PDIs.

Procedure

Activate the plugin from the All Applications page or the Developer Site.

Activation method	Procedure
All Applications page	In your PDI, navigate to **System Applications** > **All Available Applications** > **All**. Find the plugin and then activate it. For more information on using the All Applications page, see Activate a plugin.
Developer Site	**a.** Sign in to the ServiceNow® *Developer Site*. **b.** On the home page, go to the Instance Management widget and select **Manage** > **Activate plugin**. The Instance Management widget is titled **Your Instance**. It contains information such as the status, URL, and release version of your PDI. 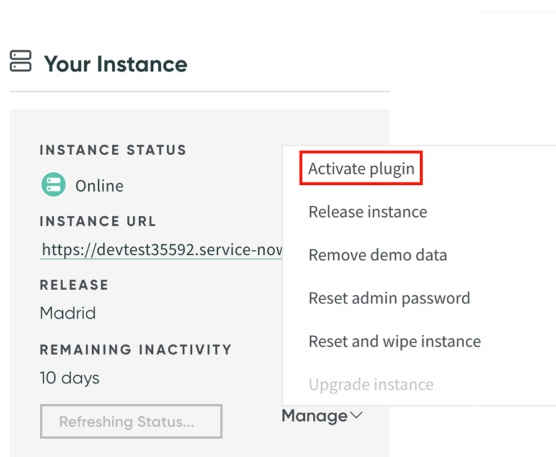 **c.** On the window, locate the plugin to activate. **d.** Next to the plugin listing, select **Activate**.

<table>
<tr><th>Activation method</th><th>Procedure</th></tr>
<tr><td></td><td>If available, you can further select to activate the plugin only or include demo data.
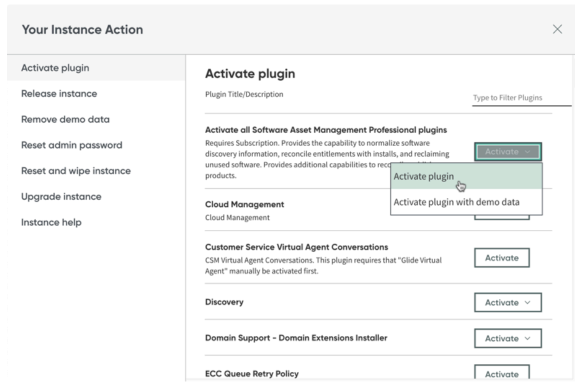
</td></tr>
</table>

Request a plugin

If a plugin does not appear on the All Applications page, you cannot activate it yourself. Request activation of the plugin by ServiceNow personnel.

Before you begin

Note:
Before requesting or activating a plugin, check whether the plugin has already been activated on your instance. For details on how to check a plugin activation status, see the Plugin Activation Status [KB0678767] article in the Now Support knowledge base.

Role required: admin

About this task

There are two ways to request a plugin:

- Access the Now Support Service Catalog on Now Support. Open the Activate Plugin page using this link: Activate Plugin.
- Access the Now Support Service Catalog through the All Applications page on your instance by following the procedure.

Procedure

1. Navigate to **System Applications** > **All Available Applications** > **All**.
2. On the All Applications page, click **Request Plugin** to open the **Activate Plugin** form on Now Support.

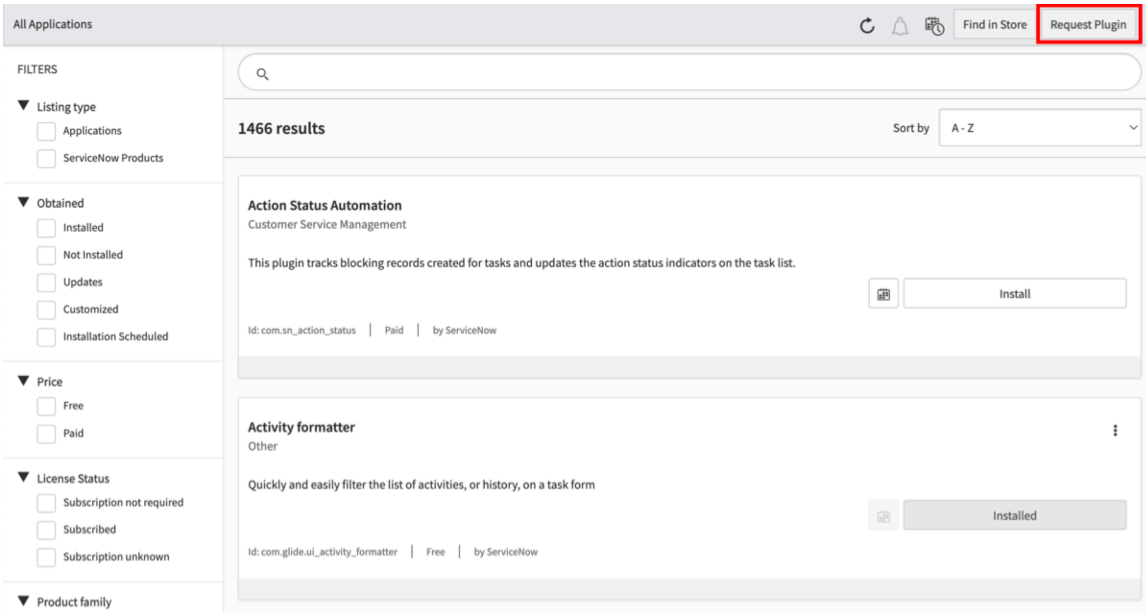

3. On the **Activate Plugin** form on Now Support, provide the following information.

Activate Plugin request form

Field	Description
What is your target instance	Instance on which to activate the plugin.
Which plugin would you like to activate	Name of the plugin to activate. **Note:** If the system does not list the plugin you want or if you are activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	The date and time to activate the plugin. **Note:** Plugins are activated in two batches, once in the morning and once in the evening, on every business day in the US Pacific time zone. If the plugin must be activated at a specific time, enter the request in the **Reason/Comments** field.

Example

For example, see the following request to activate the CSM Workspace plugin on an instance named My Instance.

servicenow.

Activate Plugin form

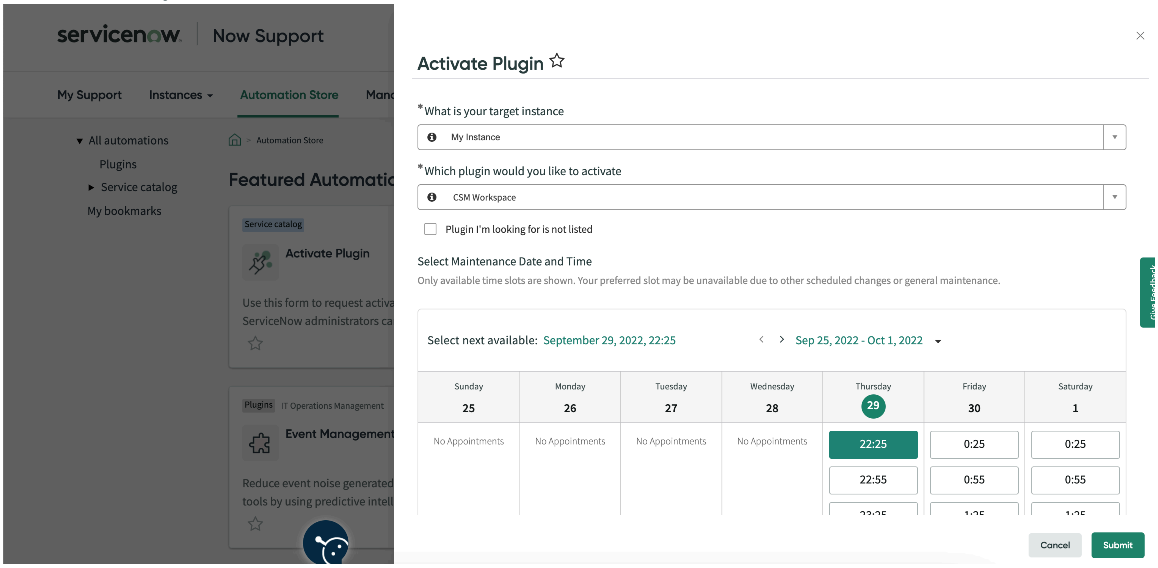

4. Click **Submit**.
   For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

List of plugins (Washington DC)

List of all plugins that you can activate if you have the admin role.

If a plugin does not appear on this list, it may require activation by ServiceNow personnel. Request these plugins through the Now Support Customer Service System instead of activating them yourself.

- For steps on activating a plugin yourself, see Activate a plugin.
- For steps on requesting a plugin that you cannot activate yourself, see Request a plugin.
- For steps on installing a ServiceNow Store application, see Install a ServiceNow Store application.

Note:
For a list of all new and changed plugins for this release, see Changes to plugins in the Washington DC release.

New plugins in Washington DC

This table lists all plugins and ServiceNow Store applications that have been introduced in Washington DC. You can activate or install these plugins or applications if you have the admin role.

Note:
By default, all ServiceNow Store applications listed in this table are compatible with Washington DC unless a note is mentioned for the application.

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
API Key and HMAC Authentication [com.glide.tokenbased_auth] (New in Washington DC)	Support API tokens for REST API endpoints. Enable API key-based authentication to securely authenticate inbound webhook URL.	Active	false	• com.glide.auth.scope • com.glide.rest.policy • com.glide.rest.auth.scope
Asset Warranty - Lenovo [com.sn_warranty_integrations_lenovo] (Available in the ServiceNow Store)	Get the warranty information of hardware assets by connecting to the Lenovo Warranty API.	Inactive	false	• sn_hamp • sn_lenovo_spoke
Capacity management [sn_cap_mgmt] (Available in the ServiceNow Store)	Enables an inventory user to view and action on the utilization of a resource. Any resource consumption will be tracked and provides a mechanism to trigger further inventory request based on capacity constraints defined.	Active	true	None
Collaborative Work Management [com.snc.cwm] (Available in the ServiceNow Store)	Provides a central hub to plan, visualize, and collaborate on work with members across the organization.	Inactive	true	• Data Grid UI Component [sn_datagrid] • Gantt UI Builder Component [sn_gantt] • Docs Component [sn_docs] • Kanban component [sn_knbn_comp] • Portfolio Planning Core [sn_align_core] • SPM Common UI Component [sn_spm_com_ui_comp]

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Contracts and Entitlement Workflows [com.sn_contract_entitl] (Available in the ServiceNow Store)	Provides workflows to manage changes on Customer contracts and entitlements.	Active	true	• Customer Contracts and Entitlements Foundation [com.com.sn_pss_core] • Lead to Cash Primitives [com.sn_l2c_core] • Customer Life Cycle Management Workflows [com.sn_l2c_cust_flows]
Creator Studio [com.app_creator_studio] (Available in the ServiceNow Store)	Provides a no-code tool that enables non-developers to build the most common types of apps for the ServiceNow AI Platform through a simplified and streamlined use-case based app creation and deployment process.	Inactive	false	[com.glide.creator_studio.global]
Customer Contracts and Entitlements [com.sn_pss_core] (Available in the ServiceNow Store)	Provides foundational objects to manage customer contracts and entitlements.	Active	true	• Customer Service Install Base Management [com.snc.install_base] • Product Catalog Management Core [com.sn.prd_pm]
Customer Life Cycle Management Workflows [com.snc.customer_lifecycle_mgmt_workflows] (Available in the ServiceNow Store)	Provides workflows to manage the life cycle of Sold Products to manage its configuration, suspend, resume and disconnect the sold product to meet customer's business needs.	Inactive	true	• Customer Service Install Base Management [com.snc.install_base] • Product Catalog Management Core [com.sn_prd_pm] • Lead to Cash Core [com.snc.l2c_core]
Desktop Assistant [sn_dex_desktop] (Available in the ServiceNow Store)	Establishes a continuous communication channel for employees, enabling them to easily access self-service options and receive timely notifications from the Service Desk. This effectively	Inactive	true	com.snc.itom.license

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	minimizes any barriers to employee productivity.			
DevOps Vulnerability Integrations [sn_devops_vul_ints] (Available in the ServiceNow Store)	Provides speed time-to-market and an organization-wide view of risk exposure for application vulnerabilities.	Inactive	false	• com.snc.ihub_spoke_util_pack • com.glide.hub.usage.dashboard • com.glide.script.vtable • com.snc.security_support.core • com.glide.hub.action_step.rest • com.glide.hub.action_type.datastream • com.glide.hub.integration.runtime • com.snc.vul_dep
DEX Content Playbook [sn_dex_content] (Available in the ServiceNow Store)	Provides the content to assist with monitoring applications and devices, as well as offer support for remediating identified experience issues.	Inactive	true	sn_dex (scoped app)
Digital End-User Experience [sn_dex] (Available in the ServiceNow Store)	Monitors applications, networks, and devices in order to detect issues before they result in downtime and reduce employee productivity.	Inactive	true	• com.sn_acc_visibility com.snc.clotho • com.glide.fenix com.sn_itom_cloud_svc • com.glideapp.itom.snac com.snc.itom.license
Digital Product Release [com.sn_dpr] (Available in the ServiceNow Store)	A dedicated release management solution designed to deliver a personalized experience for software-based release practices. Empowering central release management teams, it facilitates the definition of release policies, templates, and readiness dates.	Inactive	true	• sn_dpr_model • sn_dpr_workspace • sn_cmdb_ci_class • sn_pace • sn_pace_builder • sn_roadmap • sn_playbook_exp • now_playbook_exp • sn_devops, sn_devops_ws
Digital Product Release Data Model [sn_dpr_model] (Available	The data model for Digital Product Release. It defines the tables and	Inactive	false	com.snc.sdlc.ranking

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
in the ServiceNow Store)	relationships for release management.			
Digital Product Release Policy Content Pack [sn_dpr_content] (Available in the ServiceNow Store)	Automates and streamlines comprehensive verifications for your release process. Utilize Data Collectors and Policies aligned with the DevOps Data Model for efficient assessments of engineering completion, quality checks, and release readiness.	Inactive	false	None
Digital Product Release Workspace [sn_dpr_workspace] (Available in the ServiceNow Store)	Manage your Digital Products, Product Features and plan Versions for software release execution. The app provides capabilities required to support release execution using Digital Product Release.	Inactive	false	• com.devsnc_sn_vtb • sn_datagrid • sn_dpr_model
Dispute Rules Content Pack for Visa [com.sn_bom_visa_cp] (Available in the ServiceNow Store)	Allows issuers to easily access card network rules for initiating and investigating card dispute cases.	Inactive	false	• com.sn_ga_exp • com.sn_bom_credit_card
Document Template integration with Digital Signatures using Smart Cards [com.sn.dt-digital-signature-smart-card-integration] (Available in the ServiceNow Store)	Enables workflows for digital signing using smart cards (CAC/PIV cards) by generating and sending auto-filled documents to be signed by different participants.	Inactive	false	• com.snc.platform_document_management • com.glide.auth.mutual • com.snc.document_templates

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Employee Center Pro Kiosk	Enables all the deskless workers with simplified access to information and services for self-serve and improved productivity.	Active	true	• Employee Center Pro (sn_ex_sp_pro) • Employee Center (sn_ex_sp) • Employee Experience Foundation (sn_ex_emp_fd) • Employee Profile (sn_employee) • UKG Spoke (sn_kronos_spoke) • HR Service Delivery Integration with Ultimate Kronos Group (sn_hr_ukg) • Health and Safety Incident Management
Opportunity Marketplace [com.sn_egd_opp_market] (Available in the ServiceNow Store)	Provides a single, unified place where organizations share opportunities that are easily discoverable by employees.	Active	true	None
Email - Bounce management [com.glide.email.bounce_management] (New in Washington DC)	Enables the bounce management capability to help identify and prevent sending emails to addresses that are known to generate bounces by filtering them out while sending emails.	Active	false	com.glide.mailbox
Enterprise Asset Management for Healthcare [com.sn_eamhc] (Available in the ServiceNow Store)	Provides Enterprise Asset Management functionality targeted towards healthcare-specific roles to manage healthcare-specific models and assets and related workflows.	Inactive	true	• com.sn_eam_core • com.snc.procurement • com.sn_itam_recomm • com.snc.asset_management • sn_cmdb_ci_class • com.sn_itam_common • com.sn_ent • com.sn_risk_heatmap • com.sn_eam

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Entitlement Verification [com.sn_ent_verify] (Available in the ServiceNow Store)	Provides APIs to verify Entitlements and Characteristics.	Active	true	Customer Contracts and Entitlements [com.com.sn_pss_core]
Field Service Marketplace [com.snc.fsm_marketplace] (New in Washington DC)	Push tasks to multiple eligible resources to get work done within stipulated time period and at a better cost.	Active	true	com.snc.marketplace_core
Financial Services Operations Integration with Visa [com.sn_fso_intg_visa] (Available in the ServiceNow Store)	Enables financial institutions to integrate with Visa Resolve Online (VROL) APIs using the Visa Spoke Actions to manage various dispute lifecycle events.	Active	false	• com.sn_visa_spoke • com.sn_bom_visa_cp
Flow Designer Proactive Analytics Trigger [com.glide.hub.flow_trigger.analytics] (New in Washington DC)	Starts a flow when Proactive Analytics KPI score or KPI threshold values are met.	Active	false	• com.glide.hub.flow_trigger • com.snc.par.intelligence
GraphQL Explorer [com.glide.graphql.explorer] (New in Washington DC)	An integrated GraphQL testing tool.	Active	false	None
GRC: Management Reporting [com.sn_grc_mgmt_report] (Available in the ServiceNow Store)	Provides the Reporting capabilities for importing GRC related data to Microsoft word document using Microsoft Office 365 plugin.	Inactive	false	com.sn_esg_msoff_intg
Hardware Asset Management for TNI [com.sn_ham_tni]	Use only the Telecommunication Network Inventory (TNI) functionality	Inactive	false	san_hamp

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)	with Hardware Asset Management (HAM).			
Hardware Asset Management integration with Zero Touch Mobility [com.sn_ham_ztm] (Available in the ServiceNow Store)	Drive automation and prescriptive, standard practices for managing Mobile Device class assets throughout the end-to-end asset lifecycle from acquisition through disposition.	Inactive	false	sn_hamp
Health and Safety Risk Management [com.snc.sn_hs_rm] (Available in the ServiceNow Store)	Identify, assess, and reduce potential hazards, prioritizing safety at workplace.	Active	true	None
Health log analytics prime [com.sn_hla_gw_stream] (Available in the ServiceNow Store)	Enables transitioning to a regional data center.	Active	false	None
Hiring Core [com.sn_ta_hiring_core] (Available in the ServiceNow Store)	Record tables to store third party platform data on the ServiceNow AI Platform.	Active	false	None
IPKI Certificate Generator [com.sn_app_skyale] (New in Washington DC)	Secure your Kafka topics by generating an instance-signed certificate.	Inactive	false	com.glide.uxbuilder
ITOM Content Service [sn_smart_content] (Available in the ServiceNow Store)	Offers extensive visibility of products in your infrastructure. The classification of processes identified by Predictive Intelligence enables wider discovery and weekly updates of	Inactive	false	• Discovery [com.snc.discovery] • Normalization Data Services Client [com.glide.data_services_canonicalization.cl

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	new configuration items in the CMDB.			
I18N: Arabic Translations [com.snc.i18n.arabic] (New in Washington DC)	Provides Arabic translations of the base system UI string content in your instance.	Inactive	false	• I18N: Internationalization [com.glide.i18n] • Localization Framework Installer [com.glide.localization_framework.installer] • 18N: Knowledge Management Internationalization Plugin v2 [com.glideapp.knowledge.i18n2]
JavaScript Module Support [com.glide.module_support] (New in Washington DC)	Contains the schema for storing JavaScript modules in a file system structure.	Active	false	None
Localization Framework Hub [com.sn.localization_framework.hub] (New in Washington DC)	Enables fulfilling the translation requests received from Localization Framework Spoke.	Inactive	false	com.glide.localization_framework.installer
Localization Framework Spoke [com.sn.localization_framework.spoke] (New in Washington DC)	Enable fulfilling translation requests using Localization Framework Hub.	Inactive	false	com.glide.localization_framework.installer
Mentoring [sn_ecn] (Available in the ServiceNow Store)	Employee Connections offers improved networking opportunities and enables individuals to forge meaningful connections. It equips employees with tools such as Mentoring, facilitating the search for guidance and suitable mentors to expedite their learning and amplify their	Inactive	true	Career Conversations

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	potential for career advancement.			
Mobile Gen AI [com.glide.sg.gen_ai] (New in Washington DC)	Allows mobile to enable and integrate with generative AI capabilities.	Active	false	• com.glide.one_extend • com.glide.sg.one_extend
Opportunity Management Application [sn_l2c_oppty_mgmt] (Available in the ServiceNow Store)	Gives access to opportunity to Quote workflow and Kanban view.	Inactive	true	• com.sn_l2c_oppty_mgmt_data_model • com.snc.l2c_core
Opportunity Management Data Model [sn_l2c_oppty_mgmt_data_model] (Available in the ServiceNow Store)	Allows users to manage opportunities. It includes the Opportunity Data Model and the ability to create and manage new sales opportunities, as well as workflows to convert opportunities to quotes.	Inactive	true	• com.sn_customerservice • com.sn_prd_pm • com.sn_l2c_sales_common
Opportunity Marketplace [sn_opp_market] (Available in the ServiceNow Store)	Improve employee engagement and retention with internal mobility.	Inactive	true	• Talent Development Core • Talent Development Shared library • Hiring Core
Parallel Review and Feedback [sn_rvw_feedback] (Available in the ServiceNow Store)	Contains the basic workflow of raising a feedback and responding to feedback.	Inactive	true	None
Participant Suggestions [com.glide.sn_participant_suggestions] (Available in the ServiceNow Store)	Displays a list of users (or participants) who may be able to contribute to a Sidebar discussion.	Inactive	false	• com.glide.one_extend • com.glide.platform_ml
Product Configurator [sn_prd_config_ui]	Enables companies to configure products and services they	Inactive	false	Product Catalog Management Core [com.sn_prd_pm]

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)	market, sell, and deliver to customers.			
Quote Management Application [sn_l2c_quote_mgmt] (Available in the ServiceNow Store)	Used to create and manage product quotes	Inactive	true	• com.sn_l2c_quote_mgmt_data_model • com.sn_l2c_core
Quote Management Data Model [sn_l2c_quote_mgmt_data_model] (Available in the ServiceNow Store)	Allows users to manage the quotation process.	Inactive	true	• com.sn_customerservice • com.sn_prd_pm • com.sn.csm_pricing • com.sn_l2c_sales_common
Rack visualization component [sn_rack] (Available in the ServiceNow Store)	Capability to visualise the Rack and the equipment positioning inside the Rack. Enhances the capacity visualisation of Rack.	Active	true	None
ServiceNow Stream Connect Replication - Kafka [com.glide.hub.stream_connect.replication.kafka] (New in Washington DC)	Enables support for replicating Stream Connect topics, via a MID Server, with a customer Kafka cluster.	Inactive	false	• ServiceNow Stream Connect Replication Core [com.glide.hub.stream_connect.replication.c • Stream Connect Replication Certification [com.sn_stream_connect.replication.cert]
Site Mapping for Field Service Management [com.snc.fsm_site_mapping] (Available in the ServiceNow Store)	Navigate using wayfinding features and site maps within a native mobile application.	Active	false	• com.snc.work_management, sn_map_core • sn_map_component
Skills Industry Data [com.sn_skills_data] (Available in the	Get industry skills data.	Active	true	Skills Intelligence plugin

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Skills Intelligence Workspace [com.sn_skills_int_ws] (Available in the ServiceNow Store)	A skills intelligence workspace experience.	Active	true	Skills Intelligence plugin
Sn Topology Map [sn_topology_map] (Available in the ServiceNow Store)	Enables an inventory manager to capture and visualise a topology along with its elements.	Active	true	None
sn-4q-bubble [com.sn_4_q_bubble] (Available in the ServiceNow Store)	Application Portfolio Management bubble charts enables plotting indicator scores of your business applications. Strategize goals using these scores and create demands for your business applications.	Active	false	None
sn_supplier_tprm [sn_supplier_tprm] (Available in the ServiceNow Store)	Enables Supplier Managers to trigger risk assessments from onboarding workflows as well as allows risk data to be viewed and used within Supplier Lifecycle Operations.	Inactive	false	• Supplier Lifecycle Operations [com.snc.sn_supplier_mgmt] • Third-party Risk Management [com.sn_vdr_risk_asmt]
Strategic Portfolio Management for Telecom [sn_strg_prtf_mgmt] (Available in the ServiceNow Store)	Telecom Project templates.	Inactive	true	None
Total Cost of Ownership [sn_apm_tco] (Available in the	Helps enterprise architects to prioritize application portfolio by leveraging the application costs.	Active	true	None

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Visa Spoke [com.sn_visa_spoke] (Available in the ServiceNow Store)	The Visa spoke allows connecting with Visa's REST APIs, providing customers quick access to payment, and security data. Customers can use the spoke to search for transactions, collaborate with merchants, manage disputes and perform other functions with enhanced security.	Active	false	None
Vulnerability Response Common Workspace [sn_vul_cmn_ws] (Available in the ServiceNow Store)	Serves as a centralized hub for incorporating common workspace changes that apply across all modules or any new features that are being developed within the workspace environment.	Active	false	Vulnerability Response Common [sn_vul_cmn]

Existing plugins

This table lists all plugins that you can activate if you have the admin role.

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Access Analyzer [com.snc.access_analyzer] (New in Vancouver)	ServiceNow® Access Analyzer is an application that helps the administrators and developers to view permissions for the selected user, role, or group.	Active	false	
Account Life Cycle Events [com.sn_acct_lc] (Available in the	Allows technology industry providers to structure a repeatable, onboarding experience for internal and external customers.	Active	false	• com.snc.install_base:sys • com.sn_process_automation_designer:sys • com.glide.playbook_experience.config:sys • com.snc.csm_case_types:sys • com.snc.uib.sn_dyn_rel_rec:sys

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store gen)				• sn_csm_playbook:3.1.0 • sn_ti_core:1.0
Accounts Payable Invoice Processing sn_ap_apm (New in Vancouver)	Provides invoice automation solution for Accounts Payable teams which helps businesses reduce risk and improve productivity without additional overhead.	Active	True	• sn_ap_cm • sn_ap_ic
Accounts Payable Operations integration with Document Intelligence sn_ap_ic (New in Vancouver)	Helps automate ingesting invoices incoming from various channels, such as email, portal, integrations etc.	Active	True	com.sn_ap_ic
Additional tables for Federal Agencies [com.snc.fedtables]	Provides Federal customers additional data elements that are frequently used or required in Federal forms to build out Federal HR services.	Active	false	com.sn_hr_core
Invoice Case Management sn_ap_cm (New in Vancouver)	Allows accounts payable teams to manage invoice and payment related inquiries/ requests submitted by suppliers/employees with the objective to faster supplier issue resolution, and improved supplier relations	Active	True	• com.snc.sn_shop • Supplier Collaboration Portal • Source-to-Pay Operations
AIOPs Dashboards [com.sn_itom_aiops_dashboards] (Available in the ServiceNow Store,	One stop shop for all ITOM Health/ AIOps dashboards for multiple personas - Admin, Manager, SME.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
compatible with Vancouver, Utah)				
AIOps Experience [com.sn_sow_aiops] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	AIOps Experience is a set of new experiences composed out of the Express list, Integration Launchpad and AIOps Dashboards.	Inactive	false	Service Operations Workspace ITOM Apps
App Client UX [com.sn_app_client_ux] (New in Vancouver)	Enables admins to install, update, and manage licensed applications and plugins on their instance.	Active	false	app-client (sn_appclient)
Asset Management - Procurement Integration [com.sn_asset_proc_int] (Available in the ServiceNow Store, compatible with Vancouver, Utah)	Enables Software Asset Management (SAM) Admins to create purchase requisitions for software license requirements in the external purchasing system. Later, after the purchase order is received, the receiving transactions are used to automate the entitlement creation in Software Asset Management.	Active	false	• Procurement (com.snc.procurement) • Coupa spoke (sn_coupa_spoke)
Beans.AI Spoke [com.sn_beans_ai_spoke] (New in Vancouver)	Beans.AI is the map provider that Schedule Optimization supports for travel time estimates. The plugin enables Schedule Optimization to calculate distances between agents and task.	Active	false	Integration Hub standard installation pack

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
CMDB CI Class Models [sn_cmdb_ci_class]	Adds class models that extend the CMDB class hierarchy, including class descriptions, identification rules, identifier entries, and dependent relationships if applicable. Applications such as Discovery and Service Mapping can use these class extensions to populate configuration items (CIs) and discover various technologies and software.	Inactive	false	• Configuration Management (CMDB) (com.snc.cmdb) • Expanded Model and Asset Classes (sn_ent) • Data Foundation Model (sn_cmdb_foundation)
Customer Service Characteristics [com.snc.characteristics] (New in Vancouver)	Enables customers to use the characteristics and related features.	Active	false	None
Customer Service Install Base Characteristics [com.snc.install_base_characteristics] (New in Vancouver)	Enables customers to capture characteristics for their install bases.	Active	true	• com.snc.install_base • com.snc.characteristics
Equifax Spoke [com.sn_equifax_spoke] (Available in the ServiceNow Store)	Enables you to access important information about the credit history for a customer, fraud alerts, digital identity verification, transaction screening, and other relevant data.	Active	false	None
Events and Jobs Dashboard [sn_async_dashboard]	Monitoring dashboard for System Events and Scheduled Jobs.	Active	false	com.glide.stats.event_process_monitoring

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(New in Vancouver)				
Enterprise portfolio [enterprise_portfolio_admin] (Available in the ServiceNow Store, compatible with Utah, not compatible with Vancouver)	An application to create enterprise portfolios of types "business applications" and "application services" in Digital Portfolio Management (DPM).	Active	false	
Generative AI Controller [com.sn.generative.ai] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Natively integrates Generative AI APIs in the ServiceNow AI Platform.	Active	false	• OpenAI Generative AI Spoke (com.sn_openai) • Microsoft Azure OpenAI Generative AI Spoke (sn_azure_openai)
Glide Conversation Generative AI [com.glide.cs.genai] (New in Washington DC)	Enables Now Assist	Inactive	true	Generative AI controller
GRC Compliance Case Management Advanced [com.sn_comp_case_adv] (Available in the ServiceNow Store, compatible with	An application to manage Compliance Case Management advanced solution for IRM PRO customers.	Active	false	GRC: Compliance Case Management (com.sn_comp_case)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Vancouver, Utah, Tokyo)				
GRC Compliance Case Management Full Access [com.sn_comp_case_fa] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	An application to manage Compliance Case Management full access solution for IRM enterprise customers.	Active	false	GRC Compliance Case Management Advanced (com.sn_comp_case_adv)
GRC Employee User [com.sn_grc_emp_user] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Using this plugin, all employees in the organization can: • Report – Instances of policy/ compliance violations related to their organization. • Request – Policy modifications, clarifications, and so on, to the compliance teams.	Inactive	false	None
GRC: Third-party Due Diligence Request [com.sn_tprm_onboarding] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Enables any user to submit a request for due diligence for an engagement with a third party.	Active	true	Third-Party Risk Management GRC: Approver Configurator
Clone Admin Console [sn_instance_clone]	Enables the Clone Admin Console, which provides a unified	Active	false	High Availability Cloning (com.snc.ha)

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store, compatible with Vancouver, Utah)	admin experience and enhanced visibility for cloning data between instances, one of our most-used automations. The application includes a simplified clone request experience, a new scheduling tool to prevent timing conflicts, an on-demand backup option, ability to see all clone-related settings in one place, enhanced visibility via a dashboard to easily view current clone activity and more.			(Requires Utah patch 2 or later)
Manage Skills Configurable Page [com.snc.sn_skill_cfg_page] (Available in the ServiceNow Store)	Enables the management of skill assignments using the Next Experience user interface.	Active	true	Skills Management (com.snc.skills_management)
Map Integrations for Field Service [com.snc.app_fsm_map_integr] (New in Vancouver)	Performs intelligent travel time estimates to allocate work order tasks to agents, taking into account both the agent's location and the task's location.	Active	true	Beans.AI Spoke (com.sn_beans_ai_spoke)
Microsoft Azure DevOps Pipelines Spoke [com.sn.azure.pipln.spk]				
MID Server Experience [com.glideapp.agent.experience] (New in Vancouver)	Helps share meaningful information about your MID Servers Health, Performance, and Adoption. It’s based on Polaris	Active	false	• com.glideapp.agent • Multiple PAR plugins (All are activated automatically while activating MID Server Experience)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Experience with new look and feel.			
Platform Document Intelligence [com.glide.platform_ml_di] (New in Vancouver)	The base plugin for the Document Intelligence application.	Active	false	com.glide.platform_ml
Post-Sales Support [com.sn_pss_core] (New in Vancouver)	Provides entitlements and service contracts to enable post-sales support services to customer service users.	Active	true	• com.sn_customerservice • com.snc.install_base • com.snc.csm_case_types • com.snc.characteristics
Sustainable IT [sn_esg_sustain] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Allows you to effectively manage and monitor the emissions generated by your IT assets. Additionally, the application enables you to keep track of the energy consumption of your assets and their proper disposal after they reach the end of their lifespan.	Inactive	true	• Hardware Asset Management (HAM) Pro • ESG Management
Skills Intelligence [com.snc.skills_management] (Available in the ServiceNow Store)	An AI-driven platform that you can use in your organization to develop a workforce that is based on skills.	Inactive	true	Employee Profile
Service Operations Workspace Express List [com.sn_itom_aiops_list] (Available in the ServiceNow Store, compatible	The Express List is where operators can triage, analyze, and remediate alerts. It's a live view of alerts, with dynamic filtering, enriched information in a preview panel, and have automation embedded for the user.	Active	false	Service Operations Workspace ITOM applications

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
with Vancouver, Utah, Tokyo)				
Service Operations Workspace Integrations launchpad [com.snc.uib.sow_integrations_launchpad] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	A consolidated view for all available ITOM health integrations with easy and intuitive experience to configure event connectors. The application is activated automatically with AIOps experience.	Active	false	
Service Operations Workspace Metric Explorer [com.sn_itom_metric_exp] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	A new Metric Explorer application for SOW that replaces the earlier one. The application is installed automatically with ACC-M.	Inactive	false	
Service Operations Workspace Metric Explorer APIs [com.sn_sow_metric_exp] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	A new Metric Explorer application for SOW that replaces the earlier one. The application is installed automatically with ACC-M.	Inactive	false	
Task Grouping				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.task_grouping]				
Task Intelligence for ITSM [com.snc.itsm_ml_task] (Available in the ServiceNow Store, compatible with Vancouver, Utah)	Uses machine learning to set up, deploy, and track the solution-based models to achieve important business outcomes.		false	• Recommended Actions for ITSM - Advanced (com.snc.uib.sow_itsm_ra_advanced) • Task Intelligence Admin Console (com.sn_ti_admin)
Technology core [com.sn_ti_core] (Available in the ServiceNow Store)	Technology industry vertical CSM extension.	Active	false	com.sn_customerservice
Transporter [com.sn_transport]				
Zero Trust - Location Based Access [com.snc.zero_trust_location_access]	Location filter is a filter criteria that the admins can use while crafting the authentication policies based on the physical location of the device.	Active	false	Adaptive Authentication (com.snc.adaptive_authentication)
Zero Trust - Policy Based Session Access [com.snc.zero_trust_session_access]	ServiceNow® Zero Trust - Policy Based Session Access (Session Access) enables organizations to dynamically reduce user privilege in a web session based on a variety of factors, including IP address, location, authentication method, user's role, group, user having MFA and attributes	Active	false	Adaptive Authentication (com.snc.adaptive_authentication)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	shared by the Identity Provider (IDP).			
Vulnerability Response Integration with Claroty CTD [com.sn_clarotyctdvr] (Available in the ServiceNow Store, compatible with Vancouver, Utah)	Uses vulnerability data imported from Claroty CTD to enable risk-based action with the production process context.	Active	false	• Service Graph Connector Integration for Claroty CTD (com.sn_clarotyctdsgc) • CMDB CI Class Models (com.sn_cmdb_ci_class)
Workplace Connectors (sn_wsd_wc) (Available in the ServiceNow Store, compatible with Vancouver, Utah)	Workplace Connectors is a generic framework using which data from different kinds of workplace hardware/ sensors (such as badging systems, occupancy sensors) can be brought into WSD product via the respective spokes.	Active	false	Workplace Central (sn_wsd_central)
Threat Intelligence Security Center for Security Operations [com.snc.secops.threat.intel.security.center] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Allows you to find indicators of compromise (IoC) and enrich security incidents with threat intelligence security center data.	Inactive	false	• Security Support Common (com.snc.security_support.common) • Threat Intelligence Support Common (com.snc.threat) • Security Case Management common workspace components (com.snc.escm.ws_commons)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Network Inventory Advanced [sn_ni_adv] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Advanced capabilities of telecommunications network inventory.	Active	true	• Inventory Number Management (sn_inv_num_mgmt) • Attribute Pack (sn_attribute_pack)
Network Inventory Core [sn_ni_core] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Application to manage network inventory.	Active	true	• Inventory Number Management (sn_inv_num_mgmt) • Attribute Pack (sn_attribute_pack)
Data Separation [sn_ds] (Available in the ServiceNow Store)	Enables organizations to restrict access to sensitive data based on a lens hierarchy and its leaf node. The application supports records, related items, planning consoles, workbenches, and reports.	Active	false	• PPM Standard (com.snc.financial_planning_pmo) • Portfolio Planning (sn_align_ws)
DevOps Vulnerability Integrations [sn_devops_vul_ints] (Available in the ServiceNow Store, compatible with Vancouver, Utah, Tokyo)	Provides a new framework and data model for DevOps integration with application security tools. It's an extensible framework that also enables you to create custom integrations with any application security tool.	Active	false	com.sn_devops

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Action Status Automation [com.sn_action_status]	This plugin tracks blocking records created for tasks and updates the action status indicators on the task list.	Active	true	
Activate all Software Asset Management Professional plugins including Software Asset Workspace [com.sn_samp_master_ws] (New in Utah)	Provides the capability to normalize software discovery information, reconcile entitlements with installs, and reclaiming unused software. Provides additional capabilities to reconcile publisher products.	Inactive	true	• com.sn_samp_master • com.sn_sam_workspace
Activity formatter [com.glide.ui_activity_formatter]	This plugin lets you quickly and easily filter the list of activities, or history, on a task form.	Active	false	
Admin Center [sn_admin_center] (Available in the ServiceNow Store)	Provides a central hub for platform owners and admins to access platform capabilities easily, discover new applications, and get intelligent, actionable insights.	Active	false	sn_ace
Admin Experience Framework [com.sn.ace_framework] (New in Tokyo)	Supports new UI framework used by various applications, including Admin Center.	Active	false	None
Adoption Blueprint for Technology Excellence [sn_itsm_bus_obj] (Available in the ServiceNow Store)	Adoption blueprints, a part of Admin Center, displays a recommended set of applications and plugins available to you in an ideal order of installation and adoption to achieve business goals.	Active	false	sn_admin_center

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Advanced AI Search Management Tools [sn_ais_admin_tools]	Includes AI Search dashboards and Search Preview tool.	Inactive	false	
Advanced Work Assignment [com.glide.awa]	Automatically assigns work items to agents based on their availability, capacity, and skills.	Active	false	com.snc.skills_management
Advanced Work Assignment - Agent Affinity [com.glide.awa.agent_affinity]	Enable work assignment to the best-suited agent based on the agent's affinity to the work item.	Inactive	false	com.glide.awa
Advanced Work Assignment for HRSD [com.sn_hr_awa]	Contains out-of-box configuration data supporting routing, queuing, and assignment of HRSM cases, chats, and calls to agents.	Inactive	false	• com.glide.awa • com.sn_hr_core • com.glide.interaction.awa
Advanced Work Assignment for Incidents [com.snc.incident.awa]	Default configuration to support Advanced Work Assignment for Incident	Active	false	• com.glide.awa • com.snc.agent_workspace.itsm
Advanced Work Assignment for Source-to-Pay Operations [com.snc.sn_spend_awa] (Available in the ServiceNow Store)	Automatically assign work to your agent using advanced work assignment.	Inactive	true	• Advanced Work Assignment (com.glide.awa) • Agent Chat (com.glide.interaction.awa) • Procurement Case Management (com.sn_spend_psd)
Advanced Work Assignment for Supplier Lifecycle Operations [com.snc.sn_slm_awa]	Automatically assign supplier cases to agents based on availability and capacity.	Inactive	true	• Advanced Work Assignment (com.glide.awa) • Agent Chat (com.glide.interaction.awa)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				• Glide Virtual Agent (com.glide.cs.chatbot) • Supplier Lifecycle Operations (com.snc.sn_supplier_mgmt)
Agent Chat [com.glide.interaction.awa]	Enables Workspace Agent Chat and the Chat service channel in Advanced Work Assignment.	Active	true	• com.glide.interaction • com.glide.awa
Agent Client Collector Framework [sn_agent] (Available in the ServiceNow Store)	Management of Agent Client Collector	Active	false	• com.snc.sa.mid.webserver • com.snc.cmdb.scoped
Agent Client Collector Monitoring [sn_itmon] (Available in the ServiceNow Store)	Monitoring solution using the Agent Client Collector	Active	false	
Agent Client Collector for Investigation [sn_acc_adapter] (Available in the ServiceNow Store)	Enables you to retrieve and display the CI metrics information of the affected CIs.	Active	true	None
Agent Forecast [com.sn_agent_forecast]	Forecasts agents based on historical data.		false	com.snc.clotho
Agent Intelligence [com.glide.platform_ml]	Renamed to Predictive Intelligence.	Active	false	com.glide.platform_ml_pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Agent Intelligence Reports [com.glide.platform_ml_pa]	Renamed to Predictive Intelligence Reports.	Active	false	
Agent Schedule [com.snc.agent_schedule]	Enables customer service agents and field service technicians to see work schedules and assignments and also add personal events such as meetings or appointments.	Active	false	com.snc.app.agent_calendar_widget
Agent Workspace [com.agent_workspace]	It is a suite of tools that provide agents, case managers, help desk professionals, and managers with tools to help answer customer questions or resolve customer problems.	Active	false	• com.glide.uxbuilder • com.glide.graphql • com.glide.interaction • com.snc.agent_workspace.config • com.snc.agent_workspace.ribbon • com.snc.agent_workspace.list • com.snc.agent_workspace.form • com.snc.agent_workspace.global_search • com.snc.agent_workspace.declarative_actio
Agent Workspace - Knowledge [sn-component-workspace-knowledge]	Enables use of Knowledge Base in Agent Assist.	Active	false	com.snc.agent_workspace.config
Agent Workspace - List [com.snc.agent_workspace.list]	Workspace List Configurations	Active	false	com.snc.agent_workspace.config
Agent Workspace - Ribbon [com.snc.agent_workspace.ribbon]	Workspace Ribbon Configurations	Active	false	• com.glide.uxbuilder • com.snc.agent_workspace.config • com.sn_resolutionshaper
Aggregate Web Service	Provides SOAP Access to	Inactive	false	

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.web_service_aggregate]	GlideAggregate functionality.			
Agile Development 2.0 [com.snc.sdlc.agile.2.0]	The Agile Development 2.0 plugin provides enhanced functionality on top of Agile Development. If you already have a customized version of Agile Development, delete the customizations before activating Agile Development 2.0 to ensure that all features work properly. Refer to the documentation for detailed steps to delete the customizations.	Active	true	• com.snc.sdlc.scrum.ppm_int • com.snc.planned_task_v2 • com.snc.sdlc.agile.2.0.common • com.snc.sdlc.ranking
Agile Development — Unified Backlog [com.snc.sdlc.agile.multi_task]	Enables you to maintain a centralized backlog containing records of different task types, such as defects, problems, incident tasks, and stories. Include any tasks into your agile workflow	Active	false	• com.snc.sdlc.agile.2.0 • com.snc.sdlc.agile.multi_task.common
Agile - Scaled Agile Framework - Essential SAFe [com.snc.sdlc.safe]	Scaled Agile Framework was designed to apply Lean-Agile principles to the entire organization. Essential SAFe is the most basic configuration of the framework and it provides the minimal elements necessary to be successful with SAFe: manage your agile release train backlog, plan program increments,	Active	true	• com.snc.sdlc.ranking • com.snc.sdlc.agile.2.0
Agile - Scaled Agile Framework - Portfolio SAFe	Use Portfolio SAFe to apply lean and agile principles to your portfolio work.	Active	true	• com.snc.sdlc.safe • com.snc.portfolio_management

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.sdlc.portfolio_safe]				
Agile Development 2.0 - ATF Tests [com.snc.sdlc.agile2.0.atf]	Agile Development 2.0 - ATF Tests provides you test cases and test suites that can be run on the Agile Development 2.0 application.	Active	false	
AI Search [com.glide.ais]	Installs and enables AI Search capabilities, providing relevant, contextual, and personal search experiences across different interfaces.	Active	false	• com.glide.search.signal_data • com.glide.platform_ml • com.glide.nlu • com.snc.nlu_studio
AI Search Assist [com.snc.ai_search_assist]	Helps users deflect or quickly resolve their issues without involving the service desk operators. It can include results from Knowledge and Service catalog, allowing a user to directly order a catalog item from a search.	Active	false	com.glide.search
AI Search for Customer Portals [com.snc.csm.portal_ais]				• com.glide.ais • com.glide.ais_enabler
AI Search Index Sources [com.glide.ais.indexsources]	Includes indexed source definitions for ServiceNow AI Platform tables beyond those provided in the base system: • Catalog Task [sc_task] • Change [change_request] • Change Task [change_task] • Company [core_company]	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	• Group [sys_user_group] • Incident [incident] • Location [cmn_location] • Problem [problem] • Request [sc_request] • Requested Item [sc_req_item]			
AI Search Spoke [sn_ai_spoke] (Available in the ServiceNow Store)	Helps with Integration Hub execute AI search queries.	Active	false	com.glide.ais.external_content
Alert Rules Management [com.sn-em-arm]				
Portfolio Planning [sn_align_ws] (Available in the ServiceNow Store)	Helps the product and portfolio managers drive organizational alignment and facilitate Agile transformation by combining waterfall and agile work streams.	Active	true	com.sn_roadmap_plng
Alignment Planner Workspace integration with Azure DevOps (ADO) [com.sn_align_ado_int] (New in San Diego)	Enables import, export, and bidirectional synchronization of records between ServiceNow® Alignment Planner Workspace (APW) and Microsoft Azure DevOps by integrating the two applications.			

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Analyst workspace for Security Incident Response [com.snc.security_incident.workspace] (New in Tokyo)	Enables analysts and incident managers to prioritize their work using a dedicated workspace, monitor and view the status of incidents.	Active	false	• sn_si • sn_sirw_evam_card
Analytics Center [com.snc.pa.analytics_center]	Performance Analytics Center - Enables users to get answers about their data quickly using search-driven analytics.	Inactive	false	• com.agent_workspace • com.workspace_core • sn-component-nlq • com.snc.par.nlq.core
Analytics Task [com.snc.pa.analytics_task]	Provides the core framework for workflow related to Analytics. It consists of the analytics task table and related artifacts.	Active	false	
Angular AMB Services [com.glide.ui.ng.amb]	Angular Services for AMB	Active	false	• com.glide.ui.ng • com.glide.amb • com.glide.record_watcher
Anonymous Connect Support [com.glide.connect.anonymous_support]	Plugin to enable access and properties for Anonymous Connect Support.	Inactive	false	com.glide.connect.support
Anonymous Report Center [com.sn_anonymous_report_center]	Enables users to submit cases without disclosing their identities to agents or admins within their organizations.	Active	false	com.sn_content_delivery
API Analytics [com.glide.api_analytics]	Allows admins to view API Analytics statistics for web services.	Active	false	• com.glide.db.query_stats • com.snc.pa • com.glideapp.canvas

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
API Stats [com.glide.stats.api]	API Stats	Active	false	• com.glide.db.query_stats • com.snc.pa • com.glideapp.canvas
App APIs [com.glide.app_api]	APIs for fetching Application and Application Rollback related information.	Active	true	
App Collaboration [com.glide.app_collaboration]	Plugin that complements Delegated Developer for facilitating application collaboration. Requires App Engine license.	Active	false	com.glide.delegated_development
App Dependency Client [com.sn_dependentclient]	Plugin to provide list of store applications/ integrations based on their active plugins. Uses store's web service to create a cache of such applications on the customer instances itself.	Inactive	false	
App Engine Management Center [com.sn_aemc]				
App Engine Studio [com.snc.app-engine-studio]	Installs and enables App Engine Studio, a development tool for creators of varying skill levels to build applications that meet the immediate needs of your organization.	Inactive	true	
Application Authorization [com.glide.sys.app_authorization]		Inactive	false	
Application Creator [com.snc.apps_creator]	Retired. Replaced by the Platform as a Service plugin.	Active	false	• com.glide.ui.angular • com.glide.ui.heisenberg

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Application Creator Templates [com.snc.apps_creator_template]	Hosts system application creator templates.	Active	false	
Application Design Restrictions [com.glide.scope.design]	System for restricting and capturing tables augmented by an application	Active	false	
Application Diagnostics Tool [com.snc.app_diagnostics_tool]	Helps you to diagnose the PPM Applications, with quick diagnostic scans and resolve the issues with fix scripts.	Active	false	
Application File [com.snc.apps_file]	Associates configuration records with an application and tracks record metadata.	Active	false	
Application Insights [sn_app_insights] (Available in the ServiceNow Store)	Provides a centralized location where you can visualize and monitor system health.	Active	false	com.snc.clotho
Application Metadata [com.snc.metadata]	Adds relationships and UI actions to work with metadata.	Active	false	
Application Portfolio Management [com.snc.apm]	Helps enterprises have visibility into their business applications inventory and enables analysts to rationalize business applications. Provides Performance Analytics dashboards like Application 360, Landscape Analysis, and associated dashboards of Business Applications with Change requests, Incidents, and Problems. This plugin	Active	true	• com.snc.certification_v2 • com.snc.fiscal_calendar • com.snc.treemap • com.snc.sp_workbench_widgets • com.snc.apm.business_planner • com.snc.pa.apm • com.snc.demand_core • com.snc.apm.predictive_intelligence

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	activates the following related plugins if they are not already active: • Application Portfolio Management – Predictive Intelligence (com.snc.apm.predictive_intelligence): Predicts application category by applying algorithms like similarity on business applications-related data. • Business Planner: Enables access to the Business Planning portal.			
Application Portfolio Management – ATF Test [com.snc.apm.atf]	Validates ATF tests for Application Portfolio Management	Active	false	com.snc.apm
Application Portfolio Management - Technology Reference Model (TRM) [com.snc.apm_trm] (New in Tokyo)	Allows enterprise architects to approve or restrict the use of a software product within the organization.	Active	true	com.snc.apm
Application Portfolio Management Core [com.snc.apm_core]	Enables register a new business application	Active	false	• com.glideapp.servicecatalog.platform • com.glide.hub.designer_backend.model
Application Portfolio Management Digital Integration Management [com.snc.apm_digital_integration]	Enables APM users to manage integrations in Application Portfolio Management.	Active	true	com.snc.apm

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				
Application Portfolio Management Workspace [sn_apm_ws] (Available in the ServiceNow Store)	Enables Enterprise Architects to have a visibility in to all the important aspects of their applications in Application Portfolio Management.	Active	false	com.snc.apm
Application Service [com.snc.cmdb.it_service]	Provides functionality and view for Application Service creation, editing and history view.	Active	false	• com.snc.itom.ui • com.snc.cmdb • com.snc.cmdb.service.modeling • com.glide.ui.list_v3_components
Applications Access Control [com.snc.apps_access]	Implements file-level access for application development.	Active	false	• com.snc.apps • com.snc.apps_picker
Applications Picker [com.snc.apps_picker]	Allows users to select the desired application during application development.	Active	false	
Apply Once APIs [com.glide.system_apply_once]	APIs for altering the behavior of any 'apply_once' update in a plugin during plugin activation/upgrade	Active	false	
Appointment Booking [com.snc.appointment_booking]	Enables customers to book service appointments from the Customer and Consumer Service Portals.	Inactive	true	
Appointment Booking Demo Data [com.snc.appointment_booking_demo]	Provides demo data for the appointment booking feature.	Inactive	false	
Approvals with e-Signature	Adds a prompt for credentials when an	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.e_signature_approvals]	approver attempts to approve a request via the list context menu or Approve UI Action on the Approval form.			
Architecture Compliance [com.snc.architecture_compliance]	Manages scheduled or on-demand audits of CMDB records, to determine if configuration items (CI) match expected attributes.	Inactive	true	com.snc.certification_core
Assessment [com.snc.assessment_core]	Provides capabilities to use custom questionnaires and scripted queries to evaluate, score, and compare any records in ServiceNow.	Active	true	• com.glideapp.workflow • com.glideapp.survey • com.snc.bestpractice.task_survey • com.glide.survey_designer
Assessment Designer [com.glide.assessment_designer]	Activates assessment designer that is a drag-and-drop interface to create assessments.	Active	false	com.glide.assessment_designer.common
Assessment Designer Common [com.glide.assessment_designer.common]	Base plugin for assessment related designers like Quiz Designer, Survey Designer, Assessment designer. You must activate Assessment Designer Common plugin before using Quiz Designer, Survey Designer, or Assessment Designer. By default, this plugin is activated.	Active	false	com.glide.ui.ng.dc
Asset Management [com.snc.asset_management]	Lets you manage all your assets, consumables, and software licenses.	Active	true	• com.snc.expense_line • com.snc.model • com.snc.organization_management • com.snc.fixed_asset • com.snc.depreciation • com.snc.automation • com.glideapp.home

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Asset Management Common [com.sn_itam_common]				
Asset Management Performance Analytics [com.snc.ast_mgmt_pa]	Performance Analytics for Asset Management out-of-the-box KPIs.	Active	false	
Asset Management Workspace [com.sn_itam_workspace]	Enables ITAM capabilities in Workspace.	Active	false	• com.sn_itam_recomm • com.snc.app_shell_aw • com.snc.uib.base_agent_workspace
Asset Management Workspace - Recommendations [com.sn_itam_recomm]	Recommendations feature for workspaces.			
Assignment Workbench [com.snc.assignment_workbench]	Provides a workbench that customer service managers can use to evaluate agents based on configurable criteria, such as skills and availability, and then assign tasks to the desired agents.	Inactive	false	com.snc.matching_rule
ATF Test Generator and Cloud Runner [sn_atf_tg] (Available in the ServiceNow Store)	One-click test generation and cloud execution for ATF.	Inactive	false	• com.glide.automated_testing_framework • com.glide.auth.mutual
Availability [com.snc.availability] (New in Tokyo)	Allows organization to calculate availability for select configuration items.	Active	false	• com.snc.cmdb • com.snc.service_portfolio_core
Auto Recovery	Enables Auto Recovery functionality.	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.autorecovery]				
Automated Test Framework [com.glide.automated_testing_framework]	Testing framework and tools for creating automated tests	Active	true	• com.glide.rollback • com.glide.ui.angular • com.glide.element_mapping • com.glide.vars
Automated Test Framework - Application Navigator [com.glide.automated_testing_impl.left_nav]	Provides ATF steps that validate access to the application navigator.	Active	false	com.glide.automated_testing_framework
Automated Test Framework - Custom UI [com.glide.automated_testing_impl.custom_ui]	Adds ability to test custom UI.	Active	false	com.glide.automated_testing_framework
Automated Test Framework - Parameters [com.glide.automated_testing_impl.parameters]	Adds ability to parameterize ATF tests.	Active	false	com.glide.automated_testing_framework
Automated Test Framework - Reporting [com.glide.automated_testing_impl.report]	Automated Testing Framework Reporting test step configuration.	Active	false	• com.glide.automated_testing_framework • com.glideapp.report
Automated Test Framework - Responsive Dashboards [com.glide.automated_testing_impl.dashboards]	Automated Testing Framework Responsive Dashboard test step configuration.	Active	false	• com.glide.automated_testing_framework • com.glideapp.dashboards
Automated Test Framework - REST Inbound [com.glide.automated_testing_impl.rest_inbound]	Automated Test Framework for inbound REST requests.	Active	true	com.glide.automated_testing_framework

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Automated Test Framework - Schedule [com.glide.automated_testing_impl.schedule]	Tools for scheduling automated tests.	Active	false	com.glide.automated_testing_framework
Automated Test Framework - Allowed Errors [com.glide.automated_testing_impl.wce]	JavaScript errors found during UI test execution can be ignored or tracked as warnings. Test designers can complete their work while developers investigate and resolve scripting issues.	Active	false	com.glide.automated_testing_framework
Automated Test Framework Service Catalog [com.glide.automated_testing_impl.service_catalog]	Automated Testing Framework Service Catalog test step configuration.	Active	true	• com.glide.automated_testing_framework • com.glideapp.servicecatalog
Automated Test Framework Service Catalog Service Portal [com.glide.automated_testing_impl.service_catalog_portal]	Automated Testing Framework for Service Catalog widgets in Service Portal test step configuration.	Active	true	com.glide.automated_testing_impl.service_cat
Automated Test Framework Service Portal [com.glide.automated_testing_impl.service_portal]	Automated Testing Framework Service Portal test step configuration. This plugin can only be activated by activating the "Service Portal - Core" plugin.	Active	false	com.glide.automated_testing_framework
Automated Test Scripts For Survey [com.glide.automated_testing_impl.survey]	Automated Test Scripts For Survey.	Active	false	• com.glide.automated_testing_framework • com.snc.assessment_core
Automated Testing for Virtual Agent [com.glide.cs.atf]	Activates the Automated Testing capabilities for Virtual Agent.	Active	false	com.glide.cs

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(New in Utah)				
Automatic Assignment [com.snc.automatic_assignment]	Allows any application that uses tables that are children of Task [task] to use auto-assignment to automatically find eligible assignees for any task.	Inactive	false	
Automation Center [com.sn_coe] (Available in the ServiceNow Store)	Integrates islands of automation and delivers centralized governance and management of cross-enterprise, multi-vendor automation in a single workspace.	Active	true	sn_cmdb_ci_class (1.38.0)
Automation Center [com.sn_ac] (Available in the ServiceNow Store)	Manages and monitors your entire hyperautomation landscape.	Active	true	com.sn_cmdb_ci_class
Automation Discovery [sn_auto_discovery] (Available in the ServiceNow Store)	Creates reports from your data to show opportunities for automation.	Active	false	• com.glide.nlu • com.glide.platform_ml
AWA Channel Management [com.sn_channel_management]	Main plugin that supports various functionalities for AWA Service Channels and Queues under Manager Workspace	Active	false	• com.glide.awa • com.snc.wfo
Azure Active Directory User Mapping [sn_now_azure] (Available in the	Contains workflows and mappings for ServiceNow user to Azure Active Directory user.	Inactive	false	• com.glide.channelproxy • com.snc.msteams.app.core

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Badge Reader Integration [com.snc.badge_reader]	Badge Reader Integration Framework which allows other applications to integrate with badge reader hardware.	Inactive	false	
Base Upgrade Logger [com.glide.base_upgrade_logger]		Active	false	• com.glide.ui.mergetool • com.glide.ui.diagnostics
Basic Export Set Functionality [com.glide.system_export_set]		Active	false	
Benchmark Client [com.sn_bm_client]	Benchmark Insights for Customers.	Active	false	com.sn_bm_common
Benchmark Common [com.sn_bm_common]	Common code for Benchmark Insights.	Active	false	
Benchmarks Spoke [com.sn_bm_client.spoke]	Contains the evaluation framework for Benchmarks Recommendations.	Active	false	com.glide.hub
Best Practice - Bulk CI Changes [com.snc.bestpractice.bulkchange]	This plugin is intended to be used with the legacy Change Management state models, which were used prior to the Geneva release. It allows CMDB updates to be proposed and applied to one or more in scope configuration items as part of your change management process.	Inactive	false	
Best Practice - Change Risk Calculator [com.snc.bestpractice.change_risk]	Simple risk and impact calculations for change management.	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Best Practice - Incident Resolution Workflow [com.snc.bestpractice.incident]	Best practices for incident resolution dictate that rather than closing the incident, the incident should have a state of Resolved. This state gives the service desk a mechanism to verify that caller is satisfied with the resolution, and that the customer agrees with closing the incident.	Active	false	
Best Practice - ITIL KPI Reports [com.snc.bestpractice.itil_kpi]	Provides a series of reports that track the Key Performance Indicators (KPI) of incident management and problem management.	Active	true	com.glide.metrics
Business Location [com.snc.business_location]	This plugin supports a data model where the corporation does business with customers through physical channels such as stores and branches.	Active	true	• com.snc.service_organization • com.snc.cs_base_extension
Business Roles [com.snc.businessroles]	Activates the Business Roles module. Combined with other licensed applications, provides business roles and hierarchy for employees with integration to Okta for access provisioning during onboarding.	Inactive	false	com.sn_content_delivery
Business Rule V2 [com.glide.business_rule_v2]	Enhances business rules to support script-free conditions and behaviors.	Active	false	
Business Stakeholder [com.snc.business_stakeholder]	Contains the Business Stakeholder role which authorizes the users to approve, view/read records across the organization and view	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	reports. Note there is a fee associated with the Business Stakeholder role. Do not assign it to users without confirming your organization has the appropriate entitlement.			
Cache build stats [com.glide.stats.cache_build]	Cache build stats.	Active	false	com.glide.stats
Card Operations [com.sn_bom_credit_card]				
Case Assignment Workbench Demo [com.snc.case_assignment_workbench_demo]	Provides demo data, such as sample cases and users, so that the Assignment Workbench product features can be demonstrated on a non-production instance. Not for use on customer production instances.	Inactive	true	• com.snc.customerservice.demo • com.snc.assignment_workbench
Case Digests [com.sn_csm_case_digest]	Enables users to send case status updates and root cause analysis to customers and key internal stakeholders.	Inactive	true	• com.sn_customerservice • com.sn_csm_doc_template • com.snc.documentviewer • com.sn_publications
Case Playbook for Complaints [com.sn_csm_complaint_caseflow]	Provides a complaint case type to capture the details for a customer complaint and a playbook that provides step-by-step guidance through the lifecycle of the complaint.	Active	true	
Case Playbook for Onboarding [com.sn_csm_onboarding_caseflow]	Provides an onboarding case type to capture the details when onboarding customers for a product or service	Active	true	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	and a playbook that provides step-by-step guidance through the lifecycle of the onboarding process.			
Case Playbook for Product Support [com.sn_csm_product_support_playbook]	Case Flow for Product Support guides customer service agents through the steps needed to resolve product issues.	Active	true	
Catalog Designer Common [com.glide.ui.ng.cc]	Catalog Designer Common.	Active	false	• com.glide.ui.ng.dc • com.glideapp.workflow
CCG Content Pack [sn_itom_ccg_cp] (Available in the ServiceNow Store)	Provides the base system contents such as policies, resource collectors, configuration collectors, and remediation actions for the Cloud Configuration Governance application.	Active	false	• com.snc.cloud.core • com.glide.hub.action_step.payload • com.glide.hub.action_step.script_mid • com.snc.itom.license • com.glide.hub.action_step.get_connection_i • com.glide.hub.action_step.rest • com.glide.hub.action_step.xmlparser • com.glide.hub.action_type.datastream • com.glide.hub.integration.runtime
Central Dispatch [com.snc.central_dispatch]	Allows visually allocating tasks to agents for a logged in dispatcher.	Inactive	true	com.snc.dhtmlx.scheduler
Centralized Connection and Credential [com.snc.core.automation.connection_credential]	Centralized Connection and Credential components.	Active	false	
Certification Core [com.snc.certification_core]	Core certification structures such as filters and templates. Certification Core is activated automatically when any of these applications are activated:	Active	false	• com.glide.list_v2 • com.snc.version

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	• Desired State Certification (active by default) • Architecture Compliance • Data Certification • IT Governance Risk and Compliance			
Change Management - Approval Policy [sn_chg_pol_appr]	Change Management feature that helps in generating and configuring approvals using the Change Approval policies.	Active	false	
Change Management - ATF Tests [com.snc.change_management.atf]	This plugin loads ATF tests for Change Management when the Change Management State Model plugin is active. The demo data for this plugin is required to successfully execute these ATF tests.	Inactive	true	
Change Management - CAB Workbench [com.snc.change_management.cab]	Change Management feature which allows CAB Managers to schedule, plan and manage CAB Meetings all on the ServiceNow AI Platform.	Active	true	• com.snc.change_management.cab.form_lay • com.snc.app.calendar • com.glide.service-portal • com.snc.app_common.service_portal • com.glide.editor.tinymce
Change Management - Case Intelligence [com.snc.change_management.ml.ccbc]	Change Management feature that uses Predictive Intelligence to automatically identify cases that may have been caused by change.	Inactive	false	• com.snc.change_management.ml • com.sn_customerservice
Change Management - Change Model Foundation Data [com.snc.change_management.change_model.foundation]	Foundation data for Change Models, including State transitions and supporting flows.	Active	false	• com.snc.change_management.change_mod • com.glide.hub

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Change Management - Change Models [com.snc.change_management.change_model]	Change Models allow flexible configuration of state and transition models for Change Requests.	Active	false	• com.snc.itsm.state_transition_model • com.snc.itsm.state_transition_model.founda
Change Management - Change Models Landing [com.snc.change_management.change_model.landing]	Comprises the user interface used to create Change Requests based on Change Models.	Active	false	com.snc.change_management.change_model
Change Management - Foundation Data [com.snc.change_management.foundation]	Provides the Change Management foundation data, new states, priorities.	Active	true	com.snc.change_management
Change Management - Change Related Incidents (Maintenance mode) [com.snc.change_management.icbc]	UI and configuration for Change Related Incidents Classic environment	Active	false	com.snc.change_request
Change Management - Change Schedule [com.snc.change_management.soc]	Change Schedule.	Active	false	• com.snc.dhtmlx.gantt • com.snc.change_request • com.snc.change_management.soc.color_pic
Change Management - Change Schedule foundation [com.snc.change_management.soc.foundation]	Change Management - Change Schedule Foundation Data.	Active	false	com.snc.change_management.soc
Change Management - Change Success Score [com.snc.change_management.change_success_score]	Change Success Score is a numerical expression that represents the change success history for a particular group. It can be used to evaluate the	Active		• com.snc.change_request • com.snc.pa • com.snc.change_management.change_succ

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	likelihood of future success and help determine the level of change rigor required.			
Change Management - Collision Detector [com.snc.change.collision]	Conflict Detection is a feature of the Change Management application that identifies potential scheduling conflicts based on the configuration items in scope for the change or user or group assigned to fulfill a change.	Active	true	com.snc.maintenance_schedules
Change Management - Color Picker [com.snc.change_management.soc.color_picker]	Color Picker	Active	false	
Change Management - Core [com.snc.change_management]	Core change management plugin which both Change Management State Model and Standard Change Catalog plugins depend on. This plugin updates the **Type** field on Change Request to have the values **Normal**, **Standard**, and **Emergency**. The Type value on existing Change Requests is updated: • Routine com.glide.context_help-> standard (Standard) • Comprehensive -> normal (Normal) • Emergency -> emergency (Emergency)	Active	true	• com.snc.change_request • com.snc.change.collision

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	When creating a Change Request, an interceptor prompts you for the type of Change you want to create.			
Change Management - Incident Intelligence [com.snc.change_management.ml.icbc]	Change Management feature that uses Predictive Intelligence to automatically identify incidents that may have been caused by change.	Active	false	• com.snc.change_management.ml • com.snc.change_management.icbc
Change Management - Mass Update CI [com.snc.change_management.mass_update_ci]	Change Management feature which allows change requesters to propose and apply updates associated with a change request to one or more configuration items in scope for a change.	Inactive	false	• com.snc.change_management.state_model • com.snc.bestpractice.bulkchange
Change Management - Predictive Intelligence Core (maintenance mode) [com.snc.change_management.ml]	Core plugin containing the common components used by the Change Management Intelligence features.	Active	false	• com.snc.change_request • com.glide.platform_ml
Change Management - Risk Assessment [com.snc.change_management.risk_assessment]	Change Management feature which allows risk assessment questionnaires to be created and required to drive the assessment of risk associated with a requested change.	Inactive	true	• com.snc.bestpractice.change_risk • com.snc.assessment_core • com.glide.assessment_designer
Change Management - Standard Change Catalog [com.snc.change_management.standard_change_catalog]	Change Management feature which allows Standard change templates to be proposed, approved, and published in the ServiceNow Service Catalog.	Active	true	• com.snc.change_management • com.glideapp.servicecatalog

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Change Management - State Model [com.snc.change_management.state_model]	This plugin provides the current Change Management state models for Standard, Normal and Emergency change types, originally released in the Geneva release.	Active	true	• com.snc.change_management • com.snc.process_flow_formatter
Change Management - Success Probability [com.snc.change_management.success_probability] (New in Tokyo)	Calculates a probability to indicate what are the chances of the Change Requests completing successfully without issues.	Active	false	com.snc.change_management.change_succes
Change Management Overview Homepage [com.glideapp.report.itsm.change.overview]	Activates a Change Management Overview Dashboard in the Change Management application navigator.	Active	false	com.glideapp.report
Change Management - Risk Intelligence (maintenance only) [com.snc.change_management.ml.risk]	Change Management feature that augments existing risk assessment capabilities with Predictive Intelligence to predict change risk.	Active	false	com.snc.change_management.ml
Change Management - Standard Change Template Intelligence [com.snc.change_management.ml.sctp]	Change Management feature that uses Predictive Intelligence to identify change clusters that may be candidates for standard change templates.	Active	false	com.snc.change_management.ml
Change Management Workflows [com.glideapp.workflow.change_management]	This plugin activates the workflows which control the change processes associated with the Emergency, Standard, and Normal change types introduced in the Geneva release.	Active	true	com.glideapp.workflow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Change Request [com.snc.change_request]	The base Change Request plugin that contains the core tables for Change Management, including the change request and change task tables.	Active	true	com.glideapp.report.itsm.change.overview
Change Request Calendar [com.snc.change_request.calendar]	This plugin activates the Change Conflict Calendar introduced in the Kingston release. It enables you to view a change request and its potential conflicts once a primary configuration item, planned start date, and planned end date have been provided within a Change request.	Active	false	com.snc.app.calendar
Channel Proxy [com.glide.channelproxy]	Contains the Business Stakeholder role which authorizes the users to approve, view/read records across the organization and view reports. Note there is a fee associated with the Business Stakeholder role. Do not assign it to users without confirming your organization has the appropriate entitlement.	inactive	false	com.glide.external.app
Checklist [com.glide.ui.checklist]	Provides a simple way to track the progress of tasks without creating additional records, using checklists that can be added to any form.	Active	false	• com.glide.ui.ng • com.glide.rest.service
CI Metrics and Remediation - MECM Adapter [sn_mecm_adapter]	Enables you to retrieve and display the CI metrics information of the affected CIs.	Active	true	None

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				
CIDC Spoke [com.sn_cicd_spoke]	Provides subflows and actions to call the Continuous Integration Continuous Delivery REST API endpoints included in the com.glide.continuousdelivery plugin.			
CIWF UI Components [com.sn_ciwf_ui_cmpnt]				
Client Transaction Timings [com.glide.client_transaction]	The Client Transaction plugin provides support to track client rendering times at the server, lining the values up with the server transaction times.	Active	false	
Cloud Action Library [sn_itom_cal] (Available in the ServiceNow Store)	Use the ready-to-use actions and subflows of the Cloud Action Library application to interact with the cloud resources of the organization. The ITOM Governance features, such as Cloud Configuration Governance, use these actions to operate.	Active	false	• com.snc.cloud.core • com.glide.hub.action_step.payload • com.glide.hub.action_step.script_mid • com.snc.itom.license • com.glide.hub.action_step.get_connection_i • com.glide.hub.action_step.rest • com.glide.hub.action_step.xmlparser • com.glide.hub.action_type.datastream • com.glide.hub.integration.runtime • sn_itom_licensing
Cloud API [com.snc.cloud.api]	Cloud API Framework.	Inactive	false	• com.snc.core.automation • com.glide.ui.angular • com.glide.ui.heisenberg • com.snc.discovery.core

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Cloud Config Management [com.snc.config.mgmt]	Config Management Framework.	Inactive	false	• com.glide.rest • com.snc.cloud.core
Cloud Configuration Governance [com.sn.itom.ccg] (Available in the ServiceNow Store)	Use the application to check the configuration settings of cloud resources in your organization against a set of policies to identify violations. After identifying the violation, use remediation workflows to mitigate them.	Active	false	• com.snc.cloud.core • com.glide.hub.action_step.payload • com.glide.hub.action_step.script_mid • com.glide.hub.action_step.get_connection_i • com.snc.itom.license • com.glide.hub.action_step.rest • com.glide.hub.action_step.xmlparser • com.glide.hub.action_type.datastream • com.glide.hub.integration.runtime • sn_itom_licensing • sn.itom.cal
Cloud Configuration Governance: Content Pack [com.sn.itom.ccg.cp] (Available in the ServiceNow Store)	Provides the base-system policies, subflows, and flow designer actions for the Cloud Configuration Governance application.	Active	false	com.sn.itom.ccg
Cloud Insights Billing [sn_clin_billing] (Available in the ServiceNow Store)	Provides the tools to download the billing data for the provisioned Amazon Web Services (AWS) and Microsoft Azure resources and visualize the data in a dashboard. Use Cloud Insights Billing along with Cloud Provisioning and Governance running only on non-domain-separated instances.	Active	false	• Plugin dependencies: ◦ Orchestration ◦ com.snc.cloud.api ◦ com.snc.cmdb.scoped ◦ com.itom.jutils ◦ com.snc.clotho ◦ com.snc.ng.pattern.designer ◦ com.glide.hub.integrations • Application dependencies: ◦ sn_itom_pattern ◦ sn_cmdb_ci_class ◦ sn_cai ◦ sn_cld_intg_core

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
				◦ sn_cld_intg_aws ◦ sn_cld_intg_azure ◦ sn_cld_intg_gcp ◦ sn_cld_spend_core ◦ sn_cld_spend_aws ◦ sn_cld_spend_azure ◦ sn_cld_spend_gcp
Cloud Migration Assessment [com.sn_cloud_migration] (Available in the ServiceNow Store)	Allows you to plan, organize, and track the process of relocating your enterprise IT resources and workloads to cloud platforms.	Inactive	true	• com.sn_cloud_mig_base • com.sn_itom_licensing • com.snc.discovery
Cloud Migration Base [com.sn_cloud_mig_base] (Available in the ServiceNow Store)	A base plugin for the Cloud Migration Assessment application.	Inactive	false	None
Cloud Operations Workspace [com.cloud_operations_workspace] (Available in the ServiceNow Store)	Offers a comprehensive solution to manage the cloud operations of your organization.	Active	false	• com.snc.cloud.core • com.snc.ng.pattern.designer
Cloud Provisioning and Governance [com.snc.cloud.mgmt]	Cloud Provisioning and Governance - Integration with AWS, Azure, VMware OOB and extensible to add support for new clouds.	Inactive	false	• com.glide.service-portal • com.snc.cloud.core • com.snc.config.mgmt • com.snc.runbook_automation
Cloud Provisioning and	Cloud Provisioning and Governance Core - Discovery and Resource Blocks.	Inactive	false	• com.snc.discovery.core • com.snc.cloud.api

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Governance Core [com.snc.cloud.core]				
Cloud Provisioning and Governance: Google Cloud Connector [sn_cmp_gcp] (Available in the ServiceNow Store)	Enables integration with Google Cloud Platform.	Active	false	• com.snc.cloud.mgmt • com.snc.pattern.designer • com.snc.ng.pattern.designer
Cloud Provisioning and Governance: Terraform Connector [sn_cmp_terraform] (Available in the ServiceNow Store)	Offers ability to provision to AWS, Azure, VMware and IBM Cloud. Develop at speed, ingest your Infrastructure as Code (IaC) template and create a ServiceNow catalog item in minutes.	Active	false	com.snc.cloud.mgmt
Cloud Spend Dashboard [sn_sam_cld_spend] (Available in the ServiceNow Store)	This plugin provides Insights into SaaS, IaaS and PaaS Spend and Savings	Active	false	• com.snc.cloud.api • com.snc.cmdb.scoped • com.snc.discovery • com.itom.jutils • com.sn_itom_opt_licensing • com.snc.clotho • com.glide.hub.integration.runtime • com.glide.hub.integrations.standard • com.sn_sam_saas
CMDB-ATF Tests [com.snc.cmdb.atf]	Provides the CMDB Automated Test Framework (ATF) tests.	Inactive	false	com.snc.cmdb
CMDB and CSDM Data	Provides insights into the key foundational indicators of your	Active	false	None

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Foundations Dashboard [sn_getwell] (Available in the ServiceNow Store)	CMDB and Common Service Data Model (CSDM) as well as recommendations to ensure that the CMDB and CSDM are properly configured for optimal usage and to mitigate any potential risks.			
CMDB Dashboard [com.snc.cmdb.dashboard]	Report overall CMDB and class level aggregated CI health.	Active	false	• com.snc.cmdb • com.snc.pa • com.glideapp.canvas
CMDB Group [com.snc.cmdb.group]	Provides Configuration Item grouping capabilities.	Active	false	
CMDB Group Dashboard [com.snc.cmdb.group.dashboard]	Report the health of CIs in CMDB Groups, on a dashboard.	Active	false	• com.snc.cmdb • com.snc.pa • com.glideapp.canvas • com.snc.cmdb.group
CMDB Mainframe [com.snc.cmdb.mainframe]	One of the plugins included in the Extended CMDB plugin, used for mainframe configuration items.	Inactive	false	com.snc.cmdb
CMDB Radio Category [com.snc.cmdb.radio.category]	One of the plugins included in the Extended CMDB plugin, used for radio configuration items.	Inactive	false	com.snc.cmdb
CMDB Search [com.snc.cmdb_search]	Query-like search on CMDB CIs and relationships. Converts a free-style query with configurable synonyms and stop words, into a properly formulated query on CMDB tables.	Active	false	com.snc.cmdb
CMDB Telecom Category	One of the plugins included in the Extended CMDB plugin, used for	Inactive	false	com.snc.cmdb

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.cmdb.telecom.category]	telecom configuration items.			
CMDB Test Equipment [com.snc.cmdb.test_equipment]	One of the plugins included in the Extended CMDB plugin, used for test equipment configuration items.	Inactive	false	com.snc.cmdb
CMDB Agent Workspace [com.cmdb-workspace]	Enables CMDB capabilities in Agent Workspace.	Active	false	com.snc.agent_workspace.itsm
CMS User Interface - Service Management Core [com.snc.service_management.core.cms]	All Content Management System items (blocks, pages, and menus) used to reference core IT self-service applications are packaged in this plugin. It is also the core foundation for all Service Management applications.	Inactive	false	• com.glide.cms • com.glide.cms.extensions • com.glide.db_images
CMDB Workspace [sn_cmdb_ws] (Available in the ServiceNow Store)	A modern workspace where you can search the CMDB, understand CMDB health and CIs and access common tools	Active	true	• sn_cmdb_pg_templts • sn-cmdb-nlq-search
Coaching [com.sn_coaching]	The Coaching module facilitates the coaching of employees on their work through the use of coaching opportunities (critical moments in a process) that can be conditionally configured.	Active	true	• com.snc.planned_task_v2 • com.snc.skills_management • com.snc.organization_management
Code Search [com.glide.code-search]	A configurable Code Search API.	Active	false	
Collaboration	Retired. Replaced by the Connect plugin.	Inactive	false	com.glide.connect

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.collaboration]				
Collaboration Services [sn_tcm_collab_hook] (Available in the ServiceNow Store)	Provides the ability to communicate over Slack. The application integrates collaboration tools with ServiceNow. You can directly use a collaboration tool from task records and as a channel within communication plans. You can Initiate a conversation from a task record with the collaboration tool of your choice. Enable Collaboration Services for communication plans to pre-define Slack or Microsoft Teams as a real-time collaboration and communication channel.	Inactive	false	None
Collision APIs [com.glide.collision_detector]	APIs for predicting the behavior of the upgrade engine with regard to customizations.	Active	false	
Column Level Encryption [com.glide.now.platform.encryption]	Allows text fields and attached files to be encrypted.	Active	false	com.glide.encryption.core
Column Statistics [com.glide.column_statistics]	Column statistics used to suggest database indexes.	Active	true	
Commercial Lines Claims [com.sn_ins_claim_cml] (Available in the ServiceNow Store)	Commercial claims process digitization from First Notice of Loss (FNOL) through settlement.	Active	true	• com.snc.csm_contributor_user • sn_csm_playbook (2.0.2) • sn_bom (1.8.0) • sn_ins_siu (1.0.0)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Commercial Lines Servicing [com.sn_ins_policy_b2b]	Enables insurance carriers to be agile and responsive in resolving mid-term policy changes and requests from commercial customers. Carriers can prioritize underwriting cases by premium value with intelligent business rules and eliminate low-value manual work from underwriting queues.		Yes	• com.sn_csm_playbook • com.snc.csm_contributor_user • com.sn_bom • com.sn_bom_document • com.sn_ins_uw_b2b
Commercial Lines Underwriting [com.sn_ins_uw_b2b]	Enables insurance carriers to route complex service requests instantly to underwriters based on their existing underwriting guidelines and rules. Underwriters can collaborate with distribution and servicing teams through an efficient and transparent workflow.		No	• com.snc.csm_contributor_user • com.sn_bom
Common ITSM Service Portal Application Components [com.snc.app_common.service_portal]	Common Application Components for Service Portal.	Active	false	com.glide.service-portal
Communities [com.sn_communities]	Communities enable customers to interact with each other and share knowledge. There is no structured way to disseminate knowledge or capture knowledge from communities and make it a part of customer service and customer success.	Inactive	false	• com.glide.service-portal • com.sn_communities_global • com.snc.activity_subscriptions • com.snc.gamification • com.snc.csm_unified_theme • com.snc.assignment_workbench • com.glide.processor.ics

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Communities Demo Data [com.sn_communities_demo]	Loads demo data for the community application.	Inactive	false	• com.sn_communities • com.snc.csm_unified_theme.demo
Company extension [com.glide.company]	Adds currency columns to the Company [core_company] table.	Active	false	com.glide.currency
Configuration Compliance [snc_vulc] (Available in the ServiceNow Store)	Exposes your high-impact configuration-related security vulnerabilities, and orchestrates their remediation across frequently isolated information security, IT operations, and business process areas. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.security_support.common
Configuration Compliance Dependencies [snc_vulc_dep]	System dependencies required by Configuration Compliance. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.		false	
Configuration Management (CMDB Enterprise Edition) [com.snc.cmdb.enterprise]	A collection of modules for specialized configuration items, such as radio hardware, test equipment, and voice system hardware.	Active	false	com.snc.cmdb
Configuration Management (CMDB) [com.snc.cmdb]	Provides core functionality for the configuration management database, including enterprise hardware	Active	true	• com.glide.custom_web_service • com.sn_cmdb_content • com.snc.db.rotation • com.glide.ui.list_v3_components

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	and configuration item relationships.			
Configuration Management For Scoped Apps (CMDB) [com.snc.cmdb.scoped]	Enables scoped apps access to Identification Engine APIs.	Active	false	com.snc.cmdb.enterprise
Connect [com.glide.connect]	Provides a real-time messaging platform that connects you to your coworkers, bypassing email and static documents.	Active	false	• com.glideapp.live_feed • com.glideapp.live_common • com.glide.ui.ng • com.glide.ui.angularui • com.glide.ui.ng.amb • com.glide.notification.push • com.glide.db_audio
Connect Interaction UI [com.connect-interaction-ui]	User interface components to adapt Connect Support to use the interaction table as a source of truth.	inactive	false	• com.glide.interaction • com.glide.graphql • com.glide.connect.support • com.glide.connect.ui_components
Connect Scriptable APIs [com.glide.connect.scriptable]	Scriptable APIs for Connect.	Active	false	
Connect Service Portal Widgets [com.glide.connect.sp_widgets]	Connect Service Portal Widgets.	Inactive	false	• com.glide.service-portal • com.glide.connect • com.glide.connect.ui_components
Connect Spoke for Flow Designer [com.glide.connect_v3plus.core.ah]	The Connect Spoke for the Flow Designer provides actions that a process analyst can use when designing flows. The actions allow them to automate the creation of conversations, to add users to a conversation, and to send messages to a conversation. These actions work with	Active	false	com.glide.connect.scriptable

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Connect API version 3 and later.			
Connect Support [com.glide.connect.support]	Builds on the Connect messaging platform and enables support agents to provide real-time assistance to end users, using queues.	Inactive	true	com.glide.connect
Connect Support and Service Portal Integration [com.glide.connect.support.service-portal]	Adds Connect Support components for use in the Service Portal.	Inactive	true	• com.glide.service-portal • com.glide.connect.support
Connect Support Manager's Dashboard [com.glide.connect.managers_dashboard]	Provides a homepage and all the configuration records required to analyze Connect Support in reporting.	Inactive	false	com.glideapp.report
Connect Support Routing [com.glide.connect.support.routing]	Plugin to enable processor for Routing Connect Support request to appropriate Chat Queue.	Inactive	false	• com.snc.matching_rule • com.glide.connect.support
Consumer Service Portal [com.glide.service-portal.consumer-portal]	Enables the Consumer Service Portal, a web-based portal based on the Service Portal application that your company can use to provide information and support to consumers.	Inactive	false	• com.glide.service-portal.customer-service-base • com.glide.connect.anonymous_support
Content Analytics [com.sn_content_analytics]	Content Analytics	Active	false	com.glide.scope.access.restricted_caller
Content Automation [com.sn_content_automation]	Content automation functionality.	inactive	true	• com.sn_content_delivery • com.glide.scope.access.restricted_caller

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Content Delivery [com.sn_content_delivery]	Content Delivery	inactive	true	• com.glide.db_images • com.glide.scope.access.restricted_caller • com.glide.service-portal
Content Experiences [com.sn_content_automation]	Set up and package your content into a campaign using multiple channels to deliver the right message to the right audience at the right time.	Active	true	
Content Governance [sn_cg]	Content Governance provides communication employees a streamlined method to request content to be created for Content Publishing or Content Experiences.	Active	true	com.sn_content_publishing_store
Content Management [com.glide.cms]	Allows administrators to create custom, branded, web user interfaces on top of an existing ServiceNow instance. Use Service Portal for new development instead of CMS. Service Portal is an alternative to CMS with a refined user experience, and is active by default in the base system. For more information, see Service Portal and Content Management and Service Portal.	Inactive	false	• com.glideapp.home • com.glide.text_search • com.glide.db_images
Content Management Extended Types [com.glide.cms.types]	An extension to Content Management that adds iFrames and Flash frames. Use Service Portal for new development instead of CMS. Service Portal is an	Active	false	• com.glide.cms.type.flash • com.glide.cms.type.iframe

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	alternative to CMS with a refined user experience, and is active by default in the base system. For more information, see Service Portal and Content Management and Service Portal.			
Content Publishing [com.sn_content_publishing_store]	Content Publishing allows you to communicate with your employees using different methods.	Active	true	
Context Ranking [com.glide.sorting]	Support for drag-and-drop lists and ranking dialog. Context ranking allows a user to sort a collection of records independently of the attributes of those records. Context Ranking is activated automatically with the SDLC - Scrum Process Pack plugin.	Inactive	false	com.glide.list_v2
Context-Sensitive Help [com.glide.context_help]	Provides a context-sensitive help system, providing links to specific help pages. These help pages can be linked to the list or form view of any table, or to the form view of a specific record in a table.	Active	false	
Context-Sensitive Help Rest API [com.glide.context_help.rest_api]	REST API for the context-sensitive help system	Active	false	• com.glide.context_help • com.glide.scripted_rest_services
Contextual Search [com.snc.contextual_search]	Display related articles or service catalog items within a form or record producer.	Active	false	com.snc.application.json_service

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Contextual Search - Internal [com.snc.contextual_search.internal]	An internal plugin component for Contextual Search.	Active	false	
Contextual Search - Service Portal [com.snc.contextual_search.service-portal]	Service Portal components for Contextual Search.	Active	false	• com.snc.contextual_search • com.snc.app_common.service_portal
Contextual Security: Role Management [com.glide.role_management]	Provides the flexibility and power to protect information by controlling read/write/create/delete authorization.	Active	false	• com.glide.db_view • com.glide.ui_page
Contextual Security: Role Management V2 [com.glide.role_management.inh_count]	Prevents duplicate entries in sys_user_has_role for inherited roles, based on the value of the inh_count column.	Active	false	com.glide.role_management
Contextual Security: Role Management V2 REST API [com.glide.role_management.inh_count.rest_api]	Prevents duplicate entries in sys_user_has_role for inherited roles, based on the value of the inh_count column.	Active	false	• com.glide.role_management.inh_count • com.glide.scripted_rest_services
Continual Improvement Management (CIM) [com.sn_cim]	Core plugin for Continual Improvement Management (CIM)	Active	true	• com.snc.organization_management • com.snc.planned_task_v2 • com.snc.pa • com.snc.pa.spotlight • com.sn_coaching • com.sn_cim_atf
Continual Improvement Management Automated Tests [com.sn_cim_atf]	ATF plugin for Continual Improvement Management (CIM)	Active	true	
Contract Management	Provides the ability to manage all types of contracts.	Active	true	com.snc.asset_management

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.contract_management]				
Contract Management - Form Layouts and Behavior [com.snc.contract_management.form_layouts]	Provides the ability to customize the Contract Management form layout.	Inactive	false	
Conversation Branding Settings [com.glide.cs.branding]	Conversation Branding Settings for web chat client.	Active	false	
Conversation General Settings [com.glide.cs.settings]	Conversation General Settings for Virtual Agent.	Active	false	
Conversational Analytics [sn_ci_analytics] (Available in the ServiceNow Store)	A dashboard that helps you improve Virtual Agent (VA) interactions with users by providing deep insights into conversational data, discover which topics confuse users, and find out where users often transfer to live agents.	Active	false	• com.glide.cs • com.glide.appsee
Conversational Integration with Alexa [com.sn.va.alexa] (Available in the ServiceNow Store)	Enables Virtual Agent to interact with requesters through the Amazon Alexa's voice interface using the Virtual Agent topics.	Active	false	
Conversational Integration with LINE [sn_va_line]	The plugin helps engage with requesters who prefer the LINE messaging app to interact with your business. Requesters include customer contacts and consumers.	Active	true	• com.glide.cs.custom.adapter • com.glide.hub.integration.runtime • com.snc.ihub_spoke_util_pack • sn_agent_initiated

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Conversational Integration with Microsoft Teams [sn_va_teams] (Available in the ServiceNow Store)	Enable requesters to chat with Microsoft Teams or live agents using the Microsoft Teams application.	Active	true	
Conversational Interfaces - Diagnostics [sn_ci_diagnostics]	The CI Diagnostic Tool is a framework that allows a privileged user (admin) to perform a diagnosis of various products through scripts and powerful VA topics.	Active	false	com.glide.cs.chatbot.lite
Virtual Agent Topic Recommendations [sn_topic_recommend]	The plugin helps identify pre-built Virtual Agent topics that can be quickly implemented in your organization.	Active	false	
Conversational Integration with WhatsApp (powered by Twilio) [sn_va_whatsapp_twi]	The plugin enables customer engagement over the WhatsApp messaging app.	Active	true	• sn_twilio_spoke • com.glide.hub.integration.runtime • com.glide.cs.custom.adapter • sn_agent_initiated • com.glide.cobject • com.glide.messaging.awa
Bot Interconnect [sn_va_bot_ic] (Available in the ServiceNow Store)	Virtual Agent Bot Interconnect functions as the primary bot in a diverse chat environment. It can help you reduce complexity and create a unified chat experience for your end users. Virtual Agent Bot Interconnect gives your end users access to multiple channels and a wide variety of enterprise tasks	Active	false	com.glide.cs.chatbot

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	that are available from ServiceNow.			
Conversational integration with Slack [sn_va_slack] (Available in the ServiceNow Store)	Enable requesters to chat with Virtual Agent or live agents using the Slack application	Active	false	
Conversational integration with Workplace from Facebook sn_va_fb_workplace (Available in the ServiceNow Store)	Enable requesters to chat with Virtual Agent or live agents using the Workplace from Facebook application.	Active	false	
Conversational Integration with Facebook Messenger sn_va_fb_messenger (Available in the ServiceNow Store)	Enables users to interact with virtual and live agents using the Facebook Messenger application.	Active	true	
Conversational IVR with Amazon Connect [com.sn.va.amz.connect] (Available in the ServiceNow Store)	Enables users to conduct conversations with a bot via the phone channel, in an interactive voice response.	Active	false	Virtual Agent API (sn_va_as_service)
Virtual Agent API sn_va_as_service	Integrates any chat interface or a bot with Virtual Agent or Agent Chat.	Active	false	• com.glide.cs.custom.adapter • com.glide.cs • com.glide.cs.chatbot

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				• com.snc.ihub_spoke_util_pack • com.sn_ext_agent_management_util
Conversational Google Business Messages Adapter [com.sn.gbm.adapter] (Available in the ServiceNow Store)	This application is for mobile messaging. The application enables end users to interact with ServiceNow Virtual Agent and live agents through a Google Business Messages interface using the Google Business Messages service.	Active	false	• Plugins dependency: ◦ com.glide.cs.custom.adapter ◦ com.glide.hub.integration.runtime ◦ com.snc.ihub_spoke_util_pack ◦ com.glide.messaging.awa • App dependency: sn_agent_initiated
Conversational Interfaces Guided Setup [com.glide.ci_guided_setup]	Provides customers with a visual guidance to configure Conversational Interfaces.	Active	false	• com.snc.guided_setup • com.glide.guided_setup_config
Conversational Interfaces Landing Page [com.glide.cs.landing_experience]				
Conversational Interfaces Settings [com.glide.cs.admin_console]				
Conversational Messaging [com.glide.messaging.awa]	Conversational Messaging			
Conversational SMS Integration with Twilio [com.sn.va.sms.twilio] Maintenance mode	Integration to SMS for Advanced Work Assignment	Active		
CORS support for REST API [com.glide.rest.cors]	CORS support for REST API	Active	false	

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Cost Management [com.snc.cost_management]	Tracks operating costs for configuration items and task-related activities, allocates the costs to business consumers, and compares actual allocations to planned budgets.	Inactive	true	com.snc.expense_line
Craft.co Integration for Supplier Lifecycle Operations [com.snc.sn_supplier_craft] (Available in the ServiceNow Store)	Provides a pre-configured integration with Craft.io. Craft.io is a supplier intelligence platform that offers validated and comprehensive information about suppliers with which a company engages.	Inactive	false	com.snc.sn_supplier_mgmt
Cross-scope Access [com.glide.scope.access]		Inactive	false	
Cross-scope App Privilege Enforcement [com.glide.scope.privilege]	System for capturing and enforcing cross-scope privileges used by Scoped Apps	Active	false	
CSM Extension for Proxy Contacts [com.snc.csm_proxy_contacts]	Enables users who are internal to an organization to be the proxy case contact on behalf of customers.	Inactive	true	com.sn_customerservice
CSM Order Management [com.sn_csm_order_mgmt]				
CSS Theme support [com.glide.ui.themes]	Provides support for CSS customizations to the user interface.	Inactive	false	com.glide.ui.themes.core
CSS Theme support - UI 14 [com.glide.ui.themes.doctype]	Provides UI15 themes.	Active	false	com.glide.ui.themes.core

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
CTI Demo Data for HRSD [com.sn_hr_cti_demo]	Provides CTI Softphone application demo data for HRSD.	Inactive	false	• com.snc.cti • com.sn_hr_core
CTI Softphone [com.snc.cti]	Enables Twilio integration using Notify and OpenFrame to provide softphone functions and call center capabilities. These include make, receive phone calls, transfer, hold and mute. Applications like Customer Service and Incident Management provide demo workflows for CTI.	Active	true	• com.snc.matching_rule • com.sn_openframe • com.snc.notify • com.snc.notify.twilio_direct
Credit Assessment [com.sn_bom_credit_asmt]				
CSM and FSM Configurable Workspace Foundation [com.snc.uib.cwf_workspace]	Supports refactored version of customer workflow agent workspace.	Inactive	false	• com.snc.app_shell_aw • com.snc.uib.base_agent_workspace • com.snc.uib.cwf_workspace.integrations
CSM Configurable Workspace [com.snc.uib.csm_agent_workspace]	Supports refactored version of customer workflow customer service management agent workspace.	Active	true	• com.snc.uib.cwf_workspace • com.glideapp.servicecatalog.uxfworkspace • com.snc.uib.lookup_verify • com.devsnc_sn_response_templates
CSM Configurable Workspace [sn_csm_wfo_workspa] (Available in the ServiceNow Store)	Provides a wide range of capabilities such as Channel Management, Scheduling, Coaching, and Teams applications in CSM Manager Workspace.	Active	true	• com.snc.skills_management • com.snc.sre • com.snc.pa.customer_service
CSM Configurable Workspace	Installs Lookup and Verify artifacts for the CSM Configurable Workspace that	Inactive	true	com.snc.sn_csm_lookup_verify

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Lookup and Verify [com.snc.uib.lookup_verify]	includes pages, macroponents, data brokers etc. and can be shared in different applications.			
CSM Configurable Workspace Special Handling Notes [com.snc.uib.special_handling_notes]	Installs Special Handling Notes artifacts for the CSM Configurable Workspace that includes pages, macroponents, data brokers etc. and can be shared in different applications.	Inactive	false	
CSM Contributor User [com.snc.csm_contributor_user]	Allows you to enable contributors to report and collaborate on cases created for customers, service organizations, or themselves.	Active	true	com.snc.cs_base_extension
CSM Query Rules [com.snc.csm_query_rules]	Plugin that contains the query rules for the CustomerService tables.	Inactive	false	com.snc.query_rules
Currency support for the service catalog [com.glideapp.servicecatalog.currency]	Enables service catalog to support fully localized currencies for item prices and options.	Active	false	com.glide.currency
Custom URL [com.snc.customurl]	Activate this plugin to set up custom URL on the instance.	Inactive	false	com.snc.customurl.internal
Custom URL - Internal [com.snc.customurl.internal]	An internal plugin component to Scripted CustomURL APIs.	Inactive	false	
Customer Communities [com.sn_customer_communities]	Enables you to connect, engage, and collaborate with your employees, customers, partners, and prospects. It is only available for customers who are licensed for	Inactive	false	• com.sn_external_user_register • com.sn_communities

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Customer Services Management.			
Customer Project Management [com.snc.csm_ppm]	Provides the ability to create and manage projects for a customer account and gives end users visibility into their projects via the portal. This plugin requires the Customer Service plugin and the Project Portfolio Suite with Financials plugin.	Inactive	true	• com.sn_customerservice • com.snc.financial_planning_pmo
Customer Service [com.sn_customerservice]	Enables you to provide service and support for your external customers using several communication channels, such as email, web, and telephone. A case is created to track the issue reported or service requested, and assigned to groups or agents. Customer service agents in your organization work on the cases and resolve issues.	Inactive	false	• com.snc.asset_management • com.snc.state_flows • com.glide.connect.support • com.glide.connect.support.routing • com.snc.assignment_workbench • com.snc.cs_base • com.snc.skills_management • com.snc.assessment_core • com.snc.process_flow_formatter • com.snc.task_relations • com.snc.task_activity • com.snc.matching_rule • com.sn_resolutionshaper • com.sn_openframe • com.sn_shn • com.snc.contextual_search.dynamic_filters • com.sn_cs_social
Customer Service Base Extension Entities [com.snc.cs_base_extension]	This plugin provides granular access control through relationships, user profile attributes and roles.	Active	false	com.sn_customerservice
Customer Service Case Types [com.snc.csm_case_types]	Enables customers to manage complex case processes by defining new case types. A case type represents	Inactive	true	com.sn_customerservice

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	an individual business process and is a collection of diverse inputs and tasks that an agent performs to resolve customer requests.			
Customer Service CTI Demo Data [com.snc.customerservice_cti_demo]	Provides demo data for the CTI Softphone application.	Inactive	false	• com.sn_customerservice • com.snc.cti
Customer Service Household [com.snc.household]	Enables managing customer service for consumers belonging to the same household. Provides the ability to define the members of a household and relationships between them.	Active	true	com.snc.cs_base_extension
Customer Service Install Base Management [com.snc.install_base]	Enables customers to capture the current state of a customer's install base and establish the relationship to any downstream entities that might impact their functioning.	Inactive	false	com.sn_customerservice
Customer Service Management Demo Data [com.snc.customerservice.demo]	Provides demo data for the Customer Service Management application.	Inactive	false	• com.sn_customerservice • com.glide.service-portal.consumer-portal • com.glide.service-portal.customer-portal • com.snc.case_assignment_workbench_dem • com.snc.work_management.demo • com.snc.planned_maintenance • com.snc.kb_product_entitlements • com.snc.publications_demo • com.snc.shn_demo • com.snc.csm.order • com.snc.csm_fsm_integration

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Customer Service Management for Orders [com.snc.csm.order]	Extends the Customer Service Management application to support order-related issues by integrating with order management systems to create and resolve customer requests.	Inactive	true	com.sn_customerservice
Customer Service Mobile [com.sn_csm_mobile]	Enables the mobile user interface for Customer Service Management.	Inactive	false	com.sn_customerservice
Customer Service NLU Model for Virtual Agent Conversations [com.sn_csm.nlu]	Enables CSM Virtual Agent conversations.	Active	false	• com.sn_customerservice • com.glide.cs.chatbot • com.glide.cs.sn-va-web-client-app
Customer Service Portal [com.glide.service-portal.customer-portal]	Enables the Customer Service Portal, a web-based portal based on the Service Portal application that your company can use to provide information and support to customers.	Inactive	false	com.glide.service-portal.customer-service-base
Customer Service Social Integration [com.sn_cs_social]	Customer Service Social integration, currently adds appropriate data schema changes needed for social integration scenarios.	Inactive	false	
Customer Service Spoke [com.snc.customer_service.spoke]	Enables Flow Designer users to automate the creation of customer service cases and updates to existing cases.	Inactive	false	Flow Designer - Installer (com.glide.hub)
Customer Service Virtual Agent Conversations [com.sn_csm.virtualagent]	CSM Virtual Agent Conversations	Active	false	• com.sn_customerservice • com.glide.cs.chatbot • com.glide.cs.sn-va-web-client-app

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Customer Service with Field Service Management [com.snc.csm_fsm_integration]	Enables Account, Contact, Partner, Partner Contact, Consumer information from Customer Service to Field Service Management.	Inactive	false	• com.snc.work_management • com.sn_customerservice • com.glide.service-portal.customer-portal
Customer Service with Request Management [com.sn_cs_sm_request]	Provides an integration between Customer Service Management and the Request Management application. Enables users to create request records from a Customer Service case.	Active	true	com.sn_customerservice
Customer Service with Service Management [com.sn_cs_sm]	Provides an integration between Customer Service Management and the Incident, Problem and Change Management applications. Enables users to create incident, problem and change records from a Customer Service case.	Active	true	com.sn_customerservice
Customer Service Install Base Management [com.snc.install_base]	Enables customers to capture the current state of their install base and establish the relationship to any downstream entities that might impact their functioning.	Inactive	false	com.sn_customerservice
Dashboard for Assessment and Survey [com.snc.assessment_core.dashboards]				
Dashboard for Service Catalog [com.glideapp.servicecatalog.dashboard]				
Data Archiving [com.glide.auxdb]	Provides the ability to move a subset of data from large tables	Active	false	com.snc.db.replicate

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	into the data archive. Typical candidates include historical ITIL documents such as incidents that were closed last year, but the functionality supports archiving of non ITIL documents as well.			
Data Classification [com.glide.data_classification]		Active	true	
Data Certification [com.snc.certification_v2]	Enables field-level certification of data, either scheduled or on-demand.	Inactive	true	• com.glide.list_v2 • com.snc.certification_core
Data Discovery APIs [com.glide.data_discovery]	Provides the APIs to discover sensitive data.	Active	false	
Data Filtration [com.glide.data_filtration] (New in Tokyo)	Controls the access to tables and records based on subject attributes when performing read queries.	Inactive	false	None
Data Lookup and Record Matching Support [com.glide.data_lookup]	Allows administrators to define rules that automatically set one or more field values when certain conditions are met. This plugin completely replaces Priority Lookup. Any custom logic defined in the CalculatePriority business rule must be manually translated into the new priority data lookup definition.	Active	true	
Data Lookup and Record Matching Support for Service Catalog	Allows administrators to perform data lookups for variables on service catalog item screens, on requested items, and on catalog tasks as	Active	false	• com.glideapp.servicecatalog.platform • com.glide.data_lookup

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.data_lookup.catalog]	a user fills out the values contained in variables.			
Data Policy 2 [com.glide.data_policy2]	Defines mandatory or read-only requirements for table fields.	Active	false	
Data Structures [com.snc.datastructure]	Provides element types: DataStructure and DataObject for flyweight data that can be stored internally as JSON and utilized via the DataStructure API.	Active	false	
Database Column Offrow Migration [com.glide.db.offrow_migration]	UI/Script elements to allow for compatible columns to be moved offrow.	Inactive	false	
Database Rotation [com.snc.db.rotation]	Provides the ability to break large tables into smaller, more manageable tables â€“ these smaller tables contain a specified period of data and the table with oldest data is purged on an on-going basis.	Active	false	
Database Rotation with Default Tables [com.snc.db.rotation_default_tables]	Adds database rotation functionality to default tables.	Active	false	com.snc.db.rotation
Database Storage for Audio Files [com.glide.db_audio]	Allows audio files to be uploaded and stored in the database, and referenced in HTML.	Active	false	
Database Storage for Images [com.glide.db_images]	Allows images to be uploaded and stored in the database, and referenced in HTML.	Active	false	
Database Storage for Video Files	Allows video files to be uploaded and stored in the	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.db_video]	database, and referenced in HTML.			
Database Views [com.glide.db_view]	Allows you to define table joins for reporting purposes.	Active	false	
Database Views for Service Management [com.snc.service.db_views]	Predefined Database Views for Service Management.	Active	false	• com.glide.db_view • com.snc.service • com.snc.problem • com.snc.incident • com.glide.metrics
DataTables 1.1.0 Components [com.glide.ui.ng.datatables]	DataTables 1.1.0 Components	Inactive	false	
Default Financial Management Cost Model [com.snc.financial_management_model]	Provides a preconfigured Cost Model and artifacts to enable Financial Analysts to assess spending data and generate reports. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	com.snc.financial_management
Delegated Dev User Administration [com.sn_dd_user_admin]	Delegated Dev User Administration.	Active	false	
Delegated Development [com.glide.delegated_development]	Support for controlled development by users who do not have full administrative rights.	Active	false	
Delegated Request Experience (Maintenance mode)	Enables delegated requests for catalog items and activates the Requested For variable type and the	Active	false	• com.glideapp.servicecatalog.platform • com.glideapp.servicecatalog.requested_for

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glideapp.servicecatalog.request_management.delegated_request_experience]	Requested For field in Requested Item			
Delete Recovery [com.glide.delete_recovery]	Provides the ability to recover individual deleted items including all related cascaded deleted items.	Active	false	
Delete Recovery: Partial Undelete Support [com.glide.delete_recovery.partial_undelete]	Extends the Delete Recovery plugin to support partially undeleting records from a transaction.	Active	false	com.glide.delete_recovery
Demand Core [com.snc.demand_core]	Demand Core Plugin aids in planning, organizing, and managing Demands with basic features.	Active	false	
Demand Management [com.snc.demand_management]	Aids with capturing the demand and provides tools to screen, assess, and prioritize it. Only upgrade is allowed for this plugin. Activation should be done through PPM Standard plugin.	Inactive	true	• com.snc.process_flow_formatter • com.snc.assessment_core • com.snc.timeline_visualization • com.snc.bubblechart_workbench
Depreciation [com.snc.depreciation]	Core depreciation capabilities.	Active	false	
Designer Common [com.glide.ui.ng.dc]	Provides common components required by designers such as the form designer and the quiz designer.	Active	false	• com.glide.ui.tablet.theme • com.glide.ui.ng • com.glide.ui.themes.doctype
Desired State Certification [com.snc.certification_desired_state]	Evaluates records to see if they match a desired state, scheduled or on demand.	Active	true	com.snc.certification_core
DevOps [sn_devops] (Available in the	Provides data insights, accelerates change, and increases visibility into your DevOps environment.			

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
DevOps Config [sn_devops_config] (Available in the ServiceNow Store)	Manages and validates configuration data in your DevOps environment using a single system.	Active	false	sn_cdm
DevOps Insights [sn_devops_insights] (Available in the ServiceNow Store)	Shows the results of your overall DevOps process in a data-driven dashboard.			
DevOps Integration with Rally [sn_devops_rally] (Available in the ServiceNow Store)	Provides integration with Broadcom Rally to expose data like projects, user stories, and defects from the Rally tool in DevOps Change acceleration. You can then associate these Rally planning objects to commits and pipelines for full change traceability and acceleration.	Active	false	com.sn_devops
DevOps Util [com.snc.sn_devops_util] (New in Tokyo)	Serves as an interface to call global APIs from DevOps.	Active	false	
DHTMLX Gantt Library [com.snc.dhtmlx.gantt]	DHTMLX Gantt library.	Inactive	false	
DHTMLX Scheduler Library [com.snc.dhtmlx.scheduler]	DHTMLX Scheduler library.	Active	false	
Diagnostic Monitoring	Provides advanced diagnostic monitoring	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.monitor.diagnostics]	of each node in a ServiceNow instance.			
Dictionary overrides [com.glide.dictionary_override]	Allows specific Dictionary values to be overridden for extended table elements. For example, this plugin allows the default value for the Assigned To field in the Incident table to be different than the default value specified for the Assigned To field in the Task table.	Active	false	
Discovery and Service Mapping Patterns [com.sn_itom_pattern] (Available in the ServiceNow Store)	Deploys additional Discovery and Service Mapping Patterns.	Active	false	None
Discovery for Operational Technology [com.sn_otsm_discovery]				
Distribution Download Proxy [com.glide.dist_download_proxy]	Proxy for downloading a distribution. This plugin is activated automatically when the MID Server is activated, to handle the REST requests for file downloads that a MID Server sends to the instance.	Active	false	
Document Intelligence Administration [com.snc.docintel_admin] (Available in the	Better administration experience for Document Intelligence.	Active	false	• Document Intelligence (sn_docintel) • Predictive Intelligence (com.glide.platform_ml)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Document Service [com.sn_bom_document]				
Document Templates [com.snc.document_templates] (Available in the ServiceNow Store)	Enables you to create HTML and PDF document templates to generate standard letters or documents. Automates and simplifies the process of filling, signing, and reviewing a document online.			
Document Template Integration with AdobeSign [sn_dt_adobesign] (Available in the ServiceNow Store)	Enables you to use the Adobe Sign application for signing documents that are generated from ServiceNow Document Templates.	Active	true	None
Document Template integration with DocuSign [sn_dt_docusign] (Available in the ServiceNow Store)	Enables you to use the DocuSign application for signing documents that are generated from ServiceNow Document Templates.	Active	true	None
Documents Configuration [com.glide.doc.config]	Provides a feature to create a Table of Contents configuration (that is adding table of contents) and a document configuration such as page numbers to a PDF document that is generated from an HTML	Active	false	None

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	document. The plugin provides features to navigate easily through the PDF using Table to Contents and Document Configuration.			
Domain Support - Domain Extensions Installer [com.glide.domain.msp_extensions.installer]	Enables domain separation for Cloud Provisioning and Governance. Instances activating Domain separation should use this plugin. In addition to functionality in the Domain Support plugin, this plugin installs Domain Support - Domain Extensions that include further domain separation enhancements for all domain customers including Managed Service Providers. If the Domain Support plugin is already active, content in the Domain Support - Domain Extensions plugin will not be installed to the instance.			
Dynamic Related Records for Configurable Workspace [com.snc.uib.sn_dyn_rel_rec]	Dynamic related records provide agents with access to relational information based on the context of the current record. Admin can configure related record definitions that increase agent productivity by making data available in context to the action that the agent is performing.			

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Dynamic Scheduling [com.snc.dynamic_scheduling]	Dynamic scheduling for Service Management Applications with support for bulk task recommendations and interval based auto assignment.	Inactive	false	• com.snc.matching_rule • com.snc.service_management.core
Dynamic Translation [com.glide.dynamic_translation]	Adds ability to translate the given text from source language to target language using external machine translation providers.	Active	false	com.glide.hub.integrations
Dynamic Translation for Virtual Agent [com.glide.cs.dynamic.translation.virtual_agent] (New in Tokyo)	Activates the Dynamic Translation for Virtual Agent.	active	false	• com.glide.dynamic_translation • com.glide.dynamic_translation.va_async
Dynamic Translation Spoke [com.glide.dynamic_translation_spoke]	Provides actions that enable the ability to translate text(s) from one language to other language(s) and to detect the language of given text.	Active	false	• com.glide.hub • com.glide.hub.integration.runtime
E164 Compliant Phone Number [com.glide.phone_number]	Ensures that all necessary information for a phone number is included and properly formatted to successfully route an international call over a territory's public telephone network.	Active	true	
Email - OAUTH support for IMAP and SMTP [com.glide.email.oauth]	Support for XOAUTH and XOAUTH2 email authentication.	Inactive	false	• com.glide.email_accounts • com.snc.platform.security.oauth
Email - Support for Email Processing by Microsoft Graph API	Processes inbound emails using Microsoft Graph API.	Active	false	• com.glide.email_accounts • com.snc.platform.security.oauth • com.glide.email.oauth

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[glide.email.graph] (New in Tokyo)				
Email Access Restriction [com.glide.email_access_restriction]	Allows email table ACL conditional overrides	Active	false	
Email Accounts [com.glide.email_accounts]	Enables you to define email accounts and settings in individual records.	Active	false	
Email Address Filter [com.glide.email_address_filter]	APIs and data structures for email address filtering	Active	false	
Email Automatic User Creation [com.glide.email_createuser]	Sets glide.email.create_userid_from_email=true so that when automatic user creation is enabled, the UserID of newly-created users matches the user's email address. Also widens sys_user.user_name column to 100 bytes to accommodate longer UserIDs based on email addresses.	Active	false	
Email Client [com.glide.email_client]	Enables you to send email directly from any record	Active	false	com.glide.email_address_filter
Email Client Template [com.glide.email_client_template]	Template defines default content for the Email Client.	Active	false	
Email Digest [com.glide.email_digest]	Email Digest	Active	false	• Subscription Based Notifications • Outbound Email Notifications • com.glide.notification.preference.service
Email Filters [com.glide.email_filter]	Filters emails into different mailboxes or junk, depending on	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	headers and subject. Ignores any email that contains a VCAL invitation.			
Email Notification Preview Plugin [com.glide.email_notification_preview]	Allows you to easily preview a ServiceNow email notification without sending the notification. This allows you to preview notifications at design time.	Active	false	
Email Notification Script for Contextual Search Results [com.snc.adv_cxs_results_email_script]	Provides script for embedding contextual search results based on Predictive Intelligence additional resources and search resources in email notifications.	Inactive	false	
Email Retention [com.glide.email_retention]	Provides retention policy for email.	Active	false	• com.glide.mailbox • com.glide.auxdb
Email Service [com.glide.email.service]	Defines javascript and REST API for sending email.	Active	false	com.glide.mailbox
Email Unsubscribe [com.glide.email_unsubscribe]	Adds the ability to include an unsubscribe link in notifications.	Active	false	• com.glide.email_ordered_processing • com.glide.notification.subscription
Email/SMS Separation [com.glide.email_sms_separation]	Separation of the sending of SMS and email by adding a column to the email table. Prevents SMS messages from slowing down email message sending and vice versa. Activation of this plugin will cause email to stop sending during activation. On systems with large email tables this can take hours, and is not recommended.	Active	false	com.glide.notification
Embedded Help	Provides targeted help content to a	Active	false	com.glide.ui.hopscotch

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.embedded_help]	user in a UI page. Administrators can create customized embedded help to assist users with specific business processes within an organization.			
Embedded Help for ServiceNow Flow Designer [com.glide.hub.flow_designer_embedded_help]	This plugin is for the embedded help panel in Flow Designer. In Quebec, there is a help panel to help users easily access documentation on Flow Designer without having to navigate to the documentation site.	Active	false	
Employee Center [com.snc.employee_center]				
Employee Center - Core [com.snc.employee_center_core]				
Employee Center - Pro [com.snc.employee_center_pro]				
Employee Center - Pro [com.sn_ex_employee_center_pro]	Employee Center Pro	False	Inactive	com.sn_hr_service_portal
Employee Center integration with Zoom [com.snc.sn_ex_ec_zm] (Available in the ServiceNow Store)	Enables Employee Center integration with Zoom.	Active	true	sn_ex_sp

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Employee Document Management [com.sn_employee_document_management]	Enables you to manage large numbers of employee documents efficiently. Provides storage space and a filing system that allows you to easily retrieve documents, as well as define who can view sensitive documents, and when to purge documents. **Important:** To use Employee Document Management, the Human Resources Scoped App: Core plugin must be licensed and activated.	Inactive	true	• com.snc.documentviewer • com.glide.scope.access.restricted_caller
Employee Relations integration with NAVEX EthicsPoint [sn_hr_navex_ep] (Available in the ServiceNow Store)	Pulls employee relation cases from NAVEX EthicsPoint into ServiceNow Platform, enabling an HR Agent to work on employee relation cases from an ServiceNow instance.	Active	false	None
Employee Service Center [com.sn_hr_service_portal]	Activates a portal framework that allows administrators to build a mobile-friendly self service experience for users.	Active	true	• com.glide.connect.support • com.glide.connect.sp_widgets • com.sn_content_delivery • com.sn_content_automation • com.glide.va.sp_widgets • com.glideapp.user_criteria.scoped.api
EMR Help [sn_ind_rmt_help] (Available in the	Simplifies and streamlines the process to submit ServiceNow IT service requests related to	Inactive	true	• com.glide.encryption • com.sn_ind

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	an electronic medical record (EMR) system.			
Encrypted Workflow Scratchpad [com.snc.encrypted.scratchpad]	Encrypted Scratchpad support on the Workflow Context and Workflow Executing Activities.	Active	false	• com.glideapp.workflow • com.glide.encryption
Encryption Core [com.glide.encryption.core]	Common scripts and logic between Edge Encryption and Encryption Support.	Inactive	false	
Encrypted Workflow Scratchpad [com.snc.encrypted.scratchpad]	Encrypted Scratchpad support on the Workflow Context and Workflow Executing Activities.	Inactive	false	• com.glideapp.workflow • com.glide.encryption
Engine Based Notifications [com.glide.email_engine_notifs]	Container for engine-based email notifications. Contains a set of default email notifications. Installed only on new z-boots.	Active	false	
Enhanced Web Service Provider - Common [com.glide.web_service_provider_v2]		Active	false	com.glide.web_service_provider
Enterprise Asset Management [com.sn_eam] (Available in the ServiceNow Store)	Provides a comprehensive end-to-end solution for managing the full lifecycle of non-IT assets, providing deep visibility into the asset estate, minimizing costly downtime, and maximizing the asset usable life.	Inactive	true	• com.sn_eam_core • com.snc.procurement • com.sn_itam_recomm • com.snc.asset_management • sn_cmdb_ci_class • com.sn_itam_common • com.sn_ent • com.sn_risk_heatmap
Environmental, Social, and Governance Management [com.sn_esg] (Available in the	Provides capabilities to collect, manage, and report organizational environmental, social, and governance (ESG) data and activities to stakeholders	Active	true	• com.sn_grc • com.sn_grc_workspace • com.sn_compliance • com.snc.goal_framework • com.sn_grc_metric

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	including investors, customers, and regulators.			
E-Signature [com.snc.esign]	Allows for the electronic signing of documents on both desktop and mobile. Signatories can sign with their credentials, as an acknowledgment, or as a typed or drawn e-signature.	Active	false	• com.snc.signaturepad • com.snc.document_management • com.snc.documentviewer
ERP Customization Mining [com.snc.uib.sn_erp_mining] (Available in the ServiceNow Store)	Enables use of ERP Customization Mining to identify customized application candidates for replatforming.	Active	true	• com.snc.sn_erp_integration • com.glide.platform_ml
ERP Canvas [com.snc.sn_erp_integration] (Available in the ServiceNow Store)	Enables access to remote tables for ERP Customization Mining.	Active	false	com.glide.script.vtable
ESS Portal (implemented within Content Management) [com.glide.cms.extensions]	ESS portal content management application. Demo data includes the actual ESS portal. Use Service Portal for new development instead of CMS. Service Portal is an alternative to CMS with a refined user experience, and is active by default in the base system. For more information, see Service Portal and Content Management and Service Portal.	Active	true	com.glide.cms
Event Management Mobile	Content Of Event Management Mobile Interface Content.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Interface Content [com.em-mobile-app]				
Event Management metric explorer [com.sn-em-metric-explorer]				
Event Management Overview Homepage [com.glideapp.report.em]	Event Management Overview Homepage	Inactive	false	• com.glideapp.report • com.glideapp.itom.snac
Event stats [com.glide.stats.event]	Event stats.	Active	false	com.glide.stats
Evidence Management [com.sn_evidence_management]	Enables you to collect, categorize, retain, and secure evidence that can be used in cases that require gathering and reviewing artifacts for an investigation.	Active	true	
Execution Plan support for the service catalog [com.glideapp.servicecatalog.execution_plan]	Service Catalog execution plans.	Active	true	com.glide.execution_plan
Expense Line [com.snc.expense_line]	Core expense line table that enables cost tracking. Integrated with asset management, CMDB, cost management, and contract management.	Active	false	
Extended CMDB [com.snc.extended_cmdb]	Provides specialized configuration items, such as radio hardware, test equipment, and voice system hardware.	Inactive	false	com.snc.cmdb

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
External Content for AI Search plugin [com.glide.ais.external_content] (New in Utah)	Allows ingesting external content to AI search.	Active	false	None
External file data storage [com.sn.external.files]	Stores information about files in third-party systems and provides actions to manage the information.			
Facilities Service Management clo [com.snc.facilities_service_automation]	Manages facilities requests and enables users to report and track requests by their location. To view requests on a floor plan, the Facilities Visualization Workbench (com.snc.facilities_service_automation.fvw) plugin is required. Integration files installed when Facilities Visualization Workbench (com.snc.facilities_service_automation.fvw) is also installed. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.service_management.core
Facilities Service Management CMS Portal (Maintenance mode only) [com.snc.facilities_service_automation.cms]	Lets you launch Facilities Service Automation and other service management applications from a single CMS page. Manages facilities requests and enables users to report and track requests by their location on a floor plan. Activation of this plugin on production	Inactive	true	• com.snc.facilities_service_automation • com.snc.service_management.core.cms

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	instances may require a separate license. Contact ServiceNow for details.			
Facilities Service Management Mobile (Maintenance mode only) [com.snc.facilities_service_automation_m]	Manages facilities service management mobile components.	Inactive	true	• com.snc.facilities_service_automation • com.snc.service_management_m
Field Normalization [com.snc.field_normalization]	Provides support for cleaning up messy data through normalization and transformation.	Inactive	true	• com.glide.vars • com.glide.system_update_set
Field Service Advanced Parts Sourcing [com.snc.fsm_advanced_parts_sourcing] (New in Tokyo)	Enables Field Service agents to request and source multiple parts. Agents can receive mobile notifications when part requests are raised by their peers.	Inactive	true	com.snc.fsm_advanced_parts_sourcing
Field Service Capacity Management [com.snc.fsm_capacity_management]	Provides capabilities to manage task capacity for groups and field agents.	Active	false	com.snc.work_management
Field Service Contractor Management [com.snc.fsm_contractor_management]	Enables to outsource field service tasks. Contractor companies can re assign tasks to field agents and they can work on assigned tasks	Active	true	• com.snc.work_management • com.snc.outsourced_service_provider
Field Service Crew Operations [com.snc.fsm_crew_scheduling]	Helps organizations to assign the same set of resources repeatedly, to different tasks. This simplifies the scheduling process and allows field service agents to work together more consistently.	Active	true	com.snc.work_management

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Field Service Management [com.snc.work_management]	Provides support for scheduling and managing on-location worktasks.	Inactive	false	• com.snc.service_management.core • com.sn_shn • com.snc.assignment_workbench • com.snc.central_dispatch • com.snc.wo_signature_pad • com.snc.wm_questionnaire • com.snc.dynamic_scheduling
Field Service Management Access Hours Management [com.snc.fsm_access_hours]	Enables the management of access hours for work order tasks.	Active	false	com.snc.work_management
Field Service Management CMS Portal [com.snc.work_management.cms]	Lets you launch Field Service Automation and other service management applications from a single CMS page.	Inactive	false	• com.snc.work_management • com.snc.service_management.core.cms
Field Service Management Intelligent Task Recommendations [com.snc.fsm_task_recommendations]	Provides functionality to recommend tasks for assignment to field service technicians as well as manage criteria which can be used to make the task assignment recommendations.	Active	false	com.snc.work_management,com.snc.task_rec
Field Service Management Configurable Dispatcher Workspace [com.snc.uib.fsm_dispatcher_workspace]	Provides the FSM components and pages to support dispatcher flows on Agent Workspace.	Active	false	• com.snc.uib.fsm_agent_workspace • com.sn_fsm_components
Field Service Management Configurable Workspace [com.snc.uib.fsm_agent_workspace]	Provides the FSM components, lists, and forms to support FSM on Agent Workspace.	Active	false	• com.snc.uib.cwf_workspace • com.snc.work_management
Field Service Management - Customer Experience	Engage customers with relevant updates on reported issues	Active	false	• com.snc.csm_fsm_integration • com.snc.notify.twilio_direct

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.fsm_customer_experience]				
Field Service Management Demo Data [com.snc.work_management.demo]	Demonstration Data for Field Service Management covering the medical and telecommunication domains. NOTE: Installing this plugin adds new Configuration Item tables and relationships to the database.	Inactive	true	• com.snc.work_management • com.snc.wo_signature_pad
Field Service Management Geolocation Demo Data [com.snc.work_management_geolocation.demo]	Enables geolocation capabilities for the Field Service Management application.	Inactive	true	com.snc.work_management.demo
Field Service Management Mobile [com.snc.work_management_m]	Adds the Field Service Management menu to the mobile user interface.	Inactive	true	• com.snc.service_management_m • com.snc.work_management
Field Service Management Scheduling Flow Designer Flows [com.snc.sn_app_fsm_scheduling_flows] (New in Utah)	Updates an agent's schedule to maximize utilization if an unplanned event happens during their shift.	Inactive	false	• com.snc.work_management • com.snc.dynamic_scheduling • com.snc.agent_schedule
Field Service - Questionnaire [com.snc.wm_questionnaire]	Create Questionnaires for Work Orders or Work Orders Tasks.	Inactive	false	
Field Service Map [com.snc.fsm_map]	Displays the dispatch map for viewing agents, tasks, and agent routes.	Active		com.snc.work_management
Field Service Management Agent Workspace [com.snc.agent_workspace.fsm]	Enables field service users to manage work orders and tasks using the FSM agent workspace. This plugin activates the Agent Workspace			

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	(com.agent_workspace) plugin.			
Field Service Multi-Day Task Scheduling [com.snc.fsm_multiday_tasks]				
Field Service - Signature Pad [com.snc.wo_signature_pad]	Create PDFs for closed work orders that include the name and signature of the customer and a summary of the completed work.	Inactive	false	• com.snc.pdf_generator • com.snc.signaturepad
Field Service Resource Scheduling [com.snc.fsm_resource_scheduling] (New in Utah)	Enables setup of shared equipment like cranes, diggers to be set up as schedulable resources.	Active	true	com.snc.fsm_crew_scheduling
Field Service Spoke [com.snc.field_service_spoke]	Enables Flow Designer users to automate the creation of work orders and work order tasks and updates to existing work orders and work order tasks.	Inactive	false	Flow Designer - Installer (com.glide.hub)
Field Service Task Bundling [com.snc.fsm_task_bundle] (Available in the ServiceNow Store)	Enables the management of bundled work order tasks.	Active	false	com.snc.work_management
Field Service Territory Planning [com.snc.fsm_territory_planning] (New in Tokyo)	Enables definition and management of Territories for Field Service specific use cases. By using Field Service Territory Management with work orders and resources, customers can make sure that work order tasks are scheduled only to	Active	true	• com.snc.work_management • com.snc.territory_planning

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	field agents with a matching territory.			
Field Service with Project Management [com.snc.wm_ppm]	Enables Field Service Management integration with the Project Management application. This plugin activates the Project Management plugin, which may require additional licenses.	Inactive	true	• com.snc.work_management • com.snc.financial_planning_pmo
Field Service with Project Management Demo [com.snc.wm_ppm_demo]	Demonstration data for the Field Service Management integration with the Project Management application.	Inactive	false	• com.snc.wm_ppm • com.snc.work_management.demo • com.snc.customerservice.demo
Field Service with Service Locations support [com.snc.fsm_service_locations]	Enables to create a new service location from the WO/WOT whenever the preferred location does not exist among the stored locations.	Active	true	com.snc.work_management,com.sn_fsm_com
Finance Close Automation [sn_fcms]	Reduce close risk with intelligent task automation	Inactive	true	• sn_fin • sn_fcms_intg • com.glide.scope.access.restricted_caller
Finance Common Architecture [com.sn_fin] (Available in the ServiceNow Store)	Maintains primary data such as Enterprise Resource Planning (ERP) sources, legal entities, accounting periods, and so on.	Inactive	true	• com.snc.fiscal_calendar • com.snc.fin_dep • com.glide.scope.access.restricted_caller • com.glide.web_service_insert_multiple
Finance – ERP Integration [com.sn_fcms_integrations] (Starting with the Utah release, Finance – ERP Integration	Provides integration support between the ServiceNow platform and common ERP systems.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
is renamed as ERP Integration Framework.) (Available in the ServiceNow Store)				
Finance Service Management [com.snc.finance_service_automation] (Maintenance mode only) Planned for deprecation in March 2025 or subscription term end.	Manages finance requests and enables users to report and track those requests. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.service_management.core
Finance Service Management Mobile [com.snc.finance_service_automation_m]	Manages Finance Service Management mobile components.	Inactive	true	• com.snc.finance_service_automation • com.snc.service_management_m
Financial Charging [com.snc.service_charging]	Enables financial analysts and business unit managers to measure and show charges based on the actual consumption of services. Get insight into your services and the utilization of those services by each department. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	com.snc.itfm_core

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Modeling Engine [com.snc.financial_management]	Enables financial analysts to assemble spending data, build cost models, and generate reports to show how funds are being used. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	• com.snc.itfm_core • com.snc.financial_planning • com.snc.cost_management • com.snc.service_charging
Financial Management For APM [com.snc.financial_management_for_apm]	Enables integration Financial Management with Application Portfolio Management providing preconfigured Cost Models.Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Active	false	com.snc.financial_management
Financial Management for CSM [com.snc.financial_management_for_csm]	Plugins for Customer Service Management and Field Service Management provide integration with the Financial Management application. These plugins add cost allocations for the CSM and FMS applications as well as dashboards and reports. Financial administrators can use these cost allocations on the Financial Management workbench to allocate, track, and report on expenses.	Inactive	true	• com.snc.financial_management • com.sn_customerservice

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Financial Management For FSM [com.snc.financial_management_for_fsm]	Enables integration Financial Management with Field Service Management providing preconfigured Cost Models and a dashboard. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	• com.snc.financial_management • com.snc.work_management • com.snc.pa • com.snc.pa.fm.fsm
Financial Management for SPM [com.snc.financial_management_for_spm]	Enables integration of Financial Management with Service Portfolio Management providing preconfigured Cost Models. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Active	true	• com.snc.financial_management • com.spm_owner_workspace
Financial Planning [com.snc.financial_planning_default_definition]	Provides preconfigured Financial Plan Definition and all its artifacts which enables Financial Analysts to plan for costs, make forecasts, and evaluate actual expenses versus planned expenses. Planning take into consideration a wide range of items in your infrastructure, including assets, labor, and the configuration items in the CMDB. Activation	Inactive	true	com.snc.financial_planning

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.			
Financial Services Operations - Core [com.sn_bom]				
Financial Services Business Customers Lifecycle Operations [com.sn_bom_clo_b2b]				
Financial Services Business Deposit Operations [com.sn_bom_deposit_b2b] (New in San Diego)	Application specific to business deposit operations focused on origination, servicing, and closure of deposit accounts and is part of the Financial Services Operations product.	Inactive	true	• com.snc.csm_contributor_user • sn_csm_playbook • sn_bom • sn_bom_document
Financial Services Complaint Management [com.sn_bom_compl]				
Financial Services Know Your Customer [com.sn_bom_kyc]				
Financial Services Personal Customers Lifecycle Operations				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.sn_bom_clo_b2c]				
Financial Services Personal Deposit Operations [com.sn_bom_deposit_b2c] (New in San Diego)	Application specific to personal deposit operations focused on origination, servicing, and closure of deposit accounts and is part of the Financial Services Operations product.	Inactive	true	• com.snc.csm_contributor_user • sn_csm_playbook • sn_bom • sn_bom_document
Financial Services Remote Tables (New in San Diego)				
Financial Services Treasury Operations [com.sn_bom_treasury]				
Fiscal Calendar [com.snc.fiscal_calendar]	Enables to generate and manage different kinds of Fiscal calendar which is used in various Financial applications. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	false	com.glide.dynamic_operands.datetime
Fixed Asset [com.snc.fixed_asset]	Core fixed asset tracking	Active	false	
Fixed Costs [com.snc.fixed_costs]	Fixed Costs	Active	false	com.glide.metrics
Flow Client [com.glide.hub.flowclient]	The client side API for interacting with flows, subflows, and actions	Active	false	com.glide.hub.scriptable.plan.runners

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Flow Designer - Designer [com.glide.hub.designer]	Flow Designer - Designer	Active	false	com.glide.hub.designer_backend.suite
Flow Designer - Engine Reporting Dashboard [com.glide.hub.flow_reporting.dashboard]	Flow Engine Reporting PA Dashboard for Flow Designer.	Active	false	• com.glide.hub.flow_reporting • com.glide.dashboards
Flow Designer - Flow Engine [com.glide.hub.flow_engine]	Flow Designer runtime components.	Active	false	
Flow Designer - Flow Engine Reporting [com.glide.hub.flow_reporting]	Flow Engine Reporting components.	Active	false	• com.glide.hub.flow_engine • com.glideapp.dashboards
Flow Designer - Installer [com.glide.hub]	Suite of plugins necessary to support the Flow Designer Designer App.	Active	false	• Flow Designer - Flow Engine • Flow Designer - Flow Engine Reporting • com.glide.hub.flow_reporting.dashboard • Outbound Tracking • Flow Designer Action Step - Log • Flow Designer Action Step - Script • Flow Designer Action Step - CRUD • com.glide.hub.action_step.notification • Flow Designer Action Step - CORE • com.glide.hub.action_step.service_catalog • com.glide.hub.designer_backend.model • Flow Designer Action Step - Email • Flow Designer System Level Actions • Connect Spoke for Flow Designer • Flow Designer - Designer

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
				• com.glide.hub.action_step.email_header • com.glide.hub.action_step.look_up_email_at • App Dependency Client
Flow Designer Action Step - CORE [com.glide.hub.action_step.core]	Action Step - Core operations on GlideRecord.	Active	false	com.glide.hub.action_step.template
Flow Designer Action Step - CRUD [com.glide.hub.action_step.crud]	Action Step - CRUD operations on GlideRecord.	Active	false	com.glide.hub.action_step.template
Flow Designer Action Step - Email [com.glide.hub.action_step.email]	Action step for emails.	Active	false	• com.glide.email.service • com.glide.hub.action_step.template
Flow Designer Action Step - Log [com.glide.hub.action_step.log]	Action Step - Log.	Active	false	com.glide.hub.action_step.template
Flow Designer Action Step - Notification [com.glide.hub.action_step.notification]	Action steps for notifications.	Active	false	• com.glide.notification • com.glide.hub.action_step.template
Flow Designer Action Step - Payload Builder [com.glide.hub.action_step.payload]	Action Step - Payload Builder	Active	false	com.glide.hub.action_step.template
Flow Designer Action Step - Script [com.glide.hub.action_step.script]	Action Step - Script.	Active	false	com.glide.hub.action_step.template
Flow Designer Action Step - Service Catalog [com.glide.hub.action_step.service_catalog]	Action Step - Service Catalog.	Active	false	com.glide.hub.action_step.template

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Flow Designer Action Step Template [com.glide.hub.action_step.template]	Action Step Template Design Time components.	Active	false	
Flow Designer Action Trigger [com.glide.hub.flow_trigger]	Action plans that map a Trigger (Record, Email, REST) to an Actions.	Active	false	com.glide.hub.flow_engine
Flow Designer Designer Backend [com.glide.hub.designer_backend.suite]	Suite of plugins necessary to support the Flow Designer Designer App.	Active	false	• com.glide.hub.action_step.log • com.glide.hub.action_step.script • com.glide.hub.action_step.crud • com.glide.hub.action_step.email • com.glide.hub.action_step.notification • com.glide.hub.action_step.core • com.glide.hub.designer_backend.model • com.glide.hub.action_type.system
Flow Designer Designer Model [com.glide.hub.designer_backend.model]	Data model and REST API for representing Process Plans as Flows, Actions, and Steps.	Active	false	• com.glide.element_mapping • com.glide.hub.flow_engine • com.glide.hub.flow_trigger • com.glide.hub.flow_reporting • com.glide.hub.action_step.template
Flow Designer Scriptable Plan Runner API's [com.glide.hub.scriptable.plan.runners]	The API's needed to run flows, subflows, and actions through script.	Active	false	com.glide.hub.flow_engine
Flow Designer support for the Service Catalog [com.glideapp.servicecatalog.flow_designer]	Service Catalog Flow Designer	Active	true	com.glide.hub.flow_trigger.catalog
Flow Designer System Level Actions [com.glide.hub.action_type.system]	Action Type Definitions for low-level functions such as GlideRecord operations.	Active	false	com.glide.hub.designer_backend.model
Flow Template Builder				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.sn_flow_template] (New in San Diego)				
Form Designer [com.glide.ui.ng.fd]	Form Designer	Active	false	• com.glide.ui.ng.dc • com.glide.ui.tablet.theme • com.glide.ui.ng • com.glide.ui.themes.doctype
Form Personalization [com.glide.ui.personalize_form]	Enables users to personalize the layout for any form view.	Active	false	com.glide.ui.doctype
FSM Agent Workspace [com.snc.agent_workspace.fsm]	Enables users to manage work orders and tasks within the agent workspace.	Inactive	false	com.agent_workspace
FSM Shift Scheduling [com.snc.sn_fsm_shift_schdl] (Available in the ServiceNow Store)	Enables the management of agent shift Scheduling functionalities for Field Service Management Workforce Optimization.	Active	true	• com.snc.sn_wfo_cfg_ws • com.sn_shift_planning_host • com.snc.uib.agent_shift_planning • com.snc.skills_management • com.snc.skill_determination • com.snc.sn_wfo_skill_review • com.glide.dynamic_translation
FSM Team Performance Management [com.snc.sn_fsm_team_mgmt] (Available in the ServiceNow Store)	Manages the team performance functionalities for Field Service Management Workforce Optimization.	Active	true	• com.snc.sn_fsm_shift_schdl • com.sn_team_perf • com.snc.pa • com.snc.pa.insights • com.snc.pa.spotlight • com.snc.work_management_pa
FSO Process Optimization Content Pack [sn_bom_po] (Available in the	Provides Process Optimization capabilities to Financial Services Operations applications.	Inactive	false	• com.sn_process_optimization • sn_bom

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Fullcalendar Library [com.snc.fullcalendar]	Fullcalendar library v2.5.0.	Inactive	false	
Geolocation [com.snc.geolocation]	Uses Google Maps to track users, plan efficient routes between locations, and assist in finding accurate travel times. The system locates users from latitude and longitude information provided by their mobile devices.	Inactive	false	
Github Spoke [com.sn.github.spoke]	The GitHub Spoke for IntegrationHub provides actions that a Process Analyst can use when designing flows. The actions allow them to automate GitHub tasks.			
Gitlab Spoke [com.sn.gitlab.spoke]	The GitLab Spoke for IntegrationHub provides actions that a Process Analyst can use when designing flows. The actions allow them to automate Gitlab.			
Glide Conversation Server [com.glide.cs]	Glide Conversation Server.	Inactive	true	• com.glide.connect.support • com.glide.cs.branding • com.glide.cs.live_agent_settings • com.glide.cs.sn-va-web-client-app
Glide Conversation Server Adapters [com.glide.cs.adapter]	Glide Conversation Server Adapters	Active	false	• com.glide.cs • com.glide.external.app • com.glide.channelproxy • com.snc.bot_install_ui

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Glide Metadata [com.glide.metadata]	Core metadata support.	Active	false	
Glide Metadata Delete [com.glide.metadata_delete]	Core metadata delete support.	Active	false	• com.snc.metadata • com.snc.metadata_tree
Glide Logging Framework Extensions [com.glide.log]	Contains extensions to glide logging framework and script API.	Active	false	None
Glide Notification Translation [com.glide.notification.translation] (New in Tokyo)	Adds the ability to send multi-lingual emails based on recipient preferences.	Active	false	com.glide.localization_framework.installer
Glide Virtual Agent [com.glide.cs.chatbot]	Activates the Virtual Agent framework and the other necessary plugins.	Active	false	• com.glide.cs.chatbot.lite • com.glide.cs.custom.adapter • com.glide.cs.nlu.topics • com.glide.nlu.intent.discovery • com.glide.nlu.ibmwatson.intent.discovery • com.glide.nlu.msluis.intent.discovery
Glide Virtual Agent Lite [com.glide.cs.chatbot.lite]	Activates the lite version of the Virtual Agent bot platform and other necessary plugins.	Active	false	• com.glide.cs • com.glide.va.sp_widgets • com.glide.cs.adapter • com.glide.cs.settings • com.glide.cs.topics • com.glide.service-portal.agent-chat • com.glide.cs.topic_blocks • com.snc.conversation_builder
Global Application File Management	UI components and utilities for managing application files in the global scope.	Active	false	• com.glide.scripted_rest_services • com.glide.ui.ng

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.global_app_files]				
Goal Framework [sn_gf] (Available in the ServiceNow Store)	Enables strategic planning and helps define and construct strategic priorities and associate goals for the organization. This application facilitates personas such as strategy officers, SRO, ePMO, and portfolio managers to track strategy execution by associating work and planning items such as demand, project, and portfolio to the defined goals.	Active	false	None
Goal Framework for SPM [sn_gfa] (Available in the ServiceNow Store)	Provides the capability to monitor the goal progress in real time by automating the actual value of the targets from any source present on the ServiceNow AI Platform.	Active	true	sn_gf
Google Calendar Spoke [com.sn.gcalendar.spoke]	The Google Calendar Spoke for IntegrationHub provides actions that a Process Analyst can use when designing flows. The actions allow them to automate Google Calendar.			
Google Maps plugin [com.glideapp.google_maps]	Allows the display of Google Maps within the product as map pages.	Active	false	• com.glide.ui.ng • com.glide.rest.service
Granular Delegation [com.glide.granular_service_delegation]	Delegate specific assignments and approvals to another user for a period of time.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
GraphQL Schema for Contextual Search [com.snc.contextual_search.graphql]	Provides GraphQL schema for contextual search services	true	false	com.glide.graphql
GraphQL Schema for Knowledge in Contextual Search Application Portfolio Management [com.snc.knowledge3.graphql]	Provides GraphQL schema for contextual search sources like Knowledge	true	false	com.snc.contextual_search.graphql
GRC: Audit Management [com.sn_audit]	The GRC: Audit Management application provides a centralized process for Internal Audit teams to automate the complete audit life cycle. Project driven audits allow auditors to quickly scope engagements, conduct fieldwork, collect control evidence, and track audit observations.	Inactive	true	• com.sn_grc • com.glide.ui.ng • com.snc.common_workbench
GRC: Audit Management Workspace [sn_audit_workspace]				
GRC: Business Continuity Management – Components [sn_bcm_components]	Provides seismic components for Business Continuity Management application			
GRC: Business Continuity Management – Core [sn_bcm]	Provides capabilities to configure the core data and workspace that are required for the business continuity management applications.			

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
GRC: Business Continuity Planning [sn_bcp]	Provides capabilities for business continuity management teams to create, review, and approve business continuity and disaster recovery plans.			
GRC: Business Impact Analysis [sn_bia]	Provides structured workflows for business continuity management teams to assess the impact of a disruption to critical resources such as business processes, applications, and locations.			
GRC: Crisis Management [sn_recovery]	Provides capabilities for business continuity management teams to activate business continuity plans during plan exercises and actual crisis events.			
GRC: Compliance UCF [com.sn_comp_ucf]	The GRC: Compliance UCF plugin is an add-on that provides GRC: Policy and Compliance Management with the capability to download regulatory content and common controls from the Unified Compliance Framework (UCF).	Inactive	false	com.sn_compliance
GRC:Mobile [com.sn_grc_mobile]	There are two GRC core applications that are viewed on the GRC Mobile app: Policy and Compliance Management and Risk Management. There are two GRC core applications that are viewed on the GRC Mobile app: Policy	Inactive	false	com.glide.sg , com.glide.sg.agent_native_client

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	and Compliance Management and Risk Management.			
GRC: Performance Analytics Integration [com.sn_grc_pa]	GRC: Performance Analytics Integration enables risk and compliance managers to continuously monitor key risk and control indicators using the advanced reporting capabilities provided by Performance Analytics.	Inactive	false	com.sn_grc
GRC: Policy and Compliance Management [com.sn_compliance]	The GRC: Policy and Compliance Management application provides a centralized process for creating and managing policies, standards, and internal control procedures that are cross-mapped to external regulations and best practices. The application provides structured workflows for the identification, assessment, and continuous monitoring of control activities.	Inactive	true	• com.sn_grc • com.snc.comp_attest
GRC: Predictive Intelligence [com.sn_grc_pred_intel]	Enables customers to utilize the machine learning algorithms for a wide variety of GRC use cases such as Issue Assignment. This plugin may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.glide.platform_ml • com.sn_grc
GRC: GRC Profile Dependencies [com.snc.grc_profile_dep]	Installs all the dependent plugins required to support the GRC Profile application.	Active	false	• com.snc.grc.common • com.snc.planned_task_v2

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
GRC: Regulatory Change Management [com.sn_grc_reg_change]	The Regulatory Change Management application enables customers to check upcoming regulatory changes, assess their impact, and implement risk and compliance-related changes, ensuring overall regulatory compliance.	Active	true	• com.snc.grc_policy_dep • com.snc.grc_profile_dep • com.snc.grc_risk_dep • com.glide.hub.action_step.sftp • com.glide.hub.action_type.datastream • com.glide.hub.integration.runtime
GRC integration with Thomson Reuters Regulatory Intelligence [com.sn_grc_int_tr]	Use the GRC integration with Thomson Reuters Regulatory Intelligence application to integrate with the Thomson Reuters platform and to consume the Thomson Reuters Regulatory Intelligence (TRRI) feeds into your ServiceNow instance.	Active	true	• com.snc.grc_profile_dep • com.snc.grc_policy_dep • com.snc.grc_risk_dep
GRC: Advanced Risk [com.sn_risk_advanced]	As a decision-maker, at any level of the business, you can view up-to-date information regarding your organization's risk posture. Link intermediate risk statements to risks, as well as enterprise-level risk statements. Your risk management team can provide a rollup of risk statements and scores from the most granular risk to the highest enterprise-level risk statement. Additionally, as a risk manager, you can monitor risk events, relate them to existing risks, perform root-cause analysis, and	Active	true	• com.snc.grc_profile_dep • com.snc.grc_risk_dep

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	track all remediation tasks.			
GRC: DevOps Accelerator [com.sn_grc_devops] (Available in the ServiceNow Store)	Maps the control objectives drawn from regulations, standards, and frameworks, such as CIS controls, NIST 800-53, ISO 27002, PCI DSS etc., to DevOps Policy as a Code Engine (PaCE). These policies are provided by the DevOps Config Policy Content Pack. This allows compliance and DevOps managers to preventatively monitor control compliance, get real-time visibility into evidence of PaCE policy execution, and allow for exception management.	Active	false	None
GRC: Risk Management [com.sn_risk]	The GRC: Risk Management application provides a centralized process to identify, assess, respond to, and continuously monitor Enterprise and IT risks that may negatively impact business operations. The application also provides structured workflows for the management of risk assessments, risk indicators, and risk issues.	Inactive	true	• com.sn_grc • com.snc.risk_asmt
GRC: Risk Management Dependencies [com.snc.grc_vrm_dep]	Installs all the dependent plugins required to support the GRC Risk Management application.	Active	false	• com.snc.grc_profile_dep • com.snc.risk_asmt

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
GRC: SIG Questionnaire Integration [com.sn_sig_asmt]	The GRC: SIG Questionnaire Integration plugin installs the Shared Assessments Standardized Information Gathering (SIG) questionnaire for use with the GRC: Vendor Risk Management application.	Inactive	true	• com.snc.assessment_core • com.sn_vdr_risk_asmt • com.snc.sig_asmt_core • com.snc.document_management
GRC: Vendor Risk Management [com.sn_vdr_risk_asmt]	The GRC: Vendor Risk Management application allows risk management stakeholders to identify, assess, and manage risk across the vendor ecosystem.	Inactive	true	• com.sn_comp_asmt • com.snc.vendor_core • com.sn_grc • com.snc.vdr_risk_asmt_designer • com.snc.vendor_portal
GRC: Vendor Risk Management Dependencies [com.snc.grc_vrm_dep]	Installs all the dependent plugins required to support the GRC Vendor Risk Management application.	Active	false	• com.snc.grc_profile_dep • com.snc.vdr_risk_asmt_designer • com.snc.vendor_core • com.snc.sig_asmt_core • com.snc.assessment_core • com.snc.document_management • com.snc.vendor_portal
GRC:Virtual Agent [com.sn_grc_virtual_agent]	The GRC:Virtual Agent provides chat bot functionality for common GRC operations.			
GRC: Workbench [com.sn_grc_workbench]	The GRC Workbench utilizes CMDB information to show the upstream and downstream relationships across all applications. These relationships enable consistent risk mapping and modeling across the enterprise. The GRC Workbench does	Inactive	false	com.sn_grc

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	not work with Legacy GRC.			
Group Life Servicing [com.sn_ins_group_life] (Available in the ServiceNow Store)	Group life insurance management.	Active	true	com.snc.csm_contributor_user
Guided Application Creator [com.glide.sn-guided-app-creator]	Enables users to set up custom business applications that are immediately ready to use.	Active	false	com.sn_appcreator
Guided Decisions (Maintenance mode only) [com.snc.guided_decisions]	Enables the framework and execution engine for creating guided decisions.	Active	false	• com.glide.decision_table • com.snc.guided_decisions_guidance • com.glide.hub
Guided Decisions - Guidance (Maintenance mode only) [com.snc.guided_decisions_guidance]	Enables the framework for creating a guidance for guided decisions.	Active	false	• com.glide.hub • com.snc.guided_decisions_guidance_globa
Guided Decisions - Guidance Global Entities (Maintenance mode only) [com.snc.guided_decisions_guidance_global]	This global plugin is loaded as dependant of Guided Decisions Guidance to access global entities.	Active	false	
News Integration for Supplier Lifecycle Operations [com.snc.sn_supplier_news] (Available in the	Retrieve recent news and other articles from preferred news channels. There is a pre-configured integration with Microsoft Bing available with this application.	Inactive	false	None

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Guided Decision - Next Best Action [com.snc.next_best_action]	Recommends relevant actions like Guided Decisions for a given context. The plugin enables defining rules to recommend Next Best Actions.	Active	false	com.sn_next_best_action_list
Guided Decisions for Playbook in Workspace [com.snc.guided_decisions_playbook_experience]	Enables activity types, definitions, and UI components for the display of guided decisions in a playbook on Workspace.	Active	false	• com.snc.guided_decisions • com.snc.guided_decisions_guidance • com.glide.playbook_experience.config
Guided Decisions Experience [sn_ga_exp] (Available in the ServiceNow Store)	Enables customers to use Guided Decisions with Playbooks, Recommended Actions, and other features.	Inactive	false	com.snc.guided_decision
Guided Setup Configuration [com.glide.guided_setup_config]	Provides customers with a configuration on the Guided Setup Content.	Active	false	
Guided Setup for CMDB [com.snc.cmdb.guided_setup]	Provides customers with a visual guidance to configure CMDB. Requires Core UI.	Active	false	• com.snc.guided_setup • com.snc.cmdb
Guided Setup for Performance Analytics [com.snc.pa.guided_setup]	Provides customers with visual guidance to configure Performance Analytics. Requires Core UI.	Active	false	• com.snc.guided_setup • com.snc.pa • com.glideapp.dashboards
Guided Setup for ServiceNow Flow Designer [com.glide.hub.flow_designer_guided_tours]	A guided tour instructing users on how to start creating a flow.	Active	false	
Guided Setup Framework	Provides customers with visual guidance to configure	Active	false	• com.glideapp.ui_components • com.glide.scripted_rest_services

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.guided_setup]	ServiceNow applications.			• com.glide.embedded_help • com.glide.scoped_analytics_framework
Guided Tour Designer [com.glide.sn_tourbuilder]	ServiceNow Guided Tour Designer, aka GTD is a simplified way to create a guided tour. It provides easy to use, drag and drop interface to quickly build guided tours without writing any code. You can add, delete, modify, reorder steps as well as test the tour by previewing it.	Active	false	
Guided Tours [com.glide.guided_tours]	A guided tour provides a walk-through of a procedure in a UI page. Administrators can develop customized guided tours to assist users with specific business processes within an organization.	Active	false	• com.glide.ui.hopscotch • com.glide.embedded_help
Hardware Asset Management (Maintenance mode only) [com.sn_hamp]	Hardware Asset Management Professional focuses on driving automation and prescriptive, standard practices for managing physical hardware assets throughout the end-to-end asset lifecycle from acquisition through disposition.	Active	false	com.sn_hwnorm
Hardware Model Normalization [com.sn_hwnorm]	Hardware normalization is a net new app planned for limited access in the Orlando release. This feature will enable customers with ITAM hardware normalization plugin to received standardized content	Inactive	true	• com.snc.model • com.glide.data_services_canonicalization.cli • com.sn_glidequery

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	to normalize records for IT hardware manufacturers, models and model lifecycles.			
Health and Safety Incident Management [sn_ohs_im] (Available in the ServiceNow Store)	Enables organizations to record, investigate, and report on health and safety related employee incidents and observations.	Active	true	• Plugins dependency: ◦ com.snc.app_shell_aw ◦ com.snc.uib.base_agent_workspace ◦ com.glide.playbook_experience.config ◦ com.glide.pad.license • Store app dependency: ◦ sn_wsd_core ◦ sn_ohs_comp ◦ sn_ohs_im_pa ◦ sn_imt_core ◦ sn_ent
Health and Safety Incident Management OSHA Content Pack [com.snc.sn_hs_im_osha] (Available in the ServiceNow Store)	Adds Occupational Safety and Health Administration (OSHA) report autofill and export capability to the Health and Safety Incident Management application.	Active	true	
Health Log Analytics Core [com.glideapp.itom.occultus]	Essential components for Health Log Analytics scoped application.		false	• com.glide.indexed-docstore.es • com.snc.clotho • com.itom.jutils • com.glide.services • com.glide.auth.mutual • com.glideapp.itom.snac
Health Log Analytics Viewer [com.sn-log-viewer]				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
High Security Settings [com.glide.high_security]	Enables a high-security environment, including high-security system properties and access controls.	Active	false	com.glide.role_management
History Sets [com.glide.history_set]	Maintains sys_history_set and sys_history_line tables to view a record's audit, email, and relationship data in a table format.	Active	false	
History Walker [com.glide.history_walker]	Provides script access to History Walker	Active	false	
Homepage Deprecation Help Tool [sn_home_page_depr]				
Homepage Splash Page [com.glideapp.home_splash_page]	Provides a splash page for homepages. Upon logging in, instead of going to home.do, users go to a splash page that gives them the ability to cancel the home page transaction. This is useful when homepages take a long time to load.	Active	false	com.glideapp.home

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Important: The functionality found in homepages, arranging information from your instance to tell a story about your data, is found in dashboards on new instances. On upgraded instances with Next Experience enabled, users can view existing homepages if they have a direct URL, but they can't create or edit them. Responsive dashboards and Analytics Center dashboards take over homepage functionality. Use the Homepage deprecation help tool to convert the homepages on your instance to responsive dashboards. For more information, see: • Dashboards in the Analytics Center. • Working with responsive dashboards.			
Hopscotch [com.glide.ui.hopscotch]	Hopscotch is a framework to make it easy for developers to add product	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	tours to their pages. Hopscotch accepts a tour JSON object as input and provides an API for the developer to control rendering the tour display and managing the tour progress. To learn more about Hopscotch and the API, check out http://linkedin.github.io/hopscotch.			
HR Service Delivery integration with Microsoft Teams [sn_now_teams_hr] (Available in the ServiceNow Store)	The integration provides a broad set of service-delivery capabilities within Microsoft Teams, to help resolve employees' HR requests. Employees will be able to get answers to questions, stay up to date on the latest news and events, report issues, and discuss open HR Cases with fulfillers.	Inactive	false	• sn_hr_va • sn_now_teams
HR Service Delivery Integration with SuccessFactors [sn_hr_sf] (Available in the ServiceNow Store)	Synchronizes employee to-dos so that employees can view and complete their to-dos from either system.	Active	false	• sn_successfactors • sn_hr_integr_fw • com.sn_hr_core
HRSD Process Optimization Content Pack [com.sn_hr_process_optimization] (Available in the ServiceNow Store)	Provides pre-configured process optimization models and improvement initiatives for HRSD processes.	Inactive	false	• com.sn_process_optimization • com.sn_hr_core

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
HTML Sanitizer [com.glide.htmlsanitizer]	Automatically cleans up HTML markup in all HTML fields to remove unwanted code to protect against security concerns such as cross-site scripting attacks.	Active	false	
Human Resources Application: Core CMS [com.snc.hr.core.cms] (Maintenance mode only)	Provides case and knowledge management for HR. Standardizes the documentation, interaction, and fulfillment of employee inquires and requests while having visibility into the quantity and type of cases coming in. **Warning:** Be sure that you do not have plugin Human Resources Application: Core activated. This plugin is a new HR application that duplicates some functionality in Human Resources Application: Core. If you have any questions about this plugin, contact your ServiceNow account manager. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	• com.glide.country • com.glide.db_images • com.snc.appview • com.snc.document_management • com.snc.knowledge3 • com.snc.pdf_generator • com.snc.sc_catalog_manager • com.snc.signaturepad • com.snc.skills_management • com.snc.task_activity • com.snc.matching_rule • com.snc.hr.scoped_security
Human Resources	HR Integration provides the ability to integrate with	Inactive	true	com.glide.scope.access.restricted_caller

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Scoped App: Integrations [com.sn_hr_integrations]	multiple third-party HR systems for the scoped version of HRSM. HR Integration pushes HR profile information from HRSM to a third-party HR application and pulls HR profile information from a third-party HR application to HRSM.			
Human Resources Scoped App: Lifecycle Events [com.sn_hr_lifecycle_events]	Activates the **Lifecycle Event** module for HR. .Combined with other licensed applications, provides a full-service, employee journey experience such as onboarding or parental leave of absence for employees and those managing the process	Inactive	true	• com.sn.hr.core • com.sn_hr_service_portal
Human Resources Scoped App: Lifecycle Events for Enterprise [com.sn_hr_lifecycle_enthelene]	Human Resources Scoped App: Lifecycle Events for Enterprise enables you to automate onboarding and other employee lifecycle events that span across multiple departments such as IT, Facilities, Finance, and Legal.	Inactive	true	• com.sn_hr_lifecycle_events • com.sn_hr_core • com.sn_hr_service_portal • com.sn_hr_onboarding • com.snc.businessroles
Human Resources Scoped App: Mobile [com.sn_hr_mobile]	Activates the Mobile module for HR. Combined with Mobile Employee Experience Native Application plugin, items related to HR will be shown in the native mobile application.	Active	true	• com.glide.mobile-employee • com.sn_hr_core
Human Resources Scoped App: Mobile Employee Onboarding	Enable your new hires to complete onboarding to-dos, ask questions, and view relevant content using the new	Inactive	true	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.sn_hr_onboarding]	Mobile Onboarding application.			
Human Resources Scoped App: Parental Journey [com.sn_hr_parental_journey]	Activates the Employee Journey for Parental Leave of Absence. Combined with other licensed applications, provides a full-service, parental leave of absence experience for employees and those managing the process. **Important:** To use Parental Leave of Absence, the Human Resources Scoped App: Core and Human Resources Scoped App: Lifecycle Events plugins must be licensed and activated	Active	true	com.sn_hr_lifecycle_events
Human Resources Scoped App: Security [com.sn.hr.scoped_security]	Activation of this plugin on production instances requires a separate license. Contact ServiceNow for details. Human Resources Scoped App: Core plugin requires this plugin.	Inactive	false	com.glide.service-portal
Human Resources Scoped App: Service Portal [com.sn.hr.service_portal]	Provides a single place for employees to quickly and easily get all the HR services they need.	Inactive	true	• com.sn_hr_core • com.glide.connect.support
Human Resources Scoped App: Virtual Agent Conversations	Activates a conversational bot platform for providing user assistance through conversations	Active	false	• com.glide.cs • com.sn_hr_core • com.glide.cs.chatbot

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.sn_hr_virtual_agent]	within a messaging interface for HR. **Important:** To use Virtual Agent for HR, the Human Resources Scoped App: Core and the Glide Virtual Agent plugins must be licensed and activated.			• com.glide.cs.sn-va-web-client-app • com.sn_hr_service_portal • com.glide.va.sp_widgets
Human Resources Scoped App: Workspace [sn_hr_agent_workspace]	Activates a configurable service desk application that provides agents with an integrated and graphically intuitive user experience for HR. **Important:** To use Agent Workspace for HR, the Human Resources Scoped App: Core plugin must be licensed and activated.		false	• com.agent_workspace • com.snc.sn_lookup_and_verify_config • com.sn_hr_core
IBM License Compliance for Software Asset Management [com.sn_samp_ibm_lic] (Available in the ServiceNow Store)	Enables integration between the Software Asset Management publisher pack for IBM and the Anglepoint Elevate platform. With this integration, you can track and manage IBM licensing directly without having to integrate with the IBM License Metric Tool (ILMT) or BigFix Inventory.	Inactive	false	com.sn_samp_master_ws
Individual Life Servicing [com.sn_ins_indiv_life]	Individual life insurance policy management.	Active	true	• com.snc.csm_contributor_user • sn_csm_playbook (2.0.2)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				• sn_bom (1.8.0) • sn_ins_indiv_uw (1.0.1)
Indoor Mapping [sn_map_core] (Available in the ServiceNow Store)	Provides the capabilities to build, visualize, and integrate digital maps into various ServiceNow products and use cases.	Active	true	sn_wsd_indoor_map
Industrial Process Manager [com.sn_otsm] (Available in the ServiceNow Store)	Enables teams to map and visualize the industrial equipment models and associated production process at individual facilities.	Active	False	• com.sn_mfg_common • com.sn_cmdb_ci_class • com.sn_isa_model
Industrial Workspace Common [com.sn_mfg_common] (Available in the ServiceNow Store)	Provides common components used in Industrial Process and Operational Technology products.	Active	True	com.snc.cmdb.workspace
ISA Equipment Model [com.sn_isa_model] (Available in the ServiceNow Store)	Creates the ISA-95 Equipment Model data foundation that is required for the ServiceNow Operational Technology solution.	Active	True	com.snc.cmdb.integration_util
Issue Auto Resolution for HR [com.snc.hr.issue_auto_resolution] (New in Tokyo)	Enables organizations to provide near real-time responses to routine and non-critical employee inquiries using self-service content already defined by the organisation.	Active	false	com.sn_hr_core

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Employees can ask for help in their own natural language through the channel that is convenient to them and in the time frame that feels effective for them.			
ITAM Health Check application [com.sn_itam_health_chk] (Available in the ServiceNow Store)	Performs Software Asset Management configuration scans to determine issues with configuration, that could impact software License compliance results. Additionally, you can get recommendations to fix the issues and thereafter create tasks for successful resolution.	Active	false	com.snc.samp
IntegrationHub ETL [com.sn_int_studio] (Available in the ServiceNow Store)	User interface for configuring templates to ingest CMDB data from third party sources.	Active	true	• com.glide.integration_studio • com.snc.cmdb.integration_util
ITSM Success Dashboard [com.snc.success_dashboard_itsm] (Available in the ServiceNow Store)	Provides insights to the IT leadership teams and process owners to measure the performance and improvement of their ITSM implementation using the KPIs defined by ServiceNow.	Active	false	• com.snc.pa.premium.all_content • com.snc.self_service_analytics_core
Journey Designer [com.sn_jny] (New in Tokyo)	Enables managers and employees to engage in rich, collaborative journeys from onboarding to offboarding and everything in between.	Inactive	true	• com.sn_hr_core • com.sn_hr_lifecycle_ent • sn_ex_sp_pro • sn_ect sn_cd
I18N: Brazilian Portuguese Translations	An Internationalization plugin for language internationalization.	Inactive	false	com.glide.i18n

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.i18n.brazilian_portuguese]	Provides Brazilian Portuguese.			
I18N: Czech Translations [com.snc.i18n.czech]	An Internationalization plugin for language internationalization. Provides Czech.	Inactive	false	com.glide.i18n
I18N: Dutch Translations [com.snc.i18n.dutch]	An Internationalization plugin for language internationalization. Provides Dutch.	Inactive	false	com.glide.i18n
I18N: Finnish Translations [com.snc.i18n.finnish]	An Internationalization plugin for language internationalization. Provides Finnish.	Inactive	false	com.glide.i18n
I18N: French - Canada Translations [com.snc.i18n.french-canada]	An Internationalization plugin for language internationalization. Provides French - Canadian.	Inactive	false	com.glide.i18n
I18N: French Translations [com.snc.i18n.french]	An Internationalization plugin for language internationalization. Provides French.	Inactive	false	com.glide.i18n
I18N: German Translations [com.snc.i18n.german]	An Internationalization plugin for language internationalization. Provides German.	Inactive	false	com.glide.i18n
I18N: Hebrew Translations [com.snc.i18n.hebrew]	An Internationalization plugin for language internationalization. Provides Hebrew.	Inactive	false	com.glide.i18n
I18N: Hungarian Translations [com.snc.i18n.hungarian]	An Internationalization plugin for language internationalization. Provides Hungarian.	Inactive	false	com.glide.i18n
I18N: Internationalization [com.glide.i18n]	An Internationalization plugin for language internationalization. Provides the elements necessary for translating an instance without any translation preloaded.	Inactive	false	• com.glide.system_import_set • com.glideapp.knowledge.i18n2
I18N: Internationalization	Modules and import set maps to merge all	Inactive	false	com.glide.i18n

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Translation helper [com.glide.i18n.translation_helper]	languages into one table for translation maintenance.			
I18N: Italian Translations [com.snc.i18n.italian]	An Internationalization plugin for language internationalization. Provides Italian.	Inactive	false	com.glide.i18n
I18N: Japanese Translations [com.snc.i18n.japanese]	An Internationalization plugin for language internationalization. Provides Japanese.	Inactive	false	com.glide.i18n
I18N: Knowledge Management Internationalization Plugin v2 [com.glideapp.knowledge.i18n2]	Activating internationalization plugins for any of the available languages automatically activates this plugin. This plugin helps create and maintain translations of an article in various languages in a way that is easy to manage translations while authoring as well as viewing articles.	Inactive	false	• com.glideapp.knowledge3 • com.glide.i18n
I18N: Korean Translations [com.snc.i18n.korean]	Korean Translations	Inactive	false	com.glide.i18n
I18N: Polish Translations [com.snc.i18n.polish]	An Internationalization plugin for language internationalization. Provides Polish.	Inactive	false	com.glide.i18n
I18N: Portuguese Translations [com.snc.i18n.portuguese]	An Internationalization plugin for language internationalization. Provides Portuguese.	Inactive	false	com.glide.i18n
I18N: Russian Translations [com.snc.i18n.russian]	An Internationalization plugin for language internationalization. Provides Russian.	Inactive	false	com.glide.i18n
I18N: Simplified Chinese Translations	An Internationalization plugin for language internationalization.	Inactive	false	com.glide.i18n

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.i18n.chinese]	Provides Simplified Chinese.			
I18N: Spanish Translations [com.snc.i18n.spanish]	An Internationalization plugin for language internationalization. Provides Spanish.	Inactive	false	com.glide.i18n
I18N: Thai Translations [com.snc.i18n.thai]	An Internationalization plugin for language internationalization. Provides Thai.	Inactive	false	com.glide.i18n
I18N: Traditional Chinese Translations [com.snc.i18n.traditional_chinese]	An Internationalization plugin for language internationalization. Provides Traditional Chinese.	Inactive	false	com.glide.i18n
I18N: Turkish Translations [com.snc.i18n.turkish]	An Internationalization plugin for language internationalization. Provides Turkish.	Inactive	false	com.glide.i18n
IBM License Compliance for Software Asset Management [com.sn_samp_ibm_lic]				
IBM Watson Assistant Integration [com.glide.cs.ibm.watson.assistant.topic]	Enables the IBM Watson Assistant topic to run an IBM skill (conversation) created in IBM Watson with the Assistant V1 API. The topic runs in the Virtual Agent web client.	Inactive	false	
IBM Watson Assistant Integration V2 [com.glide.cs.ibm.watson.assistant.v2.topic]	Enables the IBM Watson Assistant Chat Integration to run a dialog skill (conversation) created in IBM Watson with the Assistant V2 API. The topics runs in the Virtual Agent web client.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
IBM Watson Translator Service Spoke [com.glide.ibm_translation_spoke]	Using IBM translation services, it adds ability to translate text from one language to another and detect the language of text.	Inactive	false	com.glide.hub.integrations
Ideation with PPM [com.snc.ppm_innovation_management]	Enables you to manage your demands, resources, portfolios and projects, and gives full visibility from idea to execution. Agile management and test management help you to improve productivity and service delivery. Enables you to manage ideas with full visibility from idea submission to execution. Collaborate with the idea submitter and other stakeholders to Identify ideas with the greatest potential and convert them into different work entities for execution.	Inactive	true	• com.snc.project_portfolio_suite • com.snc.financial_planning • com.snc.rate_model • com.snc.comments_and_feedback
Incident [com.snc.incident]	Provides base functionality for Incident Management.	Active	true	• com.snc.service • com.snc.incident_notification • com.glideapp.report.itsm.incident.overview
Incident Management - ATF Tests [com.snc.incident.atf]	Delivers ATF tests for Incident Management.	Active	true	
Innovation Management Utility [com.snc.innovation_management_util]	Provides capability to create a work entity like demand/project from an idea.	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Incident Assignment Workbench Demo [com.snc.incident_assignment_workbench_demo]	Intelligent agent recommendation through dynamic matching rules/ criteria.	Inactive	true	• com.snc.contact_management • com.glideapp.live_feed_document • com.glide.phone_number
Incident Communications Management [com.snc.iam]	Allows crisis managers to manage communications for major issues, to bring together all involved users via SMS, conference calls, email subscriptions, and live feed.	Inactive	false	com.snc.assignment_workbench
Incident Management - ATF Tests [com.snc.incident.atf]	Provides Incident Management ATF Tests.	Active	true	
Incident Management - Major Incident Management [com.snc.incident.mim]	Provides a recommended incident response process for managing incidents a business critical incident with pre-defined incident communication plans and a workbench to handle different aspects of the process including communication, collaboration, and postmortem.	Inactive	true	• com.snc.iam • com.snc.incident.updates • com.snc.task_outage
Incident Management Notification [com.snc.incident_notification]	Provides Notification functionality to Incident Management.	Active	false	
Incident Management Best Practice - San Diego [com.snc.best_practice.incident.sandiego]				
Incident Overview Homepage	Provides Incident Overview Homepage.	Active	false	com.glideapp.report

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glideapp.report.itsm.incident.overview]				
Incident Resolution Fields [com.snc.incident_resolution_fields]	Adds **Resolved** and **Resolved by** fields to the Incident table, similar to **Closed** and **Closed by**, populated with a business rule when an incident is resolved or closed.	Active	false	
Incident Updates [com.snc.incident_updates]	Incident Updates Plugin - Includes any updates to the Incident Base Plugin.	Active	true	
Index Suggestion [com.glide.index_suggestion]	Suggest database indexes for slow queries.	Active	false	Stats Tools, Column Statistics
Innovation Management [com.snc.innovation_management]	Idea Management application enables idea managers to manage ideas at organization level.	Inactive	false	com.snc.comments_and_feedback
Innovation Management - ATF Tests [com.snc.innovation_management_atf]	Innovation Management - ATF Tests provides you test cases and test suits that can be run on the Innovation Management Application.	Inactive	false	
Insert Multiple Web Service [com.glide.web_service_insert_multiple]	Enables multiple inserts for the Direct SOAP API.	Inactive	false	
Instance Security Center [com.glide.instance_sec_center]	This plugin activates the Instance Security Center Portal. Search for Instance Security Center in the left navigation menu.	Active	false	• com.snc.pa.instance_sec_dash • com.snc.pa.sp.widget • com.glide.service-portal.esm
Instance Data Replication [com.glide.idr]	Replicates records from one instance to one or more other instances.	Inactive	false	
Instance Sync Connector				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[sn_instance_sync_connector]				
Instance Sync Connector - Core [com.snc.instance_sync_connector_core]	ITSM Instance Sync Connector Core plugin that provides an ability to activate and leverage Integration Hub Process Sync Framework.	Active	false	com.glide.hub.process.sync
Integration - Common Components [com.snc.integration.common]	Provides common scripts for integrations.	Inactive	false	• com.glide.system_import_set • com.glide.web_service_provider • com.glide.system_property_categories
Integration - JDBC [com.snc.integration.jdbc]	Provides integration through JDBC.	Inactive	false	• com.glide.web_service_application • com.glideapp.agent • com.snc.integration.common
Integration - Multifactor Authentication [com.snc.integration.multifactor.authentication]	This plugin to set up Multifactor authentication on the instance.	Active	false	
Integration - Multiple Provider Single Sign-On Enhanced UI [com.snc.integration.sso.multi.ui]		Inactive	false	• com.glide.ui.ng.guided_flow • com.snc.integration.sso.multi.installer
Integration - Multiple Provider Single Sign-On Installer [com.snc.integration.sso.multi.installer]	Use this plugin instead of the Integration - Multiple Provider Single Sign-On plugin to activate the Multiple Provider Single Sign-On feature. The multiple provider single sign-on plugin enables organizations to authenticate against multiple IDPs (Identity providers) using SAML. It also supports authentication using	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	multiple digest configurations.			
Integration - Tivoli Enterprise Console (TEC) 2.0 [com.snc.integration.tivoli_tec]	Provides version 2.0 integration with Tivoli Enterprise Console (TEC).	Inactive	false	• com.snc.integration.common • com.glide.web_service_application • com.glideapp.agent
Integration - Verizon eBonding [com.snc.integration.verizon_ebonding]	Provides integration with Verizon eBonding.	Inactive	false	• com.snc.integration.common • com.glide.custom_web_service • com.glide.static_wsdl • com.glide.custom_web_service • com.glideapp.ecc_retry_policy
IntegrationHub Remote Process Sync [com.glide.hub.process.sync]			false	• com.glide.hub.process.sync.model • com.glide.hub.process.sync.jobs • com.glide.hub.process.sync.actions • com.glide.hub.action_step.rest • com.glide.hub.integration.runtime • com.glide.hub.flow_designer_introspection
IntegrationHub Remote Process Sync License [com.glide.hub.remote.process.sync.license]	A ServiceNow AI Platform Remote Process Sync subscription enables use of Remote Process Sync with custom and platform tables.		false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Note: Use of Remote Process Sync on custom apps or tables that extend from a ServiceNow created application requires that users be subscribed to both the application extended as well as ServiceNow AI Platform Remote Process Sync.			
Integrations - External Authentication Framework [com.glide.external.app]		Inactive	false	
Intent Discovery [sn_nlu_discovery] (Available in the ServiceNow Store)	Discovers user intents from your incidents, cases, or requests.	Active	false	• com.glide.nlu • com.glide.platform_ml
Interaction Logging, Routing, and Queueing [com.glide.interaction]	Interaction Logging, Routing, and Queueing.	Inactive	false	
Interactions Management [com.glide.interaction]	Activates the Interactions Management System	Active	false	
Interactive Analysis [com.glideapp.interactive_analysis]	Interactive Analysis	Active	false	• com.glideapp.home.publishers • com.glide.ui.ng

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
				• com.glide.ui.ng.amb • com.glide.ui.ng.form_elements
Interview Templates [com.sn_interview_templates]	Enables you to create reusable interview templates that are consistent and efficient for use in investigations.	Active	true	com.sn_component_interview_templates
Investment Funding [sn_invst_pln] (Available in the ServiceNow Store)	Enables you to plan and manage investments by allocating funds to investment entities, such as Business Units, Products, Teams, and so on.			
Investment Funding - ATF Tests [com.snc.investment_planning.atf]	Investment Funding - ATF Tests provides you test cases and test suits that can be run on the Investment Funding Application.	Active	false	com.snc.financial_planning_pmo.atf
Investment Planning [com.snc.investment_planning]	Enables continuous and flexible investment planning for the funding entities based on their priorities and strategic objectives. It provides the options of top-down and bottom-up funding and defunding an entity.	Active	false	• com.snc.app.grid • com.snc.app.widgets • com.snc.fiscal_calendar • com.snc.planned_task_v2
Investment Planning for PPM [com.snc.investment_planning_pmo]	Enables continuous and flexible investment planning for PPM funding entities based on their priorities and strategic objectives. It provides the options of top-down and bottom-up funding and defunding an entity.	Inactive	false	• com.snc.investment_planning • com.snc.financial_planning_pmo

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
IP Range Based Authentication [com.snc.ipauthenticator]	Controls access to an instance based on IP address.	Active	false	
ISC NLU Model for Virtual Agent Conversations [com.glide.isc_nlu]	Activates the Natural Language Understanding (NLU) content pack for the Instance Security Center.	Inactive	false	com.glide.isc_virtualagent
ISC Virtual Agent Conversations [com.glide.isc_virtualagent]	Activates the Virtual Agent content pack for the Instance Security Center.	Inactive	false	• com.glide.cs.chatbot • com.snc.nlu_studio
Issue Auto-Resolution [com.glide.cs.auto_resolution]	Intercepts incidents submitted through the Service Portal or email and gives your end users an option to have Virtual Agent help find a resolution. **Note:** Activate Glide Virtual Agent plugin to enable this feature.	Inactive	false	com.glide.cs.chatbot
IT Asset Management Mobile [com.sn_itam_mobile]	Inventory manager can receive and create assets at the dock.	Active	false	• com.glide.sg • com.glide.sg.offline
ITBM Security [com.snc.itbm_security]	Common plugin to manage security within IT Business Management plugins.	Active	false	
IT Data Mart [com.snc.it_data_mart]	Stores the information that the IT Finance application uses to allocate expenses to specific accounts and segments in the general ledger.	Inactive	true	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
IT Service Management integration with Microsoft Teams [sn_now_teams_it] (Available in the ServiceNow Store)	The integration enables employees to interact with the IT service desk, directly in Microsoft Teams. Employees can ask questions, report on issues, request services and chat with a virtual or live IT service desk agent to resolve issues.	Inactive	false	• sn_now_teams • com.glide.cs.chatbot.lite
Item Designer support for the service catalog [com.glideapp.servicecatalog.item_designer]	Service Catalog item designer.	Active	true	com.glideapp.servicecatalog
ITOM Guided Setup [com.snc.guided_setup_metadata.itom]	ITOM Guided Setup Metadata.	Active	false	com.snc.guided_setup
ITSM and PA Demo Data [com.snc.itsm_pa_demo]	Demo data for Incident, Problem, Change, Task SLAs, Business Services, Service Offering, Service Commitments, and Performance Analytics.	Inactive	true	• com.snc.incident • com.snc.problem • com.snc.change_request • com.snc.sla • com.snc.cmdb • com.snc.service_portfolio • com.snc.service_portfolio.sla
ITSM CSDM Best Practice - Quebec [com.snc.best_practice.itsm.csdm.quebec]	ITSM CSDM Best Practices for Quebec release includes best practice configurations to support Common Service Data Model (CSDM). This plugin includes a business rule that populates the assignment group on the incident, problem, and change request based on the support group field on the CI or Service Offering.	Inactive	false	
ITSM Guided Setup	Contains metadata for ITSM Guided Setup.	Active	false	com.snc.guided_setup

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.guided_setup_metadata.item]	This plugin is for ServiceNow internal usage only.			
ITSM Mobile Agent [com.sn_itsm_mobile_agent]	New redesigned ITSM Mobile experience for Incident Management and On-Call.	Active	false	
ITSM NLU Model for Virtual Agent Conversations [com.snc.itsm.nlu]	This plugin contains NLU models used in the Virtual Agent topics for ITSM-related use cases.	Active	false	com.glide.platform_ml com.glide.dev-studio, com.snc.nlu_studio
ITSM Process Optimization Content Pack (Maintenance mode only) [com.sn_itsm_process_opt_cp]	Provides pre-configured process optimization models and improvement initiatives for the ITSM Processes.	Active	false	
ITSM Roles - Change Management [com.snc.itsm.roles.change_management]	This plugin is included with the ITSM Roles plugin. This plugin adds the sn_change_read and sn_change_write roles. It also updates the existing security model for Change Management to factor these new roles.	Inactive	false	
ITSM Roles - Incident Management [com.snc.itsm.roles.incident_management]	This plugin is included with the ITSM Roles plugin.This plugin adds the sn_incident_read and sn_incident_write roles. It also updates the existing security model for Incident Management to factor these new roles.	Inactive	false	
ITSM Roles - Problem Management [com.snc.itsm.roles.problem_management]	This plugin is loaded with the ITSM Roles plugin.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ITSM Roles - Request Management [com.snc.itsm.roles.request_management]	This plugin adds sn_request_read and sn_request_wirte. It also updates the existing security model for Request Management to factor these new roles.	Inactive	false	
ITSM Shared VA Topic Blocks [com.glideapp.cs.itsm_topic_blocks]	Contains prebuilt reusable conversation topic blocks for common actions like ordering from the catalog and searching the knowledge base.			
ITSM Spoke [com.snc.itsm.spoke]	This Spoke enables Flow Designers to automate and create flows and actions associated with ITSM	Active	false	com.glide.hub
ITSM Virtual Agent Conversations [com.snc.itsm.virtualagent]	Contains prebuilt conversation topics for ITSM related use cases.	Active	true	• com.glide.cs.chatbot • com.glide.service-portal
ITSM Virtual Agent Conversation Topics Lite (Maintenance mode) [com.snc.itsm.virtualagent.lite]	Read only conversations for basic ITSM self service use cases, including opening incidents, checking status,searching the knowledge base, submitting requests, and checking for outages.	Active	false	• com.glide.cs.chatbot.lite • com.glideapp.cs.sm_topic_blocks
ITSM Workspace [com.snc.agent_workspace.itsm]	Provides you the functionality of ITSM Agent Workspace. The plugin is available by default for new and existing users. The ITSM Workspace plugin also activates the ITSM Workspace Landing Pages (com.snc.agent_workspace.itsm.landing_page) plugin.	Active	false	• com.agent_workspace • com.glideapp.servicecatalog.workspace • com.snc.agent_recommend • com.snc.agent_workspace.itsm.landing_pag

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ITSM Workspace Landing Pages [com.snc.agent_workspace.itsm.landing_page]	Delivers ITSM Agent Workspace Landing Pages - basic version.	Inactive	true	
ITSM Workspace Landing Pages - Premium [com.snc.agent_workspace.itsm.landing_page_premium]	Delivers ITSM Agent Workspace Landing Pages - premium version.	Inactive	true	Agent Workspace, Performance Analytics - Content Pack - ITSM Dashboards
JDBC Step for [com.glide.hub.action_step.jdbc] ServiceNow IntegrationHub	Activates the JDBC step in Flow Designer, enabling you to create a reusable action to send SQL commands to a relational database.	Inactive	false	com.glide.hub.action_step.template
JS Code Coverage Debug [com.glide.js.coverage]	Debug utility that collects coverage data for server-side Javascript.	Inactive	false	• com.glide.sessiondebug • com.glide.snc_code_editor
JSDebugger [com.glide.glide-js-debugger]	ServiceNow server-side scripts debugger.	Active	false	com.glide.glide-js-debugger-api
JSON Service request/response model [com.snc.application.json_service]	Core components and helpers for a JSON request/response model. Includes JSON and XML transports for NG and GlideAjax support. An extension of the processor framework.	Active	false	
KCS Integration for Incident Management [com.snc.incident.knowledge]	Provides integration of Incident Management with the Advanced Knowledge Management features.	Inactive	true	com.snc.knowledge_advanced.installer
Keylines Business Service Maps [com.snc.keylines_bsm_map]	An interactive and graphical interface to visualize Configuration Items (CIs) and their relationships. Provides filtering	Active	false	• com.glide.diagrammer • com.glideapp.bsm_map2

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	capabilities to manage data being displayed, allowing for configuration by the user to view in context to their role. Additional capabilities provide for the displaying of related tasks such as Incidents, Problems, Changes and Certification Tasks.			
Knowledge API [sn_km_api]	Provides the ability to search, view, and fetch most-viewed and featured knowledge articles.	Inactive	false	
Knowledge Blocks [com.snc.knowledge_blocks]	This plugin adds an advanced editing feature to Knowledge Management which allows user to create re-usable content blocks that can be included in existing knowledge articles. This feature will give the ability to have one large article and be able to just display pieces of that article to specific users.	Active	true	com.snc.knowledge_advanced.installer
Knowledge Capabilities in UI Builder [com.snc.uib.knowledge]	Enables configurable Knowledge Management artifacts (Pages, Data Transforms, and GraphQL queries) in UI Builder.	Active	false	
Knowledge Management Core [com.glideapp.knowledge]	Installs the core Knowledge Management items used to allow other Knowledge related plugins to work, such as Knowledge V3, Knowledge Advanced, Knowledge Service Portal. This plugin is activated by default.	Active	true	• com.glide.ui.angular • com.glide.ui.ng • com.glideapp.user_criteria • com.snc.db.rotation • com.glideapp.live_feed_v2 • com.glideapp.report.knowledge.overview

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Knowledge Document [com.snc.knowledge_document]	Adds knowledge-based functionalities to the Managed Documents plugin. You can create a knowledge article from a document, or update a knowledge document to a newer revision.	Inactive	false	com.snc.document_management
Knowledge Management - External Content Integration [com.snc.knowledge_external_integration]	Adds external content search capabilities to the knowledge management application. Once configured, this feature creates a copy of the external content on ServiceNow as knowledge articles, and then indexes the articles through Zing Search. You must have appropriate reuse and copy privileges before you configure an external source to be searchable using this feature.	Inactive	true	• com.snc.core.automation.connection_creder • com.snc.knowledge3 • com.snc.knowledge_advanced.installer
Knowledge Management - Mobile Requestor [com.glideapp.knowledge.mobile_requestor]	This is a maintenance plugin used to activate Knowledge Management Mobile Requestor features.	Inactive	false	• com.glideapp.knowledge • com.glide.script.vtable • com.glide.mobile-employee.webview • com.glide.mobile-employee.search
Knowledge Management - Service portal [com.snc.knowledge_serviceportal]	Provides knowledge management features on the service portal.	Active	false	• com.glide.service-portal.esm • com.snc.knowledge3
Knowledge Management Advanced [com.snc.knowledge_advanced]	This plugin adds advanced features to Knowledge Management such as version control and subscriptions. It requires that all	Inactive	false	com.snc.knowledge3

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Knowledge bases are Knowledge v3. Please use Knowledge Management Advanced Installer plugin to activate.			
Knowledge Management Advanced Installer [com.snc.knowledge.advanced.installer]	Use this plugin to install the Knowledge Advanced plugin. When you activate/ upgrade this plugin, it will do validations on knowledge articles and knowledge bases to make sure that the Knowledge Advanced plugin can be successfully installed. If validation fails, look at the errors in the Plugin Activation Logs and follow instructions given to fix them. Once all the issues are fixed, please re-run this plugin to install Knowledge Advanced.	Inactive	false	
Knowledge Management V3 [com.snc.knowledge3]	Activate this plugin when you upgrade from Eureka or earlier.	Active	true	• com.glide.ui.angular • com.glide.ui.ng • com.glideapp.knowledge2 • com.glideapp.user_criteria • com.snc.contextual_search • com.snc.db.rotation • com.glideapp.live_feed_v2 • com.glideapp.report.knowledge.overview
Knowledge Management Wiki Support [com.glideapp.knowledge2.wiki]	Activate this plugin to enable support for Wiki type Knowledge articles.	Active	false	• com.glide.wiki • com.glideapp.knowledge2
Knowledge Management KCS Capabilities	Activate this plugin to enable the use of KCS roles (kcs_contributor, kcs_publisher, and kcs_candidate) and	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
plugin (com.snc.knowledge_kcs_capabilities) [com.snc.knowledge_kcs_capabilities]	use the metadata fields for the KCS article state (governance and confidence).			
Knowledge Overview [com.sn_knowledge_overview] (New in Tokyo)	A dashboard to track knowledge anaytics in real time.	Active	false	• com.glideapp.report.knowledge.overview • com.glideapp.knowledge • com.glideapp.dashboards
Knowledge Product Entitlements [com.snc.kb_product_entitlements]	Knowledge Product entitlements, enables Customer Service admins to restrict access to customers based on products they are entitled to.	Inactive	false	com.sn_customerservice
Knowledge Management - Add-in for Microsoft Word [com.snc.knowledge.ms_word]	Enables authoring content in Microsoft Word and publishing it as a knowledge article using a Word add-in.	Inactive	false	• com.servicenow_now_kb_office_addin • com.glideapp.knowledge • com.sn_outlook_addin
KPI Composer [sn_kpi_composer]	KPI Composer is an application for designing KPIs and Dashboards by using performance measurement methodologies.	Active	false	
KPI Details Workspace Integration [com.snc.pa.kpi_details_ws_intgn]	Performance Analytics - KPI Details allows user to explore more for a certain KPI,this plugin used to integrate sn-kpi-details component to workspace.	Active	false	• com.agent_workspace • com.workspace_core
LDAP Support Enhanced UI [com.glide.ldap.ui]		Inactive	false	• com.glide.ui.ng.guided_flow • com.glide.ldap
Legal Request Management [sn_lg_ops] (Available in the	Enables employees in the organization to submit legal requests and to track their progress while users in the legal department with	Inactive	false	• com.glide.ui.checklist • com.snc.planned_task_v2 • com.workspace_core • sn_ex_emp_fd

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	appropriate roles can work to resolve these requests.			
Legal Matter Management [sn_lg_matter] (Available in the ServiceNow Store)	Enables you to complete legal requests that need cross-departmental tasking and a workflow with a mechanism to store supporting documents and track important milestones.	Inactive	false	sn_lg_ops
Legal Conflict of Interest [sn_lg_coi] (Available in the ServiceNow Store)	Enables you to manage the disclosure, approval, and registry of conflict of interest that might arise from employees having competing interests or loyalties.	Inactive	false	sn_lg_ops
Legal Investigations [sn_lg_investigate] (Available in the ServiceNow Store)	Enables you to streamline the process for investigating internal complaints.	Inactive	false	• sn_lg_ops • sn_lg_matter
Legal Digital Forensics [sn_lg_forensics] (Available in the ServiceNow Store)	Enables you to handle digital forensics requests for data discovery related to custodial and non-custodial data sources that are subject to investigations or litigation.	Inactive	false	• sn_lg_ops • sn_lg_matter
Legal Simple Contracts [sn_lg_contracts] (Available in the ServiceNow Store)	Enables employees to submit legal requests for getting legal support and guidance for contracts.	Inactive	false	• sn_lg_ops • sn_doc

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Legal Stock Preclearance [sn_lg_stock_cp] (Available in the ServiceNow Store)	Enables you to manage stock preclearance requests complying with the company's stock preclearance policy.	Inactive	false	• sn_lg_ops • com.snc.fiscal_calendar
Legal Virtual Agent Conversations [sn_lg_va] (Available in the ServiceNow Store)	Provides predefined conversation topics to enable automated chats with employees seeking legal services.	Inactive	false	• sn_lg_ops • com.glide.cs.chatbot
Legal Mobile [sn_lg_mobile] (Available in the ServiceNow Store)	Enables your employees to submit and track legal requests and lawyers to work on these requests on their mobile devices.	Inactive	false	• sn_lg_ops • com.glide.sg.agent_native_client • com.glide.sg
Legal Counsel Center [sn_lg_workspace] (Available in the ServiceNow Store)	Enables legal department members to categorize, prioritize, and efficiently address legal issues.	Inactive	false	• sn_lg_ops • sn_lg_ws_comps
Legal Counsel Center Components [sn_lg_ws_comps] (Available in the ServiceNow Store)	Seismic Components for Legal Counsel Center.	Inactive	false	• com.servicenow_now_button • com.servicenow_now_chart_bar
Performance Analytics Content Pack for	Provides insights into legal operations to make data-driven decisions for your Legal department.	Inactive	false	• sn_lg_ops • com.snc.pa • com.snc.pa.premium

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Legal Service Delivery [sn_lg_pa] (Available in the ServiceNow Store)				
Legal Practice Applications core [com.sn_lg_lpa_core] (Available in the ServiceNow Store)	Provides pre-built legal workflows with configurable intake forms, process flows, disposition repository, analytics, and integrations with third-party tools.	Inactive	false	Legal Request Management (sn_lg_ops)
Legal Service Management (Maintenance mode only) Planned for deprecation after Feb 1, 2023 [com.snc.legal_service_automation]	Manages legal matters and enables users to report and track matters. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.service_management.core
Legal Service Management Mobile [com.snc.legal_service_automation_m]	Manages Legal Service Management mobile components.	Inactive	true	• com.snc.legal_service_automation • com.snc.service_management_m
Limit Concurrent Sessions [com.glide.limit.concurrent.sessions]	Activate this plugin to limit concurrent interactive sessions on the instance.	Inactive	false	
Link Generator [com.snc.linkgenerator]	Activates the Link Generator that allows you to create deep links on an HR form to access information outside the application. You can generate URLs to manage content, knowledge articles, and catalogs or	Inactive	true	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	access social media or fulfillment requests			
List v2 [com.glide.list_v2]	Updates to the display of lists that include a cleanup UI, hierarchical lists, and related lists embedded in forms.	Active	false	• com.glide.tiny_url • com.glide.ui_list_edit_with_form • com.glide.db_context_menu
List v3 Components [com.glide.ui.list_v3_components]	Provides components required for List v3, including REST endpoints and Angular components.	Active	false	• com.glide.ui.list_api • com.glide.ui.ng • com.glide.ui.ng.amb • com.glide.ui.ng.form_elements • com.glide.ui.ng.filter • com.glideapp.report.charting_v2
Live Agent Conversation Settings [com.glide.cs.live_agent_settings]	Live Agent Conversation Settings	Active	false	com.glide.interaction
Live Feed [com.glideapp.live_feed_v2]	Provides an updated application to post and share content in a ServiceNow instance.	Active	true	• com.glideapp.live_feed • com.glideapp.ui_components
Live Feed [com.glideapp.live_feed]	Provides a place to post and share content in a ServiceNow instance.	Active	true	• com.glide.custom_web_service • com.glideapp.live_common • com.glide.notification
Live Feed Document - follow tasks (Incident, Change, etc.) [com.glideapp.live_feed_document]	Enables you to manage your task conversations and comments from **My Feed** or a document group.	Active	false	com.glideapp.live_feed
Localization Framework [com.glide.localization_framework]	A framework to manage localization workflows across the platform and products. To activate this plugin, install the Localization Framework Installer plugin.	Active	false	• com.glide.localization_framework.sdl_spoke • com.glide.localization_framework.xtm_spoke • com.glide.localization_framework.va • com.glide.localization_framework.service_ca • com.glide.i18n

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Localization Framework for Service Catalog [com.glide.localization_framework.service_catalog]	Enables translation of catalog items using Localization Framework. To activate this plugin, install the Localization Framework Installer plugin.	Active	false	com.glide.localization_framework
Localization Framework for Virtual Agent Topic [com.glide.localization_framework.va]	Enables translation of Virtual Agent Topic using Localization Framework. To activate this plugin, install the Localization Framework Installer plugin.	Active	false	com.glide.localization_framework
Localization Framework Installer [com.glide.localization_framework.installer]	Enhances the existing experience of translating the artifacts in an instance by automating most of the manual tasks in the localization process.	Active	false	• com.glide.hub.integrations • com.glide.localization_framework • com.glide.localization_framework.sdl_spoke • com.glide.localization_framework.xtm_spoke • com.glide.localization_framework.va • com.glide.localization_framework.service_ca
Macroponent Catalog Item for UIB [com.glideapp.servicecatalog.macroponent]	Component for UI Builder for Catalog Item and Record Producer.		false	• com.devsnc_sn_catalog_form • com.devsnc_sn_catalog_action • com.devsnc_sn_catalog_description • com.devsnc_sn_catalog_form_connected • com.devsnc_sn_catalog_price • com.devsnc_sn_catalog_quantity • com.devsnc_sn_catalog_additional_details • com.devsnc_sn_catalog_request_for
Magellan Navigator [com.glide.ui.magellan_navigator]	Provides a redesigned application navigator for Core UI. Combines standard navigation capabilities, customizable favorites, and recently accessed items in a single responsive control.	Active	false	• com.glide.ui.magellan_navigator_api • com.glide.ui.doctype • com.glide.ui.ng

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Maintenance Schedules [com.snc.maintenance_schedules]	Links configuration items to maintenance schedules. The maintenance schedules are checked against the planned dates for changes, and those that appear outside the maintenance schedule are so marked.	Active	false	• com.glide.schedules • com.snc.cmdb
Major Issue Management [com.sn_majorissue_mgt.dyn]	A set of capabilities used to manage customer-facing communications and resolution processes for common issues.	Inactive	true	• com.sn_customerservice • com.sn_publications • com.snc.task_communication_management
Managed documents [com.snc.document_management]	• Lightweight • ITIL-based solution for managing electronic documents within your ServiceNow instance. To enable the ability to publish to the knowledge base, activate the Knowledge Document plugin.	Inactive	true	
Manager Hub [sn_mh] (Available in the ServiceNow Store)	Increases managers self-service and proactive engagement with their team by providing insights and recommended actions for what's most urgent and important to drive team success. It enables managers to grow as leaders through curated and personalised resources.	Active	true	com.snc.employee_center_pro
Manager Workspace [com.snc.manager_workspace]	Manager Workspace App core shell.Install ITSM Manager Workspace or CSM	Active	true	com.snc.wfo

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Maintenance mode only)	Manager workspace to get the modules			
Marketing Service Management [com.snc.marketing_service_automation] (Maintenance mode only) Planned for deprecation after Feb 1, 2023	Manages marketing requests and enables users to report and track those requests. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.service_management.core
Marketing Service Management Mobile [com.snc.marketing_service_automation_m]	Manages Marketing Service Management mobile components.	Inactive	true	• com.snc.marketing_service_automation • com.snc.service_management_m
Meeting Extensions for Microsoft Teams [sn_now_teams_ext] (Available in the ServiceNow Store)	Enables ServiceNow to be embedded directly within a Microsoft Teams meeting kicked off from a Major Incident.	Inactive	false	sn_notify_msteams
Merge Tool [com.glide.ui.mergetool]	User interface for performing merges between two payloads.	Active	false	• com.glide.ui.angular • com.glide.snc_code_editor
Metadata Source Control [com.glide.source_control]	Source control integration applies and commits changes in an external repository.	Active	false	com.glide.system_update_set
Metadata Tree [com.snc.metadata_tree]	Hierarchical representation of metadata.	Active	false	com.snc.apps_file
Metadefender Integration [com.snc.threat.metadefender]	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.threat.intelligence

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Metric Definition [com.glide.metrics]	Provides an easy, declarative way to define metrics and allows the definitions to be tracked and stored by the system.	Active	true	com.glide.schedules
MetricBase [com.snc.clotho]	Time-series data storage and processing.	Inactive	false	• com.glide.highcharts • com.glide.hub.flow_trigger
MetricBase Demo [com.snc.clotho.demo]	Sample time-series data you can use to explore MetricBase functionality.	Inactive	true	com.snc.clotho
Microsoft AD Spoke for IntegrationHub [com.sn.ad.spoke]	The AD Spoke for IntegrationHub provides actions that a Process Analyst can use when designing flows. The actions allow them to automate the management of users, groups, computers, and objects in AD.	Inactive	false	• com.glide.hub.designer_backend.model • com.glide.hub.action_step.powershell
Microsoft AD Spoke for Password Reset [com.snc.password_reset.ms.ad.spoke]				
Microsoft Azure AD Spoke for IntegrationHub [com.sn.azure_ad.spoke]	The Azure AD Spoke for IntegrationHub provides actions that a Process Analyst can use when designing flows. The actions allow them to automate the management of users, security groups, and office groups. User management includes applying licenses that result in user provisioning into Office 365.	Inactive	false	• com.glide.hub.designer_backend.model • com.glide.hub.action_step.rest

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Microsoft Azure AD Spoke for Password Reset [com.sn.password_reset.aad.spoke]	The Password Reset Azure AD Scoped App provides a prebuilt connection to credential stores in Azure AD, allowing reset and changing of password.	Active	false	• com.sn.azure_ad.spoke • com.glideapp.password_reset • com.sn.password_reset.ah
Microsoft Azure Translation Spoke [com.glide.microsoft.translation.spoke]	Using Microsoft translation services, adds ability to translate text from one language to another and detect the language of text.	Inactive	false	com.glide.hub.integrations
Microsoft Integrations - Core [sn_now_teams] (Available in the ServiceNow Store)	Scoped application for store release of Microsoft Teams integration with ServiceNow.	Inactive	false	• com.glide.channelproxy • com.glide.cs.adapter • com.glide.hub.designer_backend.model • com.snc.msteams.app.core • com.glide.hub.action_step.rest • com.glide.hub.integration.runtime • sn_msteams_ahv2 • sn_tcm_collab_hook • sn_now_azure • sn_va_teams
Microsoft OneDrive Spoke [com.sn.onedrive.spoke]	Contains data model needed to integrate with third-party document service providers.	Active	false	
Microsoft OneDrive document services Spoke [com.sn.documents.onedrive]	Collaborate in Microsoft OneDrive and perform copy, delete, restore, and version actions on documents directly in a ServiceNow instance. This integration is bi-directional: Changes are reflected in both Microsoft OneDrive and in the instance.	Active	false	• com.snc.multiprovider_documents • com.sn.onedrive.spoke • com.sn.azure_ad.spoke

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Microsoft SharePoint Spoke [sn_sp_spoke] (Available in the ServiceNow Store)	Enables the Integration Hub to connect with SharePoint.	Active	false	• com.glide.cobject • com.glide.hub.dynamic_inputs • com.glide.hub.action_step.rest • com.glide.hub.action_type.datastream • com.glide.hub.integration.runtime
MID Server [com.glideapp.agent]	The Management, Instrumentation, and Discovery (MID) Server is a Java application that runs as a Windows service or UNIX daemon. The MID Server facilitates communication and movement of data between the ServiceNow platform and external applications, data sources, and services.	Active	false	com.glide.custom_web_service
MID Server Distributed Cluster [com.snc.agent.distributed.cluster]	Enables the MID Server distributed cluster type, for Operational Intelligence.	Inactive	false	com.snc.sa.metric
MID Server Key Management Framework Signature Configuration [com.glideapp.agent.kmf_config]	Creates KMF signature configuration records for MID Server tables.		false	
Mobile Analytics [com.glide.mobile.analytics]	Checks for new mobile applications to register, syncs Mobile Analytics instance settings to the ServiceNow server, and provides users access to the Mobile Analytics Dashboard. This plugin adds the mobile_analytics_admin and	Active	false	com.glide.serviceproxy

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	mobile_analytics_viewer roles.			
Mobile App Builder [sn_mab] (Available in the ServiceNow Store)	A ServiceNow builder for native mobile applications.	Active	false	sn_mab_api
Mobile App Builder API [sn_mab_api] (Available in the ServiceNow Store)	An API for mobile application support with access to mobile application configurations.	Active	false	None
Mobile Card Builder [sn_mobile_card_bui] (Available in the ServiceNow Store)	An interactive item view, view template, and view config builder tool integrated with Mobile Studio.	Active	false	None
Mobile Employee My Assets [com.glide.mobile-employee.myassets]	This plugin contains the configuration and records to create the my assets view that are used in Mobile Employee Experience native application. Do not activate this plugin directly. Activating com.glide.mobile-employee will activate this plugin for you.	Inactive	false	
Mobile Device ITIL and Service Management [com.snc.itil_mobile]	Applications and modules for ITIL and Service Management.	Active	false	• com.glide.labels • com.glide.mobile
Mobile Location Tracking	Enables use of action-based location tracking. This tracking option starts and	Active	false	• com.snc.geolocation • com.glide.sg

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.slide.sg.location.tracking] (New in Utah)	stops tracking based on actions the user performs. The manual tracking option isn't dependent on the installation of this plugin.			
Mobile SDK [com.snc.mobile_sdk]				
Mobile Studio [com.glide.sg-studio]	Plugin for Mobile Studio	Active	false	com.glide.sg
Mobile UI [com.glide.ui.m]	User interface for mobile devices running iOS 6+ or Android 4+ with the Chrome browser. Provides nearly full-product functionality. Adds a System Mobile UI application to the application navigator in the standard UI for configuring the mobile UI.	Active	true	• com.glide.wiki • com.glide.labels • com.glide.db_images • com.glide.ui.font_icons • com.glide.ui.ng • com.glide.ui.zepto • com.glide.ui.tablet.theme • com.glide.ui.recent_selections • com.glide.ui.table_titles • com.snc.platform.security.oauth • com.glide.ui.magellan_navigator_api • com.glide.ui.list_api • com.glide.notification.push
Mobile Publishing [com.glide.sn-mobile-whitelabel]	Mobilew plugin to handle custom app branding requests, statuses, and publishing.	Inactive	false	
Model Management [com.snc.model]	Enables you to manage and maintain model categories, models, suites, and bundled models.	Active	true	
Modeling Engine [com.snc.financial_management]	Enables financial analysts to assemble spending data, build cost models, and	Inactive	true	• com.snc.itfm_core • com.snc.financial_planning

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	generate reports to show how funds are being used. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.			• com.snc.cost_management • com.snc.service_charging
Modeling Engine – ATF Tests [com.snc.financial_management.atf]	ATF tests for Modeling Engine	Inactive	false	com.snc.financial_management_model
Multi Provider Document Services Framework [com.snc.multiprovider_documents] (New in Tokyo)	Enables your team to collaborate in real time and add versions, store, copy, delete, restore, and upload documents.	Active	false	None
Mutex stats [com.glide.stats.mutex]	Mutex stats.	Active	false	com.glide.stats
My Assets [com.snc.asset_myassets]	Provides users with self-service access to their own assets, contracts, and requests.	Active	false	com.snc.asset_management
Name-Value Pairs field editor [com.glide.ui.name_values_editor]	Adds a detailed editor for name-value pairs field types.	Active	false	com.glide.ui.ng
Natural Language Query [com.snc.nlq]	Translates natural language utterances into search queries to quickly find and filter data.	Inactive	false	com.glide.platform_ml
Next Best Action for Customer Service [com.snc.cs_next_best_action]	Provides Next Best Action and Guided Decision configurations and features specific to Customer Service.	Active	false	• com.sn_customerservice • com.snc.next_best_action
Next-Gen BSM	Next Generation BSM (NG-BSM) built on D3	Active	false	com.glide.ui.heisenberg

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.ng_bsm]	and Angular. Provides an enhanced, modern interactive graphical interface to visualize Configuration Items (CIs) and their relationships. Provides filtering capabilities to manage data being displayed and displays related information for CIs such as Events, Incidents, Problems, and Changes.			
NG Common for the Guided Flow experience [com.glide.ui.ng.guided_flow]	Common Angular Components for the Guided Flow experience.	Inactive	false	• com.glide.ui.tablet.theme • com.glide.ui.ng • com.glide.ui.themes.doctype
NG shared components [com.glide.ui.ng]	Provides libraries and services common to plugins using Angular.	Active	false	• com.glide.ui.angular • com.glide.ui.angularui • com.glide.ui.ng.amb
NLU Active Learning - Properties [com.glide.nlu.active_learning_properties]	Installs the required properties for NLU Active Learning.	Active	false	com.glide.nlu
NLU Common Model [com.glide.nlu.model]	The NLU Common Model plugin packages commonly used pattern entities that can be imported and used in any NLU model in the NLU Model Builder. Commonly used patterns like email, phone and ServiceNow specific pattern entities like INT, RITM are made available.	Inactive	false	com.glide.nlu
NLU Workbench	Enables the creation of Natural Language Understanding	Active		

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.nlu_studio]	(NLU) models. These models can understand the intent (action) and entities (details about the action) for a given user utterance/ sentence.			
NLU Workbench - Advanced Features [com.snc.nlu.workbench.advanced] (Available in the ServiceNow Store)	Adds additional features to NLU Workbench including NLU Conflict Review, NLU Batch Testing, and NLU Model Performance.	Active	false	• com.snc.nlu_studio • com.glide.nlu • com.glide.platform_ml • sn_nlu_discovery (store)
NLU Workbench - Core [com.glide.nlu]	Core NLU. Installs the required tables for persisting NLU models that are created using the NLU Model Builder.	Active	false	
Notify Connector for Microsoft Teams [sn_notify_msteams] (Available in the ServiceNow Store)	Expands the Notify providers by managing and initiating a Microsoft Teams meeting directly from any task record.	Inactive	false	• com.snc.notify • sn_now_azure • sn_msteams_com_spk
Normalization Data Services Client [com.glide.data_services_canonicalization.client]	Helps maintain consistency by ensuring that records for a given company refer to that company by the same name.	Inactive	false	com.snc.cmdb
Normalization Data Services Configuration [com.glide.data_services_canonicalization.config]	Normalization Data Services Configuration.	Active	false	com.snc.guided_setup
Notification Preference Service	Defines javascript and REST API for configuring	Active	false	com.glide.notification

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.notification.preference.service]	notification preferences.			
Notification Preference User Interface [com.glide.notification.preference.ui]	Defines the user interface for the Notifications tab in the system settings menu.	Active	false	NG shared components, com.glide.notification.subscription, Notification System Push Addon, com.glide.notification.preference.service
Notification System Push Addon [com.glide.notification.push]	Adds support for push notifications to the notification system.	Active	false	• com.glide.notification • com.glide.push
Notify [com.snc.notify]	Provides powerful platform features for workflow-driven voice calls, conference calls, and SMS messages making it possible to create flexible Interactive Voice Response (IVR) systems to do virtually anything. Requires the Twilio Driver and a separate contract with Twilio for SMS and Voice capabilities.	Inactive	true	• com.glide.phone_number • com.snc.notify.twilio
Notify-Twilio Direct driver [com.snc.notify.twilio.direct]	Provides Notify support for Twilio. Requires a separate contract with Twilio for SMS/Voice capabilities.	Inactive	false	• com.snc.sn_context • com.snc.notify
Now Experience Analytics [com.snc.now-experience-analytics]				
NowMQ v2 Messaging API [com.snc.nowmq_v2]	Provides a simple-to-use queue with features like priorities, completion FIFO, custom states, and backend reconfiguration or migration without any client changes.	Active	false	com.snc.nowmq_v2_apis

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Column Level Encryption Enterprise [com.glide.now.platform.encryption]	Column Level Encryption Enterprise enables field encryption with NIST best practice Key Management Framework (KMF) key lifecycle management and FIPS 140-2-L3 HSM key protection along with the automating Key and Data Migration jobs.		false	• com.glide.kmf.global • com.glide.encryption
OAuth 2.0 [com.snc.platform.security.oauth]	The implementation of OAuth 2.0 to support token granting and authentication.	Active	false	• com.glide.certificates • com.snc.platform.security.oauth.legacy
OAuth 2.0 legacy (Do not activate. Use 'OAuth 2.0') [com.snc.platform.security.oauth.legacy]	Legacy implementation of OAuth 2.0. Install com.snc.platform.security.oauth instead.	Active	false	
ODBC Commons [com.glide.odbc.commons]	ODBC Commons.	Active	false	
Omni Experience - Standard Feature Set [sn-oe-sfs] (Available in the ServiceNow Store)	Contains multiple Conversational Interfaces applications, including Conversational Interfaces Admin Console, diagnostic content for Advanced Work Assignment, and Sidebar.	Active	true	• com.glide.cs.collab com.glide.cs.chatbot.lite • com.sn.ace_framework
Omnichannel Callback [sn_omni_callback] (Available in the ServiceNow Store)	Adds new set of platform capabilities that enable other ServiceNow applications to display or announce callback options to the users.	Active	false	• com.glide.cs • com.glide.interaction.awa • com.glide.cs.custom.adapter • com.snc.appointment_booking

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Omnichannel Callback for Customer Service Management [com.sn_omnichannel_callback] (Available in the ServiceNow Store)	Extends Omnichannel Callback capabilities to support CSM specific use cases. This includes customer experience, agent experience and CSM data model support.	Active	false	• com.sn.omnichannel.callback • com.sn_customerservice
On-Call: Mobile Push Notifications [com.sn_on_call_m_notif]				
On-Call: PA and Reports Content Pack [com.snc.on_call.pa.content]	• Dashboards • Reports and Performance Analytics content pack for On-Call Schedules and Escalations	Active	false	• com.glide.script.vtable • com.snc.on_call_rotation
On-Call Scheduling [com.snc.on_call_rotation]	Provides the ability to create on-call schedules and escalation trees. When an incident is created, dynamically route the escalation to an on-call resource. On Call allows you to configure and build different on-call schedules per process and assignment group. When utilizing the Notify plugin, resources can use SMS and Voice escalations to interact with the escalation to acknowledge incidents, and so on.	Inactive	true	• com.glide.schedules • com.glide.notification • com.glide.ui.ng • com.snc.app.calendar • com.snc.fullcalendar
Openframe [com.sn_openframe]	An interface to integrate external communication	Inactive	false	com.glide.ui.ng

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	systems with ServiceNow. This plugin brings a UI frame that is accessible and available anywhere within ServiceNow screens.			
Openframe Integration App [com.sn_openframe_integration]	Holds records to support UIB framework.	Inactive	false	
Operational Intelligence [com.snc.sa.metric]	Operational Intelligence provides the ability to capture, and then explore and analyze operational metrics data, identifying and indicating anomalies.	Active	false	• com.glideapp.itom.snac • com.snc.sa.mid.webserver • com.snc.clotho • com.glide.highcharts • com.snc.agent.distributed.cluster • com.glide.ui.list_v3_components • com.snc.itom.ui • com.snc.sa.metric.health • com.oi-scoped-app,
Operational Intelligence - WS Scoped App [com.oi-scoped-app]	Operational Intelligence - Workspace Scoped App for different user interfaces.	Active	false	com.glide.graphq
Operational Technology Change Management [com.sn_ot_chg_mgmt] (Available in the ServiceNow Store)	Enables engineers to implement changes to Operational Technology assets and production processes.	Active	false	• CMDB CI Class Models (com.sn_cmdb_ci_class) • ISA Equipment Model (com.sn_isa_model) • Operational Technology Manager (com.sn_ot_foundation)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Operational Technology Incident Management [com.sn_ot_inc_mgmt] (Available in the ServiceNow Store)	Enables manufacturers to manage OT device incidents from open to closure.	Active	false	• CMDB CI Class Models (com.sn_cmdb_ci_class) • ISA Equipment Model (com.sn_isa_model) • Operational Technology Manager (com.sn_ot_foundation)
Operational Technology Manager [com.sn_ot_foundation] (Available in the ServiceNow Store)	Enables manufacturers to aggregate OT data from multiple sources to build a solid data foundation of OT environments.	Active	False	• com.sn_mfg_common • com.sn_cmdb_ci_class • com.sn_isa_model • com.sn_otsm_sgc • com.sn_datagrid
Order Guide Sequential Fulfillment [com.glideapp.servicecatalog.order_guide_sequencing]				
Ordered Email Processing [com.glide.email_ordered_processing]	Allows inbound email actions to be ordered, and programmatically stop processing.	Active	false	
Organization Extension [com.snc.organization_extension]	Provides Goals, Enterprise Strategy and Business Unit strategy entities.	Active	true	
Organization Management [com.snc.organization_management]		Active	false	
Outage Numbering [com.snc.outage_numbering]	This plugin will introduce a new unique identifier column for the outage table. This plugin is activated automatically for all instances except those who already have a number prefix column on the outage table. Please do not	Active		

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	activate this plugin if you have a number prefix for this table.			
Outbound Email Notifications [com.glide.email_outbound]	Enables outbound email notifications.	Active	false	
Outbound HTTP Log [com.glide.outbound_http_log]	Outbound HTTP request log	Active	false	
Outbound Tracking [com.glide.outbound_tracking]	Outbound request tracking.	Active	false	com.glide.usageanalytics
Outlook Actionable Messages [com.sn_ms_oam]	Outlook Actionable Messages	Active	false	com.snc.platform.security.oauth
Outsourced Customer Service (Maintenance mode only) [com.snc.outsourced_service_provider]	Enables onboarding Outsourced Service Provider (OSP) organizations. An enterprise can engage OSPs to provide customer service to external customers for reasons such as regional support, different languages, seasonal availability and so on.	Active	false	• com.sn_customerservice • com.snc.outsourced_service_provider
Overview Pages [com.glide.ui.overview_help]	Framework for Overview pages.	Active	false	com.glide.ui.angular
Packages Call Removal Tool [com.glide.script.packages_call_removal]	Scans scripts for Packages calls to ServiceNow Java classes, proposes changes to replace them with alternate scriptable names, and facilitates the script changes. Packages calls to ServiceNow Java	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	classes will eventually be disallowed in a future ServiceNow release, and this utility helps prepare an instance for that.			
PA Dashboard for NLU Model Builder [com.glide.nlu.pa]	The NLU Model Builder PA plugin tracks the performance of NLU models being used in Virtual Agent. The plugin packages an NLU Performance dashboard, which reports on how many predictions were correct, how many predictions were skipped by the NLU model and how many were considered wrong by Virtual Agent users.			
PA New Model [com.snc.pa.datamodel]				
Package Optimizer [com.glide.package_optimizer]	Optimizes the application upgrade process by comparing the generated file hashes from the previous install.	Active	false	None
PAR Intelligence [com.snc.par.intelligence] (New in Utah)	The core plugin for PAR Intelligence. The classpath for the Forecast Glide Scriptable API, which lets you apply Performance Analytics indicator forecasting to other data sources.	Active	false	
PAR Visualization Migration [com.snc.par.visualization.migration]	Migrates the JSON structure of seismic visualization components in pre-San Diego instances to the new format. This plugin also migrates sn-component-	Active	false	now-visualization-extensions

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	visualization json to now-visualization-extensions json structure.			
Password Reset [com.glideapp.password_reset]	Provides the ability to create self-service and service desk password reset processes for a ServiceNow instance.	Active	true	• com.snc.process_flow_formatter • com.glide.notification • com.glide.usageanalytics
Password Reset API [com.glideapp.password_reset.api]	Contains scoped APIs related to the Password Reset application.	Active	false	
Password Reset - Orchestration Add-on [com.glideapp.password_reset.addon.orchestration]	Password reset add-on to enable the use of ServiceNow Orchestration. Includes support for Active Directory and remote SOAP based credential systems.	Inactive		Password Reset, Orchestration
Password Reset Spoke [com.sn.password_reset.ah]	Provides subflows and actions in Integration Hub required for the Password Reset application.	Active	false	com.glide.hub.designer_backend.model
Password Reset Windows App [com.glideapp.password_reset_desktop]	Password Reset Windows integration.	Inactive	false	
Pattern Designer Enhancements [sn_itom_pde] (Available in the ServiceNow Store)	Offers the Command Validation Framework. Provides discovery administrators with ability to validate credential access for specific commands, including Unix and Windows commands, SNMP queries, SNMP walk, Registry queries, HTTP GET calls.	Inactive	false	None
Patient Support Services [sn_patientservice]	Streamlines the patient onboarding, education, and engagement for various patient support services such	Inactive	true	• com.snc.business_location • com.snc.csm_case_types • com.snc.household • com.snc.install_base

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)	as discount plans, adherence programs, opioid, and diabetes management.			• com.glide.encryption • com.sn_external_user_register • com.snc.pdf_generator • com.snc.signaturepad • com.glideapp.user_criteria.scoped.api • sn_hcls • sn_csm_playbook • com.sn_ind • sn_doc • sn_prd_pm • sn_ciwf_ui_cmpnt
Payment Operations [com.sn_bom_payment]				
PDF Generator [com.snc.pdf_generator]	Provides a tool to generate PDF documents. The Human Resources application uses this with various documents.	Inactive	false	com.snc.signaturepad
Performance Analytics [com.snc.pa]	Enables users to define and track key performance indicators (KPIs) and visualize these in scorecards and dashboards. Users can report and compare multiple time series, do advanced trend analysis, and compare their performance with preset targets.	Active	false	• com.glide.ui.angular • com.glideapp.report • com.glideapp.home.publishers • com.snc.core.automation • com.glide.usageanalytics
Performance Analytics - Administrator Console [com.snc.pa.administration.console]	A single console where you can: • Explore and manage Performance Analytics widgets and dashboards,	Active	false	• com.snc.pa • com.glide.ui.ng • com.glide.ui.ng.amb • com.snc.pa.diagnostics • com.snc.pa.usage.overview

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	including Platform Analytics Solutions • Diagnose and resolve configuration errors • View usage analytics • Modify advanced configuration settings • Access ServiceNow help.			
Performance Analytics - Configuration Generator [com.snc.pa.configurationgenerator]	Provides a configuration generator for creating a set of Performance Analytics indicators, breakdowns, dashboards, and widgets based on the task-derived table.	Inactive	false	com.snc.pa
Performance Analytics - GRC: Advanced Dashboards [com.sn_grc_pa_advanced]	Contains the Advanced Performance Analytics dashboards created for GRC.	Inactive	true	• com.sn_compliance • com.sn_risk_advanced • com.sn_audit
Performance Analytics - Content Pack - Advanced Work Assignment [com.snc.pa.awa]	Provides Performance Analytics indicators and visualizations for Advanced Work Assignment. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Active	false	• com.snc.pa • com.glide.awa
Performance Analytics - Content Pack - Application Portfolio Management	Application Portfolio Management dashboards developed using Performance analytics premium.	Inactive	true	• com.snc.apm • com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.pa.apm]				
Performance Analytics - Content Pack - Application Portfolio Management and Change Management [com.snc.pa.apm.change_request]	Provides integration of Application Portfolio Management with Change Management that enables performance analytics dashboards of business applications associated with Change requests. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	• com.snc.apm • com.snc.pa • com.snc.pa.change
Performance Analytics - Content Pack - Application Portfolio Management and Problem Management [com.snc.pa.apm.problem]	Provides integration of Application Portfolio Management with Problem Management which enables performance analytics dashboards of business applications. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	• com.snc.apm • com.snc.pa • com.snc.pa.problem
Performance Analytics - Content Pack - Change Management [com.snc.pa.change]	Provides Performance Analytics indicators and visualizations for Change Management. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.glideapp.report.itsm.change.overview

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Performance Analytics - Content Pack - Cloud Resources [com.sn_disco_cd]				
Performance Analytics - Content Pack - Communities [com.snc.pa.communities]	Provides Performance Analytics indicators and visualizations for Communities. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.sn_communities
Performance Analytics - Content Pack - Configuration Management (CMDB) [com.snc.pa.cmdb]	Provides Performance Analytics indicators and visualizations for Configuration Management (CMDB). Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.snc.cmdb • com.snc.pa.change
Performance Analytics - Content Pack - Content Analytics [com.snc.pa.premium.content_analytics]	Content Analytics - Performance Analytics	Active	false	• com.sn_content_analytics • com.snc.pa.premium
Performance Analytics - Content Pack - Content Automation [com.snc.pa.premium.content_automation]	Content Automation Analytics	Active	false	• com.sn_content_analytics • com.sn_content_automation • com.snc.pa.premium
Performance Analytics - Content Pack - Customer	Updates the Customer Service dashboard and incorporates	Active	false	• com.snc.pa • com.glideapp.report.customer_service

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Service Management - Advanced [com.snc.pa.customer_service_advanced]	indicators for the following plugins: Major Issue Management, Customer Service Case Action Status, Customer Service with Request Management, Customer Service with Service Management, Agent Chat, Advanced Work Assignment for CSM, and Performance Analytics - Content Pack - Advanced Work Assignment. Activate these plugins in order to view the related indicators. Without these plugins, indicators may not display correctly. Activation of the Performance Analytics â€" Content Pack â€" Customer Service Management - Advanced plugin on production instances may require a separate performance Analytics license. Contact ServiceNow for details.			• com.sn_customerservice • com.snc.pa.spotlight
Performance Analytics - Content Pack - Customer Service [com.snc.pa.customer_service]	Provides Performance Analytics indicators and visualizations for Customer Service. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.glideapp.report.customer_service • com.sn_customerservice • com.sn_customerservice_pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Performance Analytics - Content Pack - Discovery [com.snc.pa.discovery]	Performance Analytics content pack for Discovery. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.snc.pa.cmdb • com.snc.discovery
Performance Analytics - Content Pack - Event Management [com.snc.pa.em]	Performance Analytics content pack for Event Management core. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.glideapp.itom.snac • com.glideapp.report.em
Performance Analytics - Content Pack - Field Service Management [com.snc.work_management_pa]	Provides Performance Analytics content for Field Service Management. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	• com.snc.pa • com.snc.work_management
Performance Analytics - Content Pack - Financial Management [com.snc.pa.fm]	Provides Performance Analytics indicators and visualizations for Financial Management. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.financial_management • com.snc.pa
Performance Analytics - Content Pack - Financial Management	Enables Performance Analytics dashboards for Financial Management associated with	Inactive	true	• com.snc.financial_management_for_csm • com.snc.pa

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
for Customer Service [com.snc.pa.fm.csm]	Customer Service Management. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.			
Performance Analytics - Content Pack - Financial Management for Field Service Management [com.snc.pa.fm.fsm]	Enables Performance Analytics dashboards for Financial Management associated with Field Service Management. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.financial_management_for_fsm • com.snc.pa
Performance Analytics - Content Pack - Financial Management for Financial Planning [com.snc.pa.financial_planning]	Enables Performance analytics dashboards to associate to Financial Planning Process. Activation of this plugin on production instances may have licensing implications. Contact your ServiceNow account team for details.	Inactive	true	• com.snc.financial_planning • com.snc.pa
Performance Analytics - Content Pack - GRC: Audit Management [com.sn_audit_pa]	Provides Performance Analytics reports for the GRC Audit Management application. Activation of this plugin on production requires a PA Premium license. Contact ServiceNow for details.	Inactive	false	• com.sn_grc_pa • com.sn_audit

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Performance Analytics - Content Pack - GRC: Policy and Compliance Management [com.sn_compliance_pa]	Provides Performance Analytics reports for the GRC Policy and Compliance Management application. Activation of this plugin on production requires a PA Premium license. Contact ServiceNow for details.	Inactive	false	• com.sn_grc_pa • com.sn_compliance
Performance Analytics - Content Pack - GRC: Risk Management [com.sn_risk_pa]	Provides Performance Analytics reports for the GRC Risk Management application. Activation of this plugin on production requires a PA Premium license. Contact ServiceNow for details.	Inactive	false	• com.sn_grc_pa • com.sn_risk
Performance Analytics - Content Pack - Human Resources Employee Document Management Scoped App [com.sn_hr_employee_files_pa]	Provides Performance Analytics reports and dashboards for Employee Document Management.	Inactive	false	• com.sn_employee_document_management • com.sn_hr_core
Performance Analytics - Content Pack -Vulnerability Response [sn_vul_analytics] (Available in the ServiceNow Store)	Provides an integration of Vulnerability Response with Performance Analytics for trend-based reporting. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.		true	• com.snc.pa • com.snc.treemap • sn_vul
Performance Analytics - Content Pack - Human Resources Lifecycle	Enables you to define and track key performance indicators (KPIs) for Enterprise Onboarding and	Inactive	false	com.sn_hr_lifecycle_events

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Events Scoped App [com.sn_hr_lifecycle_pa]	Transitions and show these in scorecards and dashboards. Activation of this plugin on production will require a PA Premium license. Contact ServiceNow for details.			
Performance Analytics - Content Pack - Human Resources Scoped App [com.sn_hr_pa]	Enables you to define and track key performance indicators (KPIs) for HR and show these in scorecards and dashboards. Activation of this plugin on production will require a PA Premium license. Contact ServiceNow for details.	Inactive	true	com.sn_hr_core
Performance Analytics - Content Pack - Incident SLA Management [com.snc.pa.sla]	Performance Analytics content pack for Incident SLA Management. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	com.snc.pa
Performance Analytics - Content Pack - ITSM Dashboards [com.snc.pa.itsm_dashboards]	Performance Analytics content pack for ITSM Dashboards. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa.change • com.snc.pa.problem • com.snc.pa.request • com.snc.pa.request2 • com.snc.pa.knowledge • com.snc.pa.sla • com.snc.pa.spotlight.incident
Performance Analytics - Content Pack - Knowledge Management	Performance Analytics content pack for Knowledge Management. Activation of this plugin on production	Inactive	false	• com.snc.pa • com.snc.knowledge3 • com.glideapp.report.knowledge.overview

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.pa.knowledge_v2]	instances may require a separate Performance Analytics license. Contact ServiceNow for details.			
Performance Analytics - Content Pack - Operational Intelligence [com.snc.sa.metric.pa.content]	Performance Analytics content pack for Analytics and Reporting Operational Intelligence Solution Dashboard and KPIs. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.sa.metric • com.snc.pa • com.glideapp.canvas • com.snc.pa.em
Performance Analytics - Content Pack - Problem Management [com.snc.pa.problem]	Provides Performance Analytics indicators and visualizations for Problem Management. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.glideapp.report.itsm.problem.overview
Performance Analytics Content Pack for Procurement Service Management [com.sn_spend_pa] (Starting with the Utah release, Performance Analytics Content Pack for Procurement Service Management	Provides a set of pre-configured metrics and dashboards to assess spend, operational efficiency, and team performance across the Sourcing and Procurement Operations product.	Inactive	false	• com.snc.sn_shop • com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
is renamed as Performance Analytics for Sourcing and Procurement Operations.) (Available in the ServiceNow Store)				
Performance Analytics - Content Pack - PPM Standard Dashboards [com.snc.pa.pmo_dashboards]	Performance Analytics content pack for PPM Standard core. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Active	false	• com.snc.pps_dashboards • com.snc.financial_planning_pmo
Performance Analytics - Content Pack - PPM Standard [com.snc.pps_dashboards]	Performance Analytics content pack for Project Portfolio Suite Dashboards core. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.snc.project_portfolio_suite
Performance Analytics - Content Pack - Request Management (Requested Item) [com.snc.pa.request]	Performance Analytics content pack for Request Management for Requested Items. Installed with Performance Analytics - Premium. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	com.snc.pa

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Performance Analytics - Content Pack - Request Management (Requests) [com.snc.pa.request2]	Performance Analytics content pack for Request Management for Requests. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	com.snc.pa
Performance Analytics - Content Pack - Service Desk Chat [com.snc.pa.chat]	Provides Performance Analytics indicators and visualizations for Service Desk Chat. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.glide.connect • com.glide.connect.support • com.glide.connect.managers_dashboard • com.snc.pa
Performance Analytics - Content Pack - Service Mapping [com.snc.service-mapping.pa.content]	Performance Analytics content pack for Service Mapping Dashboard and KPIs. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Inactive	false	• com.snc.pa • com.snc.pa.cmdb
Performance Analytics - Content Packs for ITSM [com.snc.pa.premium.all_content]	Installing this plugin installs content packs for Incident, Problem, Change, Incident SLA, Knowledge, and Requested Items for Performance Analytics.	Inactive	false	• com.snc.pa.change • com.snc.pa.problem • com.snc.pa.request • com.snc.pa.request2 • com.snc.pa.knowledge • com.snc.pa.sla
Performance Analytics - Content Pack - Time Card Management	Performance Analytics content pack for Time card management. Activation of this plugin on production	Active	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.pa.time_card]	instances may require a separate Performance Analytics license. Contact ServiceNow for details.			
Performance Analytics - Context Sensitive Analytics for Change Management [com.snc.pa.change.context_sensitive_analytics]	Provides additional ability to open Performance Analytics context sensitive Performance Analytics dashboards in change request form based on UI actions.	Inactive	false	• com.snc.pa • com.snc.pa.change
Performance Analytics - Context Sensitive Analytics for Chat [com.snc.pa.chat.context_sensitive_analytics]	Provides additional dashboard to analyze Connect Support metrics for support queues in Performance Analytics. Adds a related link to the Chat Queue Entry [chat_queue_entry] form to quickly display Performance Analytics metrics for the associated queue.	Inactive	false	• com.snc.pa • com.snc.pa.chat
Performance Analytics - Context Sensitive Analytics for Customer Service [com.snc.pa.cs.context_sensitive_analytics]	Performance Analytics adding ability to open Performance Analytics context-sensitive dashboards in customer service form based on UI actions.	Inactive	false	• com.snc.pa • com.snc.pa.customer_service • com.sn_customerservice • com.snc.pa.customer_service
Performance Analytics - Context Sensitive Analytics for Incident [com.snc.pa.incident.context_sensitive_analytics]	Performance Analytics adding ability to open Performance Analytics context-sensitive dashboards in incident form based on UI actions.	Inactive	false	com.snc.pa
Performance Analytics - Context Sensitive Analytics	Performance Analytics adding ability to open Performance Analytics context	Inactive	false	• com.snc.pa • com.snc.pa.problem

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
for Problem Management [com.snc.pa.problem.context_sensitive_analytics]	sensitive Performance Analytics dashboards in problem form based on UI actions.			
Performance Analytics - Diagnostics [com.snc.pa.diagnostics]	Enables users to run diagnostics on Performance Analytics records and related applications.	Active	false	com.snc.pa
Performance Analytics - Domain Support [com.snc.pa.domain_support]	The Performance Analytics Domain Support plugin provides additional features to support scores collection on domain separated instances. This plugin requires the Domain Separation plugin to be installed.	Active	false	• com.snc.pa.premium • com.glideapp.dashboards • com.glide.domain.msp_extensions.installer
Performance Analytics - Example - LinkedIn [com.snc.pa.linkedin]	Automatically imports LinkedIn data historically and daily.	Inactive	false	com.snc.pa
Performance Analytics - Example - Stocks Quotes [com.snc.pa.stock]	Automatically imports stock quotes data historically and daily.	Inactive	false	com.snc.pa
Performance Analytics - Example - Twitter [com.snc.pa.twitter]	Automatically import Twitter data historically and daily.	Inactive	false	com.snc.pa
Performance Analytics - PA Solution Library [com.snc.pa.solution.library]	Allows you to upgrade the visuals of dashboards that are shipped in other content packs. Upgrade or duplicate whole dashboards or selected elements.	Active	false	com.snc.pa
Performance Analytics - Premium	Enables the functionality	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.pa.premium]	provided by the fully licensed version of Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics - Scores Migration [com.snc.pa.scores_migration]	Migrates scores from old tables to new tables. To provide optimal performance, the Scores [pa_scores] table is being migrated to Scores Level 1 [pa_scores_l1] and Scores Level 2 [pa_scores_l2]. During migration you cannot collect, modify, or delete scores. Scheduled data collection and cleanup jobs will not run during migration.	Active	false	com.snc.pa
Performance Analytics - Spotlight [com.snc.pa.spotlight]	Enables records to be ranked by multiple weighted criteria. When a record surpasses a threshold of total weights, a Spotlight is created to highlight that record.	Inactive	false	com.snc.pa

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Performance Analytics - Spotlight - change spotlight content pack [com.snc.pa.spotlight.change]	Provides preconfigured Spotlight components and a dashboard to identify and prioritize Change records according to multiple weighted criteria. Requires fully licensed Performance Analytics.	Inactive	false	• com.snc.pa.spotlight • com.snc.pa.change
Performance Analytics - Spotlight - Incident Spotlight content pack [com.snc.pa.spotlight.incident]	Provides preconfigured Spotlight components and a dashboard to identify and prioritize Incident records according to multiple weighted criteria.	Inactive	false	• com.snc.pa.spotlight • com.glideapp.canvas
Performance Analytics - Spotlight - Problem spotlight content pack [com.snc.pa.spotlight.problem]	Provides preconfigured Spotlight components and a dashboard to identify and prioritize Problem records according to multiple weighted criteria. Requires fully licensed Performance Analytics.	Inactive	false	• com.snc.pa.spotlight • com.snc.pa.problem
Performance Analytics - Spotlight - Request spotlight content pack [com.snc.pa.spotlight.request]	Provides preconfigured Spotlight components and a dashboard to identify and prioritize Request records according to multiple weighted criteria. Requires fully licensed Performance Analytics.	Inactive	false	• com.snc.pa.spotlight • com.snc.pa.request2
Performance Analytics and Reporting - Service Portal Widgets [com.snc.pa.sp.widget]	Support for Performance Analytics widgets in the Service Portal. A number of JavaScript and CSS libraries are included to be able to run Performance Analytics widgets	Active	false	com.glide.service-portal

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	independently in the Service Portal.			
Performance Analytics Content Pack for FSO [sn_bom_pa] (Available in the ServiceNow Store)	Provides Performance Analytic capabilities to Financial Services Operations applications.	Inactive	false	• com.snc.pa • sn_bom
Performance Analytics Content Pack for Major Incident Management [com.snc.pa.incident.mim]	Provides Performance Analytics indicators and visualizations for identifying and tracking major incidents. Extended functionality available with fully licensed Performance Analytics.	Inactive	false	com.snc.incident.mim
Performance Analytics Usage Overview [com.snc.pa.usage.overview]	This plugin contains dashboards, reports, and indicators to monitor Performance Analytics and Reporting usage.	Active	false	com.snc.pa
Performance Analytics Premium [com.snc.pa.premium]	Enables the functionality that you are entitled to with a subscription to Performance Analytics across all ServiceNow® products. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	• Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics and Reporting - Service Portal Widgets [com.snc.pa.sp.widget]	Support for Performance Analytics widgets in the Service Portal. A number of JavaScript and CSS libraries are included to be able to run Performance Analytics widgets independently in the Service Portal.	Inactive	false	com.glide.service-portal
Performance Analytics and Reporting for Universal Request [com.snc.universal_request.pa]	Dashboards and reports for enhanced visibility on universal request metrics.	Active	false	com.snc.pa,com.snc.universal_request.reporti
Performance Analytics Premium for APM [com.snc.pa.premium.apm]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes APM and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Performance Analytics Premium for Business Management [com.snc.pa.premium.itbm]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Business Management and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa
Performance Analytics Premium for Cloud Control Center [com.snc.pa.premium.ccc]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Cloud Control Center and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	• Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics Premium for Customer Service [com.snc.pa.premium.cs]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Customer Service and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa
Performance Analytics Premium for Financial Management [com.snc.pa.premium.fm]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Financial Management and Performance Analytics. Requires additional fee. Functionality includes the ability to:	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	• Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics Premium for Human Resource Management [com.snc.pa.premium.hr]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Human Resource Management and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa
Performance Analytics Premium for IT Operations Management [com.snc.pa.premium.itom]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes ITOM and Performance	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics Premium for IT Operations Suite [com.snc.pa.premium.itos]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes IT Operations Suite and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa
Performance Analytics Premium for Project	Enables the Performance Analytics functionality that you are entitled	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Portfolio Management [com.snc.pa.premium.ppm]	to with a subscription that includes PPM and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics Premium for Security Incident Response [com.snc.pa.premium.sir]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes SIR and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa
Performance Analytics	Enables the Performance	Active	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Premium for Security Operations [com.snc.pa.premium.sir]	Analytics functionality that you are entitled to with a subscription that includes Security Operations and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics Premium for Service Management [com.snc.pa.premium.service_management]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Service Management and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets	Inactive	false	com.snc.pa

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	• Use Performance Analytics with external data • Preserve scores beyond 180 days			
Performance Analytics Premium for Software Asset Management [com.snc.pa.premium.sam]	Enables the Performance Analytics functionality that you are entitled to with a subscription that includes Software Asset Management and Performance Analytics. Requires additional fee. Functionality includes the ability to: • Create indicators, breakdowns, and other records • Create interactive filters and use interactive analysis • Create text analytics widgets • Use Performance Analytics with external data • Preserve scores beyond 180 days	Inactive	false	com.snc.pa
Performance Dashboards [com.glide.performance_dashboards]	Performance Dashboards	Active	false	com.glide.cms
Performance Metrics for VTB [com.glide.ui.vtb_metrics]	Provides profiler page for benchmarking Java methods in Visual Task Board.	Inactive	false	com.glide.ui.vtb
Personal Lines Servicing [com.sn_ins_policy_b2c]	Enables insurance carriers to be agile and customer-centric in resolving mid-term policy changes and requests. Insurance carriers can prioritize underwriting cases		Yes	• com.sn_csm_playbook • com.snc.csm_contributor_user • com.sn_bom • com.sn_bom_document • com.sn_ins_underwrite

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	by premium value with intelligent business rules and eliminate low-value manual work from underwriting queues.			
Personal Lines Underwriting [com.sn_ins_underwrite]	Enables insurance carriers to route complex severing requests instantly to underwriters based on their existing underwriting guidelines and rules. Underwriters can collaborate with distribution and servicing teams though an efficient and transparent workflow.		No	• com.snc.csm_contributor_user • com.sn_bom
Planned Maintenance Management [com.snc.fsm_pm] (Available in the ServiceNow Store)	Manages the regular preventive maintenance of assets.	Active	true	Field Service Management (com.snc.work_management)
Planned Work Management [com.snc.fsm_planned_work_management] (Available in the ServiceNow Store)	Enables organizations to create and manage the planned activities for assets that are scheduled to be executed periodically.	Active	true	• com.snc.fsm_template_management • com.snc.planned_maintenance
Platform as a Service [com.snc.paas]	Allows the development of custom applications to meet business needs.	Active	false	com.glideapp.workflow
Playbooks for Customer Service Management [com.sn_csm_playbook]	Playbooks for Customer Service Management guides customer service agents through the various tasks and	Active	false	• com.glide.pad.core • com.glide.playbook_experience.config • com.sn_customerservice

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	their sequence to resolve customer issues and visualizes the entire lifecycle across diverse and siloed processes.			
Portal Next Experience Theme [sn_sppolaris_theme] (Available in the ServiceNow Store)	Allows the Next Experience theme to be used with the Customer Support (csm), Customer Service (csp), and Knowledge (kb) portals.	Inactive	false	com.glide.service-portal
Portfolio Management [com.snc.portfolio_management]	The Portfolio Management plugin.	Inactive	false	
PPM Mobile [com.sn_ppm.mobile]	This plugin provides Project Portfolio Suite Mobile user experience. This provides access to project status and project status report in the agent mobile application. PPM Standard is required. But if it is not active, installing PPM Mobile also activates PPM Standard plugin.	Active	false	com.snc.financial_planning_pmo
PPM Ridac [com.snc.ppm_ridac]	The PPM RIDAC flow actions plugin. Installs the flow actions needed for RIDAC conversions.	Active	false	
PPM Standard Multicurrency [com.snc.ppm_multicurrency]	Enable advanced ability to manage PPM financials in multiple currencies.	Inactive	true	• com.snc.financial_planning_pmo • com.glide.currency2
Predictive Intelligence [com.glide.platform_ml]	This is the new name for the Agent Intelligence plugin, effective in the New York release and beyond. Predictive Intelligence enables	Active	false	com.glide.platform_ml_pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	the creation of machine learning solutions using data in your instance. The plugin provides various capabilities and solution types for training the system to predict, recommend, and drive data outcomes.			
Predictive Intelligence - Enhanced UI [com.snc.ml_ui]	This plugin enhances the UI for Predictive Intelligence.	Active	false	
Predictive Intelligence for Contextual Search [com.snc.contextual_search_ml]	Enables customers to leverage machine-learning algorithms for searching with Contextual Search. For example: Similar Open Incidents. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.contextual_search • com.glide.platform_ml
Predictive Intelligence for Customer Service Management [com.snc.csm_ml]	Enables customers to leverage machine learning algorithms for searching related cases in Customer Service Management. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.sn_customerservice • com.snc.contextual_search_ml • com.snc.adv_cxs_results_email_script
Predictive intelligence for CSM - Task intelligence [com_snc_csm_ml_task]				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Predictive Intelligence for Field Service Management [com.snc.fsm_ml]	Provides various capabilities driven by machine-learning solutions. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Active	true	• com.snc.work_management • com.snc.contextual_search_ml
Predictive Intelligence for Incident [com.snc.incident.ml]	Enables customers to leverage machine-learning algorithms with application logic for predicting Open change request and Open Problem for Incident. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.contextual_search_ml
Predictive Intelligence for Incident Management [com.snc.incident.ml_solution]	Predictive Intelligence for Incident Management provides solution definitions as templates so that you can build models without writing any code.	Inactive for new customers. Active for existing ITSM Pro customers.	false	NA
Predictive Intelligence for Major Incident Management [com.snc.incident.mim.ml_solution]	Predictive Intelligence for Major Incident Management provides solution definitions as templates so that you can build models without writing any code.	Inactive for new customers. Active for existing ITSM Pro customers.	false	• com.snc.incident.mim • com.snc.incident.ml_solution
Predictive Intelligence for Knowledge Management [com.snc.knowledge_ml]	Provides various Knowledge Management capabilities driven by machine-learning solutions. For example, capabilities such as similar articles and knowledge	Inactive	false	com.snc.contextual_search_ml

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	demand insights. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details. **Note:** When you install and activate the Predictive Intelligence for Knowledge Management plugin (com.snc.knowledge_ml), the Predictive Intelligence for Contextual Search plugin (com.snc.contextual_search_ml) is also activated.			
Predictive Intelligence for PPM [com.snc.ppm_ml] (Available in the ServiceNow Store)	The Predictive Intelligence for Project Management capability uses machine-learning algorithms to search and display similar projects and demands.	Active	false	• com.snc.contextual_search_ml • com.snc.financial_planning_pmo
Predictive Intelligence for Universal Request [com.snc.universal_request.ml]	Contains the machine learning solutions for the Universal Request use cases.	Active	false	• com.snc.universal_request • com.glide.platform_ml
Predictive Intelligence Incident Estimated Time to Resolve [com.snc.incident.ml.ettr]	Enables customers to leverage machine-learning algorithms with application logic for predicting the estimated time to resolve for an incident. Activation of this plugin on production instances may require a separate license.	Inactive	false	• com.glide.platform_ml • com.sn_itsm_ettr_card

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Contact ServiceNow for details. **Important:** It is not recommended for customers to activate this plugin directly. To activate this plugin, follow the instructions given in KB0821555.			
Predictive Intelligence Workbench HRSD Content [com.sn_piwb_hrsd_content]	Enables customers to seed HRSD specific content to provide implementation guidance for use cases created through Predictive Intelligence Workbench.	Active	false	• com.sn_hr_core • com.glide.platform_ml
Predictive Intelligence Reports [com.glide.platform_ml_pa]	This is the new name for the Agent Intelligence Reports plugin.	Active	false	
Predictive Intelligence Workbench [com.sn_piwb_ml]	Predictive Intelligence Workbench core plugin provides a common framework for creating use case templates, managing use cases and dashboard to communicate value across all SN apps	Active	false	• com.glide.platform_ml • com.glide.platform_ml_pa
Pre-Visit Management [sn_previsit] (Available in the ServiceNow Store)	Streamlines the scheduling process of procedure requests for patients and increase visibility to pre-authorization approvals prior to scheduled procedures.	Inactive	true	• com.snc.business_location • com.snc.csm_case_types • com.snc.household • com.snc.install_base • com.glide.encryption • com.sn_external_user_register • com.snc.pdf_generator • com.snc.signaturepad • com.glideapp.user_criteria.scoped.api • sn_hcls

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
				• sn_csm_playbook • com.sn_ind • sn_doc • sn_prd_pm • sn_ciwf_ui_cmpnt
Proactive Customer Service Operations [com.snc.proactive_cs_ops]	Enables customers to create cases proactively from alerts either manually or through automation and track the accounts and the corresponding install base items affected by the alert.	Active	false	• com.sn_majorissue_mgt • com.snc.install_base • com.glideapp.itom.snac • com.sn_cs_sm_request • com.snc.csm_action_status
Proactive Customer Service Operations with Event Management [com.snc.proactive_cs_itom]	Provides an integration between Customer Service Management and ITOM Event Management. Enables customers to create cases proactively from alerts either manually or through automation and track the accounts and the corresponding install base items affected by the alert.	Inactive	false	• com.snc.proactive_cs_ops • com.glideapp.itom.snac
Proactive Prompts [com.sn_pp] (Available in the ServiceNow Store)	Enables contextual and proactive engagement for managers and employees by bringing actionable insights as prompts into their flow of work.	Active	true	com.glide.cs.chatbot
Problem Management [com.snc.problem]	Helps to identify the cause of a service interruption reported by a significant or recurring incidents.	Active	true	• com.snc.service • com.glideapp.report.itsm.problem.overview
Problem Management — ATF Tests	Delivers ATF tests for Problem Management.	Active	true	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.problem.atf]				
Problem Overview Homepage [com.glideapp.report.itsm.problem.overview]	Displays Problem overview homepage.	Active	false	com.glideapp.report
Problem Management Best Practice - Madrid [com.snc.best_practice.problem.madrid]	Helps to identify the cause of a service interruption reported by a significant or recurring incidents. Provides roles for problem management including a problem coordinator, problem manager and problem administrator. Provides fields to record the category, where the issue was first reported, the workaround, the cause notes and the fix notes. Searches for and attaches knowledge articles. Communicates when a workaround or fix is available.	Active	true	• com.snc.problem • com.snc.problem_task
Problem Management Best Practice - Madrid - Knowledge Integration [com.snc.best_practice.problem.madrid.knowledge]	Helps to identify the cause of a service interruption reported by a significant or recurring incidents. Create Known Error article for Incident Deflection (requires Knowledge Management Advanced plugin to be activated).	Active	true	• Problem Management Best Practice - Madrid (com.snc.best_practice.problem.madrid) • Knowledge Management Advanced Installer (com.snc.knowledge_advanced.installer)
Problem Management Best Practice - Madrid - State Model [com.snc.best_practice.problem.madrid.state_model]	Provides support for state management to control how a problem or problem task is allowed to transition through a predefined list of states.	Inactive	true	Problem Management Best Practice - Madrid (com.snc.best_practice.problem.madrid)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Problem Overview Dashboard [com.snc.pa.problem.dashboards]	Displays Problem Overview Dashboard.	Active	false	com.snc.problem
Problem Tasks [com.snc.problem_task]	Adds a Problem Task table with a reference to the Problem table.	Active	false	
Process Automation Designer for ITSM [com.snc.itsm.playbook]	Process Automation Designer enables business process owners to consolidate multiple flows across the enterprise into a single unified process. Process owners can create and manage cross-functional automated processes to provide end-users with a simplified, task-oriented interface to guide them through these processes, such as playbooks in Agent Workspace.Process Automation Designer for ITSM enables process owners to create and manage the cross-functional IT process flows.	Active	false	• com.glide.pad.core • com.playbook_experience • com.snc.incident
Process Flow Formatter [com.snc.process_flow_formatter]	Quickly summarizes multiple pieces of information about a process and displays the stages graphically at the top of a form.	Active	false	
Process Optimization [com.sn_process_optimization]	Creates business process flows from the data in audit trails. This allows in-depth analysis of business processes to improve outcomes.	Active	false	• com.agent_workspace • com.workspace_core • com.devsnc_sn_promin_workbench
Process Optimization Content Pack for CSM	Provides pre-configured process optimization models and improvement initiatives for the CSM Processes.	Inactive	false	• com.sn_customerservice • com.sn_process_optimization

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.csm_process_optimization]				
Process Optimization Content Pack for FSM [com.snc.fsm_process_optimization] (Available in the ServiceNow Store)	Creates business process flows from the work order task data in audit trails, allowing process owners to perform in-depth analysis and discover process insights to improve business outcomes.	Active	true	• sn_po (v23.0.22) • com.snc.work_management
Process Optimization for external data [com.sn_po_extdata]				
Process Optimization Workspace [com.sn_po_workspace]	Delivers the Process Optimization workspace feature for analyst users.	Inactive	false	• com.agent_workspace • com.sn_process_optimization • com.sn_po_workspace_components
Procurement [com.snc.procurement]	Allows users to create purchase orders and obtain items for fulfilling service catalog requests.	Inactive	true	
Procurement Case Management [com.sn_spend_psd] (Available in the ServiceNow Store)	Enables all employees to request services from the procurement team, and allows the procurement team to manage those requests.	Inactive	true	• com.sn_spend_sdc • com.sn_spend_workspace
Procurement File Transfer Framework [com.sn_spend_ftp_intg] (Available in the ServiceNow Store)	Extends the Procurement Integration Framework with SFTP and FTP functionality, designed for integrating with systems that do not support REST or SOAP-based integration methods.	Inactive	false	• com.sn_spend_intg • com.glide.hub.action_step.ssh • com.glide.hub.action_step.sftp • com.glide.hub.flow_trigger.rest

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Procurement Integration Framework [com.sn_spend_intg] (Available in the ServiceNow Store)	Provides a set of staging tables, transform maps, and workflows to integrate sourcing, third-party catalogs, ordering, shipments, and invoicing with Sourcing and Procurement Operations.	Inactive	false	• com.glide.explicit_roles • com.sn_pr
Procurement Service Management NLU Model [com.sn_spend_nlu] (Starting with the Utah release, Procurement Service Management NLU Model is renamed as Natural Language Understanding Models for Sourcing and Procurement Operations.) (Available in the ServiceNow Store)	Provides Natural Language Understanding (NLU) models to enhance the virtual agent conversation interface, using natural human utterances to detect the correct conversations intended by employees.	Inactive	false	• com.glide.cs.chatbot • com.sn_shop_va
Procurement with Project Management [com.sn_spend_ppm] (Starting with the Utah release, Procurement with Project Management is renamed as Project Costing for	Provides a set of capabilities to automate the calculation of planned versus actual cost by linking purchase orders to cost plans within Project Portfolio Management.	Inactive	false	• com.snc.sn_pr • com.snc.financial_planning_pmo

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Sourcing and Procurement Operations.) (Available in the ServiceNow Store)				
Product Catalog [com.snc.product_catalog]	Provides information about individual models. Models are specific versions or various configurations of an asset. Models published to the product catalog are automatically published to the service catalog.	Active	true	
Product Catalog Advanced [com.sn_prd_pm_adv]				
Product Inventory Advanced [com.sn_prd_invt]				
Program Common [com.snc.program_common]	The Program Common plugin aids in planning, organizing, and managing programs with basic features.	Active	false	• com.snc.planned_task_v2 • com.snc.itbm_security
Profanity filter for agent chat [com.sn.va.profanity]				
Project Management [com.snc.project_management_v3]	The Project Management application aids in planning, organizing, and managing projects and resources in order to setup, execute, and complete a project faster and easier. Only upgrade is	Inactive	true	• com.snc.project_management_delete • com.glide.schedules • com.snc.planned_task_v2 • com.snc.time_card • com.snc.skills_management • com.snc.process_flow_formatter • com.snc.cost_management

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	allowed for this plugin. Activation should be done through PPM Standard plugin.			• com.snc.project_management_db_views • com.snc.project_workbench • com.snc.timeline_visualization • com.glide.ui.checklist
Project Management TeamSpace 1 [com.snc.ppm_teamspace_1]	Installs a Project TeamSpace so that each team or each department like Marketing, Finance, IT-Team1, IT-Team2 can implement PPM suitable to their needs without overstepping each other.	Inactive	true	com.snc.financial_planning_pmo
Project Management TeamSpace 2 [com.snc.ppm_teamspace_2]	Installs a Project TeamSpace so that each team or each department like Marketing, Finance, IT-Team1, IT-Team2 can implement PPM suitable to their needs without overstepping each other.	Inactive	false	com.snc.financial_planning_pmo
Project Management TeamSpace 3 [com.snc.ppm_teamspace_3]	Installs a Project TeamSpace so that each team or each department like Marketing, Finance, IT-Team1, IT-Team2 can implement PPM suitable to their needs without overstepping each other.	Inactive	false	com.snc.financial_planning_pmo
Project Management TeamSpace 4 [com.snc.ppm_teamspace_4]	Installs a Project TeamSpace so that each team or each department like Marketing, Finance, IT-Team1, IT-Team2 can implement PPM suitable to their needs without overstepping each other.	Inactive	false	com.snc.financial_planning_pmo
Project Management TeamSpace 5	Installs a Project TeamSpace so that each team or each department like	Inactive	false	com.snc.financial_planning_pmo

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.ppm_teamspace_5]	Marketing, Finance, IT Team1, IT Team2 can implement PPM suitable to their needs without overstepping each other.			
Project Portfolio Suite [com.snc.project_portfolio_suite]	The Project Portfolio Suite (PPS) plugin activates an integrated set of applications for project portfolio management and IT software development. Only upgrade is allowed for this plugin. Activation should be done through PPM Standard plugin.	Inactive	true	• com.snc.project_management_v3 • com.snc.demand_management • com.snc.resource_management • com.snc.test_mgmt • com.snc.sdlc.scrum.ppm_int • com.snc.test_mgmt.ppm_int • com.snc.program_management
Project workspace [com.snc.project_workspace]				
PPM Standard [com.snc.financial_planning_pmo]	Enables you to manage your demands, resources, portfolios and projects, and gives full visibility from idea to execution. It also helps you plan, track, and manage cost and budget on projects and demands in a portfolio to strike a balance between investment and returns. Agile management and test management help you to improve productivity and service delivery.	Inactive	true	• com.snc.project_portfolio_suite • com.snc.financial_planning • com.snc.rate_model
PPM Standard - ATF Tests [com.snc.financial_planning_pmo.atf]	ATF tests for PPM Standard	Active	false	
Protocol Profile Manager	Defines properties associated to protocols such as	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.protocol_profile]	default port and keystore.			
Proxy Agent for connecting to Natural Language Understanding providers [com.glide.nlu.proxy]	The base NLU proxy agent for connecting to NLU providers. Intended for use by Virtual Agent and other clients.	Active	true	
Proxy Agent to the IBM Watson Natural Language Understanding server [com.glide.nlu.ibmwatson.intent.discovery]	Activates the IBM Watson Assistant Intent and Entity integration, which enables Virtual Agent to use intents, entities, and utterances defined in IBM Watson Assistant.	Inactive	false	com.glide.nlu.proxy
Proxy Agent to the Microsoft LUIS Natural Language Understanding server [com.glide.nlu.msluis.intent.discovery]	Proxy agent to the Microsoft LUIS Natural Language Understanding server.	Inactive	false	com.glide.nlu.intent.discovery
Proxy Agent to the ServiceNow Natural Language Understanding server [com.glide.nlu.intent.discovery]	Proxy agent for connecting to ServiceNow's NLU server for NLU intent discovery. Intended for use by Virtual Agent and other clients.	Active	false	com.glide.nlu.proxy
Purchase and Receipt Automation [com.snc.sn_pr] (Starting with the Tokyo release, Purchase and Receipt Automation is renamed as Sourcing and Purchasing Automation.)	Provides workflows and automation for sourcing requests, negotiations, and purchase requisitions.	Inactive	true	• com.glide.graphql • com.glideapp.user_criteria.scoped.api • sn_shop • sn_spend_sdc • sn_spend_psd • sn_spend_workspace

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				
Purchase Automation Integration with Risk Assessment [com.sn_spend_vrm] (Starting with the Utah release, Purchase Automation Integration with Risk Assessment is renamed as Risk Assessments Integration for Sourcing and Procurement Operations.) (Available in the ServiceNow Store)	Provides a set of capabilities to trigger risk assessments on a supplier during the sourcing or purchase requisition workflow using Vendor Risk Management.	Inactive	true	• com.snc.sn_pr • com.sn_vdr_risk_asmt
Purchase Modification Experience Pack for Procurement Case Management [com.sn_spend_cp] (Starting with the Utah release, Purchase Modification Experience Pack for Procurement Case Management is renamed as	Provides a set of pre-built playbooks, workflows, and experiences for employees, sourcing, and procurement to automate work that is typically managed through emails and spreadsheets.	Inactive	false	• com.snc.sn_shop • com.snc.sn_pr

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Playbooks for Sourcing and Procurement Operations.) (Available in the ServiceNow Store)				
Push Feedback [com.glide.push.feedback]	Collects feedback from the Apple push notification service or another feedback service. Provides a REST API for other instances to collect from this service.	Inactive	false	com.glide.push
Push Notification [com.glide.push]	Defines push notification message enqueuing and sending Responsible for processing push notifications to their next destination.	Active	false	
Push Retention [com.glide.push_retention]	Provides retention policy for push notifications.	Active	false	• com.glide.push • com.glide.auxdb
Qualys Vulnerability Integration [snc_vul_qualys] (Available in the ServiceNow Store)	Provides the ability to integrate the ServiceNow Vulnerability Response application with the Qualys vulnerability scanner. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	snc_vul
Query Rules [com.snc.query_rules]	Plugin that brings in the table and the logic needed to generate encoded queries per table for a user based on the roles the user has.	Inactive	false	
Query stats	Query stats.	Active	false	com.glide.stats

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.stats.query]				
Quick Actions [com.glide.quickactions]	Enables slash commands in Agent Workspace chat.	Active	false	
Quiz Designer [com.glide.quiz_designer]	Provides the ability to send scored questionnaires quickly and easily to one or more users. A quiz may have categories of questions that are assigned only to some users.	Inactive	false	• com.glide.assessment_designer.common • com.snc.assessment_core
Random Watermark Support [com.glide.email.random_watermark]	Generate unpredictable watermarks, and enable matching on these watermarks in inbound emails.	Active	false	com.glide.mailbox
Rapid7 Vulnerability Integration [sn_vul_t7] (Available in the ServiceNow Store)	Provides the ability to integrate the ServiceNow Vulnerability Response application with the Rapid7 Nexpose and InsightVM products. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.			sn_vul
Rate Limit for REST API [com.glide.rest.rate_limit]	Rate Limit support for REST API	Active	true	
Read only roles for Agile - Scaled Agile Framework [com.snc.sdlc.safe_read_roles]	Provides read-only roles for Agile - Scaled Agile Framework.	Active	false	com.snc.agile_read_roles
Read only roles for Agile Development 2.0	Provides the sn_agile_read role that has read-only access for all Agile	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.agile_read_roles]	Development 2.0 tables. The data from these tables can be used for generating reports.			
Read only roles for Application Portfolio Management [com.snc.apm_read_roles]	The plugin provides read only roles for Application Portfolio Management.	Active	false	
Read only roles for PPM Standard [com.snc.pmo_read_roles]	Read only roles for PPM Standard.	Active	false	
Read only roles for Release Management [com.snc.release_management_read_roles]	Provides the sn_release_read role that has read-only access for all Release Management tables. The data from these tables can be used for generating reports.	Active	false	
Read only roles for Test Management 1.0 [com.snc.tm1_read_roles]	Provides the sn_tm1_read role that has read-only access for all Test Management 1.0 tables.	Active	false	
Read only roles for Test Management 2.0 [com.snc.tm2_read_roles]	Provides the sn_tm2_read role that has read-only access for all Test Management 2.0 tables. The data from these tables can be used for generating reports.	Active	false	
Read-Only User Role [com.snc.read_only.role]	Enables Read-Only user role functionality.	Active	false	
Recommended Actions [sn_nb_action]	Recommends relevant actions, such as Guidances and Guided Decisions, for a given context.	Inactive	false	• sn_gd_guidance • com.snc.next_best_action • sn_nba_list

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)	Enables users to define contexts and strategies to recommend these actions.			
Recommended Actions Advanced [sn_nb_action_adv] (Available in the ServiceNow Store)	Enables customers to create recommendations based on AI/ML models created with Predictive Intelligence.	Inactive	false	com.snc.recommended_action
Recommended Actions for Customer Service [sn_cs_nb_action] (Available in the ServiceNow Store)	Installs the Recommended Actions capability, which recommends relevant actions for a given context. These actions appear on the Case form in the CSM Workspaces if users have defined strategies for the provided Case context.	Inactive	false	• com.sn_customerservice • com.snc.recommended_action
Redox Electronic Health Record Spoke [sn_redox_spoke] (Available in the ServiceNow Store)	An Integration Hub spoke that is used to facilitate Redox outbound calls.	Inactive	false	• com.glide.hub.integration.runtime • com.glide.hub.action_step.rest • com.glide.cobject
Redox Inbound Integration [sn_redox] (Available in the ServiceNow Store)	Improves inbound and proactive outbound service scalability and capabilities to integrate with healthcare systems that use the Redox platform.	Inactive	false	None
Relationship Layout	Enables scoped relationships to be associated to out-	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.ui.relationship_layout]	of-scope related list views.			
Release Management [com.snc.release_management_v2]	The Release Management v2 plugin is a rewrite of the original release management module. All products, releases, features, and release tasks are planned_task extensions, and much of the project management functionality (Gantt charts, timelines, time cards) is shared with release.	Active	true	• com.snc.release_management • com.snc.planned_task • com.snc.process_flow_formatter
Remote Tables [com.glide.script.vtable_demo]	Supports remote tables and associated script definitions that can be used to retrieve and optionally cache data from external sources.Tra			
Report - PDF Page Header Footer Templates [com.glideapp.report_page_hdrftr]	Allows the user to configure and store page headers and footers for PDF reports.	Active	false	• com.glideapp.report • com.glideapp.report2
Report Access Request [com.glideapp.report_access_request]	Provides the ability to request access to reports that have been restricted by report_view ACL.		false	
Report ACL Assessment [com.par_report_acl_assessment]				
Report Charting v2 [com.glideapp.report.charting_v2]	Installs V2 of ServiceNow charts with HighCharts.	Active	false	• com.glideapp.report • com.glideapp.report2
Report Engine - use summary table for reports	Causes the data from all reports, custom and standard, to be stored in the sys_report_summary	Active	false	com.glideapp.report2

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glideapp.summary_report_engine]	table and separates the data from the rendering process for all reports. Report data is periodically purged from the sys_report_summary table (approximately every two hours).			
Report Security - enforce access control checks [com.glideapp.report_security]	Enforces ACL checks when reports are created, deleted or updated.	Active	false	• com.glideapp.report • com.glideapp.report2
Reporting Statistics Reports [com.glideapp.report_statreports]	Provides reports and dashboards on reporting statistics.	Inactive	false	com.glideapp.report
Required Form Fields [com.snc.required_form_fields]	Allows an administrator to specify required fields that cannot be removed from a form.	Inactive	false	
Resource Management [com.snc.resource_management]	Enables resource requesters and resource managers to plan, organize, and manage resources for both planned and unexpected work. Activating Resource Management automatically activates the Project Management plugin if it is not already active. Only upgrade is allowed for this plugin. Activation should be done through PPM Standard plugin.	Inactive	true	• com.snc.schedule_loader • com.snc.process_flow_formatter • com.snc.cost_management • com.snc.pps_portal_common
Resource Matching Engine [com.snc.matching_rule]	Provides a Resource Matching Engine.	Inactive	false	• com.snc.skills_management • com.snc.agent_schedule

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Responsive Canvas [com.glideapp.canvas]	Enables responsive canvas. For dashboards, the responsive canvas dynamically responds to dashboard resizing and enables dragging to place and resize widgets. Easily find and preview widgets that you want to add directly from the canvas.	Active	false	• com.glideapp.home • com.glide.text_search • com.glide.db_images • com.glide.ui.angular • com.glide.ui.ng
Responsive Dashboards [com.glideapp.dashboards]	Easily create, modify, and share dashboards using responsive and dynamic widget layouts.	Active	false	• com.glideapp.canvas • com.glideapp.home • com.glide.text_search • com.glide.db_images • com.glide.ui.angular • com.glide.ui.ng • com.snc.pa
REST API Builder Backend [com.glide.rest.api_builder.backend]	REST API Builder Framework data model and crud interfaces.		false	com.glide.api_runtime
REST API for global text search [com.glide.globalsearch]	REST API for global text search.	Active	false	com.glide.rest
REST API Trigger Flow Designer [com.glide.hub.flow_trigger.rest]	Adds ability for flow designers to create a REST endpoint that triggers a flow upon request.		false	• com.glide.hub.flow_trigger • com.glide.api_runtime • com.glide.rest.api_builder.backend
REST API Provider [com.glide.rest]	Provides a REST API framework to support RESTful services.	Active	false	• com.glide.ui.angular • com.glide.ui.heisenberg
Restore Deleted Records [com.snc.undelete]	Restores deleted records from audited tables and references to those records. Also restores any records that were deleted as	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	a result of a cascade delete rule.			
Robotic Process Automation (RPA) Hub [com.sn_rpa_foundation] (Available in the ServiceNow Store)	Provides the capability to integrate the ServiceNow platform with applications that do not support APIs.	Active	true	sn_cmdb_ci_class (1.36.0)
RPA Sample Template [com.sn_rpa_template] (Available in the ServiceNow Store)	RPA templates are prebuilt automations that enable customers to kickstart their RPA initiatives. This template provides automating user password reset functionality for Oracle E-Business Suite (EBS) where APIs aren't available. It guides the ServiceNow developer to implement password reset functionality.	Active	false	com.sn_rpa_foundation (3.0.7)
Role Delegation [com.snc.role_delegation]	Allows an administrator to designate role delegators, who can delegate any role they have to members of their group.	Active	false	com.glideapp.workflow
S/MIME Email [com.glide.email.smime] (New in Tokyo)	Enables sending/ receiving signed and encrypted emails using S/MIME.	Active	false	• com.glide.certificates • com.glide.kmf
Sales Force Automation application template [com.snc.sfa2]	Provides tools to manage sales and marketing operations throughout the sales life cycle from lead generation through contract completion.	Inactive	true	com.snc.process_flow_formatter

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Scenario Planning for PPM [sn_pw_scenario] (Available in the ServiceNow Store)	Scenario Planning is a simple system to create and compare scenarios for a portfolio.	Active	true	com.snc.financial_planning_pmo
Schedule Optimization [com.snc_schedule_optimization] (New in Utah)	Agent schedules are automatically optimized daily to consolidate tasks and minimize travel time.	Inactive	false	• com.glide.platform_ml • com.snc.work_management • com.snc.app_fsm_sched_model
Schedule Optimization - Common [com.snc_schedule_optimization.common]				
Schedule Optimization - UI components [com.snc_schedule_optimization.ui]				
Schema Map v3 [com.glide.erd]	Displays the details of tables and their relationships in a visual manner, allowing administrators to view and easily access different parts of the database schema.	Active	false	com.glide.diagrammer
Scoped Analytics Framework [com.glide.scoped_analytics_framework]	Analytics Framework for Scoped Applications.	Active	false	
Scoped Application Restricted Caller Access [com.glide.scope.access.restricted_caller]	Allow scoped applications to restrict access to public tables and script includes.	Inactive	false	
Script Execution Context	Provides virtual closure mechanism.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.sn_context]				
Script stats [com.glide.stats.script]	Script stats.	Active	false	com.glide.stats.script
Script Templates [com.glide.script.templates]	Provides templates for some script fields.	Active	false	
Script Allowlist Manager [com.glide.script.whitelist]	Provides temporary support for continued direct invocation of Java Packages, Constructors, and Methods that are added to an allow list.	Active	false	
Scripted REST APIs [com.glide.scripted_rest_services]	Provides a framework for building Scripted REST APIs.	Active	false	• com.glide.rest • com.glide.scripted_rest_services.internal • com.glide.scripted_rest_services.errors
Scripted REST APIs - Error types [com.glide.scripted_rest_services.errors]	An internal plugin component to Scripted REST APIs.	Active	false	
Scripted REST APIs - Internal [com.glide.scripted_rest_services.internal]	An internal plugin component to Scripted REST APIs.	Active	false	
Scrum Program [com.snc.sdlc.scrum_program]	Plan and track the work across multiple Scrum teams that are working towards common objectives.	Inactive	true	• com.snc.program_common • com.snc.sdlc.agile.2.0
SCSS Bootstrap Theme [com.glide.ui.scss.bootstrap]	Theme assets for Bootstrap using SCSS.	Active	false	com.glide.ui.scss
SCSS Content Provider [com.glide.ui.scss]	SCSS Content Provider.	Active	false	
SDL Translation Management	It integrates Localization Framework with the	Active	false	com.glide.localization_framework

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
System Spoke for Localization Framework [com.glide.localization_framework.sdl_spoke]	SDL - Translation Management System. To activate this plugin, install the Localization Framework Installer plugin.			
SDLC - SCRUM [com.snc.sdlc.scrum]	Adds a release process specific to a Scrum development process. Sits on top of the SDLC application, adding additional agile notions like epics, stories, and sprints.	Inactive	true	com.snc.sdlc
Search Suggestions [com.glide.search.suggestions]	Provides suggestions when you enter text in a search field.	Active for new instances. Upgraded instances adopt the previous setting.		
Security Center [com.sn.vault_security_center]				
Security Dashboard [com.glide.security.dashboard]	Activates the Security Dashboard.	Inactive	false	com.snc.guided_setup
Security Incident Analytics [com.snc.security_incidentanalytics]	Provides an integration of Security Incident Response with Performance Analytics for trend based reporting. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	• com.snc.pa • com.snc.treemap • com.snc.security_incident
Security Incident Response [com.snc.security_incident] (Available in the	Provides visibility into the state of an organization's security using many of the same workflow and reporting capabilities ServiceNow is known	Inactive	true	• com.snc.service_management.core • com.snc.task_outage • com.snc.treemap • com.snc.secops.orchestration • com.snc.security_support.sir

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	for. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.			• com.snc.whtp • com.snc.threat
Security Incident Response Dependencies [com.snc.si_dep] (Available in the ServiceNow Store)	Installs all the dependent plugins required to support the Security Incident Response application.	Active	false	• com.snc.security_support.core • com.snc.security_support.sir • com.snc.runbook_automation.runtime • com.snc.task_outage • com.snc.treemap • com.snc.whtp • com.glide.scope.access.restricted_caller • com.sn_dependentclient
Security Incident Response Mobile [com.sn_sir_mobile] (Available in the ServiceNow Store)	As a security operations center (SOC) manager or a user with the ServiceNow AI Platform security analyst role (sn_si.analyst), you can log in to a ServiceNow AI Platform instance directly from your Android or iOS mobile device. With the Security Incident Response Mobile app, you can view, edit, and assign your most current and critical SIR security incidents and response tasks. Notifications inform you when critical security incidents assigned to you arrive.	Inactive		com.glide.sg.agent_native_client
Security Incident Response UI [com.app_secops_ui] (Available in the	Provides an enhanced user interface for monitoring and resolving threats to an organization's security. Activation of this plugin on production instances may require a	Inactive	true	com.snc.security_incident

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	separate license. Contact ServiceNow for details.			
Security Integration Framework [com.snc.sec.int.framework] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	
Security Jump Start (ACL Rules) [com.snc.system_security] (Available in the ServiceNow Store)	Adds ACL rules to provide a jump-start on securing many system tables, making it easier for an organization to get into production more quickly.	Active	false	
Security Operations 'Have I Been Pwned?' Integration [com.snc.secops.pwned] (Available in the ServiceNow Store)	Provides the ability to submit Whois lookups on domain names and URLs to obtain context on URL observables, and to make better determination on threats. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.threat • com.snc.security_incident
Security Operations ArcSight Logger Integration [com.snc.secops.arcsight.logger] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.security_incident

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Security Operations Carbon Black Integration [com.snc.secops.carbonblack] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.secops.orchestration • com.snc.security_incident
Security Operations CrowdStrike Host Integration [com.snc.secops.crowdstrike.host] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.security_incident
Security Operations CrowdStrike Intelligence Integration [com.snc.secops.crowdstrike.intel] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.threat • com.snc.threat.intelligence
Security Operations Elasticsearch Integration [com.snc.secops.elasticsearch] (Available in the ServiceNow Store)	Searches your Elasticsearch logs and adds relevant sighting information to your security incidents. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.security_incident
Security Operations McAfee ESM Integration [com.snc.secops.mcafee.esm]	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.security_incident

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				
Security Operations Microsoft Exchange Integration [com.snc.secops.ms-exchange] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.security_incident
Security Operations Palo Alto Networks - AutoFocus [com.snc.secops.paloalto.autofocus] (Available in the ServiceNow Store)	Provides the ability to integrate the ServiceNow Security Incident Response application with the Palo Alto Networks AutoFocus. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.security_incident • com.snc.secops.orchestration
Security Operations Palo Alto Networks - Firewall [com.snc.secops.paloalto] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.secops.orchestration • com.snc.security_incident
Security Operations Palo Alto Networks - WildFire [com.snc.secops.paloalto.wildfire] (Available in the	Provides the ability to integrate the ServiceNow Security Incident Response application with the Palo Alto Networks WildFire application. Activation of this plugin on production	Inactive	false	• com.snc.security_incident • com.snc.secops.orchestration

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	instances may require a separate license. Contact ServiceNow for details.			
Security Operations Splunk ES Event Ingestion Integration [com.snc.secops.splunkes] (Available in the ServiceNow Store)	This is the Splunk Enterprise Security Integration for Security Incident Response.			
Security Operations QRadar SIEM Integration [com.snc.secops.qradar.siem] (Available in the ServiceNow Store)	Provides the ability to integrate the ServiceNow Security Incident Response application with QRadar SIEM data. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.security_incident
Security Operations Splunk Integration [com.snc.secops.splunk] (Available in the ServiceNow Store)	Searches your Splunk logs and adds relevant sighting information to your security incidents. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.security_incident
Security Operations Spoke [com.snc.secops.spoke] (Available in the ServiceNow Store)	Provides the ability to use Security Operations flows and actions. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.security_incident • com.glide.hub

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Security Operations Tanium Integration [com.snc.secops.tanium] (Available in the ServiceNow Store)	Provides the ability to integrate the ServiceNow Security Incident Response application with Tanium data. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.security_incident • com.snc.secops.orchestration
Security Operations WHOIS Integration [com.snc.secops.whois] (Available in the ServiceNow Store)	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.threat
Security Support Common [com.snc.security_support.common] (Available in the ServiceNow Store)	Provides common functionality for use across the various security applications, such as Security Incident Response. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.security_support.core
Security Support Orchestration [com.snc.secops.orchestration]	Provides an integration of Security Operations with Orchestration to allow the facilitation of workflow activities within Security Incident Response, Threat Intelligence, or Vulnerability Response. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	• com.snc.runbook_automation.runtime • com.snc.security_support.common

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Self-Service Analytics Core [com.snc.self_service_analytics_core]	Provides the self-service analytics framework for configuring deflection contexts and activity patterns to collect the case reduction (deflection) metrics.	Active	false	com.snc.activity_subscriptions
Self-Service Analytics for Customer Service [com.snc.pa.self_service_analytics_csm]	Tracks case reduction (deflection) metrics and self-service KPIs using performance analytics dashboards. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Active	false	• com.sn_customerservice • com.snc.pa.self_service_analytics • com.snc.self_service_analytics_core
Self-Service Analytics PA [com.snc.pa.self_service_analytics]	Enables users to define and track key performance indicators (KPIs) to measure the effectiveness of self-service channels such as knowledge, community, virtual agent, and service catalog. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Active	false	com.snc.pa
Self Service Password Reset [com.snc.password_reset]	Allows locally authenticated users to request a temporary password if they forget their current password.	Active	false	com.glideapp.password_reset
Self-Service Portal for Analytics [com.snc.pa.bi_service]	Self-Service Portal for Analytics - Allows users to request services related to dashboard access.	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Sensitive Data Handling [com.glide.sensitive_data_handling]				
Sentiment Analysis [com.snc.sentiment_analysis]	This plugin is used for Sentiment Analysis	Inactive	false	com.snc.core.automation.connection_credenti
Server-side JavaScript Debugger [com.glide.debugger]	Allows application developers and administrators to efficiently debug scripts that drive the applications they develop and support.	Active	false	• com.glide.ui.angular • com.glide.ui.zepto • com.glide.ui.font_icons
Service Catalog - ATF Tests [com.glideapp.servicecatalog.atf.test]	Provides Service Catalog ATF Tests	Active	true	
Service Catalog - Channel Source Analytics [com.glideapp.servicecatalog.analytics]	Enables the channel source analytics for Catalog Items.	Active	false	com.glideapp.servicecatalog
Service Catalog - Domain Separation [com.glideapp.servicecatalog.domain_separation]	Enables service providers to separate catalog items for different domains			
Service Catalog - Workspace [com.glideapp.servicecatalog.workspace]	Service Catalog - Workspace	false		
Service Catalog Builder [com.glideapp.servicecatalog.catalog_builder]	Catalog builder for easy creation and maintenance of catalog items. Installation of this plugin allows creation and editing of catalog items as draft and each change needs to be published.	Active	true	• com.glideapp.servicecatalog.platform • com.glideapp.servicecatalog.composite_rec • com.glideapp.servicecatalog.wizard • com.glideapp.servicecatalog.catalog_templa • com.glideapp.servicecatalog.catalog_builde

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Service Catalog Builder Experience [com.glideapp.servicecatalog.catalog_builder_experience]	This plugin is for catalog builder experience in UXF.		false	• com.sn_canvas_blank • com.sn_canvas_chrome
Service Catalog Builder internal components [com.glideapp.servicecatalog.builder.internal_components]	Internal components that make up Catalog Builder. Install this plugin to edit catalog builder experience.		false	
Service Catalog CMS Extension [com.glideapp.servicecatalog.cms]	Provides the ability to define the catalog experience within CMS.	Active	false	• com.glideapp.servicecatalog.platform • com.glide.cms
Service Catalog Composite Record Producer [com.glideapp.servicecatalog.composite_record_producer]	Composite Record Producer allows records to be created and edited in multiple related tables without needing scripts.		false	com.glideapp.servicecatalog.platform
Service Catalog core applications [com.glideapp.servicecatalog]	Allows customers to order predefined, bundled goods and services from your IT organization or other departments.	Active	true	• com.glideapp.servicecatalog.platform • com.glideapp.servicecatalog.execution_plar • com.glideapp.servicecatalog.currency
Service Catalog Macroponent internal components [com.glideapp.servicecatalog.macroponent.internal_components]				
Service Catalog Manager [com.snc.sc_catalog_manager]	Provides utility for managing categories, catalog items, and Knowledge Base links for Service Catalogs.	Inactive	false	
Service Catalog Mobile Employee [com.glideapp.servicecatalog.mobile-employee]	This plugin contains the configuration, records, and catalog webviews that are used in Mobile Employee Experience native application.	Inactive		

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Note: Do not activate this plugin directly. Activating com.glide.mobile-employee activates this plugin for you.			
Service Catalog My Request for Mobile [com.glideapp.servicecatalog.mobile-request-filter]	This plugin contains the configuration, records that are used in Mobile native application for My Request. Do not activate this plugin directly. Activating com.glide.mobile-employee will activate this plugin for you.	Inactive	false	
Service Catalog REST API [com.glideapp.servicecatalog.rest.api]	Service Catalog REST API.	Active	false	com.glideapp.servicecatalog.scoped.api
Service Catalog Scoped API [com.glideapp.servicecatalog.scoped.api]	Service Catalog Scoped API to support application creation on Service Catalog platform.	Active	true	com.glideapp.servicecatalog
Service Catalog Template [com.glideapp.servicecatalog.catalog_template]	Framework to create and maintain catalog templates.		false	com.glideapp.servicecatalog.platform
Service Catalog Wizard [com.glideapp.servicecatalog.wizard]	Activating this plugin creates the data model for wizards.		false	com.glideapp.servicecatalog.platform
Service Creator [com.glide.service_creator]	Enables a department to offer custom services through the service catalog, such as the HR department offering tuition reimbursement for further education.	Active	true	• com.glide.ui.ng.cc • com.glideapp.servicecatalog.item_designer

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Service Credits [com.sn_service_credits]	Vendor service credits.		false	• com.snc.service_portfolio • com.snc.service_portfolio.sla_commitment • com.snc.task_outage
Service Delivery Common [com.sn_spend_sdc] (Starting with the Utah release, Service Delivery Common is renamed as Common Service Delivery.) (Available in the ServiceNow Store)	Contains Service Task and Service Request tables, as well as other infrastructure that forms the basis of Finance and Supply Chain Workflows products.	Inactive	false	• com.glide.pad.core • com.glide.pad.license • com.playbook_experience
Service Fulfillment Steps [com.glideapp.servicecatalog.service_fulfillment_steps]	Adds support for Service Fulfillment using Data Driven Workflows.	Active	false	• com.glideapp.servicecatalog.platform • com.glideapp.servicecatalog.catalog_builder
Service Graph Connector for Microsoft Defender for IoT (Azure) [com.sn_msftd4iotazsgc] (Available in the ServiceNow Store)	Integrates Microsoft Defender for IoT (Azure) with the ServiceNow Operational Technology Manager application to import devices and sensor appliances.	Active	False	• com.sn_cmdb_ci_class • com.snc.cmdb.integration_util • com.glide.hub.action_type.datastream • com.snc.itom.discovery.license • com.snc.itom.license
Service Graph Connector Integration for Claroty CTD [com.sn_clarotyctdsgc] (Available in the	Integrates Claroty CTD with the ServiceNow Operational Technology Manager application to import sites, detected	Active	False	• com.sn_cmdb_ci_class • com.snc.cmdb.integration_util • com.glide.hub.action_type.datastream • com.snc.itom.discovery.license • com.snc.itom.license

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	devices by each site, connection (or baselines), and installed programs.			
Service Graph Connector for Operational Technology [com.sn_otsm_sgc] (Available in the ServiceNow Store)	Enables you to import your existing Operational Technology data from a populated Microsoft Excel flat-file spreadsheet. You use it in the Integration Hub Extract Transform Load (ETL) to upload this data to the Configuration Management Database (CMDB).	Active	false	• com.sn_cmdb_ci_class • com.snc.cmdb.integration_util
Service Integration and Management [com.snc.siam_core]	SIAM Core plugin that provides support for Global scope access, required by applications like VMW and Service Onboarding.	Inactive	false	
Service Level Management [com.snc.sla]	Provides the core SLA functionality. SLA Definitions provide conditions to start, pause, stop, cancel and reset Task SLAs against any Task type. In addition, you can specify a schedule on the definition to define the working hours and also a workflow to run against each Task SLA, which is typically used to generate notifications.	Active	true	• com.glideapp.workflow • com.glide.schedules • com.glide.relative_duration
Service Level Management - Contract Management Integration [com.snc.sla.contract2]	Extends the existing SLA functionality, by utilizing a contract as the primary document that houses all appropriate data needed to drive the SLA processing of your task.			

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	The attach_sla field is added.			
Service Level Management Dashboard [com.snc.sla.overview]	This plugin provides the My SLAs Dashboard for Service Level Management	Active	false	• com.glideapp.canvas • com.snc.sla
Service Level Management Guided Tour [com.snc.sla.guided_tour]	Provides the Guided tour of SLA functionality.	Active	false	
Service Level Management PA Dashboard [com.snc.pa.sla.overview]	This plugin provides the SLA Overview (Premium) Dashboard. Activation of this plugin on production instances may require a separate Performance Analytics license. Contact ServiceNow for details.	Active	false	• com.glideapp.canvas • com.snc.pa • com.snc.sla • com.snc.sla.breakdowns
Service Operations Workspace Service Dashboard [com.sn_itom_service_dashboard] (Available in the ServiceNow Store)	Monitors the health of services in a dashboard.		false	Service Operations Workspace ITOM Applications
Service Management Geolocation [com.snc.service_management.geolocation]	Provides Service Management geolocation capabilities.	Inactive	true	com.snc.geolocation
Service Management Geolocation Mobile [com.snc.service_management_m]	Adds a menu in the new mobile UI for Service Management Geolocation.	Inactive	false	• com.glide.ui.m • com.snc.service_management.geolocation
Service Management	Provides core functionality for	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Virtual Agent Core [com.glideapp.sm_va_core]	Service Management Virtual Agent			
Service Management Virtual Agent Topic Blocks [com.glideapp.cs.sm_topic_blocks]	This plugin contains prebuilt re-usable conversation topic blocks for common actions like ordering from the catalog and searching the knowledge base.	Inactive	false	• com.glide.cs.chatbot • com.snc.contextual_search • com.glideapp.servicecatalog.platform
Service Mapping Plus [sn_sm_scoped_app] (Available in the ServiceNow Store)	Offers the service mapping functionality enhanced with Machine Learning.	Inactive	false	com.snc.service-mapping
Service Modeling [com.snc.cmdb.service_modeling]	Core infrastructure for Service Modeling used in Service Mapping and Delivery.	Inactive	false	• com.glide.vars • com.snc.cmdb.enterprise
Service Portal Analytics [com.glide.service-portal.analytics]	User Experience Analytics for Service Portal provides dashboard views for monitoring the key performance indicators (KPIs) of web applications built on Service Portal. You can use these insights to optimize your portal.	Active	false	com.glide.appsee
Service Portfolio Management Foundation Demo Data [com.snc.service_portfolio.demo_data]	Adds Demo Data for Service Portfolio Management Foundation and SLA Commitments. If you installed the SPM SLA Commitments plugin before New York, this activates the newer version.	Active	true	• com.snc.service_portfolio • com.snc.service_portfolio.sla_commitment
Service Request Criteria				

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.sn_req_criteria]				
ServiceNow Add-Ins for Microsoft Office [com.sn_outlook_addin]	Enables users to interact with ServiceNow from within Microsoft Office	Inactive	false	• com.snc.platform.security.oauth • com.glide.ui.vtb • com.glideapp.servicecatalog • com.glide.service-poortal
ServiceNow Application Repository [com.snc.apprepo]				
ServiceNow Antivirus Program [com.glide.snap]	ServiceNow Antivirus Program	Active	false	
ServiceNow Certificate Inventory and Management [com.sn_disco_certmgmt]	Automatically track certificates, provide awareness of upcoming expirations, and drive the workflow to renew them.			
ServiceNow Cloud Encryption [com.glide.platform.cloud_encryption]				
ServiceNow Document Viewer [com.snc.documentviewer]	Document viewer is a platform feature that will enable users to view enterprise class documents inline within the platform attachment instead of downloading it to the local device and then opening the documents with a locally installed viewer.	Active	false	com.snc.apppdfgenerator
ServiceNow Flow Designer - Dynamic Inputs [com.glide.hub.dynamic_inputs]	Enable flow designers to configure an action using input values that have been dynamically created. During flow design,	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	when you add an action with dynamic inputs to a flow, it calls a data gathering action to collect data and dynamically shows choices based on that returned data. Use dynamic inputs to look up and display choices dynamically rather than using hard-coded options. For example, you can use a dynamic input to generate choices for a spoke integration.			
ServiceNow Flow Designer - Dynamic Outputs [com.glide.hub.dynamic_outputs]	Enable flow designers to configure an action using output values that have been dynamically created. During flow design, when you add an action with dynamic outputs to a flow, it calls a data gathering action to create a complex object. Use dynamic outputs to create a complex object with a dynamic structure rather than creating a hard-coded object with a static structure. For example, you can use a dynamic output to generate a complex object for a spoke integration.	Inactive	false	
ServiceNow Flow Designer - Introspection [com.glide.hub.flow_designer_introspection]	Enables the use of Dynamic Inputs and Dynamic Outputs for Actions, Subflows and Flows in Flow Designer.	Inactive	false	• com.glide.hub.dynamic_inputs • com.glide.hub.dynamic_outputs
ServiceNow IntegrationHub Action Step - MFT	Provides capabilities to copy directory using the SFTP step.			

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.hub.action_step.mft]				
ServiceNow IntegrationHub Action Template - Data Streaming [com.glide.hub.action_type.datastream]	Activates Data Stream actions in Flow Designer, enabling you to send REST or SOAP requests from Flow Designer to APIs that return a stream of response data larger than 10 MB. Parse stream data into a series of complex object outputs and use the data pills in other actions in a flow.	Inactive	false	com.glide.hub.action_step.template
ServiceNow IntegrationHub Action Step - RTE [com.glide.hub.action_step.rte]	Action Step - RTE.		false	com.glide.hub.action_step.template
ServiceNow IntegrationHub Action Step - SFTP [com.glide.hub.action_step.sftp]	Provides an SSH File Transfer Protocol step to create a reusable action that can access, transfer, or manage files.	Inactive	false	com.glide.hub.action_step.template
ServiceNow IntegrationHub Enterprise Pack Installer [com.glide.hub.integrations.enterprise]	Installs IntegrationHub enterprise pack to automate human resources, customer relationship management, enterprise resource planning, and more. Includes IntegrationHub professional pack, Microsoft SCCM spoke, and Data Stream actions.		false	• com.glide.hub.integrations.professional • com.sn.sccm.spoke
ServiceNow IntegrationHub Professional Pack Installer	Installs IntegrationHub professional pack plugins to automate IT operations. Includes		false	• com.glide.hub.integrations.standard • com.glide.hub.action_step.powershell • com.sn.ad.spoke

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.hub.integrations.professional]	IntegrationHub standard pack, Microsoft AD spoke, Microsoft Azure AD spoke, and SSH and Powershell steps.			• com.sn.azure_ad.spoke • com.glide.hub.action_step.ssh
ServiceNow IntegrationHub Standard Pack Installer [com.glide.hub.integrations.standard]	Installs IntegrationHub standard pack plugins to automate developer operations. Includes IntegrationHub starter pack, and JDBC and XML parser steps.			
ServiceNow IntegrationHub Starter Pack Installer [com.glide.hub.integrations.starter]	Installs IntegrationHub starter pack plugins to build your own integrations. Includes HipChat spoke, Microsoft Teams spoke, Slack spoke, Slack webhooks spoke, eBonding spoke, legacy IntegrationHub usage dashboard, and REST and SOAP steps.			
ServiceNow ITOM Licensing [com.sn_itom_licensing]				
ServiceNow Language Detection Service Spoke [com.glide.language_detection_spoke]	Using ServiceNow Language Detection service, it adds ability to detect the language of given text.	Active	false	• com.glide.hub • com.glide.hub.integration.runtime
ServiceNow Mobile Request - My Request Filter [com.glideapp.servicecatalog.mobile-request-filter]	This plugin contains the configuration and records that are used in the Now Mobile app. Do not activate this plugin directly. Activating com.glide.mobile-employee will activate this plugin for you.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Subscription Management [com.snc.usage_admin.snc]	ServiceNow Subscription Management.	Active	false	com.snc.usage_admin.base
Service Owner Workspace (Maintenance mode only) [com.spm_owner_workspace]	Provides a premium Service Portfolio Management experience. Portfolio managers and service owners access an integrated and graphically intuitive user interface to manage and monitor portfolios and services.	Inactive	false	• com.snc.pa.premium • com.snc.spm • com.snc.spm.spend • com.snc.service_workspace
Service Portal Agent Chat [com.glide.service-portal.agent-chat]	This plugin brings Agent Chat to Service Portal.	Inactive	false	• com.glide.cs • com.glide.cs.sn-va-web-client-app
Service Portal - Core [com.glide.service-portal]	Service Portal is a portal framework that allows administrators to build a mobile-friendly self service experience for users.	Active	false	• com.glide.ui.web_db_types • com.glide.ui.angularui • com.glide.ui.ng.filter • com.glide.ui.spectrum • com.glide.ui.m • com.glide.connect • com.glide.ui.scss • com.glide.ui.scss.bootstrap • com.glide.service-portal.designer • com.glide.automated_testing_impl.service_p
Service Portal - Knowledge Base [com.glide.service-portal.knowledge-base]	Knowledge base for SP	Active	false	
Service Portal La Jolla Brand Update	Updates your /sp and /sp_config portals with the ServiceNow La Jolla branding. The	Active	false	com.glide.service-portal

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.service-portal.themes.la-jolla]	logo, favicon, login page background image, hero banner, theme, and widget instance options are updated when the portal record has not been modified. When the portal record is modified no changes are made and an administrator can manually apply changes by downloading the La Jolla images from the db_image table and applying as well as updating the theme setting of the portal record.			
Service Portal - Service Catalog [com.glide.service-portal.service-catalog]	Enables Service Catalog widgets for Service Portal.	Active	false	No
Service Portal - Service Catalog v2 [com.glideapp.servicecatalog.portal]	Enables Service Catalog widgets for Service Portal.	Active	false	No
Service Organization [com.snc.service_organization]	Provides the base framework to support any internal or external entity involved in the customer service value chain. E.g. Branches, Stores, Franchises, Institutions.	Active	false	
Service Portal - Service Status [com.glide.service-portal.service-status]	Service Status gives information on current, planned, and historical outages for Business Services so they do not have to call your Service Desk. Loading demo data	Active	true	com.glide.service-portal

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	randomly generates 90 days of outages for 19 demo data Business Services.			
Service Portal Announcements [com.glide.service-portal.announcements]	Enables administrators to broadcast announcements to Service Portal users.	Active	true	
Service Portal Configuration Pages [com.glide.service-portal.config]	Service Portal Configuration Pages.	Active	false	com.glide.service-portal
Service Portal Designer [com.glide.service-portal.designer]	Drag-and-drop wysiwyg portal designer	Active	false	com.glide.service-portal
Service Portal for Enterprise Service Management [com.glide.service-portal.esm]	Provides a default service portal with easy configuration and designer. End-user resources include Knowledge Base, Service Catalog, and Services Status.	Active	true	• com.glide.service-portal • com.glide.service-portal.service-status • com.glide.service-portal.sqanda
Service Portal La Jolla Brand Update [com.glide.service-portal.themes.la-jolla]	This plugin updates your /sp and /sp_config portals with the ServiceNow La Jolla branding. If the portal record has not been modified, the logo, favicon, login page background image, hero banner, theme, and widget instance options are updated.			
Service Portal Surveys [com.glide.service-portal.survey]	Service Portal Surveys and Assessments.	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Service Portal User Criteria Support [com.glide.service-portal.user-criteria]	Enables Service Portal User Criteria support. Installing this plugin will create new user criteria records for each of your widgets and pages that already have role-based permissions on them.	Inactive	false	• com.glide.service-portal • com.glideapp.user_criteria
Service Portfolio Management Foundation [com.snc.service_portfolio]	Allows an organization to document the business services it provides using a standardized, structured format. Performance against availability commitments is calculated and can be displayed in a homepage.	Inactive	true	• com.snc.service_portfolio_core • com.glideapp.servicecatalog • com.glideapp.summary_report_engine
Service Portfolio Management Core [com.snc.service_portfolio_core]	This plugin contains core functionality for Service Portfolio Management available out of the box by default.	Inactive	false	com.snc.cmdb
Service Portfolio Management Premium [com.snc.spm]	Collects service offering metrics and allows for roll-up calculations to parent services and taxonomy nodes for performance scores and other metrics as viewed via the Service Owner Workspace plugin.	Inactive	false	• com.snc.pa.premium • com.snc.service_portfolio • com.snc.service_portfolio.sla_commitment • com.snc.task_outage
Service Portfolio Management Estimated Spend [com.snc.spm.spend]	Collects service offering cost and allows for selection between local model and Financial Management model.	Inactive	false	com.snc.spm
Service Portfolio Management SLA Commitments	Allows commitments to be defined by an SLA, so that staff can track how efficiently the service desk	Active	true	• com.snc.service_portfolio • com.snc.sla

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.snc.service_portfolio.sla_commitment]	meets commitments for a service offering.			
Service Workspace [com.snc.service_workspace]	Service Workspace application.	Inactive	false	com.workspace_core
SessionDebug [com.glide.sessiondebug]	Provides SessionDebug statements and filtering.	Active	false	
Shift Planning (Maintenance mode only) [com.sn_shift_planning]	Shift Planning plugin. Plan team schedules, shifts and workflows for managing timeoffs and shift swaps requests.	Active	false	
Shift Planning [com.sn_shift_planning_host]				
Shodan Exploit Integration [sn_vul_shodan] (Available in the ServiceNow Store)	Provides the ability to integrate the ServiceNow Vulnerability Response application with the Shodan product. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.		false	sn_vul
ShoppingHub [com.snc.sn_shop] (Starting with the Tokyo release, ShoppingHub is renamed as Procurement Common Architecture.) (Starting with the Utah release, Procurement Common Architecture	Provides an architecture to store purchase orders, requisitions, sourcing requests, and other objects that are commonly used across the source-to-pay business processes.	Inactive	true	• com.sn_fin • com.sn_spend_sdc

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
is renamed as Source-to-Pay Common Architecture.) (Available in the ServiceNow Store)				
Shopping Hub [com.snc.sn_spend_uib] (Available in the ServiceNow Store)	Provides a streamlined, e-commerce like experience for employees to self-service requests, and source or purchase products and services through the procurement organization. This is an add-on experience to Employee Center.	Inactive	false	• com.sn_pr • com.sn_cmn_uib_comp
ShoppingHub Mobile [com.sn_shop_mobile] (Starting with the Utah release, ShoppingHub Mobile is renamed as Shopping Hub Mobile.) (Available in the ServiceNow Store)	Enables Shopping Hub and other employee experiences for procurement to be enabled within the Now Mobile application.	Inactive	false	• com.glide.sg • com.snc.sn_pr • com.glide.sg.agent_native_client
Sidebar Server [com.glide.cs.collab] (New in Tokyo)	A server for the Sidebar feature.	active	false	• com.glide.cs • com.glide.cs.custom.adapter • com.glide.quickactions
Signature Pad [com.snc.signaturepad]	Provides a tool to allow a digital signature in a .pdf document. The Human Resources application	Inactive	false	com.glide.db_images

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	uses this with various documents. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.			
Sitemap Generator [sn_ux_seo_sitemap] (Available in the ServiceNow Store)	Helps you to define and automatically generate XML sitemaps to improve search engine optimization of your public portal pages.	Inactive	false	• com.glide.scope.access.restricted_caller • com.glide.hub.integration.runtime • com.glide.hub.action_step.rest • com.glide.hub.action_step.script
Skill Determination [com.snc.skill_determination]	Automatically determine and assign skills for work items.	Active	false	com.snc.skills_management
Skills Management [com.snc.skills_management]	The Skills Management application enables you to organize skills data of service agents and employees in a centralized location. You can use this data to assign tasks based on skills, analyze skill gaps to identify coaching and training needs, optimize your workforce, create career growth plans, and perform skill-based hiring.	Inactive	true	
Skill Recommendation [com.snc.sre] (Maintenance mode only)	Skill Recommendation ML prediction model and workflow to recommend and assign skills to agents	Active	false	• com.snc.skills_management • com.glide.platform_ml
Skills Library Data for Skills Management [com.snc.skills_management.seed_data]	Provides a library of skill categories and skills that you can start using for quick on-boarding.	Active	true	com.snc.skills_management

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
SLA Breakdowns [com.snc.sla.breakdowns]	Provides the ability to generate breakdown data for each Task SLA record. For example, you can generate breakdown data on the basis of the **Assignment group** and **Assigned to** fields. **Note:** You need to have up-to-date versions of script includes *`TaskSLAController`* and *`RepairTaskSLAController`* before activating this plugin to ensure that breakdown data is generated correctly. If you have customized versions of the script includes *`TaskSLAController`* and *`RepairTaskSLAController`*, you need to incorporate all customizations into the versions of these files from the most recent upgrade.	Active	false	com.snc.sla
SLA Timeline [com.snc.sla.timeline]	Provides the ability to view an SLA in a timeline.	Active	false	
SM Planned Maintenance [com.snc.planned_maintenance]	Allows setup and configuration for repeating and triggered requests. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.service_management.core

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
SMS Preferences [com.snc.sms_pref]	Provides ability to set SMS preferences for receiving messages from different providers.	Inactive	false	
sn-custom-tinymce-scoped-app [com.sn-custom-tinymce-scoped-app]	Scoped app providing custom tinymce build	Inactive	false	
sn-openframe-uxb [com.devsnc_sn_openframe_uxb]	Openframe component.	Active	false	• com.glide.uxbuilder • com.servicenow_behavior_key_binding • com.servicenow_sass_generic • com.servicenow_sass_global • com.servicenow_now_legacy_icon • com.servicenow_ui_core • com.servicenow_ui_effect_graphql • com.servicenow_ui_effect_amb • com.servicenow_ui_effect_http • com.servicenow_ui_effect_update_state • com.servicenow_ui_renderer_snabbdom • com.sn_seismic_post_message • com.sn_uxpage_presource • com.sn_translate • com.devsnc_library_enhanced_test • com.devsnc_library_uxf • com.servicenow_behavior_rtl
SNC Code Editor [com.glide.snc_code_editor]		Active	false	
sn-process-automation-designer [com.sn_process_automation_designer]	This plugin is the Process Automation Designer application	Active	false	• com.glide.uxbuilder • com.devsnc_sn_glide_form_controls • com.devsnc_sn_record_input • com.servicenow_now_alert • com.servicenow_now_button • com.servicenow_now_card

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
				• com.servicenow_now_dropdown • com.servicenow_now_heading • com.servicenow_now_icon • com.servicenow_now_label_value • com.servicenow_now_modal • com.servicenow_now_record_list_connecte • com.servicenow_now_text_link • com.servicenow_now_radio_buttons • com.servicenow_now_visual_board • com.servicenow_sass_generic • com.servicenow_sass_global • com.servicenow_ui_core • com.servicenow_ui_effect_graphql • com.servicenow_ui_effect_helpers • com.servicenow_ui_effect_http • com.servicenow_ui_enzyme_adapter • com.servicenow_ui_renderer_snabbdom • com.sn_component_error_handler • com.sn_pd_picker • com.sn_uxpage_presource
sn-ui-builder [com.sn_ui_builder]	UI Builder is an enjoyable, easy to use builder for creating UIs at ServiceNow.	Active	false	• com.glide.uxbuilder • com.servicenow_behavior_focus • @servicenow/now-alert • @servicenow/now-modal • @servicenow/now-text-link • sn-devx-components • com.sn_http_request • com.sn_translate • sn-ui-builder-components • sn-uxf-builder-middleware
Software Asset Management Machine Learning Normalization [com.sn_sam_ml_normalization]	Provides core capabilities to normalize unrecognized discovered software using machine learning. Requires		false	• com.snc.samp • com.glide.platform_ml

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	Software Asset Management Professional.			
Software Asset Workspace [com.sn_sam_workspace]	Enables the Software Asset Management (SAM) capabilities in Workspace.		false	• com.snc.app_shell_aw • com.snc.uib.base_agent_workspace • com.sn_itam_card • com.snc.sams
Software Development Lifecycle (SDLC) [com.snc.sdlc]	Extends the Release Management v2 plugin by adding some new structures to accommodate the software development life cycle. This plugin is designed to accommodate most non-agile development methodologies, including the common waterfall method of development.	Inactive	true	com.snc.release_management_v2
Source-to-Pay Workspace [com.sn_spend_workspace] (Available in the ServiceNow Store)	Provides a single environment for Procurement Specialists to work on purchase requisitions, sourcing requests, negotiations, procurement requests, and more.	Inactive	false	com.snc.uib.sn_dyn_rel_rec
Special Handling Notes [com.sn_shn]	Special Handling Notes.	Inactive	false	
Special Handling Notes Demo Data [com.snc.shn_demo]	Special Handling Notes application demo data.	Inactive	false	• com.sn_shn • com.snc.customerservice.demo
State Flows [com.snc.state_flows]	Enables advanced users to customize the state flow of any task table that uses states.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
State Transition Models [com.snc.itsm.state_transition_model]	State Transition Models allow flexible configuration of state and transition models.	Active	false	
State Transition Models - Foundation Data [com.snc.itsm.state_transition_model.foundation]	Foundation data for State Transition Models, including State transition conditions.	Active	false	com.snc.itsm.state_transition_model
Stats Tools [com.glide.db.query_stats]	Records statistics for system activities that affect performance. Query, Script, and Transaction stats, database performance tuning tools.	Active	false	com.glide.monitor.round_robin_database
Strategic Planning [sn_apw_advanced]	Helps to prioritize all work, align goals, visually roadmap, and track progress when using Agile, waterfall, or hybrid approaches.			
Strategic Spend Tracking for PPM [sn_ppm_sst] (Available in the ServiceNow Store)	Evaluate strategic value of your projects and demands.	Active	True	com.snc.financial_planning_pmo
Studio [com.glide.dev-studio]	Allows developers to add and update application files.	Active	false	• com.glide.code-search • com.sn_appcreator
Subscription Administration Base [com.snc.usage_admin.base]	Subscription Administration Base.	Active	false	com.glide.subscription_framework

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Subscription Based Notifications [com.glide.notification]	Allows users to subscribe to notifications on a task or CI without being on the watchlist or being one of the assigned users.	Active	false	• com.glide.ui_policy • com.glide.email_notification_preview
Subscription Based Notifications 2.0 [com.glide.notification.subscription]	The notification subscription model - improved and simplified.	Active	false	com.glide.notification
Subscription Management and Enforcement Framework [com.glide.subscription_framework]	Subscription Management and Enforcement Framework.	Active	false	com.glide.usageanalytics
Subscriptions and Activity Feed Framework [com.snc.activity_subscriptions]	This plugin provides a generic set of artifacts to handle subscriptions for any defined subscribable object. Any entity can be defined as a subscribable object and a set of subscribers can subscribe to the objects. When a event occurs related to the subscribable object, activities can be tracked and subscribers can be notified.	Inactive	true	
Success Indicators [com.snc.vendor.insights]	Leverages odds ratio and ML capabilities for apps like Vendor Manager Workspace, etc.	Active	false	
Supplier Collaboration Portal [com.snc.sn_supplier_sp]	Provides a single, one-stop experience for suppliers to get self-service, complete tasks and make requests into an organization.	Inactive	true	• Supplier Common Architecture (com.snc.sn_slm) • Service Portal (com.glide.service-portal)

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
(Available in the ServiceNow Store)				• E-signature (com.snc.esign) • Employee Center (com.snc.employee_center)
Supplier Common Architecture [com.snc.sn_slm] (Available in the ServiceNow Store)	Provides a common architecture to track data objects related to a supplier used in both Supplier Lifecycle Operations and Supplier Collaboration Portal.	Inactive	false	• Finance Common Architecture (com.sn_fin) • Service Delivery Common (com.sn_spend_sdc) • Document Management (com.snc.platform_document_management) • External User Registration (com.snc.external_user_self_registration) • Common Vendor Core (com.snc.sn_vendor_core) and (com.snc.vendor_core)
Supplier Lifecycle Operations [com.snc.sn_supplier_mgmt] (Available in the ServiceNow Store)	Collaborate with suppliers, manage supplier relationships, monitor risk, compliance, and performance across the supplier life cycle.	Inactive	true	• Supplier Common Architecture (com.snc.sn_slm) • Playbook Experience (com.playbook_experience) • Employee Experience Foundation (com.snc.sn_ex_emp_fd) • News Integration for Supplier Lifecycle Operations (com.snc.sn_supplier_news) • Map UI Component for threat and alert data feeds (com.sn_fam_map)
Supplier Lifecycle Operations Integration with SAP [sn_slm_sap_int] (Available in the ServiceNow Store)	Send suppliers created in Supplier Lifecycle Operations to SAP ECC and SAP S4 HANA.	Inactive	false	• Finance Common Architecture (com.sn_fin) • Finance – ERP Integration (com.sn_fcms_integrations) • Primary Data Integration with SAP (sn_sap_data_int)
Survey Designer	Activates survey designer that is	Active	false	com.glide.assessment_designer.common

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.survey_designer]	a drag-and-drop interface to create surveys.			
Syntax Editor [com.glide.syntax_editor]		Active	false	com.glide.snc_code_editor
System Applications Core [com.snc.apps]	Core applications development.	Active	false	com.snc.apps_picker
System Applications Support [com.snc.apps_complete]	Provides support for creating and managing applications.	Active	false	• com.snc.apps • com.snc.apps_picker • com.snc.apps_access • com.snc.apps_file • com.snc.apps_hub • com.snc.apps_legacy • com.snc.metadata • com.snc.metadata_tree • com.snc.apps_creator • com.glide.sessiondebug • com.glide.autorecovery • com.glide.transaction_scope • com.glide.ui.mergetool • com.glide.ui.relationship_layout
System Import Sets [com.glide.system_import_set]	Provides the functionality for import sets.	Active	false	• com.snc.automation • com.glide.ui_policy • com.glide.system_import_data_source
System Update Set Picker [com.glide.system_update_set_picker]	Allows users to choose an update set for tracking customizations.	Active	false	com.glide.system_update_set
System Update Sets (viewer) [com.glide.local_update_set]	Provides useful descriptions for update set entries and supports viewing contents of update sets.	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
System Update Sets (with remote update set support) [com.glide.system_update_set]	Facilitates moving customizations between systems. Supports viewing contents of update sets.	Active	false	• com.snc.apps_hub • com.glide.system_update_set_picker • com.glide.system_update_set_preview
System Update Sets Preview [com.glide.system_update_set_preview]	Allows users to preview the changes that are performed by an update set and predict whether there will be any collisions in attempting to apply the update set.	Active	false	com.glide.system_update_set
System Web Services [com.glide.web_service_application]	Provides a series of web service import sets.	Active	false	• com.glide.web_service_provider • com.glide.web_service_import_sets • com.snc.web_service_import_set_tables
Tablet Device Support - iPad with iOS 6+ [com.glide.ui.tablet]	Provides a UI supporting nearly full-product functionality on the Apple iPad.	Inactive	false	• com.glide.context_help • com.glide.ui.font_icons • com.glide.ui.tablet.theme
Tag Data Governance [com.sn_itom_tag] (Maintenance mode only)	This store application(beta) provides customers a way to get their cloud resources tagged per enterprise norms, in order to improve return on investments.	Active	false	
Tag Governance [sn_itom_tag] (Available in the ServiceNow Store)	Use the application to identify on-premises or cloud resources that are inconsistent and don't comply with the tag policies of your organization. When covered by the ITOM Governance subscription, you can also fix tag compliance issues. Use remediation scripts to add or modify cloud tags.	Active	false	• com.snc.discovery • com.snc.itom.vis.license • com.glide.hub.integrations • sn_cai
Targeted Communications	Provides a way to publish and send out	Inactive	true	com.snc.cs_base

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.sn_publications]	newsletter like articles to targeted internal/ external customers.			
Targeted Communications Demo Data [com.snc.publications_demo]	Targeted Communications application demo data.	Inactive	true	com.sn_publications
Task Activities [com.snc.task_activity]	Enables support for activities on task tables.	Inactive	false	
Task Intelligence for Customer Service [com.snc.csm_ml_task] (Available in the ServiceNow Store)	Offers several AI capabilities including language detection, multi-lingual email and case categorization, case sentiment analysis, and document intelligence. These features automate several routine tasks across the case lifecycle, enabling agents to focus on case resolution.	Active	false	• sn_ti_admin.com.sn_ti_admin • sn_docintel.com.snc.docintel • com.sn_customerservice • com.glide.platform_ml_task • com.glide.i18n • com.glide.dynamic_translation
Task-Outage Relationship [com.snc.task_outage]	Allows users to create an outage from an incident and a problem form. Incidents and problems have many to many relationship with outages. For new instances from Jakarta only, this feature is also available on the Change form.	Inactive	false	
Team Development [com.snc.apps_hub]	Supports parallel development on multiple, non-production ServiceNow instances by providing branching operations, the ability to compare a development instance to other development instances, and a central dashboard for	Active	false	• com.snc.apps • com.glide.ui.angular

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	all team development activities.			
Team Performance [com.sn_team_perf]	Team Performance module for Workforce Optimization	Active	false	com.snc.wfo
Template Management for Field Service [com.snc.fsm_template_management] (New in Tokyo)	Enables administrators to configure work order templates to process data dynamically rather than using the static information described in the work order templates. It fetches information from the source table to populate fields in the work order form and creates relevant tasks for a work order.	Active	false	com.snc.work_management
Templated Snippets [com.sn_templated_snip]	Activated with Human Resources Scoped App: Core [com.sn_hr_core]. Creates pre-defined and reusable responses that can be added to any table extending the task table including the HR Case table.	Inactive	true	com.glide.scope.access.restricted_caller
Test Management [com.snc.test_mgmt]	Provides a tool for manual software testing.	Inactive	true	com.snc.planned_task_v2
Test Management 2.0 [com.snc.test_management.2.0]	Enables you to manage testing processes and deliver software products more efficiently, with fewer errors.	Inactive	false	• com.snc.planned_task_v2 • com.snc.sdlc.agile.2.0.common
Test Migration 2.0 - Data Migration [com.snc.test_migration_v1_v2]	The Test Migration plugin provides functionality for migrating Test Case, Test and Test Suite to Test Management 2.0 Test version, Test Step and Test Set.	Active	false	• com.snc.test_mgmt • com.snc.test_management.2.0

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Text Index Attachments [com.glide.text_index_attachments]	Text Index Attachments - natural language search using Lucene syntax.	Active	false	com.glide.text_index
Text Search [com.glide.text_search]	Text Searching across multiple tables.	Active	false	com.glide.text_index
Threat Core [com.snc.threat]	Integrates Threat Intelligence with other Security Operations applications. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.secops.orchestration
Threat Intelligence [com.snc.threat.intelligence]	Provides a point of reference for your company's Structured Threat Information Expression (STIX) data. Included in Threat Intelligence is the Security Case Management application. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.threat
Time card management [com.snc.time_card]	Works with the Task table to record time worked on projects, incidents, problems, and change requests. Task assignees can record time worked in the **Time worked** field on a task record or enter hours directly onto a time card.	Inactive	false	com.snc.time_card_portal
Time Recording for Customer Service [com.snc.csm_time_recording]	Enables customer service ages to record time on cases, case tasks, and other activities. Recording time automatically generates time cards	Inactive	false	• com.snc.wm_time_recording • com.sn_customerservice

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	and time sheets for approval by customer service managers.			
Time Recording for Field Service [com.snc.time_recording_fsm]	Extends the functionality of the Time Card Management and Cost Management applications to Field Service Management. Field service agents record time worked on tasks and other activities. These time worked entries automatically create time cards and weekly time sheets. Managers can review and approve time sheets as well as view and create labor rate cards.	Inactive	true	• com.snc.work_management • com.snc.wm_time_recording
Timeline Visualization [com.snc.timeline_visualization]	Enables graphical representation of activities over time to provide a high-level view of strategic and operational activities in your organization such as incidents, problems, changes, and projects. A base system visualization provided by this plugin is the CIO Roadmap. This roadmap shows projects grouped by portfolios. Organization leaders can use the CIO roadmap to monitor and evaluate the status of current and upcoming projects.	Inactive	true	com.snc.project_management_v3
Tiny URL Support [com.glide.tiny_url]	Enables support for generating shortened URLs to eliminate problems with very long URLs in Internet Explorer. New	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	properties will be added to the System properties page.			
TinyMCE HTML Field Editor [com.glide.editor.tinymce]	Enables users to edit HMTL fields with the TinyMCE editor instead of the legacy (htmlArea) editor.	Active	false	
Transaction Design Scope [com.glide.transaction_scope]	Handles transaction design scope management.	Active	false	
Transaction Quotas [com.glide.quota]	Allows definition of quota policies for different types of transactions. A transaction quota cancels any transaction in violation of the policy and notifies the user of the cancellation.	Active	false	
Transform Definitions [com.glide.transform.definitions]	Activates transform functions which enables flow designers to transform a data pill or field's data type without writing a script.	Active	false	com.glide.transform
Transform Service [com.glide.transform]	Supports use of Transform APIs associated with the Remote Tables functionality.			
Translated Email Retention [com.glide.email_retention.translation] (New in Tokyo)	Provides retention policy for Email translations.	Inactive	false	com.glide.notification.translation
Transaction stats [com.glide.stats.transaction]	Transaction stats.	Active	false	com.glide.stats
Tree map [com.snc.treemap]	Enables support for treemap view on any applications.	Inactive	false	com.snc.pa

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Trusted Security Circles Client [com.snc.intel_sharing.client]	Provides the ability to set up communication channels that connect sets of Trusted Security Circles customers who have some kind of underlying relationship. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.threat
Trusted Security Circles Client (Advanced) [com.snc.intel_sharing.clnt.adv]	Provides the capabilities of the basic level, along with the ability to join any available trusted circle and initiate an unlimited number of threat shares per day. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	false	com.snc.intel_sharing.client
UI 15 [com.glide.ui.doctype]	Enables UI15, a user interface for modern browsers with refined colors and usability enhancements.	Active	false	• com.glide.ui11 • com.glide.ui.recent_selections • com.glide.ui.themes.doctype • com.glide.ui.font_icons • com.glide.ui.ng • com.glide.editor.tinymce • com.glide.ui.personalize_form • com.glide.ui.overview_help
UI Common [com.glide.ui.common]	Glide UI Common	Inactive	false	
UI Components [com.glideapp.ui_components]	Provides all common angular components for apps.	Active	false	• com.glide.ui.ng • com.glide.ui.angularui • com.glide.ui.heisenberg

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
UI16 [com.glide.ui.ui16]	Enables Core UI, a user interface that provides an updated look and usability improvements. Notable features include real-time form updates, user presence, a redesigned application navigator with tabs for favorites and history, and enhanced activity streams.	Active	false	• com.glide.ui.doctype • com.glide.ui.concourse • com.glide.ui.form_presence • com.glide.ui.snippets
Universal Request Integration for Incident management [com.snc.incident.universal_request]	This plugin delivers Universal Request Integration for Incident Management. It is a developer plugin.	False	True	com.snc.universal_request
Universal Request Integration with Microsoft Teams [com.snc.universal_request.ms_teams]				
Universal Request: Advanced Work Assignment [com.snc.universal_request.awa]	Enables to use Advanced Work Assignment and automatically assign work items to your agents, based on their availability in the Universal Request application.	False	True	com.snc.universal_request
Universal Request: NLU Model for Virtual Agent Conversations [com.snc.universal_request.nlu]	Allows the use of NLU for virtual agent conversations in Universal Request.	False	True	
Universal Request: Reporting [com.snc.universal_request.reporting]	Enables to use reports in Universal Request.	False	True	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Universal Request: Virtual Agent Conversations [com.snc.universal_request.va]	Helps to set up and use Virtual Agent in Universal Request.	False	True	
Predictive Intelligence for Universal Request [com.snc.universal_request.ml]	Facilitates to use machine-learning solutions in the Universal Request application, if you have the admin role.	False	True	• com.snc.universal_request • com.glide.platform_ml
Performance Analytics and Reporting for Universal Request [com.snc.universal_request.pa]	Contains Performance Analytic and Reporting dashboard created for Universal Request.	False	False	
Universal Request integration with Microsoft Teams [sn_uni_req_msteams]	Enables to use Microsoft Teams in the Universal Request application.	False	True	• com.glide.cs.chatbot • com.snc.universal_request.va • sn_now_teams
Universal Request [com.snc.universal_request]	Install this plugin to use the Universal Request application.	False	True	
Update Set Batching [com.glide.system_hierarchy_update_set]	Deploy update sets in a group.	Active	false	com.glide.system_update_set
Upgrade Blame Tool [com.glide.upgrade_blame]	Tracks the affected table records touched (insert/update/delete) by the load files during system upgrade, zboot, or plugin activation/upgrade. This plugin adds application modules to allow ServiceNow employees to access the upgrade blame log and to turn on/off the feature. It	Active	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	also adds security access control rules to prevent user from modifying the upgrade blame logs table.			
Upgrade Metric [com.glide.upgrade_metric]	This plugin uses the Upgrade Metric table to track the changes that take a long time to load during upgrades.	Active	false	
User Criteria Scoped API [com.glideapp.user_criteria.scoped.api]	User Criteria Scoped API to support CRUD operations on User Criteria.	Inactive	false	com.glideapp.user_criteria
User Guide [com.glide.user_guide]	Provides the ability to create end-user help documentation that is specific to the policies and procedures of your organization. A default help page is provided in the base system that displays Core UI help documents for system navigation and other basic operations.	Active	false	
Usage Insights [com.glide.appsee]	Checks for new web and mobile applications to register and provides users access to the Dashboard.	Active	false	com.glide.serviceproxy
User Registration Request [com.snc.user_registration]	Provides the ability for unregistered users to request access to a ServiceNow instance.	Inactive	false	
Vaccine Administration Management [com.sn_vaccine_sm] (Available in the ServiceNow Store)	Enables governments and healthcare providers to deliver vaccines, on a deadline and with finite resources.	Inactive	true	• com.snc.appointment_booking • com.snc.business_location • com.snc.csm_case_types • com.snc.household • com.snc.install_base • com.glide.encryption • com.sn_external_user_register • com.snc.pdf_generator

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
				• com.snc.signaturepad • com.glideapp.user_criteria.scoped.api • sn_hcls • sn_csm_playbook • com.sn_ind • sn_doc • sn_prd_pm • sn_ciwf_ui_cmpnt
Vendor Manager Workspace [com.snc.vlm.vmw]	Access to the Vendor Manager Workspace.	Inactive	false	• Performance Analytics - Premium • Service Workspace, Vendor Manager Workspace - Core • Service Owner Workspace, UX Framework • GraphQL Plugin • Interactions Management • Agent Workspace - Main Configuration • Agent Workspace - Ribbon • Agent Workspace - List • Agent Workspace - Form • com.snc.agent_workspace.global_search • com.snc.agent_workspace.declarative_actio • com.snc.agent_workspace.highlighted_value
Vendor Manager Workspace - Core [com.snc.vlm]	Configure the Vendor Manager Workspace.	Inactive	false	
Vendor Success Indicators [com.snc.vendor.insights]	Vendor Success Indicators to utilize Odds Ratio and other ML capabilities for Vendor Manager Workspace.	Active	false	com.snc.app.ml.insights
Version Management [com.snc.version]	Provides the ability to track, compare, and revert to multiple	Active	false	

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	versions of table records.			
Version Support [com.glideapp.version]	Supports tracking versions of files that are stored in update sets, including the ability to compare and revert to previous versions.	Active	false	
Virtual Agent Experience Pack for Procurement Service Management [com.sn_shop_va] (Starting with the Utah release, Virtual Agent Experience Pack for Procurement Service Management is renamed as Virtual Agent for Sourcing and Procurement Operations.) (Available in the ServiceNow Store)	Provides pre-configured topics for sourcing and procurement within a conversational interface, for employees to raise requests, find knowledge, and complete tasks.	Inactive	false	• com.snc.sn_pr • com.glide.cs.chatbot
Virtual Agent integration with actionable notifications [com.glide.cs.actionable.notification]	Enables ServiceNow actionable notifications on Virtual Agent chat channels.	Active	false	com.glide.cs.notification
Virtual Agent integration with notifications [com.glide.cs.notification]	Enable ServiceNow notifications on Virtual Agent chat channels.	Inactive	false	• com.glide.cs • com.glide.notification.provider

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Virtual Agent Language Detection and Translation [com.glide.cs.runtime_language_detection_translation]				
Virtual Agent Platform Topics [com.glide.cs.topics]	Core conversation topics that support Virtual Agent platform conversation flows.	Inactive	false	com.glide.cs
Virtual Agent Platform Topic Blocks [com.glide.cs.topic_blocks]	Topic blocks required to support core Virtual Agent platform conversation topics	Inactive	false	com.glide.cs
Virtual Agent Service Portal Widgets [com.glide.va.sp_widgets]	Virtual Agent Service Portal Widgets	Inactive	true	• com.glide.service-portal • com.glide.cs • com.glide.cs.sn-va-web-client-app
Virtual Agent Spoke [com.glide.cs.va_spoke]	Virtual Agent Integration Hub actions and flows.	Active	false	
Virtual Agent Web Client [com.glide.cs.sn-va-web-client-app]	Scoped app for the virtual agent web client.	Inactive	false	
VirusTotal Integration [com.snc.threat.virustotal]	Enables the VirusTotal scanner in Threat Intelligence. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	com.snc.threat
Visibility Content [sn_pattern_design] (Available in the	Supplies the updated version of the patterns that were part of the family releases up until Tokyo.	Active	false	com.sn_itom_pattern

servicenow

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)				
Visual Task Board Flow Designer Spoke [com.glide.ui.vtb.ah]	The Visual Task Board (VTB) Spoke for the Flow Designer provides Actions enable process analysts to compose flows that manipulate task boards, cards, board members, and assignees on given tasks without having to write code. The Visual Task Board Plugin (com.glide.ui.vtb.ah) should be installed to use these actions.	Active	false	• com.glide.ui.vtb • com.glide.hub.designer_backend.model
Visual Task Boards [com.glide.ui.vtb]	Allows users to organize, modify, and track progress of multiple tasks from an intuitive, Kanban-inspired interface.	Active	false	• com.glide.ui.ng • com.glide.ui.font_icons • com.glide.task • com.glide.ui.checklist • com.glide.ui.ng.amb • com.glide.rest.service • com.glide.ui.ng.filter • com.glide.ui.magellan_navigator_api
Vulnerability Response [com.snc.vulnerability] (Available in the ServiceNow Store)	Allows security users to compare security data pulled from internal and external sources and, if CIs or software are found to be vulnerable, changes and security incidents can be created using Vulnerability Groups. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Inactive	true	• com.snc.sam.core • com.snc.secops.orchestration • com.snc.security_support.vul
Vulnerability Response Common	Framework that provides the necessary	Active	false	Security Support Common [sn_sec_cmn]

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[sn_vul_cmn] (Available in the ServiceNow Store, compatible with Vancouver, Utah)	functionality to support the Vulnerability Response applications like Vulnerability Response, Application Vulnerability Response, Container Vulnerability Response and Configuration Compliance.			
Vulnerability Response Dependencies [com.snc.vul_dep]	System dependencies required by Vulnerability Response. Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.		false	
Vulnerability Response Mobile [com.snc.vul_mobile] (Available in the ServiceNow Store)	As a remediation owner, you can access the Vulnerability Response (VR) application on your ServiceNow AI Platform instance with your Android or iOS mobile device.	Active	false	com.glide.sg.agent_native_client
Vulnerability Solution Management [com.snc.vulnerability.solution] (Available in the ServiceNow Store)	Vulnerability Solution Management automatically correlates solutions with vulnerabilities, identifies preferred solutions to reduce risk, and tracks deployment progress. The Vulnerability Solution Management application enables the Vulnerability Solution Management feature in Vulnerability Response.	Active		

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Walk-up Experience [com.snc.walkup]	Enables your IT organization to set up a contact channel to support both online check-in and onsite check-in to a pre-established walk-up service center.	Inactive	true	• com.snc.agent_workspace.itsm • com.glide.awa • com.glide.interaction • com.snc.asset_management • com.glide.service-portal • com.snc.appointment_booking • com.glide.db_audio
Walk-Up for CSM (Maintenance mode only) [com.snc.walkup_for_csm]	Walk-up experience for Customer Service enables in-store support to assist onsite check-ins, managing queues, and customer interactions	Active	false	• com.snc.walkup • com.sn_customerservice
Web Service Consumer [com.glide.web_service_consumer]	Provides a SOAP Message module for developing, prototyping, and saving outbound SOAP messages that can be reused in business rules and scripts.	Active	false	• com.glide.vars • com.glideapp.ecc
Web Service Import Set Tables [com.snc.web_service_import_set_tables]	Complements direct web services and scripted web services in providing a web service interface to import set tables.	Active	false	• com.glide.web_service_import_sets • com.glide.web_service_application
Web Service Import Sets [com.glide.web_service_import_sets]	Complements direct web services and scripted web services in providing a web service interface to import sets.	Active	false	• com.glide.table_editor • com.glide.web_service_provider
Web Service Provider - Common [com.glide.web_service_provider]	Provides scripted web service and SOAP message resources.	Active	false	• com.glide.vars • com.glide.system_property_categories • com.glide.system_import_set
Web Service Provider - Custom WSDL	Provides the ability to create a scripted web service to accept any WSDL format.	Inactive	false	

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
[com.glide.static_wsdl]				
Web Service Provider - Scripted [com.glide.custom_web_service]	Enables users to create web services that are not addressed by the system. Allows a user to define input and output parameters and use JavaScript to do everything in between.	Active	false	• com.glide.web_service_provider • com.glide.web_service_consumer
WebKit HTML to PDF [com.snc.whtp]	Enables the instance to use the service WebKit HTML to PDF.	Inactive	false	com.snc.platform.security.oauth
WFO Mobile [com.snc.wfo.mobile]				
Widgets [com.snc.app.widgets]	Provides widget infrastructure along with Angular UI carousel module.	Inactive	false	
Workday ESG integration [sn_esg_workday] (Available in the ServiceNow Store)	Enables ESG users to pull critical HR data from Workday in to ESG instances.	Active	true	ESG Management
Workflow Authoring Tools [com.glideapp.workflow.authoring]	Allows you to define and modify workflows by arranging and connecting activities with transitions.	Active	true	com.glideapp.workflow
Workflow Pause Utility [com.glideapp.workflow.pause]	Allows you to pause and resume workflows.	Inactive	false	
Workflow Runtime Engine [com.glideapp.workflow]	Enables the creation of workflows that drive automated processes. This may entail generating tasks based on conditions, running scripts, generating	Active	false	• com.glide.diagrammer • com.glide.vars • com.glide.schedules • com.glide.relative_duration • com.glide.web_service_application

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
	approvals, or other actions. Satisfies the same need as the Execution Plans plugin but with greater control and an easier interface.			• com.glide.service_api • com.snc.datastructure • com.glideapp.live_feed
Workflow/ Stages element support [com.glide.stages]	Support for Workflow-type elements and stage displays.	Active	false	
Workforce Optimization (Maintenance mode only) [com.snc.wfo]	Activation of this plugin on production instances may require a separate license. Contact ServiceNow for details.	Active	false	
Workforce Optimization Common [com.sn_wfo_common]				
Workforce Optimization Dependencies [com.sn_wfo_dependency]				
Workforce Optimization for Customer Service (Maintenance mode only) [com.snc.wfo.csm]	The Workforce Optimization for Customer Service application provides Channel Management, Team Performance, Scheduling, Skill Determination and Coaching capabilities for managers, supervisors, and team leads to improve quality and efficiency of their customer service teams and enhance team satisfaction.	Active	false	• com.snc.manager_workspace • com.sn_customerservice • com.glide.interaction.awa • com.sn_channel_management • com.snc.wfo • com.sn_shift_planning • com.snc.skills_management • com.manager_workspace_components • com.sn_coaching • com.snc.sre • com.sn_team_perf

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
Workforce Optimization for Field Service Management [sn_fsm_wfo] (Available in the ServiceNow Store)	Helps to manage and maintain the productivity of your workforce from a single application using Workforce Optimization for Field Service.	Active	true	• FSM Team Performance Management (com.snc.sn_fsm_team_mgmt) • FSM Shift Scheduling (com.snc.sn_fsm_shift_schdl) • Skill Recommendation Engine (com.snc.sre) • Coaching (com.sn_coach_learning, com.sn_coaching)
Workforce Optimization for HR [com.sn_hr_wfo]				
Workforce Optimization for ITSM [com.snc.wfo_itsm] (Maintenance mode only)	Workforce Optimization bundle for IT Service Management	Active	false	• com.snc.wfo • com.sn_shift_planning • com.snc.skills_management • com.snc.manager_workspace • com.manager_workspace_components • com.sn_coaching • com.snc.sre • com.sn_team_perf • com.sn_channel_management • com.snc.incident.awa • com.snc.walkup • com.glide.interaction.awa
Workspace Core [com.workspace_core]	Core Workspace app for enabling CSM/ITSM agents to provide world-class service at light speed.	Active	false	• com.glide.uxbuilder • com.snc.agent_workspace.config • com.snc.agent_workspace.global_search
Workplace Indoor Mapping [sn_wsd_indoor_map] (Available in the	Enables organizations to digitalize their workplace floor plans. Using the Map Studio, map admins can create and update maps that can be used across the	Active	true	• sn_wsd_suite • sn_wsd_space_map

Plugin [ID]	Description	State	Has Demo Data?	Dependencies
ServiceNow Store)	Workplace Delivery Suite.			
XML Parser for ServiceNow IntegrationHub [com.glide.hub.action_step.xmlparser]	Activates the XML parser step in Flow Designer, enabling you to identify structured data from an XML payload without having to write script.	Inactive	false	com.glide.hub.action_step.template
XTM Translation Management System Spoke for Localization Framework [com.glide.localization_framework.xtm_spoke]	Integrates Localization Framework with the XTM - Translation Management System. To activate this plugin, install the Localization Framework Installer plugin.	Active	false	com.glide.localization_framework
Zoom extension for Omnichannel Callback [com.sn.zoom.callback] (Available in the ServiceNow Store)	Enables customers to request for a visual engagement session using Zoom.	Active	false	• com.snc.notify.zoom • com.sn.omnichannel.callback

Find components installed with an application

Activating a plugin installs an application on your instance. Each application consists of components such as tables, user roles, and scheduled jobs. To view all components that are installed with an application, see the Application Files table.

Before you begin

- The application plugin must be activated.
- Roles required: admin or the application admin role

Note:
The application admin role is only required for applications that use the Application administration feature. For example, Vulnerability Response.

About this task

In the Application Files table, look up the components that are installed with the plugin for the application. If the application comprises multiple plugins, look up each plugin individually.

Procedure

1. In the navigation filter, enter: `sys_metadata.list`
   The list of Application Files records is displayed.
2. Click the personalize list icon ( ), and add the **Package** column to the list.
3. Filter the list with these conditions to show only the components for the plugin:
   [Package] [is] [*<plugin name>***]**
4. In the **Class** column heading, click the menu icon ( ) and select **Group by Class**.
5. Select the class of information that you want to view.

Available system properties

Some properties are available on a system properties form, but some lesser-used properties are available only from the System Property [sys_properties] table. Sometimes, the property does not exist in a base instance, but can be added if you change the value.

Note:
The Instance Security Hardening Settings content contains detailed descriptions, and compliance values, for the security-related system properties and plugins in the ServiceNow AI Platform. To learn more about each of these properties, see Instance Security Hardening Settings.

cdu.record.watcher.timeout

Maximum waiting period for the Cloud Discovery schedule configuration wizard to stop watching the record watchers while validating account, member account Discovery, and logical database Discovery.

- Type: integer
- Default value: 120,000 milliseconds
- Location: System properties

com.cmdb.baseline.max_changes

Maximum number of changes and relationships for a CI that can appear in the baseline diff for the CI.

- Type: integer
- Default value: 100
- Location: **Configuration** > **CMDB Properties** > **Baseline Properties**

com.glide.attachment.max_size

Sets the maximum file attachment size in megabytes. An empty value in this field allows attachments up to a maximum of 1 KB.

- Type: integer
- Default value: 1024
- Location: **System Properties** > **Security**

com.glide.cs.collab.event_queue.enabled

Switch to enable the event queue for badge count calculation and publish.

Note:
If you set com.glide.cs.collab.event_queue.enabled to **true**, you must manually enable the Sidebar notification events process schedule record ("sys_trigger") by changing its *trigger_ type* value.

com.glide.cs.collab.event_queue.threshold

Calculates the unread count for each member of a chat. If the chat has too many members, this responder thread will be slow so offload this to the event queue. Default value: 50.

com.glide.cs.collab.notify_all_members

If Sidebar is unable to deliver messages to a Microsoft Teams user, this property determines who a "message not sent" notification is sent to. If this property is **true**, then a notification is sent to all Sidebar users in the discussion. If **false**, then the notification is sent only to the sender of the message.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.glide.cs.collab.teams_outage_timer

Length of time the "message not sent" notification (see com.glide.cs.collab.notify_all_members) displays.

- Type: integer
- Default value: 5
- Location: System Property [sys_properties] table

com.glide.cs.conversation_history_cross_channel.enabled

Enables (**true**) or disables (**false**) the display of cross-channel conversation history in the chat panel when an agent is having a conversation with a requester.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

com.glide.cs.field.boolean.show.choice_picker

Enables (**false**) or disables (**true**) a Boolean option to requesters whenever a dynamic choice node only has a single value associated. When set to **true**, the dynamic choice node displays to requesters with only the single value selection option and no Boolean option.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

com.glide.cs.gen_ai.enable_mid_topic_ai_search

Enables or disables AI Search result in Mid Topic Discovery using this flag.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.glide.cs.genai.input_collector_mismatch_reprompts

Sets the number of times the user is asked to answer the original question before triggering topic discovery.

- Type: integer
- Default value: 2
- Location: System Property [sys_properties] table

com.glide.cs.instance_uri_override

Enables Virtual Agent to use this URL value to override the instance URL for link generation. The value entered must be a direct, full URL such as http://www.example.com to work correctly.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

com.glide.cs.notification.create_llm_actions_after_topic_migration

Enables (`true`) or disables (`false`) whether actions for Virtual Agent notifications and Proactive Triggers are created after LLM topics are published during the NLU to LLM topic migration workflow.

- Type: true | false
- Default: true
- Location: System Property [sys_properties] table

com.glide.cs.proactive_trigger.create_llm_actions_after_topic_migration

Enables (true) or disables (false) whether actions for Proactive Triggers are created after LLM topics are published during the NLU to LLM topic migration workflow.

- Type: true | false
- Default: true
- Location: System Property [sys_properties] table

com.glide.cs.unified.prompt_enabled_mid_skill_switch

Enable (true) or disable (false) the mid-topic switch confirmation message.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.glide.csv.loader.ignore_non_parseable_lines

Enables (**true**) or disables (**false**) ignoring one or more lines (rows) in an import set that contains bad data, such as a row that is missing a column of data.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

com.glide.csv.loader.max_errors_allowed

Sets the maximum number of lines (rows) that an import can ignore before failing. If the import succeeds, the import lists the number of rows ignored due to errors.

- Type: integer
- Default value: 100
- Location: Add the property to the System Property [sys_properties] table.

com.glide.email.max_read

Specifies the maximum number of emails a POP3 reader should process concurrently.

- Type: integer
- Default value: 20
- Location: Add the property to the System Property [sys_properties] table.

com.glide.event_handler.CollabNotificationEventHandler

Specifies the event handler for the sidebar_notification queue.

com.glide.event_manager.sidebar_notification.claim_limit

Specifies the number of sidebar_notification events the job on each node polls each time.

com.glide.interaction.autocreate_relationships

Specifies whether auto-logging of records in the interaction_related_record table is turned on or off during active interactions.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

com.glide.loader.verify_target_field_size

Enables dynamic resizing of import set fields. With the default setting of **false**, up to 20 records of the source data are sampled to determine the import set field length. If the field is empty in all the sampled records, the default length of 40 is used. Any data loaded that exceeds the import set table field length is truncated. Set this property to **true** to allow any import set field to increase the column size by 100 to match the length of the data. For example, if the incoming source data has a length of 60, the system sets the length of the column to 160.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

com.glide.scripting.module.use_module_sandbox

Separates the JavaScript module runtime environment from scripts. Setting this property to false may improve performance.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

com.glide.soap.include_non_db_fields

Controls whether non-database fields, such as the sys_tags field, are included in WSDLs and SOAP responses (**true**) or not (**false**).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

com.glide.ssl.read.timeout

Sets the time-out value in seconds for SSL connections during read operations. Typically, this setting is used as part of LDAPS. If you enter timeout values for both this system property and the LDAP **Read timeout** field, the lowest timeout value takes precedence.

- Type: integer
- Default value: 10
- Location: Add the property to the System Property [sys_properties] table.

com.glide.ui.virtual_host

Specifies Guided Tour Designer callout icons. When set, callout icons appear in the designer instead of letters.

- Type: String
- Value: pa,html,script,include scripts,glide,custprefix,_sn
- Location: System Property [sys_properties] table

com.glideapp.canvas

Enables (**true**) or disables (**false**) responsive canvas for dashboards. When enabled, all new dashboards use responsive canvas.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.apps.publish.include_inactive_choices

Specifies whether choice table selections marked as inactive (`sys_choice inactive attribute = true`) should be included when publishing an application (**true**) or not (**false**).

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.apps.publish.maxrows

Defines the maximum number of data records to include when publishing an application.

- Type: integer
- Default value: 1,000
- Location: System Property [sys_properties] table

com.snc.assessment.decision_matrix_filter_max_entries

Defines the maximum number of items to show for a decision matrix field filter.

- Type: integer
- Default value: 1,000
- Location: System Property [sys_properties] table

com.snc.assessment.signature_authentication

Requires authentication for a user signature.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.dd.publish_to_app_repo_enabled

Enables (**true**) or disables (**false**) display of the **Publish To App Repo** permission in Manage Developers.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.dd.publish_to_app_store_enabled

Enables (**true**) or disables (**false**) display of the **Publish To App Store** permission in Manage Developers.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.dd.publish_to_update_set_enabled

Enables (**true**) or disables (**false**) display of the **Publish To Update Set** permission in Manage Developers.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

com.snc.dd.upgrade_app_enabled

Enables (**true**) or disables (**false**) display of the **Update App** permission in Manage Developers.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.hr.core.impersonateCheck

Allows a user to have identical permissions/access as the person they are impersonating. If **true**, access and restrictions apply when impersonating another user. If **false**, the user has identical permissions/access as the person they are impersonating.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.iam.log_level

Logging level for the business rule MapUpstreamImpactedCI. Debug is the most detailed option with full trace of how the Impacted CI List is calculated. Error is the minimal logging option with only severe errors being logged.

- Type: string
- Default value: info
- Possible values: debug, info, error
- Location: System Property [sys_properties] table

com.snc.on_call_rotation.reminders.showtz

Specifies whether to show the user time zone (**true**) or not (**false**).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

com.snc.pa.activate_kpi_signals

Create this property and set to 'false' to deactivate the KPI Signals feature.

- Type: Boolean true | false
- Default value: Does not exist but is treated as 'true'
- Location: Add the property to the System Property [sys_properties] table.

com.snc.pa.correlation.enable_insights

Enables the generation of key contributor insight cards.

- Type: true/false (Boolean)
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

com.snc.pa.dc.batch_score_insert

When true, Performance Analytics data collection jobs insert indicator scores to the database in batches. When false, indicator scores are inserted one by one. Even when true, batch score insertion is performed only as a fallback when nested data collection (com.snc.pa.dc.nested_collection) has been disabled.

- Type: Boolean true | false
- Default value: true
- Location: System Properties [sys_properties] table

com.snc.pa.dc.flush_snapshot_count

Maximum number of Performance Analytics snapshots to store in the memory during data collection. When this number is reached, the snapshots are offloaded to the database, freeing memory.

- Type: integer
- Default value: 1000
- Location: System Properties [sys_properties] table

com.snc.pa.dc.nested_collection

Activates a hierarchical approach to score collection, where Indicator > level 1 breakdown > level 2 breakdown. This approach is more resource efficient than a flat hierarchical approach, which creates redundant objects.

- Type: Boolean true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

com.snc.pa.dc.max_breakdown_elements_limit

Maximum number of breakdown elements retrieved by data collection for each breakdown source. This property is handled differently depending on whether nested collection is enabled. If nested collection is enabled, only those elements that have non-null scores count against this limit. If nested collection is disabled, such as under optimized data collection, all elements on the breakdown source count against the limit. When a breakdown source exceeds this limit, breakdown scores are not collected from that source.

- Type: integer
- Default value: 10000
- Location: System Property [sys_properties] table

com.snc.pa.dc.max_error_count

Maximum errors that may occur before data collection is stopped.

- Type: integer
- Default value: 500
- Location: System Property [sys_properties] table

com.snc.pa.dc.max_records

Maximum number of records that are stored during a data collection.

- Type: integer
- Default value: 5000
- Location: System Property [sys_properties] table

com.snc.pa.dc.max_row_count_indicator_source

Maximum number of rows that are allowed to be fetched from an Indicator Source. This property only applies to jobs that do not use Optimized data collection.

- Type: integer
- Default value: 50000
- Location: System Property [sys_properties] table

com.snc.pa.dc.hsql.max_row_count_indicator_source

Maximum number of rows that are allowed to be fetched from an Indicator Source. This property only applies to jobs that use Optimized data collection.

- Type: integer
- Default value: 1 million
- Location: System Property [sys_properties] table

com.snc.pa.dc.script_timeout

Maximum time in seconds that a script is allowed to run during a data collection cycle.

- Type: integer
- Default value: 30
- Location: System Property [sys_properties] table

com.snc.pa.dc.use_hashcode_for_distinct

Sets whether Count Distinct aggregations are stored as separate values in memory or as a HashCode. When the value of this property is True, a HashCode is used, which consumes less memory than storing individual values.

- Type: Boolean true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

com.snc.pa.insight.max_days_in_past

By default, retrieve insights created within the number of days in the past that this property specifies.

- Type: Integer
- Default value: 14
- Location: Add the property to the System Property [sys_properties] table.

com.snc.pa.insights.query_limit

Limit the number of rows when querying the par_computed_insight table in the GET response to the Insights API.

- Type: Integer
- Default value: 1000
- Location: The System Properties [sys_properties] table

com.snc.pa.target_batch_operation_limit

Applies only to KPI Details: The maximum number of targets that can be inserted or updated in one bulk action on the Targets Configuration page.

- Type: integer
- Default value: 100
- Location: System Property [sys_properties] table

com.snc.pa.dc.script.use.sandbox

Controls whether Performance Analytics data collection scripts use the sandbox. Only users with the security_admin role can modify this property.

Important:
Exercise caution when turning off this property. When set to false, unauthorized users may be able to execute malicious actions on your instance.

- Type: Boolean true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.pa.formula_indicator_valid_breakdown

Controls whether users can apply a breakdown to a formula indicator when not all contributing indicators in the formula use that breakdown. When true, all contributing indicators must use the formula. When false, a breakdown can be applied when at least one of the contributing indicators uses it.

- Type: Boolean true | false
- Default value: true
- Location: System Property [sys_properties] table

com.snc.pa.xmr.processes.limit

The maximum number of KPI Signals Configurations (pa_xmr_processes) that can be evaluated by the KPI Signals Insight Job. After this number of processes have been evaluated, the remaining processes are ignored.

- Type: Integer
- Default: 1000
- Maximum: 10,000

com.snc.project.default_schedule

Stores the sys_id of the default schedule attached to projects.

- Type: string
- Location: System Property [sys_properties] table

com.snc.project.loglevel

Automatically resubmits timed-out Ajax requests.

- Type: choice list
- Default value: 0 (debugging disabled)
- Location: System Property [sys_properties] table

com.snc.project.wbs_gantt

Automatically resubmits timed-out Ajax requests.

- Type: true | false
- Default value: false
- Location: **Project** > **Administration** > **Properties**

com.snc.task.associate_ci

List of all task types where user wants to associate CIs using a List system.

- Type: string
- Default: incident,problem,change_request
- Location: System Property [sys_properties] table

com.snc.task.refresh_impacted_services

List of all task types where **Refresh Impacted Services** UI action is enabled.

- Type: string
- Default value: incident,change_request
- Location: System Property [sys_properties] table

com.snc.time_card.default_rate

(Cost Management plugin) Sets a default hourly rate to use if no labor rate cards apply to the user.

- Type: integer
- Default value: 50
- Location: System Property [sys_properties] table

com.snc.time_worked.update_task_timer

Enables (**true**) or disables (**false**) updating the task timer value based on changes to the task time worked records. Uses the Update task time business rule.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

csm_auto_account_domain_generation

When new account in Customer Service application is created, domain is created and placed under TOP domain. If parent field on account form is populated, and new record is inserted, it creates that account as a subdomain of the parent.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

css.assessment.caption.background.color

Sets the background color of the caption on assessment and survey questionnaires.

- Type: color
- Default value: #eee
- Location:
  - **Assessments** > **Admin** > **Assessment Properties**
  - **Survey Management** > **Administration** > **Properties**

css.assessment.caption.font.color

Sets the font color of the caption text on assessment and survey questionnaires.

- Type: color
- Default value: #ffffff
- Location:
  - **Assessments** > **Admin** > **Assessment Properties**
  - **Survey Management** > **Administration** > **Properties**

css.assessment.question.header.background.color

Sets the background color of question headers on assessment and survey questionnaires.

- Type: color
- Default value: #767676
- Location:
  - **Assessments** > **Admin** > **Assessment Properties**
  - **Survey Management** > **Administration** > **Properties**

css.$nav-highlight-main

Sets the color for the navigation expanded items highlight background.

- Type: color
- Default Value: #3D4853
- Location: System Property [sys_properties] table

css.$nav-hr-color

Sets the color for the navigation separator.

- Type: color
- Default Value: #303A46
- Location: System Property [sys_properties] table

css.$navpage-header-bg

Sets the color for the header background.

- Type: color
- Default Value: #303A46
- Location: System Property [sys_properties] table

css.$navpage-header-color

Sets the color for the banner text.

- Type: color
- Default Value: #FFFFFF
- Location: System Property [sys_properties] table

css.$navpage-nav-bg-sub

Sets the color for the background for navigator and sidebars.

- Type: color
- Default Value: #455464
- Location: System Property [sys_properties] table

css.$navpage-nav-border

Sets the color for the border for Core UI.

- Type: color
- Default Value: #DDD
- Location: System Property [sys_properties] table

css.$navpage-nav-color-sub

Sets the color for module text in Core UI.

- Type: color
- Default Value: #BEC1C6
- Location: System Property [sys_properties] table

css.$navpage-nav-selected-bg

Sets the color for navigator selected tab background.

- Type: color
- Default Value: #4B545F
- Location: System Property [sys_properties] table

css.$navpage-nav-selected-color

Sets the color for the currently selected navigation tab icon color for Core UI.

- Type: color
- Default Value: #FFFFFF
- Location: System Property [sys_properties] table

css.tablet.gradient.start

Start color of the gradient for the tablet UI header.

- Type: color
- Default value: #666
- Location: **System Properties** > **Tablet UI Properties**

css.tablet.gradient.end

End color of the gradient for the tablet UI header.

- Type: color
- Default value: #111
- Location: **System Properties** > **Tablet UI Properties**

css.tablet.headerfooter.text.color

Color of the text and icons in the tablet UI header and footer.

- Type: color
- Default value: #ffffff
- Location: **System Properties** > **Tablet UI Properties**

glide.allow.new.cert_follow_on_task

Allow Desired State audit to create a new follow-on task for the same failure at each audit run.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.amb.session.logout.overlay.style

Specifies a cascading style sheet (css) that defines an overlay. The overlay masks all onscreen data on timeout, and the login screen appears on top of the overlay.

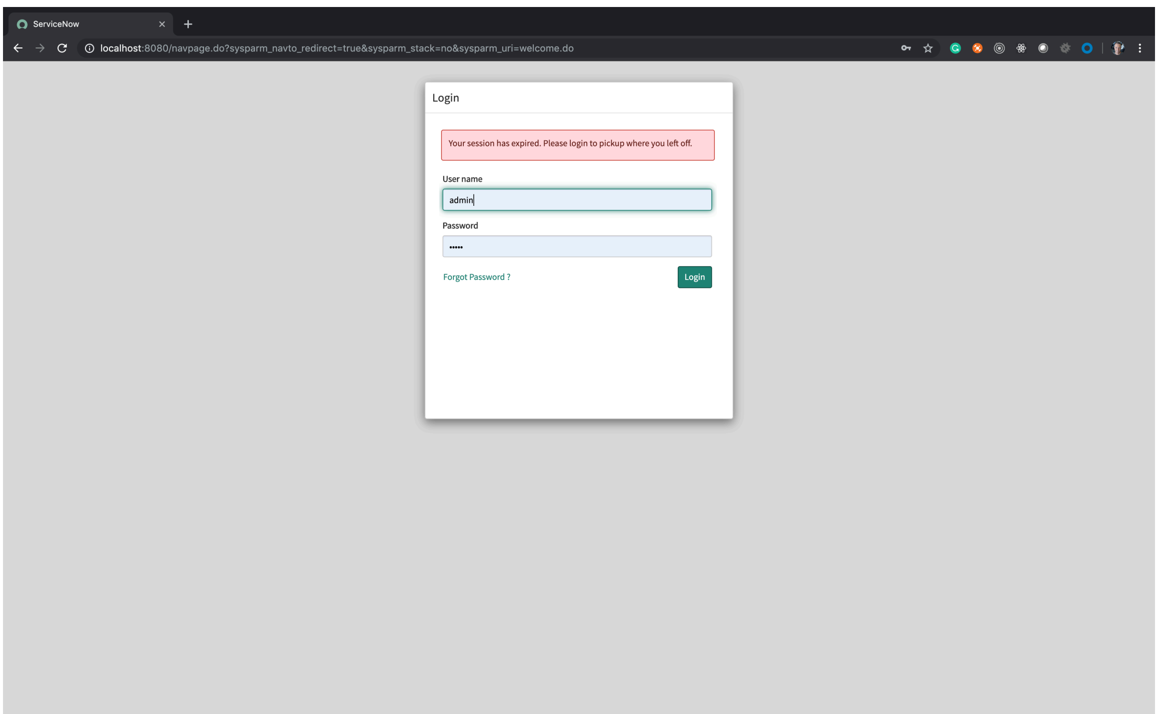

Limited to core ServiceNow AI Platform functions, not applications.

- Type: string

  Example: `background-color: #000; opacity: 0.5`
- Default value: empty (no overlay)
- Location: System Property [sys_properties] table

glide.analytics.cache.enabled

When true, caching capability is turned on for data visualizations and Platform Analytics dashboards. Caching still needs to be turned on for individual dashboard and data visualizations in UI Builder and the in-line dashboard editor, respectively.

- Type: true | false (Boolean)
- Default value: false. Turned on automatically when the Data Visualizations 3.0 app from the ServiceNow Store is installed.
- Location: System property [sys_properties] table

glide.analytics.cache.expiration.age.in.seconds

Value in seconds of the expiration_time field in analytics cache records. Applies only when glide.analytics.cache.enabled is true.

- Type: integer
- Default value: 86400 (24 hours)
- Minimum value: 3600 (1 hour)
- Location: System property [sys_properties] table

glide.angular.legacy

When true, enables legacy AngularJS behavior. When false, integrates AngularJS sanitizeUri and $http security patches. If you set this property to false, existing JSONP requests with cross-domain endpoints will fail. You can address this issue in your scripts by doing one of the following:

1. Add the domains of any newly failing requests to the *`angular.jsonp.inclusion_list.urls`* system property.
2. Opt out of this behavior by changing *`angular.jsonp.inclusion_list.enabled`* to false.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.apps.hub.current

URL of the team development parent instance.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.attachment.extensions

Comma-separated list of file extensions that can be attached. No value means that there are no restrictions.

- Type: string
- Default value: none
- Location: **System Properties** > **Security**

glide.attachment.role

Comma-separated list of roles that can create attachments.

- Type: string
- Default value: public
- Location: **System Properties** > **Security**

glide.authenticate.multisso.login_locate.user_field

Identifies a common login identifier.

- Type: string
- Default value: user_name
- Location: Add the property to the System Property [sys_properties] table.

glide.authenticate.sso.saml2.require_signed_authnrequest

Enables the Identity Provider Single sign-on service to receive a signed AuthnRequest.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.banner.image.url

URL used when clicking the banner image.

- Type: string
- Default value: home.do
- Location: System Property [sys_properties] table

glide.banner.image.url_target

Target frame used when clicking the banner image. Use `gsft_main` for the main frame, `_top` to replace the current browser window,`_blank` for a new window or tab.

- Type: string
- Default value: gsft_main
- Location: System Property [sys_properties] table

glide.basicauth.update_last_login_time

Controls whether the Last login and Last login time fields in a user record [sys_user] are updated when an integration user sends a web request to an instance using basic authentication. If true, the Last login and Last login time fields in the user's record are updated with web requests.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.bsm.color.affect_neighbors

Color of an affected neighbor node. When a node has a service issue, all nodes that depend on that node are considered affected nodes. In the map, the affected nodes are parents or grandparents of the node with the service issue.

- Type: color
- Default value: Beige
- Location: **BSM Map** > **Map Properties**

glide.bsm.map.style.font_family

Font family name used in the map text. If you designate a font that is not on the user system, the browser substitutes another font and the text may not render as you expect.

- Type: font name
- Default value: Arial
- Location: **BSM Map** > **Map Properties**

glide.bsm.map.style.font_size

Font size of the text that appears with a CI node. The default size is magnified for nodes with more connections and reduced for downstream nodes.

- Type: integer
- Default value: 14
- Location: **BSM Map** > **Map Properties**

glide.bsm.map.style.selection_background_color

Background color of a selected CI node. This color is also used with a node's Highlight Hierarchy option.

- Type: color
- Default value: RoyalBlue
- Location: **BSM Map** > **Map Properties**

glide.bsm.map.style.selection_text_color

Color of the text that appears under a selected CI node.

- Type: color
- Default value: White
- Location: **BSM Map** > **Map Properties**

glide.bsm.map.style.text_color

Color of the text that appears under an unselected CI node.

- Type: color
- Default value: Black
- Location: **BSM Map** > **Map Properties**

glide.bsm.max_levels

Maximum level depth from the root CI that can be initially displayed in business service maps. Level depth is the graph distance between the root CI and a node. This value must be an integer.

- Type: integer (1-10)
- Default value: 5
- Location: **BSM Map** > **Map Properties**

glide.bsm.new_node_color

Color for nodes that became viewable from the last expand operation.

- Type: color
- Default value: PaleGreen
- Location: **BSM Map** > **Map Properties**

glide.bsm.max_nodes

Maximum number of downstream nodes to retrieve from the database for a CI. If more nodes exist in the database, they are not displayed in the map.

- Type: integer
- Default value: 1000
- Location: **BSM Map** > **Map Properties**

glide.bsm.refresh_interval

Seconds between each automatic reloading of troubles and tasks.

- Type: integer
- Default value: 30
- Range of possible values: 1 to 3600
- Location: **BSM Map** > **Map Properties**

glide.bsm.task_threshold

Change the CI glyph color from orange to red when the number of tasks reaches this threshold.

- Type: integer
- Default value: 3
- Location: **BSM Map** > **Map Properties**

glide.bsm.too_many_children

Maximum number of child nodes to display. Nodes are collapsed for the map to meet this limit.

- Type: integer, valid values 1 or greater
- Default value: 10
- Location: **BSM Map** > **Map Properties**

glide.businessrule.async_condition_check

Specifies whether the instance checks the condition statement of async business rules (the **When** field is set to async) prior to running the business rule. If **true**, the instance evaluates the async business rule conditions a second time before running the rule. If **false** or missing, the instance does not evaluate the conditions a second time.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.businessrule.callstack

Specifies whether business rule executions are logged in the daily log when they start and finish (**true**) or are not logged (**false**). Logging is useful for troubleshooting

a problem where you must know which business rules are running and in which order.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.cache.size.service_cache

- Type: integer
- Default value: 20
- Location: Add the property to the System Property [sys_properties] table.

glide.canvas.grid.widget_render_concurrent_max

Defines the maximum number of widgets that can render simultaneously on a dashboard. With smaller values, individual widgets load more quickly. With larger values, fewer Ajax requests to the server are needed. Set the value of this property to half of the number of widgets that are visible when your most-used dashboard loads. For example, if six widgets are visible on the dashboard set the value to 3.

Note:
This system property applies to responsive dashboards only.

Two is the minimum value, and 1 is interpreted as 2. For values of zero or lower, all widgets load simultaneously.

- Type: integer
- Default value: none
- Location: System Property [sys_properties] table

glide.cmdb.query.batch_time_limit_in_sec

Time limit (in seconds) for running one batch to get one batch of query results (100 results).

- Type: integer
- Default value: 300
- Location: **Configuration** > **CMDB Properties** > **Query Builder Properties**

glide.cmdb.query.max_results_limit

Limits the number of results for a scheduled query and in the results section in the Query Builder when you click **Load All Results**.

- Type: integer
- Default value: 10,000
- Location: **Configuration** > **CMDB Properties** > **Query Builder Properties**

glide.cmdb.query.non_cmdb.black_listed_tables

Exclusion list of non-CMDB tables that appear in the CMDB Query Builder when creating a CMDB query.

- Type: string
- Default value: empty
- Location: **Configuration** > **CMDB Properties** > **Query Builder Properties**

glide.cmdb.query.query_time_limit_in_sec

Time limit (in seconds) for running an entire query to get all results.

- Type: integer
- Default value: 1,800
- Location: **Configuration** > **CMDB Properties** > **Query Builder Properties**

glide.cmdb.statemgmt.max_bulk_count

Maximum number of CIs that CI Lifecycle Management can process in a bulk update operation.

- Type: integer
- Default value: 1000
- Location: System Property [sys_properties] table

glide.cmdb.statemgmt.max_lease_expired_days

Maximum number of days that lease expiration can be set with for CI Actions.

- Type: integer
- Default value: 15
- Location: System Property [sys_properties] table

glide.cms.dashboards.sharing_with_secure_search

When enabled, the list of users, user groups, and roles that are visible in the sharing panel is restricted. The restriction is based on the configuration of the sys_user, sys_user_group, and sys_user_role ACLs. For more information, see Access control list rules.

There is a performance penalty associated with enabling this property. Performance degradation may be based on the number and complexity of business rules and ACLs on your instance.

- Type: true | false
- Default value: false
- Location: **System Properties** > **Dashboard Properties**

glide.cms.share_dashboards.role

When this property is empty, all users can share their own dashboards.

Note:
This property does not apply to users with the admin and dashboard_admin roles. Users with these roles can always share any dashboard.

- Type: string
- Default value: empty
- Location: **System Properties** > **Dashboard Properties**

glide.chart.decimal.precision

Controls the rounding precision of non-currency numeric values displayed on reports. A **Decimal Precision** value specified in a report's style options overrides this property. Currency values always have a precision of 2.

- Type: integer
- Default value: 2
- Maximum value: 4
- Location: Add the property to the System Property [sys_properties] table.

glide.chat.invite_fields

(Chat plugin) Comma-separated list of fields (glide_list or references to sys_user or sys_group) used to generate the invitations when creating a chat room from a task. Users can select a check box for each of the specified fields to invite to the chat room users that are referenced by the fields.

- Type: string
- Default value: none
- Location: **Social IT** > **Chat Administration** > **Properties**

glide.chat_room.create_roles

(Chat plugin) Comma-separated list of roles that are allowed to create chat rooms.

- Type: string
- Default value: none
- Location: **Social IT** > **Chat Administration** > **Properties**

glide.chat.show_emoticons

(Chat plugin) Enables (**true**) or disables (**false**) rendering text emoticons as images.

- Type: true | false
- Default value: true
- Location: **Social IT** > **Chat Administration** > **Properties**

glide.class.downgrade.enabled

Flag for allowing class downgrade during identification and reconciliation.

- Type: true | false
- Default value: true
- Location: System Properties [sys_properties] table

glide.class.switch.enabled

Flag for allowing class switching during identification and reconciliation.

- Type: true | false
- Default value: true
- Location: System Properties [sys_properties] table

glide.class.upgrade.enabled

Flag for allowing class update during identification and reconciliation.

- Type: true | false
- Default value: true
- Location: System Properties [sys_properties] table

glide.cmdb.health.metricProcessor.maxRunningTime

Max time in minutes for which individual metric processor runs in each scheduled cycle.

- Type: integer
- Default value: 120
- Location: Navigate to **All** > **Configuration** > **Health Preference**, and then click **System Properties** in the right hand-side navigator.

glide.cmdb.health.src.cmdb_health_audit_only

When set to **true**, disables health results from sources other than CMDB Health audit (such as cloud discovery). Only results generated by CMDB Health audit appear in the CMDB dashboard.

For example, by default, if a CI is determined to be stale by Discovery, then that CI appears as stale in the CMDB dashboard even though CMDB Health audit did not determine that CI to be stale. To disable these stale CI health results, set the property to **true**.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.cmdb.health.staleness_exclude_dependent_cis

Exclude dependent CIs for the staleness CMDB Health metric.

When enabled, dependent CIs are not checked for staleness, regardless of any staleness or inclusion rules that are defined for the respective CI types.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.cmdb.logger.source.identification_engine

Enable and configure what type of details the system logs when using IRE outside the scope of identification simulation. For example, when using an API, ECC queue or scheduled jobs.

- Type: string
- Default value: info, warn, error, debug, or debugVerbose
- Location: Add the property to the System Property [sys_properties] table.

Note:
Depending on the setting, the system can generate large amounts of data that might affect overall system performance. Set the value with caution, limiting the level of details and use time to the minimum necessary for testing or debugging.

glide.cmdb.logger.use_syslog.CMDBHealth

A comma-separated list that controls the level of logging of CMDB Health jobs. Logging creates entries in the system logs to capture messages generated by the health auditing process each time they run. This logging helps debugging if there is a failure.

- Type: string
- Default value: error
- Location: System Property [sys_properties] table

glide.cmdb_model.display_name.shorten

If **true**, generates shorter display names for product models if the name of the product model contains the manufacturer name.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.cookies.http_only

Enables (**true**) or disables (**false**) the generation of HTTP-only cookies. Set this property to **false** to use Approval with E-Signature.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.cost_mgmt.calc_actual_cost

(Cost Management plugin) When an expense line is created for any task with a **Type** of **Planned task**, sums all task expense lines and adds the total to the **Work cost** field on the task record.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.cost_mgmt.debug

(Cost Management plugin) Enables debugging of cost management processing. All logging events are recorded in the Financial Management Log [fm_log] table.

Note:
This feature can generate a large number of log records. Enable the feature only during initial testing or when troubleshooting.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.cost_mgmt.process_task_cis

(Cost Management plugin) Creates expense lines affected configuration items when creating a task expense line.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.cost_mgmt.service_allocation.method

(Cost Management plugin) Defines whether business service to cost center allocation costs should be calculated based on total units or allocated units.

For example, if the service supports 100 units but only has 50 units allocated, and cost center ABC is allocated 25 units. Using the **All units** method results in an allocation cost of 25/100, or 25% of the total cost of the service. In this case, we track unallocated costs separately. Using **Allocation units** method results in an allocation cost of 25/50, or 50% of the total cost of the service. In this case, all costs are allocated (recovered).

- Type: choice list
- Default value: all_units
- Location: System Property [sys_properties] table

glide.csv.export.line_break

Enables the user to control how exported CSV data appears in Notepad. Valid values are LF for a line feed between records and CRLF for a carriage return followed by a line feed.

- Type: string
- Default value: LF
- Location: System Property [sys_properties] table

glide.csv.use_row_currency

If false, exports currency and price field values in the user's currency. If true, exports currency and price field values in the currency of the row.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.custom.ip.authenticate.allow

Comma-separated list or dash-separated range of IP addresses that are allowed access to view the stats.do, threads.do, and replication.do pages.

- Type: string
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.
- Instance Security Hardening Settings: Performance monitoring IP restriction (instance security hardening)

glide.db.aggregates.trend.use_iso_week

If set to true, reports weekly trend by data according to ISO weeks. If false, data trended per week reports according to the first day of the week set by glide.ui.date_format.first_day_of_week.

- Type: true | false
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.

glide.db.audit.ignore.delete

Specifies the tables where the sys_audit_delete file is not updated when records are deleted.

- Type: string — comma-separated list of tables
- Location: System Property [sys_properties] table

glide.db.forced.chunk.threshold

Sets the threshold above which to forces chunk record deletions and updates. Chunking helps to prevent non-primary key deletions on huge tables from causing replication problems. Use this property to set a force chunk threshold when using the GlideRecord – DeleteMultiple and GlideRecord – UpdateMultiple methods.

- Type: integer
- Default value: 100000000 (100 million)
- Location: Add the property to the System Property [sys_properties] table.

glide.db.clone.allow_clone_target

Enables (**true**) or disables (**false**) use of a non-production instance as the target for a system clone.

- Type: true | false
- Default value: **false** for production instances, **true** for non-production instances
- Location: System Property [sys_properties] table

glide.db.impex.XMLLoader.max.file.size.mb

Controls the maximum file size allowed when importing an XML file, in megabytes. Attempting to import an XML file larger than this limit results in an error.

> **Warning:**
> Use caution when modifying this property. Uploading a large XML file may impact performance and can cause an instance outage.

- Type: integer
- Default value: 100
- Location: **System Properties** > **Import Export**

glide.db.large.threshold

Sets the number of rows above which a table is considered large and uses a different method of querying for results. When this property is absent, the instance uses the default query method.

- Type: integer
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.

glide.db.loguser

Display (**true**) or suppress display (**false**) of database errors (such as `Unique Key violation detected by database (Duplicate entry 'ABC' for key 'name') to user.`) to logged-in user.

- Type: true | false
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.

glide.db.max.aggregate.size

Sets the maximum number of groups that a grouped report or list renders. Larger values may affect system performance.

- Type: integer
- Default value: 100
- Location: Add the property to the System Property [sys_properties] table.

glide.db.max_view_records

Sets the maximum number of records returned when running a GlideRecord query in a script. Do not use values larger than the default because they may cause queries to consume excessive memory on the application server and can, in extreme cases, cause a system outage. This property does not control the maximum number of records that appear in a list, report, or exported file.

- Type: integer
- Default value: 10000
- Location: Add the property to the System Property [sys_properties] table.

glide.db.table.update_inactive_choices_enabled

Specifies whether choice table selections marked as inactive (`sys_choice inactive attribute = true`) should be loaded into client databases when applications are installed (**true**) or not (**false**).

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.debug.log_point

Specifies whether the log points feature should be enabled or not.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.debugger.log.ui

Specifies whether logs should be displayed under forms and lists and in the Session Log tab, or only in the Session Log tab.

- Type: choice list
- Default value: Page
- Location: System Property [sys_properties] table

glide.discovery.application_mapping

Enables or disables the Application Dependency Mapping (ADM) feature.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.discovery.log_debug_info

An optional property to populate debug information in Discovery log. If this property is set to **true**, the Discovery sensor extracts the debug information from ECC input message and writes this information to the Discovery log table, so that it is visible when inspecting discovery status.

- Type: true | false
- Default: false

glide.discovery.log_message_chars

When a log message is longer than this value, ServiceNow creates a preview of the message with an ellipsis at the end of the message to indicate that there is content that is not shown. The preview size prevents any one list row from taking up the entire screen.

- Type: integer
- Default value: 200 (characters)
- Location: **Discovery Definition** > **Properties**

glide.discovery.multi_page_serial_mode

This property controls the processing of multi-page pattern input records. **false** sets the pages to process in parallel and **true** sets the pages to process serially.

- Type: true | false
- Default value: false

glide.discovery.sam_batch_update_size

Enables batching the updates to table cmdb_sam_sw_install when only the field last_scanned in the table has changed. This optimizes the overall discovery execution as it reduces the multiple single update queries on the database. The default value is 100 which means 100 records at a time in the table are batched and updated. By setting this property to 1, this switches to sequential updates.

- Type: integer
- Default value: 100

glide.discovery.use_cmdb_identifiers

Controls whether Discovery uses the CMDB Identification and Reconciliation Framework (introduced with the Geneva release) or the legacy identifiers from previous releases.

- Type: true | false
- Default value: true
- Location: **Discovery Definition** > **Properties**

glide.domain.notify_change

Displays a notification message telling the user that the domain picker automatically changed.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.domain.notify_record_change

Displays a notification message telling the user that the domain picker automatically changed because the record that the user is viewing changed the domain of the user.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.duplicate_ci_remediator.dry_run

Determines whether the Duplicate CI Remediator actually remediates CI duplication by updating records in the CMDB, or not.

- Type: true | false
- Default value: false
- Location: **Configuration** > **CMDB Properties** > **Duplicate CI Remediator Properties**

glide.duplicate_ci_remediator.max.cis

Threshold for the number of duplicate CIs, which if exceeded, support for reconciliation in the Duplicate CI Remediator is limited (1,000 by default).

- Type: integer
- Default value: 1000
- Location: **Configuration** > **CMDB Properties** > **Duplicate CI Remediator Properties**

Note:
This threshold never exceeds 5,000, even if you set the property to a value greater than 5,000.

glide.duplicate_ci_remediator.max.field_length

Attributes in which max_length exceeds this property value (4000 by default) are excluded from the **Select Main CI**, **Merge Attribute Values**, and **Determine Duplicate CI Actions** tabs in the Duplicate CI Remediator wizard.

- Type: integer
- Default value: 4000
- Location: **Configuration** > **CMDB Properties** > **Duplicate CI Remediator Properties**

Note:
This property impacts the performance of de-duplication tasks, therefore be cautious about setting this value.

glide.eccprobe.longrunner.class

Identifies the name of a script include that manages long running commands.

- Type: string
- Default value: LongRunner
- Location: System Property [sys_properties] table

glide.eccprobe.longrunner.retry_minutes

Allows bursts of polling errors up to specified minutes in duration.

- Type: integer
- Default value: 10
- Location: System Property [sys_properties] table

glide.eccprobe.longrunner.interval.initial_seconds

Initial polling period for long running commands in seconds.

- Type: integer
- Default value: 20
- Location: System Property [sys_properties] table

glide.eccprobe.longrunner.interval.backoff_percent

Percentage that polling period grows per poll.

- Type: integer
- Default value: 15
- Location: System Property [sys_properties] table

glide.eccprobe.longrunner.interval.max_seconds

Maximum polling period. As the polling period decays, it gets longer until it reaches this maximum.

- Type: integer
- Default value: 300
- Location: System Property [sys_properties] table

glide.eccprobe.longrunner.debug

Allows LongRunner logs status on successful polls in addition to the usual messages on unsuccessful ones when set to **true**.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.eccprobe.longrunner.max_poll_count

Maximum number of polls that is executed per 5- second polling period. This number is generally not relevant, but may be used to contain polling under extreme loading conditions.

- Type: integer
- Default value: unlimited
- Location: System Property [sys_properties] table

glide.ecmdb.all_relationship_role

An example value is: itil,asset,configuration.

- Type: string
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.

`glide.email_address_filter.max_domains`

Sets the maximum number of domains that can be associated with a single email address filter.

- Type: integer
- Default value: 100

`glide.email_address_filter.max_exceptions`

Sets the maximum number of exceptions that can be associated with a single email address filter.

- Type: integer
- Default value: 1000

glide.email.append.timezone

Specifies whether to append the time zone to all dates and times in outbound emails (**true**) or not (**false**).

- Type: true | false
- Default value: true
- Location: **System Properties** > **Email**

`glide.email_client.quick_message.insert`

Controls how quick messages are inserted in the email client. If the property is set to **true**, quick message content is inserted at the place of the cursor. If **false**, quick message content replaces existing content in an email draft.

- Type: true | false
- Default value: true

glide.email.forward_subject_prefix

Specifies the comma-separated list of prefixes in the subject line that identify a forwarded email.

- Type: string
- Default value: fw:,fwd:
- Location: Add the property to the System Property [sys_properties] table.

glide.email.journal.lines

Specifies the number of entries from a journal field (such as Additional comments and Work notes) included in email notifications. A value of **-1** includes all journal entries.

- Type: integer
- Default value: 3
- Location: **System Properties** > **Email**

glide.email.mail_to

Specifies the email address to send notifications that use the *`${mailto:}`* variable.

- Type: string
- Default value: SMTP email address that is active by default
- Location: Add the property to the System Property [sys_properties] table.

glide.email.name_split

Delimiter between the first and last name in an email address to identify users from incoming emails. For example, a delimiter of "." in the email address john.smith@company.com tells the system to look for a user record for John Smith.

- Type: string
- Default value: period (.)
- Location: **System Properties** > **Email**

glide.email.notification.save_when_no_recipients

Controls whether (**true**) or not (**false**) a notification-generated sys_mail record is saved even if there are no recipients. Used in conjunction with other notification recipient logging properties, this property enables troubleshooting problems with notifications.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.email.override.url

Sets the URL to use in emailed links in place of the instance URL. The URL should end with nav_to.do. An example value is: https://servicenow.customerdomain.com/production/nav_to.do.

- Type: string
- Default value: Instance URL
- Location: Add the property to the System Property [sys_properties] table.

glide.email.read.active

Enables (**true**) or disables (**false**) the inbound mail server.

- Type: true | false
- Default value: true
- Location: **System Properties** > **Email**

glide.email.reply_subject_prefix

Specifies the comma-separated list of prefixes in the subject line that identify an email reply.

- Type: string
- Default value: re:aw:,r:
- Location: System Property [sys_properties] table

glide.email.smtp.active

Specifies whether to enable (**true**) or disable (**false**) the outgoing mail server.

- Type: true | false
- Default value: true
- Location: **System Properties** > **Email**

glide.email.smtp.max_recipients

Specifies the maximum number of recipients the instance can list in the To: line for a single email notification. Notifications that would exceed this limit instead create duplicate email notifications addressed to a subset of the recipient list. Each email notification has the same maximum number of recipients.

- Type: integer
- Default value: 100
- Location: Add the property to the System Property [sys_properties] table.

glide.email.smtp.max_send

Specifies how many emails to send through each new SMTP connection. The instance establishes a new SMTP connection if there are more emails to send than the specified value.

- Type: integer
- Default value: 100
- Location: Add the property to the System Property [sys_properties] table.

`glide.email_system_address_filter.max_address_filters`

Sets the maximum number of email filters that can be associated with a single system address filter.

- Type: integer
- Default value: 100

glide.email.test.user

Specifies the comma-separated list of email addresses to which the instance sends all email messages. Typically used in non-production instances for testing purposes.

- Type: string
- Default value: none
- Location: **System Properties** > **Email**

glide.email.text_plain.strip_xhtml

Indicates whether both outbound and inbound emails that are shown in comments convert the XML to plain text (**true**) or preserve the XML (**false**).

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.email.watermark.visible

Indicates whether the watermark in email notifications is visible (**true**) or is wrapped in a hidden div tag (**false**).

Note:
Email clients that use the plain text version of the email show the watermark.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.email_client.show_sms_option

Specifies whether a check box appears in the email client for sending the message to the user's SMS device (**true**) or not (**false**). If no SMS device exists, the email client sends the message to the primary email device.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.entry.loggedin.page_ess

Specifies the page that an ESS user, who by definition has no roles, is redirected to when the user logs in.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.excel.boolean.display_value

If true, translates columns of type boolean to the user's language when exporting records in Excel format. If false, exports the values in English, regardless of the user language.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.excel.max_cells

Sets the maximum number of cells in an Excel export.

- Type: integer
- Default value: 500000
- Location: Add the property to the System Property [sys_properties] table.

glide.excel.use_user_date_format

Specifies whether Excel exports use the date/time format specified in a user's profile (**true**) or not (**false**). If **false**, exports the instance date/time format defined by the `glide.sys.date_format` property.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.export.csv.charset

Specifies the character set used to export CSV files. See Supported Character Encodings for a list of supported character encoding options.

- Type: string
- Default value: windows-1252
- Location: Add the property to the System Property [sys_properties] table.

glide.export.csv.raw.value

When **true**, raw database values are exported instead of the display values when you export to CSV. When **false**, display values are exported.

- Type: true | false
- Default value: false
- Location: **System Properties** > **Import Export**

glide.export.escape_formulas

When **true**, string values that start with the characters +, -, =, or @ are prepended with a single apostrophe when you export to CSV, XLS, or XLSX files.

- Type: true | false
- Default value: **true** for new instances starting with Istanbul
- Location: System Property [sys_properties] table for new instances starting with Istanbul. For upgraded instances, Add the property to the System Property [sys_properties] table.

glide.export.excel.general_formatting

When **true**, values in cells in exported Excel files have the general format. When **false**, the cells in exported Excel files have the text format.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.export.excel.wrap_cells

When **true**, values in cells in exported Excel files are wrapped automatically. When **false**, the width of exported Excel columns is resized to fit 256 characters and values are not wrapped.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.fdih.retry.max_count

Specifies the maximum number of retry attempts when no value is specified in **Count** for a retry policy.

- Type: Integer
- Default value: 0
- Location: System Property [sys_properties] table

glide.glidesoap.proxy_host

Specifies the proxy server hostname or IP address for SOAP clients.

- Type: string
- Default value: value of `glide.http.proxy_host`
- Location: System Property [sys_properties] table

glide.glidesoap.proxy_port

Specifies the port number for the proxy server for SOAP clients.

- Type: string
- Default value: value of `glide.http.proxy_port`
- Location: System Property [sys_properties] table

glide.guest.active.session.life_span

Sets the maximum session time for a guest session regardless of user activity, in minutes.

- Type: integer
- Default value: 0
- Location: System Property [sys_properties] table

glide.guest.session_timeout

Sets the inactive session timeout for a guest session, in minutes.

Note:
This setting decouples a guest session timeout from the `glide.ui.session_timeout` property. Any changes to the `glide.ui.session_timeout` property do not affect a guest session timeout.

- Type: integer
- Default value: 30
- Location: System Property [sys_properties] table

glide.hd.payload_multi_page.device_history.mutex_expires

Total time the thread waits to update the Discovery history before giving up.

- Type: integer
- Default value: 1 minutes
- Location: System Properties

glide.hd.payload_multi_page.device_history.mutex_spin_wait

The time between when we check if the Discovery history database is free.

- Type: integer
- Default value: 100 ms
- Location: System Properties

glide.hd.payload_multi_page.device_history.mutex_max_spins

The number of times the Discovery history is checked before giving up.

- Type: integer
- Default value: 600
- Location: System Properties

glide.hd.payload_multi_page.discolog.mutex_expires

Total time the thread waits to update the Discovery log before giving up.

- Type: integer
- Default value: 1 minutes
- Location: System Properties

glide.hd.payload_multi_page.discolog.mutex_spin_wait

The time between when we check if the Discovery log database is free.

- Type: integer
- Default value: 100 ms
- Location: System Properties

glide.hd.payload_multi_page.discolog.mutex_max_spins

The number of times the Discovery log is checked before giving up.

- Type: integer
- Default value: 600
- Location: System Properties

glide.help.default.page

Sets the overall help URL for the system if you are using context-sensitive help. This URL is used when there is not any context-sensitive help available for the form, list, or record.

- Type: string
- Default value: http://wiki.servicenow.com/
- Location: System Property [sys_properties] table

glide.history.max_entries

Sets the number of characters to display as a preview of journal input fields.

- Type: integer
- Default value: 250
- Location: System Property [sys_properties] table

glide.home.page

Specifies which page to load when a user selects a homepage from the banner.

- Type: string
- Default value: `home_splash.do?sysparm_direct=true`
- Location: System Property [sys_properties] table

glide.home.refresh_disabled

Enables (**true**) or disables (**false**) homepage refresh.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.home.refresh_intervals

Comma-separated list of refresh intervals available on homepages.

- Type: string
- Default value: 300,900,1800,3600
- Location: Add the property to the System Property [sys_properties] table.

Important:

The functionality found in homepages, arranging information from your instance to tell a story about your data, is found in dashboards on new instances. On upgraded instances with Next Experience enabled, users can view existing homepages if they have a direct URL, but they can't create or edit them. Responsive dashboards and Analytics Center dashboards take over homepage functionality.

Use the Homepage deprecation help tool to convert the homepages on your instance to responsive dashboards.

For more information, see:

- Dashboards in the Analytics Center.
- Working with responsive dashboards.

glide.html.enable_media_sites

Enables a URL option for media in the HTML Editor. The media attaches, but doesn't load in the HTML Editor.

- Type: string
- Default value: youtube.com,player.vimeo.com,vimeo.com
- Location: System Property [sys_properties] table

glide.html.escape_script

Enables (**true**) or disables (**false**) JavaScript tags in HTML fields.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.html.sanitize_all_fields

Specifies whether all HTML fields are sanitized to remove unwanted code (**true**) or not (**false**).

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table
- Instance Security Hardening Settings: HTML sanitizer (instance security hardening)

glide.http.connection_timeout

Specifies the maximum number of milliseconds an outbound HTTP request (such as Web Services) waits to establish a connection.

- Type: integer
- Default value: 10000 (10 seconds)
- Location: System Property [sys_properties] table

glide.http.outbound.max_timeout

Specifies the number of seconds that RESTMessageV2 and SOAPMessageV2 APIs wait for a response from a synchronous call. The maximum value is 30 seconds.

To define a timeout longer than 30 seconds, set the `glide.http.outbound.max_timeout.enabled` system property to **false** and use the `waitForResponse()` method to set the timeout. For example, see Asynchronous SOAPMessageV2 example. If `glide.http.outbound.max_timeout.enabled` is set to true and a value is passed in the `waitForResponse()` method, the system uses the smallest value from either the `waitForResponse()` method or the `glide.http.outbound.max_timeout` system property.

- Type: integer
- Default value: 30
- Location: Add the property to the System Property [sys_properties] table.

glide.http.outbound.max_timeout.enabled

Enables setting a timeout in the `glide.http.outbound.max_timeout` system property.

To define a timeout longer than 30 seconds, set the *glide.http.outbound.max_timeout.enabled* system property to **false** and use the *waitForResponse()* method to set the timeout. For example, see Asynchronous SOAPMessageV2 example. If *glide.http.outbound.max_timeout.enabled* is set to true and a value is passed in the *waitForResponse()* method, the system uses the smallest value from either the *waitForResponse()* method or the *glide.http.outbound.max_timeout* system property.

- Type: boolean
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.http.proxy_bypass_list

Specifies the semicolon-separated list of addresses that bypass the proxy server. Use an asterisk (*) as a wildcard character to specify all or part of an address.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_host

Specifies the proxy server hostname or IP address.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_ntdomain

Specifies the domain used to authenticate the proxy server with NTLM authentication.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_nthost

Specifies the hostname used to authenticate the proxy server with NTLM authentication.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_ntpassword

Specifies the password used to authenticate the proxy server with NTLM authentication.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_ntusername

Specifies the user name used to authenticate the proxy server with NTLM authentication.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_port

Specifies the port number for the proxy server.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_password

Specifies the password used to authenticate the proxy server.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.proxy_username

Specifies the username used to authenticate the proxy server.

- Type: string
- Default value: none
- Location: System Property [sys_properties] table

glide.http.timeout

(Web service Consumer Plugin) Specifies the maximum number of milliseconds to wait before an outbound transaction times out.

- Type: integer
- Default value: 175000 (175 seconds)
- Location: Add the property to the System Property [sys_properties] table.

glide.i18n.force_index

Enables (**true**) or disables (**false**) indexing all translated fields, regardless of the value of the table attribute text_index_translations.

- Type: true | false
- Default value: true

glide.identification_engine.batch_update_last_discovered

Controls batch update of last_discoveredfield in CIs that are being processed by the identification engine.

Set to **false** if there are business rules that apply to last_discovered field, and you want to trigger these rules when calling Identification and Reconciliation API.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.identification_engine.dependent_items_local_cache_count

For optimization, a custom number of locally cached query result entries of dependent CIs.

- Type: integer
- Default value: 10000
- Location: Add the property to the System Property [sys_properties] table.

Note:
If there is a memory issue due to optimization related to using local cache, set the *glide.identification_engine.related_items_local_cache_count* and the *glide.identification_engine.dependent_items_local_cache_count* properties to 0.

glide.identification_engine.granular_insert_locking

Determines whether to use multiple granular insert locks or single global insert lock.

Set to **false** if there are performance issues associated with the usage of multiple granular insert locks.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.identification_engine.independent_items_local_cache_count

For optimization, a custom number of locally cached query result entries of independent CIs.

- Type: integer
- Default value: 100000
- Location: Add the property to the System Property [sys_properties] table.

Note:
Setting the value to 0 avoids using local cache for independent CIs and might affect performance.

glide.identification_engine.related_items_local_cache_count

For optimization, a custom number of locally cached query result entries of related/lookup items.

- Type: integer
- Default value: 15000
- Location: Add the property to the System Property [sys_properties] table.

Note:
If there is a memory issue due to optimization related to using local cache, set the *glide.identification_engine.related_items_local_cache_count* and the *glide.identification_engine.dependent_items_local_cache_count* properties to 0.

glide.identification_engine.skip_duplicates

Controls how identification processes a small set of duplicate CIs.

- When **true**: If the number of duplicate CIs is less than the threshold specified by *glide.identification_engine.skip_duplicates.threshold*, then one of the duplicate CIs is picked as a match and gets updated. The rest of the duplicate CIs are tagged as duplicates.
- When **false**: Matching a CI fails, and an error is logged.
- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.identification_engine.skip_duplicates.threshold

Maximum number of CIs that can be in a set of duplicate CIs to allow identification to process the duplicate CIs according to the setting of *glide.identification_engine.skip_duplicates*.

If the number of duplicate CIs exceeds the threshold, then identification processes the duplicate CIs as if *glide.identification_engine.skip_duplicates* is set to **false**, regardless of the setting.

- Type: integer
- Default value: 5
- Location: Add the property to the System Property [sys_properties] table.

glide.identification_logs.max_run_ids

Maximum number of log runs that can be displayed when navigating to **Configuration** > **Identification Logs**.

- Type: integer
- Default value: 1000
- Location: **Configuration** > **CMDB Properties** > **Identification/Reconciliation Properties**.

glide.image_provider.security_enabled

Controls the security settings for images. If **true**, images are visible only to authenticated and authorized users. If **false**, images are visible to anyone with a URL to the attachment.

- Type: true | false
- Default value: true

  - New/zbooted instances: property is present and set to **true**
  - Upgraded instances: **false** if property is not present, unchanged if property is present
- Location: System Property [sys_properties] table

glide.imap.secure

Enables (**true**) or disables (**false**) SSL encryption for connections to the IMAP server.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.imap.secure_port

Specifies the communications port for IMAP secure connections.

- Type: string
- Default value: 995
- Location: Add the property to the System Property [sys_properties] table.

glide.imap.tls

Enables (**true**) or disables (**false**) starting the IMAP server in Transport Layer Security (TLS) mode.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.import.debug

Enables (**true**) or disables (**false**) debug logging for all import processes.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.import.error_message.generic

Enables (**true**) or disables display (**false**), failed imports to display a generic error instead of a verbose SQL message. This property should be enabled.

- Type: true | false
- Default value:
  - New/zbooted instances: **true** for starting with Istanbul
  - Upgraded instances: false
- Location: System Property [sys_properties] table

glide.import_excel.use_only_user_session_date_format

If true, any Date or Date/Time value in an Excel spreadsheet is converted to a user session date format in the staging table upon Excel file import.

If false, any Date/Time explicit cell type in Excel is imported into a staging table with the user session date format, and any date value that is not explicitly in a Date/Time cell type in Excel is imported into a staging table using the system date format.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table if you want to change the value to false.

glide.import.sftp.debug

Enables (**true**) or disables (**false**) additional debug logging for SFTP imports. Enabling this property causes the instance to log all outgoing and incoming messages during the SSH session.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.import.scp.debug

Enables (**true**) or disables (**false**) additional debug logging for SCP imports. Enabling this property causes the instance to log all outgoing and incoming messages during the SSH session.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.import_set.preserve.leading.spaces

Specifies whether the import process preserves leading spaces in Excel data cells (**true**) or not (**false**). When **false**, the import process removes leading spaces from Excel data cells.

Note:
The import process always removes trailing spaces from Excel data cells.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.import_set_row.dynamically_add_fields

Specifies whether an import set can add new columns to the staging table (**true**) or not (**false**). Instances that contain large numbers of import sets can sometimes become unresponsive when an import adds a column because the instance must alter every row in the staging table. Sometimes the database alter table action causes an outage. Setting this property to **false** prevents an import set from adding columns to the staging table and produces a log message. As a workaround, administrators can manually add a column to the staging table by creating a new dictionary entry and then reimporting the import set.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ir.query_method

Sets the query method for global text search. Only the **Simple query** method is supported.

- Type: choice list
- Default value: simple
- Location: **System Properties** > **Global Text Search**

glide.active.session.timeout.invalidate.session

Specifies whether integration sessions can be invalidated by configuring a maximum active session time.

- Type: Boolean true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.integrations.active.session.life_span

Sets the maximum session time for a integration session regardless of activity, in minutes.

- Type: integer
- Default value: 0
- Location: System Property [sys_properties] table

glide.integration.session_timeout

Sets the inactive session timeout for integration sessions, in minutes.

- Type: integer
- Default value: 1
- Location: Add the property to the System Property [sys_properties] table.

glide.invalid_query.returns_no_rows

Controls how invalid GlideRecord queries are handled with respect to invalid or undefined field names. When this property is **true**, invalid queries containing invalid or undefined field names always return no rows. When this property is **false** (default), if a query is incorrect, and includes invalid or undefined field names, the invalid part of the query condition is ignored and results are based on the valid part of the query.

To override this logic at the session level, execute `gs.getSession().setStrictQuery(false)`. To restore strict query, execute `gs.getSession().setStrictQuery(true)`. To use this property, change the line `gr.addQuery('table', arguments.length == 1 ? record.getRecordClassName() : tableName);` in the Workflow script include to `gr.addQuery('table', (tableName) ? tableName : record.getRecordClassName() );`.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.itil.assign.number.on.insert

Controls whether (**true**) or not (**false**) a task number is generated and assigned on load (**Create New**) or when the task is submitted. This feature helps prevent unused task numbers.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.jdbcprobeloader.retry

Sets the number of times a JDBC probe attempts to connect to a JDBC data source.

- Type: integer
- Default value: 60
- Location: Add the property to the System Property [sys_properties] table.

glide.jdbcprobeloader.retry_millis

Sets the number of milliseconds a JDBC probe waits between retry attempts to a JDBC data source.

- Type: integer
- Default value: 5000
- Location: Add the property to the System Property [sys_properties] table.

glide.jquery.legacy

When true, enables legacy JQuery behavior. When false, integrates the JQuery 1.12.3 and 2.2.3 security patches. If you set this property to false, UI pages with incorrect HTML markup may render incorrectly. You can address this issue in your scripts by closing any html element tags that are self-closing and not in the valid list. For example, change `<div />` to `<div></div>`.

- Valid self-closing elements that do not need to be updated include:
  - `<area />`
  - `<base />`
  - `<br />`
  - `<col />`
  - `<embed />`
  - `<hr />`
  - `<img />`
  - `<input />`
  - `<link />`
  - `<meta />`
  - `<param />`
  - `<source />`
  - `<track />`
  - `<wbr />`
- Type: true | false

- Default value: true
- Location: System Property [sys_properties] table

glide.knowman.search.apply_role_based_security

If **true**, honors read access of knowledge bases or articles specified for roles. If **false**, enables specified user criteria to override read access specified for roles.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.knowman.search_character_limit

Minimum number of characters required for knowledge search. Search terms with fewer than this number of characters return no results.

- Type: integer
- Default value: 3
- Location: System Property [sys_properties] table

glide.knowman.search.default_language

(Knowledge Management Internationalization Plugin v2) Default language for knowledge articles. If empty, defaults to the language of the logged-in user.

- Type: string
- Default value: empty
- Location: System Property [sys_properties] table

glide.knowman.serviceportal.seo_portals

Value contains a comma-separated list of service portal record sys_ids. The SEO indexing of knowledge articles is enabled on the portals mentioned in the list and is blocked by the noindex metatag for other portals. If at least one portal is listed, SEO suggestions display when users author or edit knowledge articles.

- Type: string
- Default value: empty
- Location: System Property [sys_properties] table

Note:
If the property value is empty, the indexing is enabled for all portals.

glide.knowman.show_language_option

(Knowledge Management Internationalization Plugin v2) Specifies whether the **Language** box appears (**true**) or is hidden (**false**) on articles with multiple translations.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.knowman.show_language_option.roles

(Knowledge Management Internationalization Plugin v2) Comma-separated list of roles that can see the Language box on articles with multiple translations.

- Type: string
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.

glide.knowman.show_links

Specifies whether the **Link** line appears (**true**) or is hidden (**false**) in a knowledge article.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ldap.allow_empty_group

Specifies whether all members can be removed from an Active Directory security group (**true**) or not (**false**).

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ldap.binary_attributes

Comma-separated list of LDAP attributes that should be converted from binary format to encoded64 strings. If you set this property, only the values listed are converted. The most common attributes are *objectSID* and *objectGUID*. These converted values are unique and can be used as the coalesce field on the LDAP import mapping. If this property is blank, ServiceNow tries to map these binary attributes without the conversion and they are not guaranteed to be unique since they are not properly converted to string values.

- Type: string
- Default value: objectsid,objectguid
- Location: System Property [sys_properties] table

glide.ldap.paging

Enables (**true**) or disables (**false**) LDAP paging query support. LDAP paging is a more efficient LDAP querying method for environments with more than 1000 users.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.legacy.excel.export

Enables (**true**) or disables (**false**) exporting to XLS format when exporting to an Excel file. By default, only XLSX export is enabled. This property does not affect the Excel web service. When **true**, this property also allows users to select XLS or XLSX as the Easy Import template format.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.list.filter_max_length

Sets a maximum character limit for the condition builder query.

- Type: integer
- Default value: 0
- Location: System Property [sys_properties] table

glide.lists.live_list_enabled

Enables a list refresh prompt.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.live_feed.company_feed_exclude_groups

Controls whether messages posted to a public group appear (**true**) or are omitted (**false**) on the Company Feed.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.live_services

(Chat plugin) Enables (**true**) or disables (**false**) Live Services, such as chat support.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.login.no_blank_password

Prevents (**true**) or allows (**false**) logins from users with blank passwords. Often, importing lists of users creates many users with blank passwords. By default, this property is set to **true** on production instances.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.login.home

Sets the default homepage that users see after login. If blank, the last page visited is used. The format is *<page>.do*.

- Type: string
- Default value: home.do
- Location: System Property [sys_properties] table

glide.max_journal_list_size

Sets the maximum size, in megabytes, of journal input fields.

- Type: integer
- Default value: 10
- Location: Add the property to the System Property [sys_properties] table.

glide.messaging.reassign.enabled

Sets property that automatically reassigns messaging if the current agent is unavailable.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.notification.recipient.include_logging

Master switch to enable/disable logging all reasons a recipient was included. If **false**, no include logging is performed.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.notification.recipient.exclude_logging

Master switch to enable/disable logging all reasons a recipient was excluded. If **false**, no exclude logging is performed.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.notification.recipient.exclude_logging.device_inactive

Logs recipients who are excluded because their chosen notification device record is marked as inactive.

- Type: true | false
- Default value: true

The `glide.notification.recipient.exclude_logging` property in the **System Properties** > **Email** module must be enabled to modify this property.

glide.notification.recipient.exclude_logging.device_schedule

Logs recipients who are excluded because the chosen notification device record's schedule field excludes it.

- Type: true | false
- Default value: true

  The `glide.notification.recipient.exclude_logging` property in the **System Properties** > **Email** module must be enabled to modify this property.

glide.notification.recipient.exclude_logging.event_creator

Logs recipients who are excluded because they initiated the notification event, such as updating an incident record, and the **Send to Event Creator** check box is cleared on the notification record.

- Type: true | false
- Default value: true

The *glide.notification.recipient.exclude_logging* property in the **System Properties** > **Email** module must be enabled to modify this property.

glide.notification.recipient.exclude_logging.invalid_email

Logs recipients who are excluded because the email address for that user is invalid, for example the @ is missing, or empty.

To modify this property, you must enable the *glide.notification.recipient.exclude_logging* property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.user_calendar_integration_disabled

Logs recipients of calendar invitations who are excluded because the **Calendar Integration** field is set to **None** on the user record.

- Type: true | false
- Default value: true

The *glide.notification.recipient.exclude_logging* property in the **System Properties** > **Email** module must be enabled to modify this property.

glide.notification.recipient.exclude_logging.user_inactive

Logs recipients who are excluded because the **Active** check box is cleared on the user record.

- Type: true | false
- Default value: true

The *glide.notification.recipient.exclude_logging* property in the **System Properties** > **Email** module must be enabled to modify this property.

glide.notification.recipient.exclude_logging.user_notification_disabled

Logs recipients who are excluded because the **Notification** field is set to **Disabled** on the user record.

To modify this property, you must enable the *glide.notification.recipient.exclude_logging* property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging

Enables or disables logging all reasons that a recipient was included. This property is a main switch. If **true**, the subsequent properties dealing with the inclusion of

logging are enabled. If **false**, none of the subsequent properties relating to the inclusion of logging are enabled.

glide.notification.recipient.include_logging.delegate

Logs recipients who are included because they are delegates of another user.

- Type: true | false
- Default value: true

The `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module must be enabled to modify this property.

glide.notification.recipient.include_logging.event_parm

Logs recipients who are included because they are in the **parm1** or **parm2** fields of the event record.

To modify this property, you muse enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_fields

Logs recipients who are included via a notification target record, such as an incident record, specified in the **Users/Groups in Field** field for the notification record. The recipient_fields are fields in the target record that contain a recipient to add. For example, if the record that triggered the notification is an incident, and the **assigned_to** field for the incident is listed in recipient_fields, that user is included as a recipient.

To modify this property, you must enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_groups.group_email

Logs recipients who are included in a group email for any group provided in the notification record's recipient_groups or the event **parm1** or **parm2** field.

To modify this property, you must enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_groups.manager

Logs recipients who are included because they manage any group provided in the notification record's recipient_groups or the event **parm1** or **parm2** field.

To modify this property, you must enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_groups.membership

Logs recipients who are included via membership in any group provided in the notification record recipient_groups or the event **parm1** or **parm2** field.

To modify this property, you must enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_users

Logs recipients who are included via notification record's **Users** field (recipient_users).

To modify this property, you must enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.subscription

Logs recipients because they are subscribed via User Notification Preferences.

To modify this property, you must enable the `glide.notification.recipient.include_logging` property in the **System Properties** > **Email** module.

- Type: true | false
- Default value: true

glide.oauth.update_last_login_time

Controls whether the Last login and Last login time fields in a user record [sys_user] are updated when an integration user sends a web request to an instance using OAuth authentication. If true, the Last login and Last login time fields in the user's record are updated with web requests.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.pdf_export_from_form_list.show_report_attrs

Enables or disables displaying the PDF page header for all PDFs generated from a list.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.pdf.max_rows

Sets the maximum number of rows in an exported PDF file.

- Type: integer
- Default value: 1000
- Location: Add the property to the System Property [sys_properties] table.

glide.pdf.font.size

Sets the font size for exported PDF files.

- Type: integer
- Default value: 8
- Location: Add the property to the System Property [sys_properties] table.

glide.pg.any_rejection_rejects

Controls the default process guide rejection handling. If **true**, the first rejection rejects the entity. If **false**, all users must reject the approval.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.phone_number_e164.allow_national_entry

When **true**, users can enter phone numbers in the local format listed in the territory selector. When **false**, users must enter phone numbers in the international format listed in the territory selector.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.phone_number_e164.display_national

When set to **true** or **form**, a Phone Number (E164) field displays phone numbers in a local format on forms but as an international format on lists. When set to **all**, a Phone Number (E164) field always displays phone numbers in a local format. When set to user, a Phone Number (E164) field only displays phone numbers in a local format when the phone number matches the locale setting of the current user.

- Type: string
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.phone_number_e164.display_territory_selector

Specifies whether to display (**true**) or hide (**false**) the territory selector. Hiding the territory selector restricts users to entering only local or national phone numbers.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.phone_number_e164.display_territory_text

When set to all, a Phone Number (E164) always displays the territory label. When set to national, a Phone Number (E164) displays the territory label only if the phone number is in local format. When set to read-only, a Phone Number (E164) displays the territory label in read-only mode regardless of whether the number is in local or global format. When set to read-only-national, a Phone Number (E164) displays the territory label in read-only mode only if the number is in local format. When set to list, a Phone Number (E164) displays the territory label in a list. When set to list-national, a Phone Number (E164) displays territory label in a list if the number is in national format. When set to none, a Phone Number (E164) does not display the territory label.

- Type: string
- Default value: read-only
- Location: Add the property to the System Property [sys_properties] table.

glide.phone_number_e164.display_users_idd

Specifies whether to display the international direct dialing prefix between the territory selector and the input box on forms (**true**) or not (**false**).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.phone_number_e164.strict

Specifies whether all phone number fields must match the display format of the field's select territory. When **true**, the phone number input box displays a red line underneath phone numbers that do not match the territory format listed in the territory selector. Users cannot save an invalid phone number. When **false**, the phone number input box displays a green line underneath phone numbers that do not match the territory format listed in the territory selector. Users can save an invalid phone number, and the territory selector offers the option to select an **Other / Unknown** territory format.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.pop3.ignore_headers

Specifies the comma-separated list of email headers that cause the instance to ignore an email message. Use the format name:value to specify email header types and values. You can use a wildcard (*) for the subtype. For example, `Content-Type:multipart/*; report-type=delivery-status;` ignores emails containing a type of multipart and a parameter of *`report-type=delivery-status`*. For syntax specifications, see http://www.w3.org/Protocols/rfc1341/4_Content-Type.html.

- Type: string
- Default value: Auto-Submitted:auto-replied,X-FC-MachineGenerated:true,Content-Type:multipart/report; report-type=delivery-status;
- Location: **System Properties** > **Email**

glide.pop3.ignore_senders

Specifies the comma-separated list of senders that cause the instance to ignore an email message. Enter only the name before the at (@) sign.

- Type: string
- Default value: mailer-daemon, postmaster
- Location: **System Properties** > **Email**

glide.pop3.ignore_subjects

Specifies the comma-separated list of strings that cause the instance to ignore an email message if they are present at the start of a subject line. Values are case-insensitive.

- glide
- Type: string
- Default value: out of office autoreply, undeliverable:, delivery failure:,returned mail:,autoreply
- Location: **System Properties** > **Email**

glide.pop3.parse_end

[Legacy] Text indicating the end of the email body section where the instance should parse name:value pairs to change field values when processing inbound email actions. This property is no longer required to set field values from the email body.

- Type: string
- Default value: none
- Location: **System Properties** > **Email**

glide.pop3.parse_start

[Legacy] Text indicating the beginning of the email body section where the instance should parse name:value pairs to change field values when processing inbound email actions. This property is no longer required to set field values from the email body.

- Type: string
- Default value: none
- Location: **System Properties** > **Email**

glide.pop3.process_locked_out

Enables (**true**) or disables (**false**) the ability for locked out users to trigger inbound actions.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.pop3.reply_separators

Comma-separated list of separators that cause the instance to disregard everything below the text string in the message body. This list is case-sensitive.

- Type: string
- Default value: \n\n-----Original Message-----,\n\n _____ \n\nFrom:
- Location: **System Properties** > **Email**

glide.pop3readerjob.create_caller

Controls the behavior when an instance receives an email from an email address not associated with a user record. If **true**, the instance creates a new user record for the email address and places that new user in the **Caller** field of any created tickets. If **false**, the instance associates the new ticket to the **Guest** user record.

- Type: true | false
- Default value: false
- Location: **System Properties** > **Email**

glide.processor.json.row_limit

Specifies the maximum number of rows a JSON query returns.

- Type: integer
- Default value: 250
- Location: Add the property to the System Property [sys_properties] table.

glide.product.description

The value of this property is displayed as text in the banner next to the company logo.

- Type: string
- Default value: Service Management
- Location: **System Properties** > **Basic Configuration UI16**

glide.product.help_url

Global setting for the URL that the help icon goes to. Value must be empty for context-sensitive help to work.

glide.product.help_show

Controls whether the help icon, help.gifx, appears in (**true**) or is omitted from (**false**) the welcome banner.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.product.image.light

The Banner image displayed for Core UI Interface.

- Type: uploaded image
- Default value: none
- Location: **System Properties** > **Basic Configuration UI16**

glide.product.icon

Stores the favicon image displayed in bookmarks, tabs, and the browser address bar.

- Type: image
- Default value: favicon.ico?v=4
- Location: **System Properties** > **System**

glide.product.name

Specifies text to use in place of Service-now.com in the browser top title bar (and in browser tabs). This text is duplicated in the banner to the right of the logo unless you add display: none; to the end of the value field within the `glide.product.name.style` property.

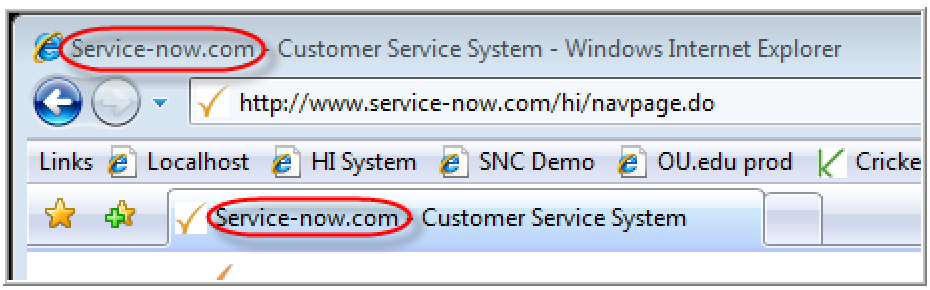

- Type: string
- Default value: ServiceNow
- Location: System Property [sys_properties] table

glide.product.name.style

Specifies the CSS properties used to display the `glide.product.name` text in the banner to the right of the logo. To not display the text, add `display: none;` to the end of the value field in this property.

- Type: string
- Default value: padding-bottom: 0px; padding-top: 0px;
- Location: System Property [sys_properties] table

glide.quota.manager.debug

Controls whether to display (**true**) or hide (**false**) additional information related to the Quota Manager, such as running transactions, canceled transactions, and which quotas are matched to transactions.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.quota.manager.heartbeat

Sets the number of seconds between the start of each Quota Manager heartbeat. This value determines how often the Quota Manager checks for transactions exceeding a quota and how often the Quota Manager writes status in the log file.

- Type: integer
- Default value: 1
- Location: Add the property to the System Property [sys_properties] table.

glide.quota.manager.minimum_transaction_time

Sets the minimum number of seconds that a transaction must run before the Quota Manager matches it to a transaction quota. You should set this value to at least 1 second because smaller values decrease performance, and because transactions shorter than 1 second are probably not worth canceling. For optimal performance, set this value to the value of your most restrictive quota. For example, if your most restrictive quota cancels transactions longer than 1 minute, set the minimum transaction time to 60 seconds.

- Type: integer
- Default value: 1
- Location: Add the property to the System Property [sys_properties] table.

glide.reconciliation.override.null

Flag for allowing or disallowing the update of an empty field by a lower priority data source.

- Type: true | false
- Default value: true
- Location: System Properties [sys_properties] table

glide.remote_glide_record.max_count

Control the maximum number of records that the GlideRecord query method returns when using the SOAP web service. The instance primarily uses this property to control the records returned when using a Perl API GlideRecord query.

> **Warning:**
> If you set this property to a large value, such as 10000, socket timeout errors might occur with MID Servers. Keep the value of this property less than 1000.

- Type: integer
- Default value: 250
- Location: Add the property to the System Property [sys_properties] table.

glide.report.add_to_list_supported

Sets whether to support the add_to_list ACL for reporting on specified columns from the Available columns list when creating or sharing a list report.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.report.calendar.max_more_events_per_day

Specifies the maximum number of calendar events that can appear in the **+ <number>** pop-up for:

- A calendar day when calendar is in month or year view
- The top 'full day' section of a calendar day when a calendar is in day or week view

When this number is exceeded, a **+ many** link appears, which opens a list of events instead of a pop-up. For more information about the maximum number of events that can be displayed in a calendar day, see system property *glide.report.calendar.max_events_displayed_per_cell*.

- Type: integer
- Default value: 30
- Location: Add the property to the System Property [sys_properties] table.

glide.report.calendar.max_events_displayed_per_cell

Defines the maximum number of events that can appear in calendar report for:

- A calendar day when calendar is in month or year view
- The top 'full day' section of a calendar day when a calendar is in day or week view

Events that exceed this value are visible via a link in the calendar cell. See *glide.report.calendar.max_more_events_per_day* for more information.

- Type: integer
- Default value: 3
- Location: Add the property to the System Property [sys_properties] table.

glide.report.new_calendar

Enables (**true**) or disables (**false**) new calendar reports. Internet Explorer 7 and 8 do not support new calendars. If you open a calendar report in one of these browsers, the old version of calendar reports is always used.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.report_home.group_report.show_usr_grp

Enables (**true**) or disables (**false**) the Reporting preferences link in the user profile.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.required.attribute.enabled

Flag for enforcing required attributes (cannot be null) during identification and reconciliation.

- Type: true | false
- Default value: true
- Location: System Properties [sys_properties] table.

glide.rest.apis.disabled

Controls which REST APIs are available on the instance, along with glide.rest.apis.enabled. The value for this property is a comma-separated list of API names, such as Table API or Aggregate API. If neither this property nor glide.rest.apis.enabled is set, all REST APIs are available. If a particular API is specified in both properties, that API is disabled. If you disable specific APIs without explicitly enabling any APIs, all REST APIs except the disabled APIs are available.

- Type: string
- Default value: all REST APIs are enabled by default
- Location: System Property [sys_properties] table

glide.rest.apis.enabled

Controls which REST APIs are available on the instance, along with glide.rest.apis.enabled. The value for this property is a comma-separated list of API names, such as Table API or Aggregate API. If neither this property nor glide.rest.apis.disabled is set, all REST APIs are available. If a particular API is specified in both properties, that API is disabled. If you enable specific APIs, only those APIs are available.

- Type: string
- Default value: all REST APIs are enabled by default
- Location: Add the property to the System Property [sys_properties] table.

glide.rest.choice.allow_non_existing_value

Enables passing non-existing values to a choice field. Does not apply to Change REST API /api/sn_chg_rest/change.

- Type: true | false
- Default value: false

  If **true**, the choice field is set to a passed value, even if the passed value does not exist in the choice field.

- Location: Add the property to the System Property [sys_properties] table.

glide.rest.debug

Logs all stages of REST processing, including processing times.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.rest.replace_request_host_with_property

Ensures that generated links in an API payload, such as links to more pages and reference fields, are based off of the instance address not the request host address.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.rollback.version

Controls whether rollback behavior is used (**true**) or not (**false**).

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.rss.max_rows

Controls the maximum number of records returned by the RSS Feed Generator.

- Type: integer
- Default value: 1000
- Location: Add the property to the System Property [sys_properties] table.

glide.sc.reset_cascade

If **true**, forces variable cascading when navigating between pages in an order guide.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.schedules.repeat_nth

Controls how a schedule entry with **Repeats** set to **Monthly** and **Monthly Type** set to **Day of the Week** is defined. Choices are **Week** or **Day**. The **Week** option is defined as choosing a day of the month in the nth week selects the nth day of the month. The **Day** option is defined as choosing the nth day of the month selects the nth day of the month.

- Type: choice list
- Default value: day
- Location: System Property [sys_properties] table

glide.schedules.fifth

Controls how a schedule entry that selects the fifth occurrence of a day in month behaves in a month containing only four occurrences of the day. Choices are **Last**, **Next**, and **Strict**. The **Last** option selects the last (fourth) occurrence of the day. The **Next** option selects the first day of the next month. The **Strict** option skips the day completely. This property is valid only when the `glide.schedules.repeat_nth` property is set to **Day**.

- Type: choice list
- Default value: last
- Location: System Property [sys_properties] table

glide.script.ccsi.ispublic

Provides privacy control over client-callable script includes that are accessed by public pages. If **false**, all client-callable script includes are private.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.script.log_level

Controls the behavior of gs.log. Set this property to **none** to disable Glide script logging, or to **print** to save log data to the file system instead of the database. Use the default value **all** to save all Glide script logs to the database.

- Type: string
- Default value: all
- Location: Add the property to the System Property [sys_properties] table.

glide.script_processor.admin

Specifies the user role necessary to access the Scripts - Background module. To require administrators to elevate privileges to access the module, set the value to **security_admin**.

- Type: string
- Default value: admin
- Location: Add the property to the System Property [sys_properties] table.

glide.script.vtable.log.debug

Enables use of debug logs for remote table script definitions. Set this property to **true**. Prolonged use of this property can affect performance, so it is best to set its value to **false** when you finish a debugging session.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.search.suggestions.enabled

Set it to **true** (default) to enable search suggestions, or set it to **false** to disable search suggestions everywhere. For more information about search suggestions, see Enable search suggestions.

glide.secondary.query.sysid

Controls whether a secondary sort using sys_id as a sort key is automatically applied when sorting requested database records on a list. This property is useful when sorting records that have multiple records found for a single value of the primary sort key. Use this property to ensure that the **Next / Previous** buttons on forms display the proper record. If **false**, there is no secondary sort, therefore records with the same primary sort value are returned as they are found in the database and may be inconsistent.

Note:
Setting this property to **true** may impact the load and sort processes. It is not supported by the Elasticsearch document store plugin.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.security.auto.resubmit.ajax

Automatically resubmits timed-out Ajax requests.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.security.checkacl.before.setvalue

Controls how the high security setting glide.security.strict.updates processes ACLs. When **true**, the instance checks ACL rules for all columns before applying any updates to a row. When **false**, the instance applies ACL rules as each column is processed in alphabetical order. This behavior may result in one update preventing another update from occurring. For example, suppose that you create a custom ACL rule to prevent updates after a record is closed. If you attempt to close an incident and also add information to a custom resolution code field in the same update, the close action is applied first because it is controlled by the [incident.state] column and the custom field by the [incident.u_resolution_code] column. After the instance applies the close action, the custom ACL rule prevents further write updates to the row.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.security.csrf.handle.ajax.timeout

Handles errors for timed out Ajax requests.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.security.csrf.strict.validation.mode

Enforces strict validation on CSRF tokens so that users cannot resubmit a request if the CSRF token does not match.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table
- Instance Security Hardening Settings: CSRF strict validation (instance security hardening)

glide.security.diag_txns_acl

Controls who can view the stats.do, threads.do, and replication.do pages. If **true**, only administrators or users from a known IP address are allowed to view the pages. If **false**, all users have access to the pages.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.
- Instance Security Hardening Settings: Performance monitoring (ACL) (instance security hardening)

glide.security.disable_ui_pages_sysparm_client_script

The system does not evaluate client scripts that are passed in URL parameters in UI pages. If you use URL parameters to load client scripts, you can add and disable this system property. Use this property only while you change the implementation to no longer depend on clients being passed in by URL parameters. Do not keep the system property disabled.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.security.mime_type.aliasset

Creates customized mime type alias sets. For example, `image/png=image/x-png`.

- Type: string
- Location: Add the property to the System Property [sys_properties] table.

glide.security.granular.create

Requires users to have write access on all individual fields on a table before they can create a record in that table.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.security.file.mime_type.validation

Enables (**true**) or disables (**false**) MIME type validation for file attachments.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table
- Instance Security Hardening Settings: Upload MIME type restriction (instance security hardening)

glide.security.strict_elevate_privilege

Forces all elevated roles to be treated equally for users with the administrator role. When enabled, administrators must explicitly, manually elevate themselves to any roles that are marked as elevated.

When disabled, administrators need only to manually elevate to the security_admin role. Other roles are automatically granted to administrators.

- Type: true | false
- Default value: **true** for new instances, **false** for upgraded instances
- Location: For new instances, this property is available on the System Property [sys_properties] table. For upgrades, Add the property to the System Property [sys_properties] table.

glide.service_portal.search_as_you_type_behavior

Select search suggestions (Suggestions) or type-ahead functionality (Typeahead). For more information about search suggestions, see Enable search suggestions.

glide.set_x_frame_options

Enables (**true**) or disables (**false**) the X-Frame-Options response header to SAMEORIGIN for all UI pages. The X-Frame-Options HTTP response header can be used to indicate whether a browser should be allowed to render a page in a <frame> or <iframe>. Set this property to **true** to avoid clickjacking attacks, by ensuring that CMS content cannot be embedded into other sites.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table
- Instance Security Hardening Settings: X-Frame-Options: SAMEORIGIN (instance security hardening)

glide.short_poll_delay

(Chat plugin) Sets the short polling delay, in milliseconds, for XMPP requests. Polling is the method by which the browser gets information from the server to send instant messages in chat.

- Type: integer
- Default value: 1000 (1 second)
- Location: **Social IT** > **Chat Administration** > **Properties**

glide.shortened_journal_length

Sets the number of characters to display as a preview of journal input fields.

- Type: integer
- Default value: 512000
- Location: Add the property to the System Property [sys_properties] table.

glide.smtp.dateformat

Specifies the date format to use for outgoing email notifications

- Type: string
- Default value: date format [sys_user.date_format] listed in the user record of the email sender.
- Location: Add the property to the System Property [sys_properties] table.

glide.smtp.default_retry

Enables (**true**) or disables (**false**) resending email when an unknown SMTP error code is encountered. The instance recognizes only the SMTP error codes defined in the *glide.smtp.defer_retry_ids* property.

- Type: true | false
- Default value: true
- Location: **System Properties** > **Email**

glide.smtp.defer_retry_ids

Specifies the comma-separated list of SMTP error codes that force the instance to resend email.

- Type: string
- Default value: 421,450,451,452
- Location: **System Properties** > **Email**

glide.smtp.fail_message_ids

Specifies the comma-separated list of SMTP error codes that prevent the instance from resending email.

- Type: string
- Default value: 500,501,502,503,504,550,551,552,553,554
- Location: **System Properties** > **Email**

glide.smtp.timeformat

Specifies the time format to use for outgoing email notifications.

- Type: string
- Default value: time format listed in email sender's user record [sys_user.time_format].
- Location: Add the property to the System Property [sys_properties] table.

glide.soap.allow_null_numeric_output

When **true**, SOAP XML responses return an empty element for numeric fields with no value. When **false**, numeric fields with no value return a value of 0 in the SOAP response.

Note:
The property `glide.wsdl.show_nillable` must be **true** to allow null values in SOAP responses.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.soap.default_security_policy

Specifies the name of SOAP security policy the instance uses when enforcing Web Services Security (WSS) for inbound requests.

- Type: string
- Default value: default security policy
- Location: Add the property to the System Property [sys_properties] table.

glide.soap.import_set_insert_serialized

Controls the processing of web service inserts. If **true**, the instance processes multiple simultaneous inserts one at a time (serially across nodes) to ensure an accurate transform. Serialized processing slows the speed at which the instance processes inserts. If **false**, multiple simultaneous inserts into an import set table result in simultaneous transforms that may produce duplicate target records due to the coalesce value being created at the same time.

Note:
Set this value to **false** only to optimize performance when the related transform map does not have a coalesce value that may be present simultaneously.

servicenow

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.soap.request_processing_timeout

Sets the maximum number of seconds that a SOAP request has to finish processing before the connection times out. This property computes a default value from the value of the property `glide.http.timeout` divided by 1000. There might be network infrastructure (such as proxy servers) in place that implements a shorter timeout. In this case, a socket timeout may occur unless this property is set to a shorter value. In general, you should set this property to a value several seconds less than the shortest socket inactivity timeout in effect anywhere in the network path between the client application and the ServiceNow instance.

- Type: integer
- Default value: 175 (value of `glide.http.timeout` divided by 1000)
- Location: System Property [sys_properties] table

glide.source_control.checksum_required

This property allows you to enable optional checksum validations and sanitizations. When set to **true**, continues to throw an error if the checksum validation of application files fails. This action prevents application developers on your instance from importing application files from source control that have been modified outside of an instance.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.source_control.checksum_quick_install

When set to **true**, bypasses the sanitization step(s) if the checksum matches. This action prevents an application developer from waiting for validation/sanitization steps to be processed if application files linked to a source control repository outside of your instance weren't modified.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.source_control.default_branch_name

Sets the default GIT repository branch to work in for application files linked to source control.

- Type: string
- Default value: sn_instances/<instance_name>
- Location: System Property [sys_properties] table

glide.spell.dictionary.en

Sets the spell checker dictionary used in the system for English users. Dictionaries are available for Brazilian Portuguese (pb.dic), Czech (cs.dic), Dutch (nl.dic), English US (en.dic), English UK (en_uk.dic), Estonian (et.dic), Finnish (fi.dic), French (fr.dic),

servicenow

German (de.dic), Hebrew (he.dic), Hungarian (hu.dic), Italian (it.dic), Polish (pl.dic), Portuguese (pt.dic), Russian (ru.dic), Spanish (es.dic), and Thai (th.dic).

- Type: choice list
- Default value: en.dic (English US)
- Location: System Property [sys_properties] table

glide.spell.dictionary.max_matches

Specifies the maximum number of spelling errors that the spell-check should detect. By default, spell-check finds only 10 spelling errors. The instance ignores any additional errors after reaching the maximum value.

- Type: integer
- Default value: 10
- Location: Add the property to the System Property [sys_properties] table.

glide.sys.activity_using_audit_direct

[Not Supported] Controls whether the record's history is generated using the Audit table (**true**) or not (**false**).

[Required] Set the value to **false** to generate history with History Sets.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.sys.audit_inserts

Specifies whether the Audit table audits inserts (**true**) or not (**false**).

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.sys.date_format

System date format for all users unless overridden in the user record.

- Type: date format
- Default value: yyyy-MM-dd
- Location: **System Properties** > **Basic Configuration U116**

glide.sys.default.tz

System timezone for all users unless overridden in the user's record.

- Type: timezone
- Default Value: none
- Location: **System Properties** > **Basic Configuration UI16**

glide.sys.domain.domain_change_notify

Domain legacy

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.sys.domain.no_change_roles

Domain legacy

- Type: string
- Default value: This field should be left empty.
- Location: System Property [sys_properties] table

glide.sys.domain.use_record_domain

Domain legacy

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.sys.domain.use_record_domain_for_client_scripts

Domain legacy

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.sys.time_format

System time format for all users unless overridden in the user's record.

- Type: time format
- Default Value: HH:mm:ss
- Location: **System Properties** > **Basic Configuration UI16**

glide.sys_reference_row_check

Specifies whether the script conditions of Access Control Rules apply to a table's reference fields (**true**) or not (**false**).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

Note:
If the `glide.sys_reference_row_check` system property is not present, or has been set to false, script conditions for Access Control Rules are not applied. This means an ACL containing scripted conditions will pass it's check as long as the other ACL criteria are met (such as role requirements).

glide.template.max_context

Specifies the maximum number of templates displayed in a form's context menu. If more than this number are available, users can click **Apply Template** on the context menu to open the reference list of templates.

- Type: integer
- Default value: 15
- Location: System Property [sys_properties] table

glide.transform.reuse_coalesce_field_value

If **true**, avoid re-evaluation of coalesce scripts to get the values.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ts.index.attachment.debug

If **true**, enables log messages for exceptions that occur when indexing attachments. You can leave this property enabled during normal operations to capture stack trace information about any exceptions.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ts.index.attachment.list_terms.debug

If **true**, the system logs all indexed terms when an attachment is indexed.

For optimal performance, set this property to **false** during normal operations. Only enable this property when you are actively debugging an issue.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ts.reindex.sys_metadata.after.upgrade

If **true**, the system automatically reindexes the sys_metadata table after an upgrade, without noticeable performance impact. To disable automatic reindexing, set the value to **false**.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.activity.displayname

Specifies whether the activity formatter shows **name** values (true) or **user_name** values (false).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.activity.email_roles

Specifies the list of roles (comma-separated) that can view email in the Activity Formatter.

- Type: string
- Default value: itil
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.activity.email.use_display

Specifies whether to display the email address or the user IDs (display value of the User table) in email headers (**true**) or not (**false**). If **true**, the instance searches for a user record with a matching email address. If the instance cannot find a matching user record, it displays the email address.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.activity_stream.form_button

Removes the activity stream button from forms.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.activity_stream.list_button

Removes the activity stream button from lists.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.activity_stream.page_size

Size of pages for Core UI form activity stream. A value of `0` means no paging.

- Type: integer
- Default value: 0
- Location: Still appears in the System Properties [sys_properties] table, but is no longer available.

glide.ui.activity_stream.scale_animated_gifs

Creates a thumbnail of animated `.gif` files. The maximum dimensions are 525 px width and 350 px height. The animation is not preserved in the thumbnail, but previewing the image displays the animation at its original size.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.activity_stream.scale_images

Creates a thumbnail of large images to display in the activity stream. The maximum dimensions are 525 pixels wide and 350 pixels high. If the user clicks the image or attachment, the preview is at full size.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.activity_stream.style.comments

Changes the color of the left bar in activity stream comments in Core UI.

- Type: color entry, either a name, such as **blue** or a code, such as **#0000FF**.
- Default value: transparent
- Location: System Property [sys_properties] table

glide.ui.activity_stream.style.work_notes

Changes the color of the left bar in activity stream work notes in Core UI.

- Type: color entry, either a name, such as **blue** or a code, such as **#0000FF**.
- Default value: gold
- Location: System Property [sys_properties] table

glide.ui.activity.style.comments

Changes the background color of the activity stream comments in UI15.

- Type: string
- Default value: background-color: background-color: WhiteSmoke
- Location: System Property [sys_properties] table

glide.ui.activity.style.work_notes

Changes the background color of the activity stream work notes in UI15.

- Type: string
- Default value: background-color: LightGoldenRodYellow
- Location: System Property [sys_properties] table

glide.ui.allow.field.dependency.for.templates

Checks dependency for choice fields at the time of applying template, preventing invalid values to be set to the dependent field. For example, a subcategory field which depends on category field.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.allow_deep_html_validation

Allows administrators to prevent users from saving invalid HTML in a journal field.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.attachment.force_download_all_mime_types

Forces download of all MIME attachment files.

- Type: true | false
- Default value: true for new instances, false for upgraded instances
- Location: System Property [sys_properties] table
- Instance Security Hardening Settings: Force Download MIME types (instance security hardening)

glide.ui.audit_deleted_tables

Comma-separated list of system tables for which the audit history tracks deletions.

- Type: string
- Default value: sys_user, sys_user_group, sys_user_role, sys_user_has_role, sys_user_grmember, sys_group_has_role, sys_security_acl_role
- Location: System Property [sys_properties] table

glide.ui.auto.recovery

Allows users to recover unsaved changes while working in the Studio.

- Type: true | false
- Default value: true
- Location: **Auto Recovery** > **Properties**

glide.ui.auto.recovery.exclude.field.types

Comma-separated list of field types to exclude from automatic recovery.

- Type: string
- Default value: none
- Location: **Auto Recovery** > **Properties**

glide.ui.auto.recovery.unsupported.field.types

Comma-separated list of field types no supported for automatic recovery.

- Type: string
- Default value: password,password2,glide_encrypted,video,user_image,image

glide.ui.auto.recovery.unsupported.tables

Comma-separated list of tables to exclude from automatic recovery.

- Type: string
- Default value: v_ws_editor
- Location: **Auto Recovery** > **Properties**

glide.ui.auto_req.extend.session

If **true**, enables the user to extend their user session by selecting a homepage refresh time. If **false**, it enforces session timeout. The session timeout value is ignored when the user specifies an automatic refresh value.

For example, if the user selects **5 minutes** for automatic homepage refresh, the session is renewed every five minutes. By adding this property and setting the value to **false**, administrators can force the user session to time out even if the user page refreshes every **x** number of minutes. The user session times out after the value specified in the session timeout, plus the selected refresh value. This property takes

effect when the **Remember me** check box is not selected for the user. Tablet and mobile devices do not support this feature.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.breadcrumb_max_entries

Specifies how many choices will be shown in the condition builder breadcrumb.

- Type: integer
- Default value: 10
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.buttons_bottom

Controls whether UI actions appear at both the bottom and top of the form (**true**) or only at the top (**false**). This property only works in a deprecated version of the UI.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.cert_task_activity.fields

Defines which journal field is the task activity field.

- Type: string
- Default value: work_notes
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.chart.bar.horiz.max_col_slant_labels

Sets the maximum number of columns in a horizontal bar chart before slanting (angling) the labels.

- Type: integer
- Default value: 5
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.chart.height

Specifies the height of a chart in pixels.

- Type: integer
- Default value: 300
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.chart.pie.labels

Enables (**true**) or disables (**false**) labels on pie chart slices.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.chart.pie.labels.max_items

Sets the maximum number of pie chart slice values that can be returned to display their labels.

- Type: integer
- Default value: 8
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.clickthrough.popup

- For reference fields, enables (**true**) or disables (**false**) display of the pop-up diamond icon for reference fields, and opening a new window when clicking the icon.
- For document ID fields, enables (**true**) or disables (**false**) displaying the information icon and opening a pop-up window with the document's form. To learn more Document ID field.

> **Note:**
> Related lists do not appear on forms opened in the pop-up window.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.clickthrough.replace

Enables (**true**) or disables (**false**) both the pop-up and click-through icons for reference fields.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.date_format.first_day_of_week

Specifies which day of the week that weeks start on for calendar reports (1=Sunday, 2=Monday...).

- Type: integer
- Default value: 1
- Location: System Property [sys_properties] table

glide.ui.date_picker.first_day_of_week

Specifies the first (leftmost) day of the week for the date and date/time picker (1=Sunday, 2=Monday...).

- Type: integer
- Default value: 1
- Location: System Property [sys_properties] table

glide.ui.default.applications

Comma-separated list of application names that open by default in the navigation pane when nothing is opened via user preferences. If the property is specified and

is blank, no applications are opened in the navigation pane when no applications are opened via user preferences. If the property is not specified, the first application that is authorized for the user opens if no applications are opened via user preferences.

- Type: string
- Default value: first authorized application for the user
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.dirty_form_support

Enables (**true**) or disables (**false**) display of a confirmation message when a form has unsaved changes and the user leaves the form through any means except a submit (such as using the green back arrow, any form button, or other). This property is not supported in Safari.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.encode_module_uri

Enables (**true**, the default) or disables (**false**) encoding of module URIs in the navigation pane. For information about how ServiceNow encodes module URIs, see Encoding module URIs.

glide.ui.escape_form_message

Specifies whether strings in `GlideForm` functions are sanitized to remove unwanted code (true) or not (false).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.escape_gwindow_message

Specifies whether strings in `GlideWindow` functions are sanitized to remove unwanted code (true) or not (false).

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.export.choice_list_max_characters

Sets the maximum number of characters that will be included from a condition field type in an export to Microsoft Excel, as well as the maximum number of characters displayed in the list view condition builder in Core UI. (Note: This property has no effect on the length of condition field values in tables.)

- Type: integer
- Default value: 80
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.filter.first_day_of_week

Identifies the first day of the calendar week for the company. By default, the start of the week is Monday, meaning that the calendar week begins with

Monday and ends with Sunday. To change this behavior, add the property `glide.ui.filter.first_day_of_week` to the instance as an integer property. Set the value to the integer corresponding with the day of the week that the calendar begins on, where 1 is Sunday, 2 is Monday, and so on. The function impacts all charts and calculations where the day of the week is used as a parameter.

- Type: integer
- Default value: 2
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.first.field.reference

Enables (**true**) or disables (**false**) having the first column in a list always link to the underlying record, even if it is a reference field. For example, if the first column on an Incident list is **Assigned to**, that value links to the Incident if this property is set to **true**.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.form_annotations

Enables (**true**) or disables (**false**) form annotations, which allow you to add Custom, Section Separator, and Line Separator to a form.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.form_multiple_splits

Enables (**true**) or disables (**false**) multiple splits and end splits in the form layout configuration slushbucket.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.gauge.view

Defines which view you want to be the default for homepage gauges.

- Type: string
- Default value: portal
- Location: System Property [sys_properties] table

glide.ui.glide_list.start.locked

Controls whether a glide_list (like the watch list) starts out locked (**true**) or unlocked (**false**) on a form.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.goto_use_contains

Controls whether the **Go to** navigation performs a "contains" query (**true**) or a "greater than" query (**false**) by default.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.homepage.parallel

Enables (**true**) or disables (**false**) use of parallel rendering.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.homepage.parallelism

Sets the maximum number of threads that should cooperate on rendering any given homepage.

- Type: integer
- Default value: 2
- Location: System Property [sys_properties] table

glide.ui.homepage.preview

Enables (**true**) or disables (**false**) displaying the preview icon for lists on the homepage.

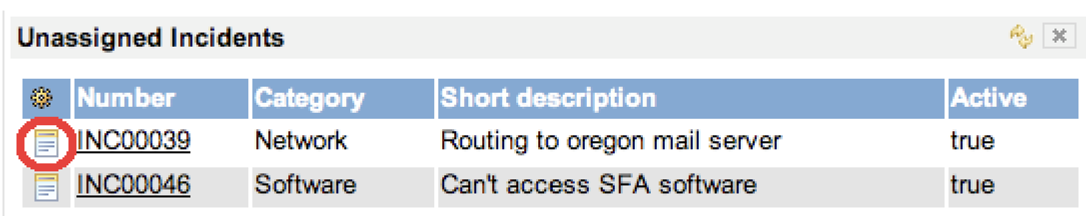

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.html.editor

Specifies which HTML field editor to use, TinyMCE or htmlArea (legacy).

- Type: string
- Default value: tinymce
- Location: **System Properties** > **UI Properties**

glide.ui.html.editor.default_target_link

Enables you to set a default target value for links in the TinyMCE.

- Type: string
- p

- Default value: ''
- Location: System Property [sys_properties] table

glide.ui.html.editor.extended_valid_elements

Defines which elements remain in the TinyMCE text when the editor saves. This functionality can be useful to add or override specific elements that should be kept.

- Type: string
- Default value: blank
- Location: System property [sys_properties] table

glide.ui.html.editor.toolbar.line1

Configures the editing toolbar (first line) for HTML fields when TinyMCE is enabled. The **spellchecker** tool is not supported in Core UI or UI15.

- Type: string
- Location: **System Properties** > **UI Properties**

glide.ui.html.editor.toolbar.line2

Configures the editing toolbar (second line) for HTML fields when TinyMCE is enabled. The **spellchecker** tool is not supported in Core UI or UI15.

- Type: string
- Location: **System Properties** > **UI Properties**

glide.ui.html.editor.v4.paste.html_import

Controls how content being pasted from sources other than Microsoft Word is filtered. This includes content copied from TinyMCE itself.

- Type: choice list
- Default value: clean
- Location: System Property [sys_properties] table

glide.ui.html.image.allow_url

Enables (**true**) or disables (**false**) uploading an image via URL from HTML Fields.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.html.editor.remove_trailing_brs

Specifies whether an empty line containing a `<br>` tag is converted to a `<div> </div>` tag for TinyMCE HTML fields. If **true**, the tag is converted. If **false**, the tag is left as `<br>`.

- Type: true | false
- Default value: true
- Location: **System Properties** > **UI Properties**

glide.ui.html.toolbar

Configures the editing toolbar for HTML fields.

- Type: string
- Default value: list of buttons
- Location: **System Properties** > **UI Properties**

glide.ui.incident_activity.fields

Defines what fields are visible in the activity formatter. If the activities are customized, the system updates this property automatically.

- Type: string
- Default value: list of fields
- Location: **System Properties** > **UI Properties**

glide.ui.incident_activity.max_addresses

Specifies the maximum number of addresses to list in an email audit record. If the number of addresses exceeds this limit, the instance truncates the list after the maximum value and displays an ellipsis character (...).

- Type: string
- Default value: 5
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.incident_alert_activity.fields

Incident communication plan activity formatter fields — the list of fields tracked from the Incident Communication Plan form in the activity formatter.

- Type: string
- Default value: opened_by, work_notes, comments, severity, estd_distruption_time, actual_disruption_time
- Location: System Property [sys_properties] table

glide.ui.journal.use_html

Specifies whether TinyMCE is available for journal fields (**true**) or not (**false**).

- Type: true | false
- Default value: false

glide.ui.js_error_notify

Displays client script errors to users with the **client_script_admin** role. Also displays a generic error message to other users who encounter a client script error.

- Type: true | false
- Default value: true
- Location: **System Properties** > **All Properties**

glide.ui.label.enable

Enables (**true**) or disables (**false**) using labels, such as Most Active, Most Recent, or user-created.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.m.clear_pasteboard_when_backgrounded

Clears the copy/paste clipboard when the application enters the background.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table. In Categories, add the mobile_post_auth category.

glide.ui.max_calendar_duration

Maximum number of days that a single calendar report entry can display.

- Type: integer
- Default value: 90
- Range of possible values: 90 to 400
- Location: System Property [sys_properties] table

glide.ui.max_calendar_records

Maximum number of records saved in a calendar report.

- Type: integer
- Default value if the property is not configured: 10,000
- Location: System Property [sys_properties] table

glide.ui.m_agents

Comma-separated list of browser agents considered as mobile browsers for the Smartphone Interface. These browsers are directed to the mobile pages instead of the full browser pages.

- Type: string
- Default value: iphone,android_phone,IEMobile,Windows Phone,iPod,Windows CE,BlackBerry,BB10
- Location: System Property [sys_properties] table

glide.ui.max_ref_dropdown

Maximum number of records for a reference field that is displayed as a choice list. If the number of available records is greater than this value, the field appears as a reference field, not as a choice list.

- Type: integer
- Default value: 25
- Location: System Property [sys_properties] table

glide.ui.mentions.default_limit

Number of users with same name displayed in Activity Stream Compose editor. When set to true, the default limit takes in a number which determines how many users with the same name are displayed.

- Type: string
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.mobile_agents

Comma-separated list of browser agents considered as mobile browsers for the Legacy Smartphone Interface. These browsers are directed to the mobile pages instead of the full browser pages.

- Type: string
- Default value: iPod,Windows CE,BlackBerry,Android,Opera Mini,IEMobile,Windows Phone,iphone
- Location: System Property [sys_properties] table

glide.ui.nav.stripe.select.maxchars

Numerical character limit for list menu choices within the nav stripe. The nav stripe appears at the top of the page when using UI11.

- Type: integer
- Default value: none
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.polaris.experience

Enables the Next Experience user interface, landing page, and Unified Navigation. For more information, see Next Experience UI.

This property does not affect anything in workspaces, such as Platform Analytics Components. Workspaces, UI Builder, and related features are always enabled.

- Type: true | false
- Default value: true on new instances, false on upgrades from Tokyo or earlier.
- Location: The System Property [sys_properties] table

glide.ui.polaris.global_search

Toggles the Next Experience search function.

- Type: true | false
- Default value: true when `glide.ui.polaris.experience` is true.
- Location: The System Property [sys_properties] table

glide.ui.polaris.history.url.param_blocklist

URL parameters that are excluded when the system identifies and clears duplicate URLs from the history. This property applies only to the Next Experience Unified Navigation, which is enabled when `glide.ui.polaris.experience` is true.

- Type: string
- Default value: ["tinyId"]
- Location: The System Property [sys_properties] table

glide.ui.polaris.list_style.enable_highlighted_value_style

Enables highlighted value style on list cells. This property applies only when `glide.ui.polaris.experience` is true.

- Type: true | false
- Default value: true when *glide.ui.polaris.experience* is true.
- Location: The System Property [sys_properties] table

glide.ui.polaris.login.show_illustrations

Enables display of illustrative graphics on the stylized login page. This property applies only when *glide.ui.polaris.experience* is true.

- Type: string
- Default value: true when *glide.ui.polaris.experience* is true.
- Location: The System Property [sys_properties] table

glide.ui.polaris.menus

Toggles the Unified Navigation menus in the header. This property applies only when *glide.ui.polaris.experience* is true, meaning that Next Experience Unified Navigation is enabled.

- Type: true | false
- Default value: true when *glide.ui.polaris.experience* is true.
- Location: The System Property [sys_properties] table

glide.ui.reference.readonly.clickthrough

Enables (**true**) or disables (**false**) reference pop-ups on read-only reference fields. Reference pop-ups and click-throughs are hidden by default if a client script, UI policy, or ACL makes the field read-only. Being able to see or click through to the target record is unrelated to the reference field being writable.

If **false**, the administrator can override the system setting for a specific field by adding the *readonly_clickthrough=true* attribute to the dictionary entry.

- Type: true | false
- Default value: false
- Location: **System Properties** > **UI Properties**

glide.ui.remember_view

Enables (**true**) or disables (**false**) use of user preferences for the last view.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.remember.me.default

Controls whether the **Remember me** check box is selected (**true**) or cleared (**false**) by default.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.report.extend_calendar_choices

Controls which field styles are applied during calendar highlighting. If **false**, field styles in only the Task table are used. If **true**, the calendar first uses field styles from the table that the report is based on. If no applicable styles exist in that table, the calendar uses field styles from the Task table.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.report.old_report_builder

Enables (**true**) or disables (**false**) use of the old report builder UI.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.section508

Enables (**true**) or disables (**false**) rendering of alternate text in place of images.

- Type: true | false
- Default value: false
- Location: **System Properties** > **UI Properties**

glide.ui.show_live_feed_activity

Enables (**true**) or disables (**false**) live feed for a record in the activity formatter.

- Type: true | false
- Default value: false
- Location: **System Properties** > **UI Properties**

glide.ui.show_template_bar

Controls whether the template bar can be displayed on the specified table. For example: `glide.ui.show_template_bar.incident = false`.

- Type: true | false
- Default value: true
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.session_timeleft

Sets how much time, in minutes, users have to extend a session before it times out due to inactivity.

- Type: integer
- Default value: 2
- Location: System Property [sys_properties] table

glide.ui.active.session.life_span

Sets the maximum session time for sessions regardless of activity, in minutes.

- Type: integer
- Default value: 0
- Location: System Property [sys_properties] table

glide.ui.session_timeout

Sets the inactive session timeout, in minutes.

- Type: integer
- Default value: 30
- Location: System Property [sys_properties] table

glide.ui.stream_icon

Toggles the displaying of the activity stream button on the specified table. For example: `glide.ui.stream_icon.hr_case = true`.

- Type: true | false
- Default value: false

glide.ui.syntax_editor.context_menu

Enables disables the context menu in script editor.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.syntax_editor.linter.eslint_config

Defines linting configurations using the ESLint utility.

- Type: string
- Default value: All recommended ESLint rules are enabled, except no-undef, no-unused-vars, no-empty, and semi.
- Location: System Property [sys_properties] table

glide.ui.table.labels

Controls whether the system uses verbose labels for table names (**true**) or literal table names (**false**).

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.ui.tablet_enabled

Enables (**true**) or disables (**false**) the tablet UI.

- Type: true | false
- Default value: true
- Location: **System Properties** > **Tablet UI Properties**

glide.ui.tablet_agents

Uses the tablet UI if one of the specified strings (comma-separated) appears in the browser user_agent header.

- Type: string
- Default value: ipad,android_tablet
- Location: **System Properties** > **Tablet UI Properties**

glide.ui.tablet.title

Displays the specified text as a brief page title for tablet UI.

- Type: string
- Default value: ServiceNow
- Location: **System Properties** > **Tablet UI Properties**

glide.ui.tablet.title.ios_webapp

Displays the specified text as the default home screen icon label used in iOS version 6 and later.

- Type: string
- Default value: ServiceNow
- Location: **System Properties** > **Tablet UI Properties**

glide.ui.task.insert

Enables (**true**) or disables (**false**) the use of **Insert** and **Insert and Stay** options on tables derived from Task (such as Incident, Change, and Problem).

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.textarea.character_counter

When **true**, displays a count of available characters for journal and multi-line text fields.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.textarea_initial_rows

Sets the number of rows initially displayed for multiline form elements. When the element is selected for editing, it is expanded. Set this value to 0 or blank to ignore this property. This functionality is used to conserve space on a form when multiline elements take up too much space.

- Type: integer
- Default value: 0
- Location: System Property [sys_properties] table

glide.ui.update_on_iterate

Controls whether updates are saved (**true**) or discarded (**false**) when a user clicks the blue arrows on a form.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.ui.ui_policy_debug

Enables (**true**) or disables (**false**) logging of UI policy processing in the JavaScript debug window

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ui.user_cookie.life_span_in_days

Sets the expiration time period for the 'remember me' cookie.

> **Note:**
> When a user accesses an instance with 'remember me' enabled, the access resets the cookie expiration period until the maximum (glide.ui.user_cookie.max_life_span_in_days) life span limit is reached.

- Type: integer
- Default value: 15 (days)

  > **Note:**
  > Maximum value 30.

- Location: System Property [sys_properties] table

glide.ui.user_cookie.max_life_span_in_days

Sets a maximum cap for the `glide.ui.user_cookie.life_span_in_days` cookie to expire before forcing a user to re-authenticate into the instance.

- Type: integer
- Default value: 30 (days)

  > **Note:**
  > Maximum value 365 days.

- Location: System Property [sys_properties] table

glide.ui.welcome.profile_link

Enables (**true**) or disables (**false**) allowing users to click their name in the welcome message and see their user profile.

- Type: true | false
- Default value: In new instances **true**. In existing instances **false**.
- Location: System Property [sys_properties] table

glide.ui11.show_switch_link

This property only affects a deprecated version of the UI. It is no longer supported. Displays (**true**) or hides (**false**) a banner link for switching between the UI11 and legacy interfaces.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

glide.update.suppress_update_version

Comma-separated list of tables for which updates are not tracked in the Versions [sys_update_version] table. You cannot compare and revert versions for tables in this list.

- Type: string
- Default value: sys_user,sys_import_set_row
- Location: Add the property to the System Property [sys_properties] table.

glide.user.default_password

Default password for new users created from incoming email. Users must reset the password at first login.

- Type: string
- Default value: password
- Location: **System Properties** > **Email**
- Instance Security Hardening Settings: Configure passwords for automatically created user accounts

glide.userauthgate.extauth.check

Enables (**true**) or disables (**false**) the *UserAuthenticationGate* checks for both external and internal authentication mechanisms.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.user.trusted_domain

Comma-separated list of trusted domains for which the instance automatically creates a guest user based on incoming emails. Use an asterisk (*) to trust all domains. If an email is not from a trusted domain, the instance processes the inbound email as a "guest user" but does not create a guest user in the instance.

- Type: string
- Default value: *
- Location: **System Properties** > **Email**

glide.workflow.model.cache.max

Maximum number of models held in the workflow cache. You must restart the instance after changing this property to apply the change.

- Type: integer
- Default value: 300
- Location: **Workflow** > **Properties**

glide.wsdl.definition.use_unique_namespace

Enables (**true**) or disables (**false**) use of a unique WSDL namespace value when publishing a ServiceNow table through web services. If **true**, the WSDL target namespace is www.service-now.com/<table name>.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.wsdl.maximum_hierarchy

Maximum number of hierarchical levels returned in a WSDL when hierarchical SOAP web service is enabled. This property has a maximum value of 3.

- Type: integer
- Default value: 3
- Location: Add the property to the System Property [sys_properties] table.

glide.wsdl.show_nillable

Enables (**true**) or disables (**false**) allowing a WSDL element to be valid if it has no element content despite a {content type} which would otherwise require content.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.ws.use_row_lock

If **true**, avoid using the current web service lock that applies before running the transform. Use the row lock that applies after the OnStart method.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

glide.xmlhttp.excessive

Sets the number of items visible in the Available half of a many-to-many or one-to-many collection box.

- Type: integer
- Default value: 100
- Location: Add the property to the System Property [sys_properties] table.

glide.xmlprocessor.use_unload_format

Enables (true) or disables (false) unloading the XML using the display_value format. To unload using this format, add the parameter `useUnloadFormat=true` to the request URL.

- Type: true | false
- Default value: false
- Location: Add the property to the System Property [sys_properties] table.

google.maps.auto_close

If **true**, automatically closes the current info window before opening a new one. Clear the check box to enable multiple pop-up windows to remain open at once.

- Type: true | false
- Default value: true
- Location: **System Properties** > **Google Maps**

google.maps.client

Client ID for Google Maps API for Business. This client ID authorizes production use of Google Maps API for Business. By default, this ID is provided by ServiceNow. ServiceNow may require you to purchase a separate Client ID based on your usage.

- Type: string
- Default value: gme-servicenow
- Location: **System Properties** > **Google Maps**

google.maps.key

Map key from Google, tied to the URL of the server. This key is the private key authorizing development use of Google Maps.

- Type: string
- Default value: empty
- Location: **System Properties** > **Google Maps**

google.maps.latitude

Starting latitude of the map. This value determines the starting latitude displayed in Google Maps.

- Type: string
- Default value: 36.008522
- Location: **System Properties** > **Google Maps**

google.maps.longitude

Starting longitude of the map. This value determines the starting longitude displayed in Google Maps.

- Type: string
- Default value: -95.221764
- Location: **System Properties** > **Google Maps**

google.maps.max_items

Maximum number of items to display on the map. This setting determines how many icons can be displayed on a map.

- Type: integer
- Default value: 500
- Location: **System Properties** > **Google Maps**

google.maps.table

Table used by the map. This setting names the table containing the name, longitude, and latitude fields.

- Type: string
- Default value: cmn_location
- Location: **System Properties** > **Google Maps**

google.maps.zoom

Starting zoom level of the map (1 is the lowest).

- Type: integer
- Default value: 4
- Location: **System Properties** > **Google Maps**

help.base.default

(Context-Sensitive Help plugin) Sets the base URL for help contexts in which ServiceNow wiki is **false** and an absolute URL is not specified.

- Type: string
- Default value: http://servicenow.com/docs/?context=
- Location: System Property [sys_properties] table

help.base.servicenow

(Context-Sensitive Help plugin) Sets the base URL for help contexts in which ServiceNow wiki is **true**

- Type: string
- Default value: http://servicenow.com/docs/?context=
- Location: System Property [sys_properties] table

mid.server.rba_debug_powershell

Enables probe-level debugging for PowerShell probes. With this property enabled, PowerShell probes write detailed activity information to the MID Server log.

- Type: true | false
- Default value: false
- Location: **Orchestration** > **MID Server Properties**

mid.version.override

Sets an override condition for the current version for all MID Servers in your environment. For more information on version control, see MID Server version control properties.

- Type: string
- Default value: none (If this property is left empty, the MID Servers get their version information from the mid.buildstamp property.)
- Location: Add the property to the System Property [sys_properties] table.

par_vis_config.data_source.can_select_indicator

Specifies roles (comma-separated) which can select indicators as data sources from the Data Visualization configuration panel. If empty, all users can select the indicator sources that they have access to.

- Type: string
- Default value: empty
- Location: System Property [sys_properties] table

par_vis_config.live_refresh_rate_min_value

Specifies the minimum interval in seconds for the **Live refresh rate** setting in the Data Visualization configuration. If set, a user can still set an empty or 0 value.

- Type: integer
- Default value: 30 (seconds)
- Location: System Property [sys_properties] table

par_viz.table_data_max_data_points

Maximum number of data points for data visualization charts based on table sources.

- Type: integer
- Default value: 10000
- Location: System Property [sys_properties] table

par_viz.table_data_max_groups

Maximum number of groups to be displayed for data visualization charts based on table sources.

- Type: integer
- Default value: 50
- Location: System Property [sys_properties] table

password_reset.activity_monitor.incident_threshold

Specifies the maximum number of lockouts before an alert is issued.

- Type: integer
- Default value: 10 (lockouts)
- Location: System Property [sys_properties] table

password_reset.activity_monitor.incident_window

Specifies the amount of time that is used for recording and counting the number of user lockouts.

- Type: integer
- Default value: 60 (minutes)
- Location: **Password Reset** > **Properties**

password_reset.captcha.ignore

Enables and disables CAPTCHA functionality.

The password reset application uses Google re-CAPTCHA as the default CAPTCHA service. To use the base system CAPTCHA, change the `password_reset.captcha.google.enabled` system property to `false`.

- Type: true | false
- Default value: false
- Location: **Password Reset** > **Properties**

password_reset.qa.num_enroll

Specifies the number of questions a user must select and answer to be enrolled in the password reset program.

Note:
This security question property can be overridden by adding a value for the *num_enroll* parameter in a security question verification.

- Type: integer
- Default value: 5 (questions)
- Location: **Password Reset** > **Properties**

password_reset.qa.num_reset

Specifies the number of questions a user must answer to verify their identity during the password reset process.

Note:
This security question property can be overridden by adding a value for the *num_reset* parameter in a security question verification.

- Type: integer
- Default value: 3 (questions)
- Possible values: integers that are less than the number specified for the *num_enroll* property.
- Location: **Password Reset** > **Properties**

password_reset.request.max_attempt

Specifies the number of password reset attempts a user has before they are locked out for a period determined by the value in *max_attempt_window*.

- Type: integer
- Default value: 3 (attempts)
- Location: **Password Reset** > Properties

password_reset.request.max_attempt_window

Specifies how long a user is locked out or prevented from changing their password after trying the maximum number of times.

- Type: integer
- Default value: 1440 (minutes)
- Location: **Password Reset** > **Properties**

password_reset.request.retry_window

Specifies the length of time before the count for password reset attempts refreshes.

- Type: integer
- Default value: 1440 (minutes)
- Location: **Password Reset** > **Properties**

password_reset.request.success_window

Specifies how long a user is locked out after they have successfully reset their password.

- Type: integer
- Default value: 10 (minutes)
- Location: **Password Reset** > **Properties**

password_reset.sms.default_complexity

Specifies the number of characters in the SMS code that are required for a user to reset their password.

Note:
This SMS code property can be overridden by adding a value for the complexity parameter in an SMS code verification.

- Type: integer
- Default value: 4 (digits)
- Location: **Password Reset** > **Properties**

password_reset.sms.expiry

Specifies the amount of time, in minutes, until the SMS code sent to the user expires.

Note:
This SMS code property can be overridden by adding a value for the expiry parameter in an SMS code verification.

- Type: integer
- Default value: 5 (minutes)
- Location: **Password Reset** > **Properties**

password_reset.sms.max_per_day

Specifies the maximum number of SMS codes that are sent to a user within one 24-hour period. When a user clicks the **Send Verification Code** button, the 24-hour period begins.

Note:
This SMS code property can be overridden by adding a value for the `max_per_day` parameter in an SMS code verification.

- Type: integer
- Default value: 10 (per day)
- Location: **Password Reset** > **Properties**

password_reset.sms.pause_window

Specifies the amount of time to pass before another SMS code can be sent to a user.

> **Note:**
> This SMS code property can be overridden by adding a value for the *pause_window* parameter in an SMS code verification.

- Type: integer
- Default value: 2 (minutes)
- Location: **Password Reset** > **Properties**

password_reset.wf.refresh_rate

Specifies how often to check the status of the workflow. Represented in the password reset process progress bar.

- Type: integer
- Default value: 90000 (milliseconds)
- Location: **Password Reset** > **Properties**

password_reset.wf.timeout

Specifies the maximum wait time, in milliseconds, for the workflow to execute. The workflow is triggered during the password reset request when the user clicks **Submit**.

- Type: integer
- Default value: 500 (milliseconds)
- Location: **Password Reset** > **Properties**

promin.show_records.limit

Specifies the maximum number of records that can be returned for an arc or node for the **Show Records** link to display. If more records are returned than the set limit, the link does not display.

- Type: integer
- Default value: 100,000 (records)
- Location: Add the property to the System Property [sys_properties] table.

remember_me.max_duration_in_days

Sets a maximum cap for the *remember_me_cookie.duration_in_days* cookie to expire before forcing a user to reauthenticate into the instance.

- Type: integer
- Default value: 30 (days)

  > **Note:**
  > Maximum value 365 days.

- Location: System Property [sys_properties] table

remember_me_cookie.duration_in_days

Sets the default expiration for the "remember me" cookie. When a user logs in with "remember me" enabled, each login within the expiration limit resets the expiration period.

- Type: integer
- Default value: 15 (days)

Note:
Maximum value 30 days.

- Location: System Property [sys_properties] table

sam.install_deletion_deadline

Defines the number of days after which a software install is deleted if not discovered with the configuration item. Use a value that is greater than the number of days between consecutive Discovery runs.

- Type: integer
- Default value: 7
- Location: System Property [sys_properties] table

security.list.internal.domains

Other domains (excluding the ones in the comma-separated list in this property) are classified as external email domains. The counts of incoming emails from these external domains display in the External Incoming Emails metric in the Instance Security Center.

- Type: choice string
- Default value: None
- Other possible values: Comma-separated list of all internal email domains in the organization.
- Location: System Property [sys_properties] table

sn_acc_vis_content.set_assigned_to

Determines if Assigned To attribute is overridden by ACC-V or not.

- Type: Boolean
- Default value: false

sn_acc_vis_content.assigned_to_user_order

Defines the source of Assigned To user and its priority for ACC-V. The order of the username sources defined is used to define the priority based on which the Assigned To user is set for a computer CI. For example, computer_system_username, logged_in_users. These two different sources decide which user name derived from which source takes preference in deciding the attribute Assigned To for a computer CI.

computer_system_username: The source is derived from the native command for Windows machine wmic computersystem get user username.

Logged_in_users: Source is OsQuery which gets logged in users from the OsQuery table.

- Type: choice string
- Default value: computer_system_username, logged_in_users

sn_acc_vis_content.persist_os_user_type

Persists the type of users that you are interested in for ACC-V.

- Type: Comma separated parameter
- Default value: local,system

sn_agent.disco_minimum_threshold_for_rediscovery_minutes

The minimum number of minutes since the last Enhanced Discovery on a given IP before a rediscovery can be triggered. This helps to prevent too frequent Enhanced Discovery in cases where a network intermittently has a break or a laptop is rebooted repeatedly.

- Type: Integer
- Default value: 60

sn_agent.disco_disable_ci_clobber_of_agentless_disco

If true, checks whether CI already exists discovered with matching criteria (discovery_source=ServiceNow) and instead of re-discovering, links the agent record to existing CI.

- Type: Boolean
- Default value: true

sn_agent.disco_ci_clobber_of_agentless_disco_threshold_days

The maximum number of days since an existing CI has been discovered using Agent-less Discovery where the existing CI links instead of re-discovering it using the Agent.#This takes effect when disco_disable_ci_clobber_of_agentless_disco is true and it has been an extended period of time (over 14 days, by default) since a non-ACC discovery_source (like ServiceNow) has discovered the CI. Once this threshold is exceeded, Enhanced Discovery is allowed to execute and clobber an existing CI regardless of whether a CI has been previously discovered by a non-ACC discovery_source. This enables the most recent discovered data to be present in the CMDB.

- Type: Integer
- Default value: 14

sn_agent.host_data_collection.disable_when_container

If true, when host data collection occurs and it is found that the Agent Client Collector is running in a container, the incoming data will not be processed and the host data collection status will be set to disabled.

- Type: Boolean
- Default value: true

sn_acc_vis_content.persist_sam_usage_metrics

Enable to leverage the SAM Basic metering functionality. If this flag is set to True, after fetching details about all the software installed, including the last accessed time via osquery, the information goes into the Software Usage [samp_sw_usage] table. The SAM plugin also needs to be enabled to persist the information.

- Type: Boolean
- Default value: true

sn_app_insights.minutes_between_triggers

How many minutes to wait before sending out a notification for an identical metric trigger.

- Type: Integer
- Default value: 30

sn_app_insights.p1_predict_factor.semaphores

The normalizing factor for data for the semaphores metric in the p1 prediction model.

- Type: Integer
- Default value: 21

sn_app_insights.p1_predict_factor.sys_load

The normalizing factor for data for the sys_load metric in the p1 prediction model.

- Type: Integer
- Default value: 923

sn_app_insights.p1_predict_max_cooldown

The number of consecutive "no p1 predicted" minutes required to exit the p1 alert state.

- Type: String
- Default value: 5

sn_app_insights.p1_predict_threshold

The minimum confidence required for the p1 predict model to predict a p1 alert state.

- Type: Integer
- Default value: 90

sn_bm_client.dashboard_display_unit

Time unit to use for KPI values on the Benchmarks dashboard. Options are **Hours** or **Days**.

- Type: choice list
- Default value: Days
- Location: System Property [sys_properties] table

sn_chg_soc.change_soc_initial_limit

Defines the number of change_request records to display on load of the Change Schedule.

- Type: Integer
- Default value: 40
- Location: System Property [sys_properties] table

sn_chg_soc.change_soc_scroll_load_limit

Defines the number of change_request records to display as the Change Schedule is scrolled.

- Type: Integer
- Default value: 20
- Location: System Property [sys_properties] table

sn_chg_soc.change_soc_total_limit

Defines the total number of change_request records that can be displayed on a Change Schedule.

- Type: Integer
- Default value: 1000
- Location: System Property [sys_properties] table

sn_chg_soc.landing_page.pinned_notification

Shows a message for pin or unpin a change schedule.

- Type: true | false
- Default value: true
- Location: System Property [sys_properties] table

sn_chg_soc.schedule_window_days

Specifies the number of days to be factored before and after the respective start/end of a change_request record when displaying blackout and maintenance window spans on the Change Schedule page.

- Type: Integer
- Default value: 30
- Location: System Property [sys_properties] table

sn_disco_cd.scheduled_job_time_out_in_mins

This property sets the value for the maximum run time of the scheduled job for Cloud Resources in minutes.

- Type: integer
- Default value: 120

sn_disco_certmgmt.cert_task_default_approval_group

This property is used to set default Approval Group name for Certificate Inventory and Management for manual approval processes for requesting new, renewing, or revoking certificates for automated workflows only. You can add more than one approval group, separated by commas. The task chooses the current domain group first and then the global domain group. If no group is found, the task will fail.

- Type: string
- Default value: none

sn_disco_certmgmt.default_cert_order_validity_period

This property is used to set the validity period of the certificate order being placed in Certificate Inventory and Management.

- Type: integer
- Default value: 1000 days

sn_disco_certmgmt.default_cert_task_priority

This property is used to configure New and Renew task priorities in Certificate Inventory and Management. Accepts values 1, 2, 3, 4, 5. If any invalid value (<1 || >5 || other strings) is provided, the value reverts to P3.

- Type: integer
- Default value: 3 (if property doesn't exist)

sn_disco_certmgmt.default_revoke_cert_task_priority

This property is used to configure Revoke task priorities in Certificate Inventory and Management. Accepts values 1, 2, 3, 4, 5. If any invalid value (<1 || >5 || other strings) is provided, the value reverts to P1.

- Type: integer
- Default value: 1 (if property doesn't exist)

sn_generative_ai.disable_dynamic_translation

Disables Dynamic Translation for large language model (LLM) conversations in Virtual Agent.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

sn_global_searchui.hide_results_count

Hides (**true**) or displays (**false**) the number of matching records for search groups and tables on the Global search page.

To modify this property, select **Enhanced Global Search UI** from the application picker before opening the System Properties table.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

sn_nowassist_va.com.glide.cs.gen_ai.enable_mid_topic_ai_search_conversational_catalogs

In Mid Topic Discovery AI Search results, change the value to true to show catalog results. By default, these Mid Topic AI Search results only show knowledge base results.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

sn_nowassist_va.synthesized_autostart_items

When synthesized response only returns a singular action, configure whether to automatically launch the action. By default, only Virtual Agent topics automatically

launch. You can configure this for whenever synthesized response returns a single conversational catalog item, a single Virtual Agent topic along with Knowledge Base information, or a single conversational catalog item along with Knowledge Base information.

- Type: string
- Default value: topic
- Location: System Property [sys_properties] table

subscription.custom_table.enforce_entitlement

You can enable guardrails around custom tables when using certain applications, such as Mobile Studio, in a production or non-production environment.

When set to **true**, the property limits the list of custom tables available in applications such as Mobile Studio to those you are entitled to use, based on the active subscriptions present on your instance. If you want to remove these guardrails, setting the property to **false** enables you to access to all custom tables, regardless of your active subscription entitlements.

- Type: Boolean
- Default value:
- Location: System Property [sys_properties] table

Note:
Setting this property to **false** enables you to use applications with custom tables that may not be entitled. See the service description for your active subscriptions to determine which custom tables (if any) you are entitled to use with selected applications.

sys_par_insights_enabled

Enables the Proactive Analytics feature.

- Type: true/false (Boolean)
- Default value: true
- Location: To change the value, create this property in the System Properties table.
- More information: Add a system property

tablet.footer.text

Text that appears on the footer of the tablet UI.

- Type: string
- Default value: Copyright ServiceNow 2012
- Location: **System Properties** > **Tablet UI Properties**

tablet.header.text

Text that appears on the right-hand side of the tablet UI header.

- Type: string
- Default value: ServiceNow
- Location: **System Properties** > **Tablet UI Properties**

Add a system property

Add or create a property to control system behavior.

Add a system property

Before you begin
Role required: admin

For more information on creating system properties for your own applications, take the training on the ServiceNow® Developer Site.

About this task
Some properties in the system aren't visible in an instance by default and must be added to the System Property [sys_properties] table. If a feature requires the addition of the property, you can add a system property.

Important:
System properties store configuration information that rarely or never changes. Each time you change or add a system property, the system flushes the cache to keep all nodes in the cluster in sync. This cache flush has a very high performance cost for one to 10 minutes, which can potentially cause an outage if done excessively. To prevent such outages, don't use a system property to store configuration information that changes more than once or twice a month. Instead, use a custom table to store regularly changing configuration information.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
   The entire list of properties in the System Properties [sys_properties] table appears.
2. Verify that the property doesn't exist by searching for the property name.
3. Select **New**.
4. Complete the System Property form using the database name of the property.

   Make sure to specify the correct data **Type** and add the new value that you want the property to use.

   Properties that you add already contain default values. You add properties to change this value.

Field	Description
Name	Name of the property you're creating.
Description	Type a brief, descriptive phrase describing the function of the property.
Choices	Comma-separated values for a choice list. If you need a different choice list label and value, use an equal sign (=) to separate the label from the value. For example, `Blue=0000FF`, `Red=FF0000`, `Green=00FF00` displays **Blue**, **Red**, and **Green** in the list, and saves the corresponding hex value in the property value field.

servicenow

Field	Description
Type	Select the appropriate data type from the list (for example, **integer**, **string**, or **true|false)**.
Value	Set the desired value for the property. All property values are stored as strings. When retrieving properties via the `gs.getProperty()` method, treat the results as strings. For example, a true|false property returns 'true' or 'false' (strings), not the Boolean equivalent.
Ignore cache	The system stores system property values in server-side caches to avoid querying the database for configuration settings. When you change a system property value, the system always flushes the cache for the sys_properties table. Use this field to determine whether to flush this property's value from all other server-side caches. The default cleared check box means that the system will flush all server-side caches and retrieve the current property value from the database. Leave the check box cleared when you want to ensure all caches have the current property value. Select the check box to ignore flushing some server-side caches, thus flushing only the cache for the sys_properties table and preserving the prior property value in all other caches. This option avoids the performance cost of flushing all caches and retrieving new property values. Typically, you should only select the check box and enable ignoring the cache when you have a system property that changes more frequently than once a month, and the property value is only stored in the sys_properties table.
Private	Set this property to true to exclude this property from being imported via update sets. Keeping system properties private prevents settings in one instance from overwriting values in another instance. For example, you may not want a system property in a development instance to use the same value as a production instance.
Read roles	Define the roles that have read access to this property.
Write roles	Defines the roles that have write access to this property.

5. Select **Submit**.
   Depending on the property name, an administrator might be able to change its value only through a new module. It may also appear in one of the Properties pages in System Properties.

   Note:
   If the **Ignore cache** check box is selected, the system flushes the server cache when the parameter is changed.

Create a system properties module

You can add a module in the application navigator to access the list of system properties. This module makes it easy to add properties to the System Properties table.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Application Menus**.
2. Search for the application you want to add the properties table to, for example **System Properties**.
   Select an application that is restricted to the admin role so that non-admin users cannot access it.
3. From the Modules related list, click **New**.
4. Complete the form fields.

Field	Description
Title	Module name. For example, **All Properties**.
Application Menu	Specifies the name of the application menu the module appears under. This field should automatically be populated with the name of the application you accessed the Modules related list from.
Link type	Specifies the type of link this modules opens. For a list of system properties, select **List of Records**.
Table	Specifies the table used by the module. Select **System Properties [sys_properties]**.

5. Click **Submit**.
6. Verify that the module was created.
   For example, navigate to **System Properties** > **All Properties**.

What to do next

If you want to include additional parameters for the list of system properties module, see Create a module.

Query join and complexity size limits

The platform uses a relational database to store data. Retrieving data can involve multiple joins to create a single result set. While these joins are usually simple, in certain cases the system may issue very large joins to bring together large numbers (>20) of tables.

Database engines normally handle multiple joins quite well. The relational model assumes joins are cheap and efficient, and this is usually true. Rarely, however, exceedingly large joins may cause a database performance issue.

To mitigate this potential complexity, use the following property to limit join complexity and size.

Navigate to **System Properties > System** and find the property with this description:

Max number of database joins per query. Smaller values cause the system to issue a larger number of less complex queries. Larger values reduce the number of queries at the cost of additional complexity per query. In the absence of known database issues stemming from large join counts, this property should remain unchanged.

This property specifies the maximum number of joins the system prefers to do. In certain edge cases, more complex queries might need to be issued, but generally no more than this number of joins will be issued. e.g., a value of 10 will result in no more than 10 tables being joined together in any given query.

The system must compensate for the fewer joins by issuing more queries to retrieve necessary data, so tuning this number down will result in more queries being sent to the database. In most cases, tuning this parameter is counterproductive.

Note:
In the absence of known database issues stemming from large join counts, this parameter should remain unchanged.

Web proxy

Several properties support Web proxy configuration.

Note:
Set these properties in the System Properties `[sys_properties]` table.

Proxy setup properties

Property	Description	Examples
glide.http.proxy_host	Specify the proxy server hostname or IP address. • Type: string • Default value: none	proxy.company.com, 192.168.34.54
glide.http.proxy_port	Specify the port number for the proxy server. • Type: string • Default value: none	8080, 9100
glide.http.proxy_username	Specify the username used to authenticate the proxy server. • Type: string • Default value: none	proxyuser

Proxy setup properties (continued)

Property	Description	Examples
glide.http.proxy_password	Specify the password used to authenticate the proxy server. • Type: string • Default value: none	password
glide.email.override.url	Set the URL to use in emailed links in place of the instance URL. The URL should end with `nav_to.do`. • Type: string • Default value: instance URL	https://servicenow.customerdomain.com/production/nav_to.do

NTLM authentication

NTLM is the most complex of the authentication protocols supported by a basic web server.

NTLM is the most complex of the authentication protocols supported by a basic web server such as HttpClient. It is a proprietary protocol designed by Microsoft with no publicly available specification. Early versions of NTLM were less secure than Digest authentication due to faults in the design. However, these were fixed in a service pack for Windows NT 4 and the protocol is now considered more secure than Digest authentication.

NTLM authentication requires that an instance of NTCredentials be available for the domain name of the server or the default credentials. Since NTLM does not use the notion of realms, HttpClient uses the domain name of the server as the name of the realm. Also, the username provided to the NTCredentials should not be prefixed with the domain:

- Correct: adrian
- Incorrect: DOMAIN\adrian

There are some significant differences in the way NTLM works compared with basic and digest authentication. These differences are generally handled by HttpClient. However, having an understanding of these differences can help you avoid problems when using NTLM authentication.

- NTLM authentication works almost exactly the same as any other form of authentication in terms of the HttpClient API. The only difference is that you need to supply NTCredentials instead of UsernamePasswordCredentials (NTCredentials actually extends UsernamePasswordCredentials so you can use NTCredentials right throughout your application, if needed).
- The realm for NTLM authentication is the domain name of the computer being connected. This can be troublesome because servers often have multiple domain names. Only the domain name that HttpClient connects to, as specified by the HostConfiguration, is used to look up the credentials. While initially testing NTLM authentication, it is best to pass the realm in as null, which is used as the default.
- NTLM authenticates a connection and not a request. So you need to authenticate every time a new connection is made and keeping the connection open during authentication is vital. For

this reason, NTLM cannot be used to authenticate with both a proxy server and the web server, nor can NTLM be used with HTTP 1.0 connections or web servers that do not support HTTP keep-alives.

Note:
Set these properties from the System Properties `[sys_properties]` table.

NTLM authentication

Property	Description	Examples
glide.http.proxy_ntusername	Specify the username used to authenticate the proxy server with NTLM authentication. • Type: string • Default value: none	username
glide.http.proxy_ntpassword	Specify the password used to authenticate the proxy server with NTLM authentication. • Type: string • Default value: none	password
glide.http.proxy_nthost	Specify the hostname used to authenticate the proxy server with NTLM authentication. • Type: string • Default value: none	nthost
glide.http.proxy_ntdomain	Specify the domain used to authenticate the proxy server with NTLM authentication. • Type: string • Default value: none	DOMAIN

Proxy servers for SOAP clients

Administrators can specify separate proxy settings for SOAP clients, such as the MID Server or ODBC Driver.

To specify a proxy server for a MID Server, see MID Server configuration

To specify a proxy server for the ODBC driver, see Configure ODBC to use proxy servers

Bypass the proxy server

Administrators can configure the instance to bypass the proxy server for specific URLs or URL patterns.

Typically, internal addresses do not need a proxy server for SOAP communications.

Bypass the proxy server

Property	Description	Examples
glide.http.proxy_bypass_list	Specify the semicolon-separated list of addresses that bypass the proxy server. Use an asterisk as a wildcard character to specify all or part of an address. • Type: string • Default value: none • Location: System Properties `[sys_properties]` table	127.0.0.1;*.internal.com;localhost

Admin Center

Admin Center provides a central hub for platform owners and admins to access platform capabilities, discover new applications, and get intelligent, actionable insights.

Use the Admin Center application to discover the capability of the available applications and how they can be implemented to work towards achieving your business goal.

Get started

Explore

Learn about Admin Center.

Discover

Discover Adoption Blueprints, recommended adoption journeys organized by business outcome.

Operate

Learn about the configurable Admin Home page features.

Configure

Learn more about how to configure applications with Configuration Hub

.

Admin Center overview

Admin Center gives a centralized control of the instance to the admins with easy access to the platform capabilities, applications, and configurable admin dashboards.

The first stop experience within Admin Center is the Admin Home page. It helps you manage and track all your administrative work within a single page experience. You can prioritize your work based on the instance data. See Admin Home for more information about Admin Home features.

Admin Center offers easy access to a number of key applications administrators need. You can access these applications directly from the **Admin** menu on the Admin Home page.

- Adoption Blueprints: Adoption Blueprints lists the available applications and features that are recommended to achieve business goals. You can target the strategic business improvements with a set of apps and features. See Adoption blueprints for more information.

  Note:
  Adoption Blueprints platform feature is available by default with instance upgrade to Tokyo and above.

- Now Assist Admin: 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 features and skills. See Now Assist Admin console for more information.
- Application Manager: As an admin, you have the ability to install, update, and manage licensed applications and plugins on your instance using the Application Manager application. See Application Manager for more information.

  Note:
  Application Manager store application is available by default with the instance upgrade of Vancouver and above.

- Security Center: Use the Security Center application to help your organization maintain the security of your ServiceNow deployments. See Security Center for more information.

  Note:
  Security Center is available by default from the Vancouver release and above with the instance upgrade. For the previous releases, you can install it from the ServiceNow store.

- Subscription Management: Proactively manage your subscriptions and monitor subscription usage on your instances using the Subscription Management application. See Subscription Management for more information.

  Note:
  The legacy version Subscription Management is available by default with the Vancouver release. You can install the new version of Subscription Management on your Vancouver instance from the ServiceNow® Store. You can also have both the legacy and new version on your Vancouver instance.

- Configuration Hub: Access and zoom in all the relevant details related to the selected application(s) using the Configuration Hub application. You can also track and view delta changes, over-allocated licenses, and customizations on a single screen. See Configuration Hub for more details.

  Note:
  Configuration Hub store application is available by default with instance upgrade of Vancouver patch 2 and above.

You can also access the Admin Experience Product Hub on the ServiceNow Communityfor additional resources about Admin Center and the applications.

Admin Home

Manage and track your daily administrative work and tools on your instance by quickly accessing the cards on the dashboard. It helps you stay updated on any high priority incidents or problems on the instance that might need immediate attention.

Watch the detailed video for more information.

Use the advanced Admin Home page to accomplish the following:

- Get notified of the critical and high priority incidents
- Stay updated about the problems and changes on the instance
- Create and switch to your personalized dashboard based on what is most important to you
- Stay updated on the current version, scheduled upgrade, and licensed applications status on your instance
- Get a comprehensive, single-page view of the delta changes on the instance

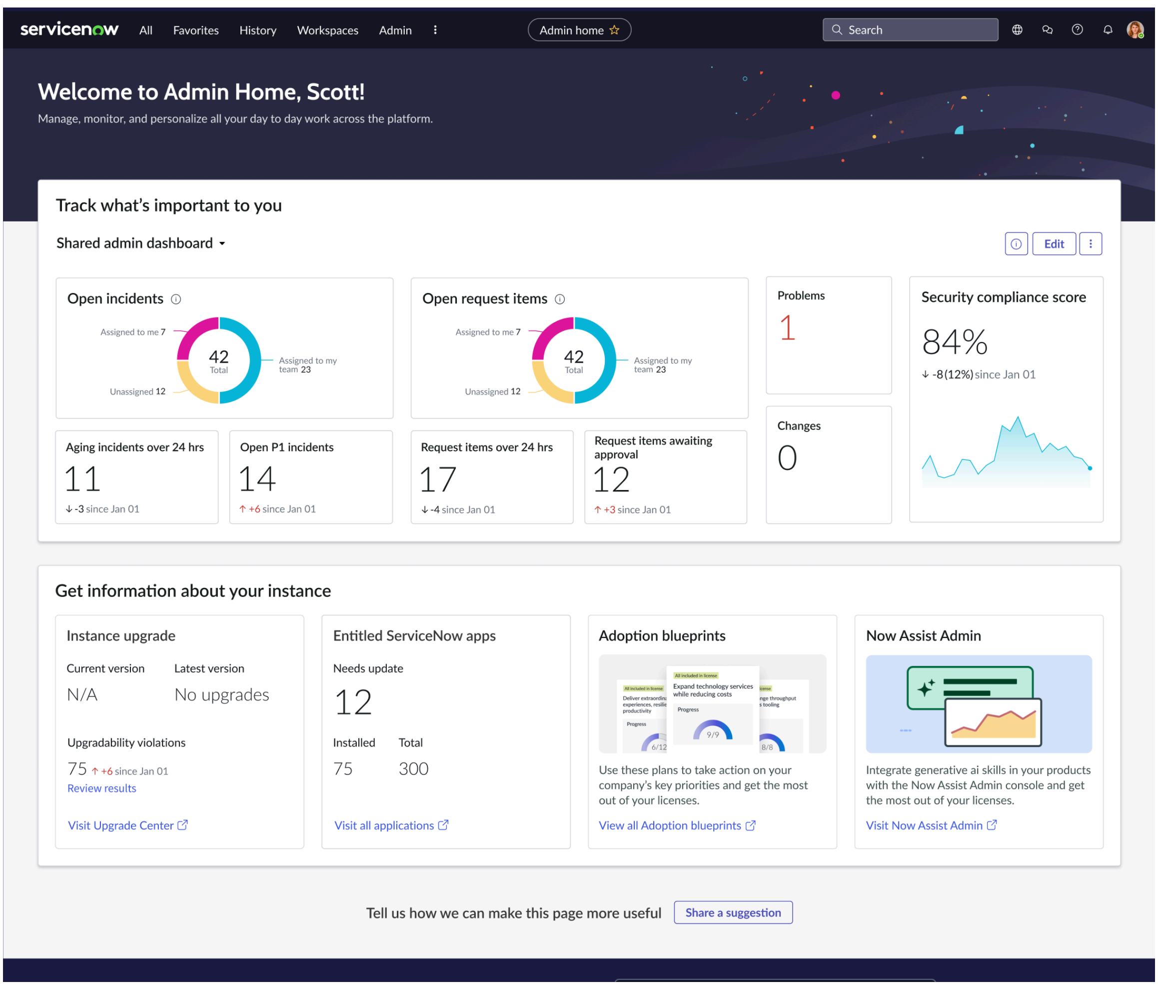

Track important information

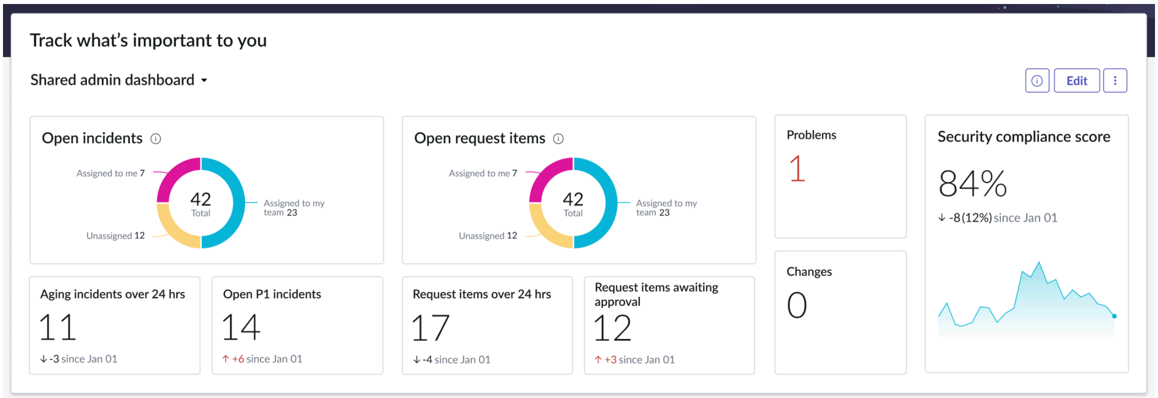

The **Track what's important to you** section of the Admin Home page provides insights into key information, such as open incidents, open requests, etc. This information is available on a default shared dashboard. You can click on each card to view more information.

Note:
You can't add or remove any cards from the Shared admin dashboard. Instead, to modify the cards, you must first create a new dashboard by selecting **Create new** dashboard. Then, modify the new dashboard as needed and share it with other administrators using the **Share** option.

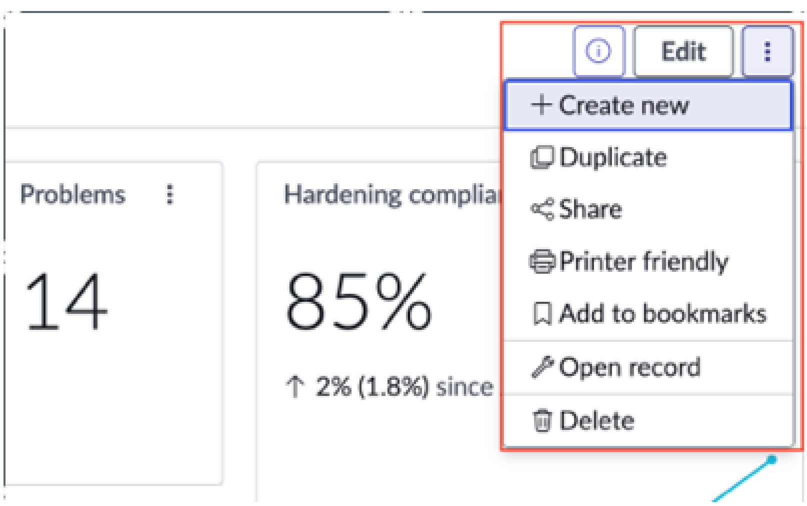

Get information about your instance

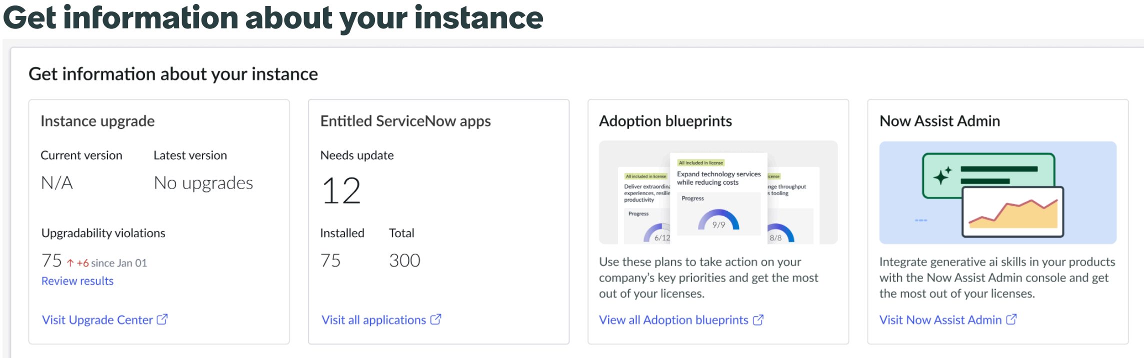

The **Get information about your instance** section of the Admin Home page shows information about your instance's current version, violations, and scheduled upgrade. It also shows the status of licensed applications as a list of installed applications on your instance and the applications that need to be updated, as well as a card with easy access to Adoption blueprints.

If your instance has one or more Now Assist plugins installed, a card displaying the status of your instance's Now Assist skills will be shown.

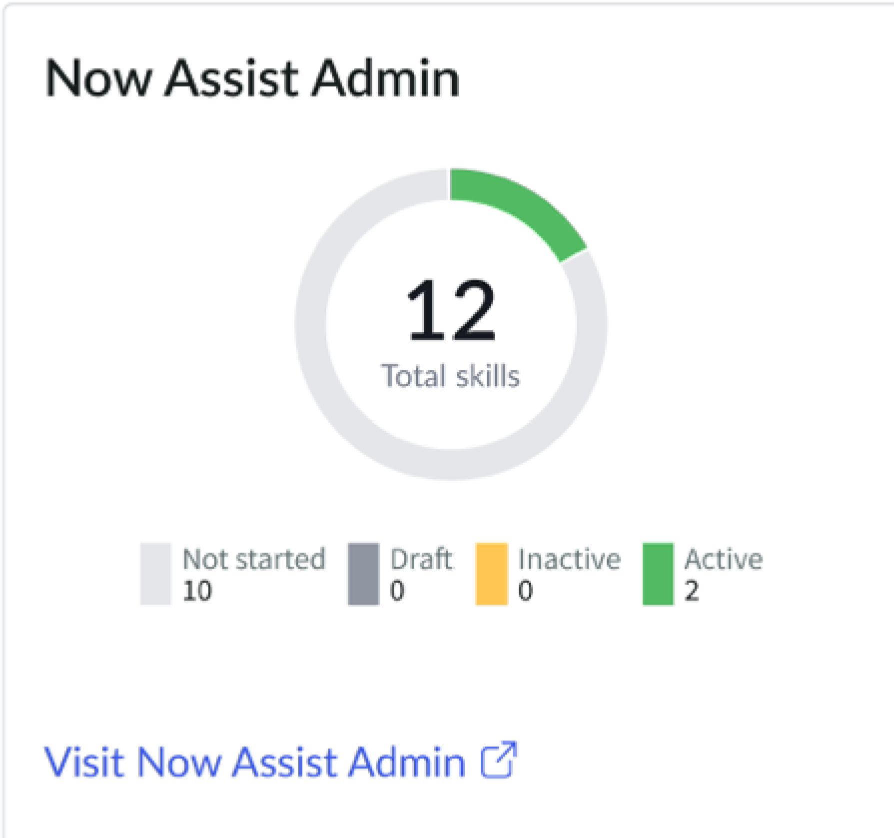

Adoption blueprints

Adoption blueprints display the applications and features available to you and the recommended solutions to achieve business goals.

Adoption blueprints overview

Adoption blueprints are the specific and measurable results that an organization needs to accomplish their growth. You can target the strategic business improvements with a set of apps and features.

Some of the goals based on which adoption blueprints are recommended:

- Optimizing operation to efficiently meet the goals of your organization
  - Reduced service operations cost for premium quality services
  - Focused on simplifying and automating transactions
- Managing risk more efficiently by improving prioritization, efficiency, and operational reporting
- Delivering seamless employee experience to increase engagement and productivity with proper utilization of resources

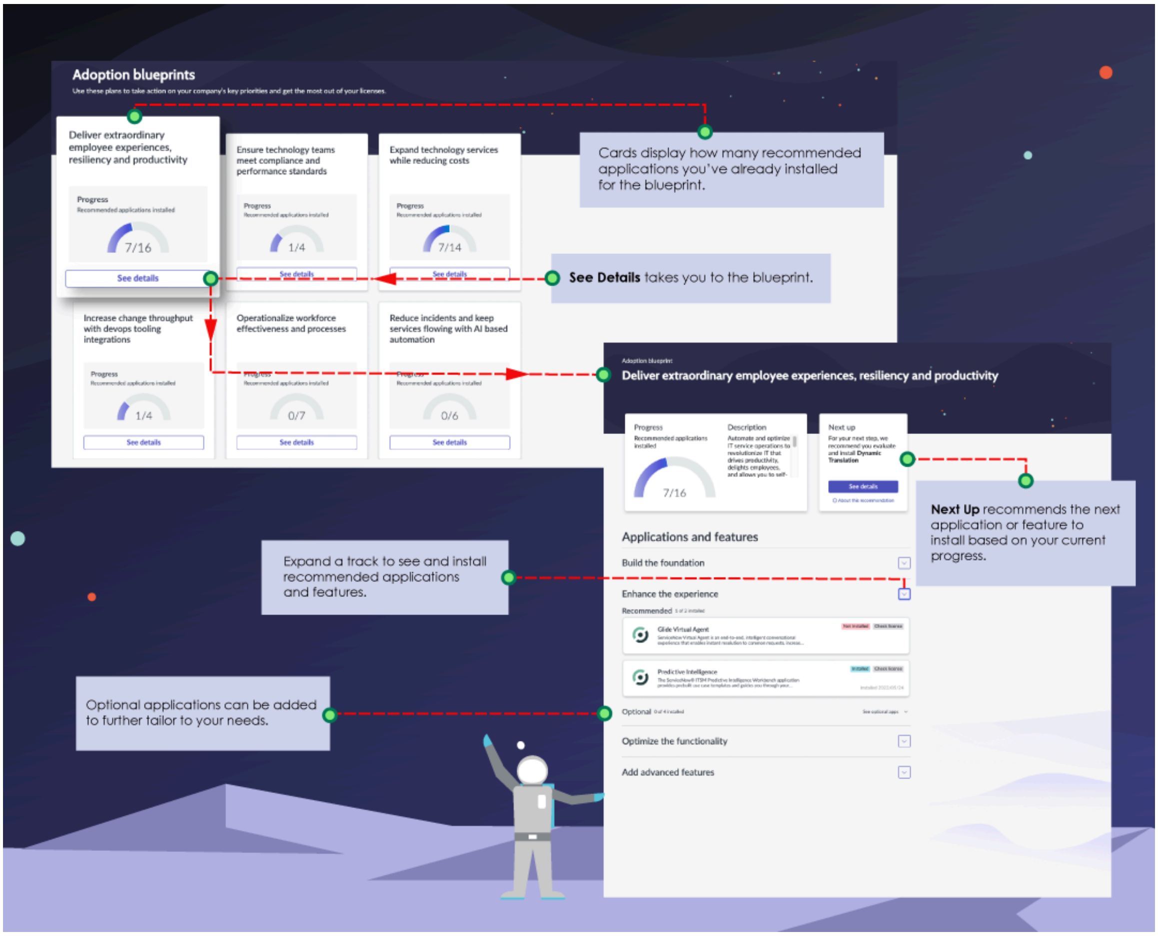

Each adoption blueprint displays the progress you have made in working towards your strategic goal. The number of applications that have been installed includes only the recommended applications.

See Select your adoption blueprint for more information.

Select **Share** if you want to share information about adoption blueprints, recommended applications and features. You can share the information with more than one recipient. The recipients are notified via email. You can also view the people with whom the details have been shared previously.

Note:
If the information is shared with a non-admin user, the share functionality provides a read-only access to the non-admin users.

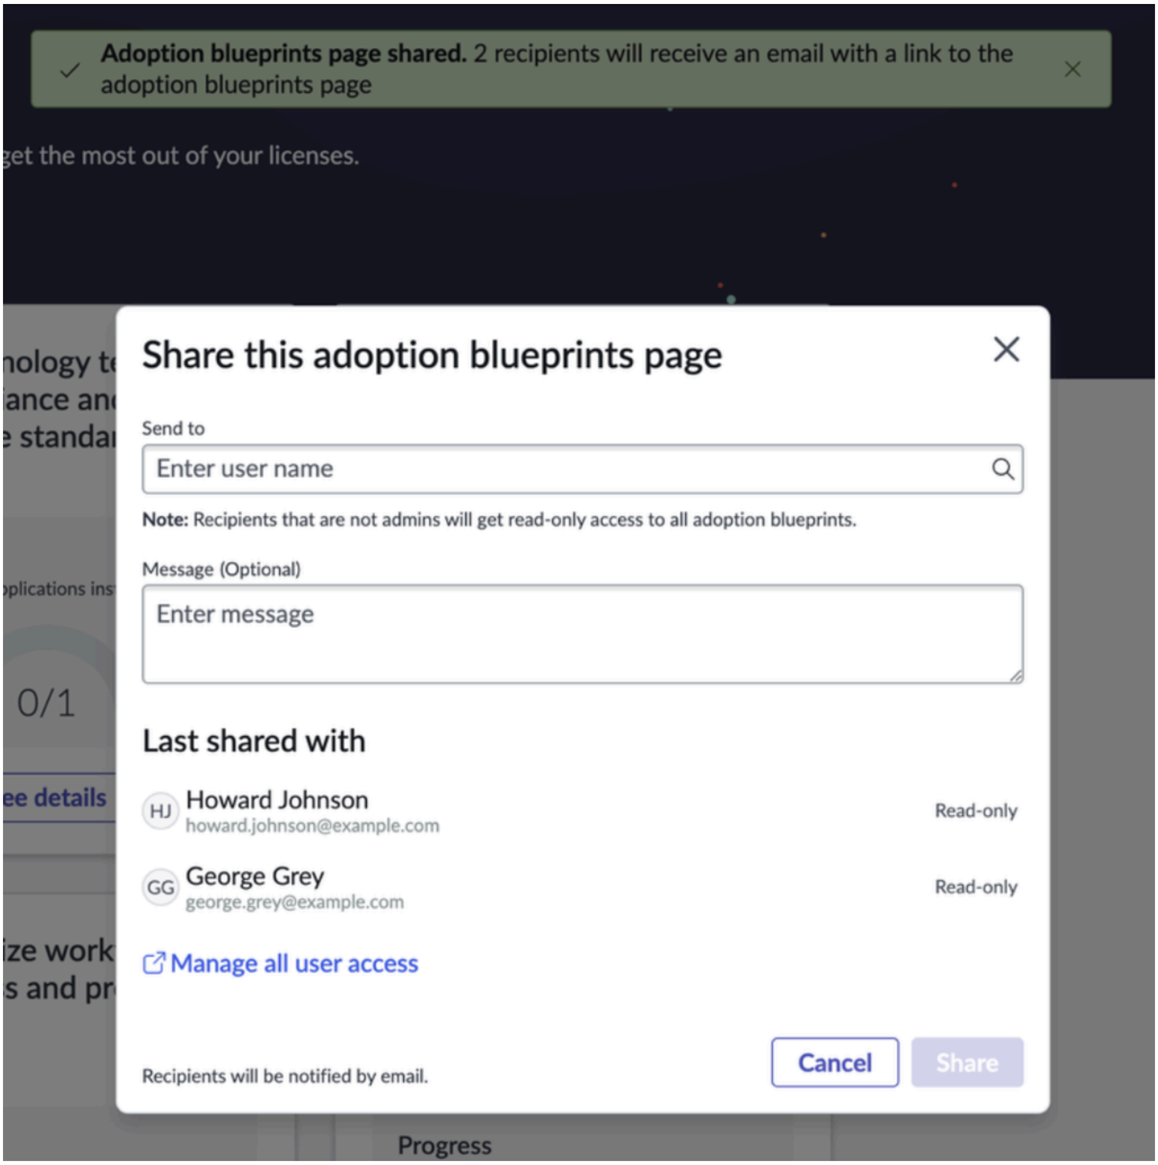

When you select the adoption blueprint, the following tracks for the recommended applications and features within the selected adoption blueprint show up.

- Build a foundation
- Enhance the experience
- Optimize the functionality
- Add advanced features

The recommended applications and features have been arranged in the order of impact in the process of giving you a measurable result. They can also be referred as crawl, walk, run, and fly phases respectively.

Note:
The applications and features mentioned in different tracks are only recommended to be installed. Some of the recommended applications and features in each track may already be installed in your instance and not require any additional action.

In each of the tracks, the applications are mentioned in the following order:

1. Recommended applications which are not yet installed
2. Recommended applications which are already installed but need to be updated
3. Recommended applications which are already installed in your instance.

Note:
The recommended applications don't need to be installed at all or in the order listed.

Select your adoption blueprint

Select an adoption blueprint that matches one of the strategic goals of your organizations. It also gives you a list of recommended applications that can be installed in the right combination to realize your business goal.

Before you begin
Role required: admin

Procedure

1. Access Adoption blueprint in one of the following ways.
2. Select an adoption blueprint from the Adoption blueprint page.
3. Select **See Details** to view the details of the selected adoption blueprint.
   The adoption blueprint details page shows the progress of the installation of the recommended applications.
4. Scroll down to see all the tracks for the selected adoption blueprint.
   See Adoption blueprints for more information.
5. **Optional:** Expand the Optional list to view the optional list of applications and features that may be beneficial for you to install.
   The total number of recommended applications and features mentioned in the progress section doesn't include the optional applications.

   Note:
   The installation of each of the recommended and optional applications and features is not mandatory to achieve the business goal.
6. Review the Next Up card to view the application recommended to install next.

   Note:
   You can skip this step if you don't want to install the application mentioned in the Next Up card. You are also required to consider the current needs of your organization along with the suggested applications.

Expand technology services while reducing costs blueprint

Select this adoption blueprint to create actionable plans to achieve your organization's business goals for automating core service and operations processes.

Note:
The applications recommended within a blueprint may include, but aren’t limited to, the following applications. Factors like entitlements and other instance dependencies can cause the number of recommended applications to vary.

Track	Recommended applications	Description	Roles installed	Additional resources
Build the foundation	Configuration Management	CMDB is a centralized source that gives you full visibility into your IT environment. By storing information about your organization's infrastructure and how it's configured, this system enables you to monitor your network and ensure stability and best performance.	cmdb_import_api_admin, cmdb_inst_admin, cmdb_ms_admin, cmdb_ms_editor, cmdb_query_builder, cmdb_query_builder_read, cmdb_read, ecmdb_admin, sn_cmdb_admin, sn_cmdb_editor, tracked_file_reader	• CMDB and Discovery deployment playbooks • CMDB 101 - What is a configuration management database and why do you need one? (ServiceNow Community) • CMDB Resources Page (KB0546686)
Enhance the experience	Service Level Management	Service Level Management (SLM) helps to gather service requirements and monitors and reports information with regard to agreed service levels (SLAs).	Roles installed with Service Level Management	• Service Level Management (SLA) Resources (KB0623502) • How SLAs Really Work (KB0529411)
Enhance the experience	CSDM and the CMDB Data Foundations Dashboards	CSDM and CMDB Data Foundations Dashboards provide insights into the key foundational metrics of your CMDB and Common Service Data Model (CSDM). It provides recommendations to ensure that the CMDB and CSDM are properly configured for optimal usage and to mitigate	Uses bases system roles. See Base system roles.	Data Foundations Dashboards for CSDM and CMDB (video)

Track	Recommended applications	Description	Roles installed	Additional resources
		any potential risks.		
Enhance the experience	Discovery	Discovery finds applications and devices on your network, and then updates the CMDB 360 with the information it finds, contributing toward self-resolution of incidents.	discovery_admin	• Discovery Welcome Guide (Community) • Discovery Best Practices (ServiceNow Community) (ServiceNow Community) • Discovery and MID Server Resources (KB0540193)
Optimize the functionality	Service Mapping	Service Mapping discovers all application services in your organization and builds a comprehensive map of all devices, applications, and configuration profiles used in these application services.	Roles installed with Service Mapping	• What is service mapping? • Service Mapping Overview (video)
Optimize the functionality	Change Management	Change Management provides a systematic approach to control the life cycle of all changes, facilitating beneficial changes to be made with minimum disruption to IT services.	Uses bases system roles. See Base system roles.	• Automate change and incident management Playbook • Change Management (ServiceNow Community) • Change Management Overview (video)

Track	Recommended applications	Description	Roles installed	Additional resources
Optimize the functionality	Incident Management	Incident Management restores normal service operation while minimizing impact to business operations and maintaining quality.	Uses bases system roles. See Base system roles.	• Incident Management limitations and support • Automate change and incident management Playbook • Incident Management Overview (video)
Optimize the functionality	Knowledge Management	Knowledge Management (KM) enables the sharing of information in knowledge bases. These knowledge bases contain articles that provide users with information such as self-help, troubleshooting, and task resolution.	Roles installed with Knowledge Management	• Knowledge Management limitations and support • Knowledge Management Quick Start Guide (ServiceNow Community)
Optimize the functionality	Problem Management	Problem Management helps to identify the cause of an error in the IT infrastructure, reported as occurrences of related incidents.	Uses bases system roles. See Base system roles.	Problem Management and 10 steps to making it successful (ServiceNow Community)
Optimize the functionality	Service Catalog	Service Catalog provides self-service opportunities for customers to request catalog items such as service and product offerings.	Roles installed with Service Catalog	Service Catalog design playbook
Optimize the functionality	Service Operations	Service Operations	Roles installed with Service	Configuration activities

Track	Recommended applications	Description	Roles installed	Additional resources
	Workspace for ITSM	Workspace is a configurable workspace that provides a unified experience for multiple IT Service Management and IT Operations Management workflows. Configure your agent experience using the easy-to-navigate interface of Service Operations Workspace for ITSM.	Operations Workspace for ITSM	for Service Operations Workspace in ITSM and ITOM (ServiceNow Community)
Optimize the functionality	Certificate Inventory and Management	Certificate Inventory and Management discovers, takes inventory of, and proactively manages all your TLS certificates. IPv6 is supported for Certificate Inventory and Management.	Roles installed with Certificate Inventory and Management	What is certificate management?
Optimize the functionality	Firewall Audits and Reporting	Firewall Audits and Reporting discovers and takes inventory of your firewall security policies, devices, device groups, and manager information.	Roles installed with Firewall Audits and Reporting	What is a firewall audit?
Add advanced features	Digital Portfolio Management	Digital Portfolio Management (DPM) manages and maintains all your services, applications, and	Roles installed with Digital Portfolio Management	• Digital Portfolio Management Resources (KB1123710)

Track	Recommended applications	Description	Roles installed	Additional resources
		products from a single location.		• Digital Portfolio Management Overview (video)
Add advanced features	Health Log Analytics	Health Log Analytics helps prevent IT issues before your users are affected. It helps you identify the root cause of an issue by enabling you to triage-related logs and analyze the raw data.	Roles installed with Health Log Analytics	• Health Log Analytics (KB1064762) • ITOM Talks Session #3 - Health Log Analytics (video)

Deliver extraordinary employee experiences, resiliency and productivity blueprint

Select this adoption blueprint to enhance your employees' IT experience by automating and optimizing IT service operations that enable self-resolution of incidents, improved productivity, and recovering from complicated IT scenarios.

Note:
The applications recommended within a blueprint may include, but aren't limited to, the following applications. Factors like entitlements and other instance dependencies can cause the number of recommended applications to vary.

Track	Recommended applications	Description	Roles installed	Additional resources
Build the foundation	Dynamic Translation	Use Dynamic Translation to enhance your experience of dynamically translating text entered in an application for a seamless localization experience.	Uses bases system roles. See Base system roles.	• Dynamic Translation limitations • Dynamic Translation overview
Build the foundation	ITSM Success Dashboard indicators	Use ITSM Success Dashboard to get an insight to the IT leadership team and process owners to measure the IT-related	Roles installed with ITSM Success Dashboard indicators	Measuring success in your ITSM deployment and how Success Dashboards can help (ServiceNow Community)

servicenow

Track	Recommended applications	Description	Roles installed	Additional resources
		performance of their implementation.		
Build the foundation	Employee Center	Use the Employee Center portal to experience a standard multi-department, dynamic portal for service delivery and employee engagement.	Roles installed with Employee Center	• Employee Center limitations and support • Employee Center (ServiceNow Community • Employee Center success story • Employee Center implementation success story
Build the foundation	Configure Walk-up Experience for Customer Service Management	Use the Walk-up Experience application to create and manage an on-site IT walk-up venue, where the requests and issues are fulfilled and solved by experienced IT technicians in person and in real time.	Roles installed with Walk-up Experience	Walk-up Experience limitations and support
Build the foundation	Service Operations Workspace for ITSM	Use the Service Operations Workspace interface to unify and configure your agent experience for multiple IT Service Management and IT Operations Management workflows.	Roles installed with Service Operations Workspace for ITSM	• Service Operations Workspace limitations and support • Service Operations Workspace overview

Track	Recommended applications	Description	Roles installed	Additional resources
Build the foundation	Survey Management	Use the Assessments and Surveys application to create, send, and collect responses for surveys in order to improve the productivity and efficiency of your processes.	Roles installed with Survey Management	• Survey Management limitations and support • Survey Management overview
Build the foundation	Benchmarks	Use the Benchmarks application to compare the performance of your organization with other recognized industry standards and implement recommendations for improved performance.	Roles installed with Benchmarks	Benchmarks limitations and support
Build the foundation	Incident Management	Use the Incident Management application to enrich IT experience by restoring normal service operation, while working toward minimizing business operations impact and maintaining quality.	incident_manager and Roles installed with Incident Management	• Incident Management limitations and support • Automate change and incident management Playbook • Incident Management Overview (video)
Build the foundation	Knowledge Management	Use the Knowledge Management application to share information in knowledge bases, contributing toward self-	Roles installed with Knowledge Management	• Knowledge Management (ServiceNow Community) • Knowledge Management limitations and support

servicenow

Track	Recommended applications	Description	Roles installed	Additional resources
		resolution of incidents.		
Build the foundation	Service Catalog	Use the Service Catalog application to self-create service catalogs contributing toward self-resolution of incidents. It will also help you to self-request a product or a service.	Roles installed with Service Catalog	• Service Catalog limitations and support • Service Catalog success story
Build the foundation	AI Search	Use the AI Search application for a modern consumer-grade search engine for Service Portal, Now Mobile, and Virtual Agent. Intelligent query features enable you to find the answers you need quickly.	Roles installed with AI Search	• AI Search (ServiceNow Community) • AI Search limitations and support
Build the foundation	Performance Analytics ITSM Dashboards	Use the Performance Analytics ITSM Dashboards content pack to access the preconfigured dashboards. These dashboards contain actionable data visualizations that help you improve your business processes and practices.	Roles installed with Performance Analytics ITSM Dashboard	Performance Analytics ITSM Dashboard (ServiceNow Community)

Track	Recommended applications	Description	Roles installed	Additional resources
Build the foundation	Service Portal	Use Service Portal to build a mobile-friendly self-service portal experience for your employees or customers.	Roles installed with Service Portal	• Service Portal limitations and support • Service Portal (ServiceNow Community)
Enhance the experience	Virtual Agent	Unlock your enterprise productivity with ServiceNow® Virtual Agent. Provide your employees and customers with a friendly messaging interface, featuring prebuilt conversations powered by artificial intelligence.	Roles installed with Virtual Agent	• Virtual Agent limitations and support • Virtual Agent success story
Enhance the experience	Predictive Intelligence Workbench ITSM content	Use the Predictive Intelligence interface to train machine learning models and can improve performance, efficiency, and flexibility to your systems across multiple business units.	Roles installed with Predictive Intelligence Workbench ITSM content	• Predictive Intelligence limitations and support • Predictive Intelligence success story
Enhance the experience	Issue auto-resolution tab	Use the **Issue auto-resolution** tab to understand how well your Virtual Agent (VA) chatbot anticipates user needs. It displays information about the number of user issues intercepted by the auto-resolution	chat_analytics_admin and chat_analytics_viewer	Issue auto-resolution (ServiceNow Community)

Track	Recommended applications	Description	Roles installed	Additional resources
		service and resolved by VA. It contributes toward auto-resolution of IT issues, enhancing employee and customer experience.		
Enhance the experience	ITSM Virtual Agent Conversations	Users can reset any Citrix desktop or application session using Virtual Agent conversation flows. Users can also provision a Citrix desktop or application.	Uses bases system roles. See Base system roles.	ITSM Virtual Agent Conversations support
Optimize the functionality	Implement NLU in Virtual Agent	Create and maintain NLU model groups and intents that map to your conversation topics in Virtual Agent. You can create and edit model groups in NLU Workbench or from topics in Virtual Agent Designer.	virtual_agent_admin	NLU support
Add advanced features	Discovery	Discovery finds applications and devices on your network, and then updates the CMDB 360 with the information it finds, contributing toward self-resolution of incidents.	discovery_admin	Discovery limitations and support

Track	Recommended applications	Description	Roles installed	Additional resources
Add advanced features	Automation Discovery	Automation Discovery helps you identify automation opportunities for your workflows. Use the discovery reports to implement or improve automation solutions like Virtual Agent (VA), and Agent assist. It helps you toward improved productivity and efficiency.	Roles installed with Automation Discovery and nlu_admin	• Automation Discovery limitations and support • Automation Discovery overview
Add advanced features	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.	nlu_admin	Intent Discovery overview

Increase change throughput with devops tooling integrations blueprint

Select this adoption blueprint to expedite change without increasing the levels of risk and overhead by using AI-based change policies and DevOps tools.

Note:
The applications recommended within a blueprint may include, but aren't limited to, the following applications. Factors like entitlements and other instance dependencies can cause the number of recommended applications to vary.

Track	Recommended applications	Description	Roles installed	Additional resources
Optimize the functionality		Use DevOps Change Velocity to connect data from your DevOps tool chain with the ServiceNow	Uses bases system roles. See Base system roles.	• DevOps Change Velocity support • DevOps Change

Track	Recommended applications	Description	Roles installed	Additional resources
		platform to deliver changes faster without sacrificing compliance.		Velocity overview (video)
Add advanced features		Use DevOps Config to validate and manage the configuration data of your enterprise applications across every stage of the DevOps pipeline.	Roles installed with DevOps Config	DevOps Config overview (video)
Build the foundation		Use DevOps Insights to plan and implement updates to your DevOps processes.	Uses bases system roles. See Base system roles.	DevOps Insights overview (video)
Build the foundation	Integration Hub	Use Integration Hub to provide both inbound and outbound integration with third-party applications, for an expedited throughput.	Uses bases system roles. See Base system roles.	• Integration Hub support • Integration Hub overview (video)
Enhance the experience		Use Site Reliability Operations as an integral part of DevOps for reliable digital services to deliver modern customer experiences.	Roles installed with Site Reliability Operations	• Site Reliability Operations support • Site Reliability Operations overview (video)

Reduce incidents and keep services flowing with AI based automation blueprint

Select this adoption blueprint to leverage AI-based automation to keep the business running by reducing and prioritizing incidents and issues.

Note:
The applications recommended within a blueprint may include, but aren't limited to, the following applications. Factors like entitlements and other instance dependencies can cause the number of recommended applications to vary.

Track	Recommended applications	Description	Roles installed	Additional resources
Build a foundation	Event Management	Use the Event Management application as a single management console to identify health issues across the datacenter.	Roles installed with Event Management	• Event Management support • Event Management overview (video)
Optimize the functionality	Health Log Analytics	Health Log Analytics helps prevent IT issues before your users are affected. It helps you identify the root cause of an issue by enabling you to triage-related logs and analyze the raw data.	Roles installed with Health Log Analytics	• Health Log Analytics (KB1064762) • ITOM Talks Session #3 - Health Log Analytics (video)
Add advanced features	Managing major incidents	Define an incident as a major incident to demand a response beyond the routine incident management process.	Roles installed with Major Incident Management	Major Incident Management overview (video)
Optimize the functionality	Service Mapping	Service Mapping discovers all application services in your organization and builds a comprehensive map of all devices, applications, and configuration profiles used in these application services.	Roles installed with Service Mapping	• What is service mapping? • Service Mapping Overview (video)
Build a foundation	Service Operations Workspace for ITSM	Use the Service Operations Workspace interface to unify and configure your agent experience	Roles installed with Service Operations Workspace for ITSM	• Service Operations Workspace limitations and support

Track	Recommended applications	Description	Roles installed	Additional resources
		for multiple IT Service Management and IT Operations Management workflows.		• Service Operations Workspace overview
Enhance the experience		Use Site Reliability Operations as an integral part of DevOps for reliable digital services to deliver modern customer experiences.	Roles installed with Site Reliability Operations	• Site Reliability Operations support • Site Reliability Operations overview (video)

Ensure technology teams meet compliance and performance standards blueprint

Select this adoption blueprint to ensure that your technology teams meet all of the required compliance standards for vendors, security, and performance.

Note:
The applications recommended within a blueprint may include, but aren't limited to, the following applications. Factors like entitlements and other instance dependencies can cause the number of recommended applications to vary.

Track	Recommended applications	Description	Roles installed	Additional resources
Enhance the experience	Certificate Inventory and Management	Certificate Inventory and Management discovers, takes inventory of, and proactively manages all your TLS certificates. IPv6 is supported for Certificate Inventory and Management.	Roles installed with Certificate Inventory and Management	What is certificate management?
Add advanced features	Cloud Provisioning and Governance	Use the Cloud Provisioning and Governance application to provide a single interface to access cloud resources, publish cloud offerings to a catalog, and	Uses bases system roles. See Base system roles.	Cloud Provisioning and Governance overview (ServiceNow community)

Track	Recommended applications	Description	Roles installed	Additional resources
		manage the usage of those resources.		
Enhance the experience	Firewall Audits and Reporting	Firewall Audits and Reporting discovers and takes inventory of your firewall security policies, devices, device groups, and manager information.	Roles installed with Firewall Audits and Reporting	What is a firewall audit?
Build a foundation	Service Level Management	Service Level Management (SLM) helps to gather service requirements and monitors and reports information with regard to agreed service levels (SLAs).	Roles installed with Service Level Management	• Service Level Management (SLA) Resources (KB0623502) • How SLAs Really Work (KB0529411)
Optimize the functionality		Use the Vendor Management Workspace to monitor your company's vendors and manage all vendor-related information.	pa_data_collector	Vendor Management Workspace overview (video)

Operationalize workforce effectiveness and processes blueprint

Select this adoption blueprint to increase the effectiveness and efficiency of your workforce processes.

Note:
The applications recommended within a blueprint may include, but aren't limited to, the following applications. Factors like entitlements and other instance dependencies can cause the number of recommended applications to vary.

Track	Recommended applications	Description	Roles installed	Additional resources
Enhance the experience	Advanced Work Assignment	Use Advanced Work Assignment to automatically assign work items to your	Roles installed with Advanced Work Assignment	Advanced Work Assignment (AWA) overview (video)

Track	Recommended applications	Description	Roles installed	Additional resources
		qualified agents, based on their availability, capacity, and optionally, skills.		
Enhance the experience	Agent Forecast	Set the minimum or maximum number of agents required per hour so that you always have the desired staffing coverage.	sn_agent_forecast_admin	Agent Forecast overview (video)
Optimize the functionality	Coaching	Use Coaching to monitor the performance of your team and coach your employees to improve their skill set.	Roles installed with Coaching	Coaching support
Build a foundation	Continual Improvement Management	Use Continual Improvement Management to request improvement opportunities, and implement phases and tasks to meet performance goals, track progress, and measure success.	Roles installed with Continual Improvement Management (CIM)	Continual Improvement Management (CIM) support
Add advanced features	Process Mining content pack	Use ITSM Process Mining to quickly analyze and optimize your business processes.	sm_cim_requestor or admin	• Process Mining success story • ITSM Process Mining Content Pack overview (video)
Build a foundation	Skill Determination	Define conditions for incidents or interactions and add skills required for those conditions using skill	skill_admin or admin	Skill Determination overview (video)

Track	Recommended applications	Description	Roles installed	Additional resources
		determination rules.		
Enhance the experience	Skills Management	Use Skills Management to associate skills with individual users or groups and assign them to tasks or projects. You can also assess the skills needed for your organization, identify gaps, and implement effective plans for hiring or training of your teams.	Roles installed with Skills Management	Skills Management (ServiceNow Community)
Add advanced features		Use Workforce Optimization for ITSM to manage and maintain the productivity of your workforce from a single location.	Roles installed with Workforce Optimization for ITSM	Workforce Optimization for ITSM overview (video)

Install applications for an adoption blueprint

Install applications from the selected adoption blueprint to work towards achieving a certain business goal of your organization.

Before you begin

Role required: admin

Procedure

1. Select the application you want to install from the Adoption Blueprints details page.
   An application details page with the key features about the selected application appears.
2. Start the installation of the selected application.

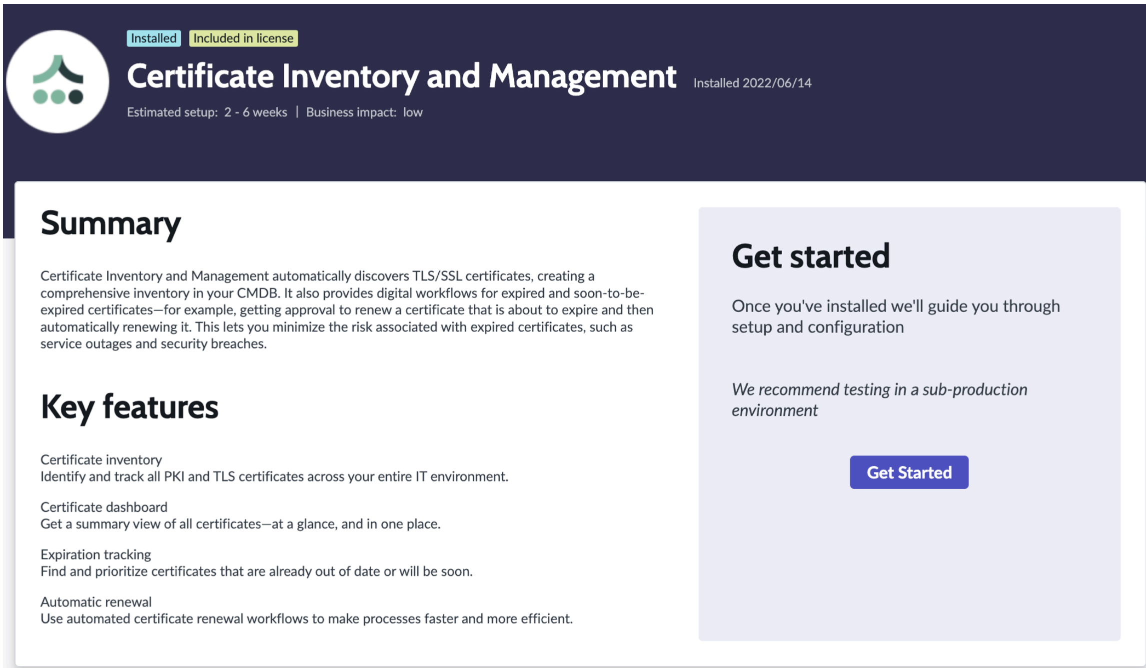

The button to start the installation or activation may vary depending on the selected application.

- **Activate**: If you have selected a plugin
- **Install**: If you have selected an application
- **Update**: If the installed application needs to be updated
- **Get started**: If the application or the plugin is already installed

  Note:
  You can view the configuration page of the application by clicking **Get Started**.

Note:
Some applications and plugins show dependency applications and roles that need to be installed.

Configuration Hub

Experience the ability to zoom into an application or a group of applications and view all the relevant tables and records on a single page using the Configuration Hub. As an admin you can access all relevant details related to the selected application(s).

You can achieve the following by using Configuration Hub.

- One-stop experience to access all relevant information
- Enhanced ability to save your preferences for all tables and records
- Ability to track and view delta changes, over-allocated licenses, and customizations on a single screen
- Enhanced use of multiple filters simultaneously to drill down to a particular application
- Ability to open and edit a record within Configuration Hub. You don't need to navigate to the record using the conventional ways.

Note:
Configuration Hub store application is available by default with any instance upgrade of Vancouver patch 2 and above. You can navigate to Configuration Hub from the **Admin** menu.

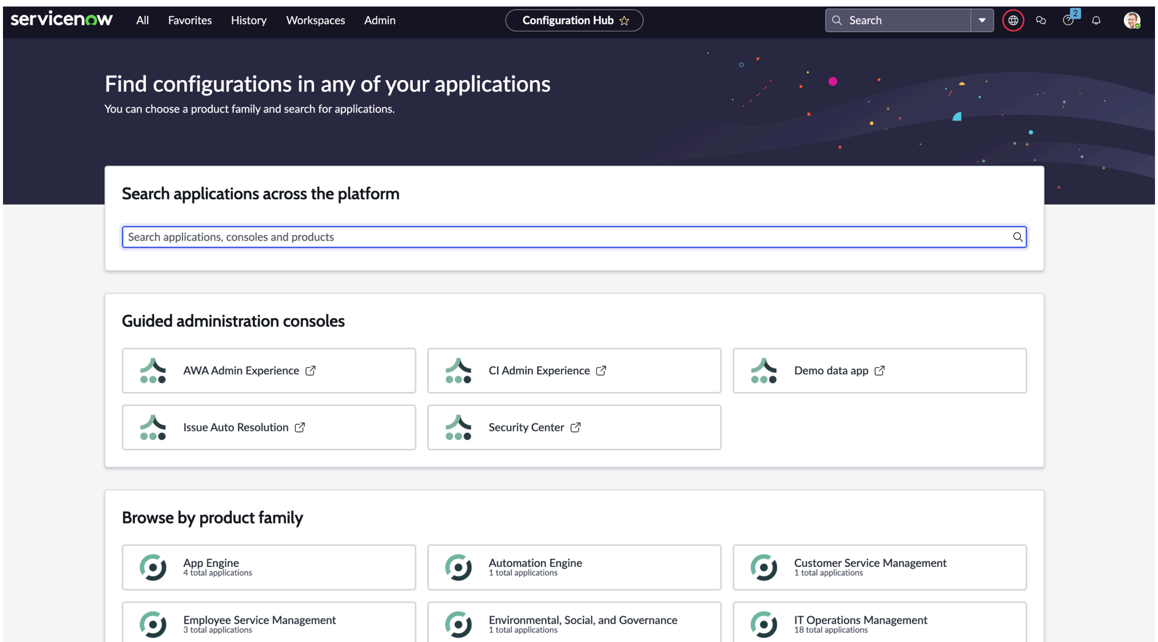

You can search through all the application, console or product names that have been installed on the instance within a certain scope.

Note:
The application name doesn’t populate in the search bar if it’s not yet installed on the instance. The number under the family name shows the number of applications in the selected family have been installed on the instance. If there is a product family that doesn’t have any app installed, it won’t show up on the list.

It can also help you search applications that have been grouped by SN store, taxonomy or product families. When you select a product family, it takes

to the product home page. See Explore the product home page for more information.

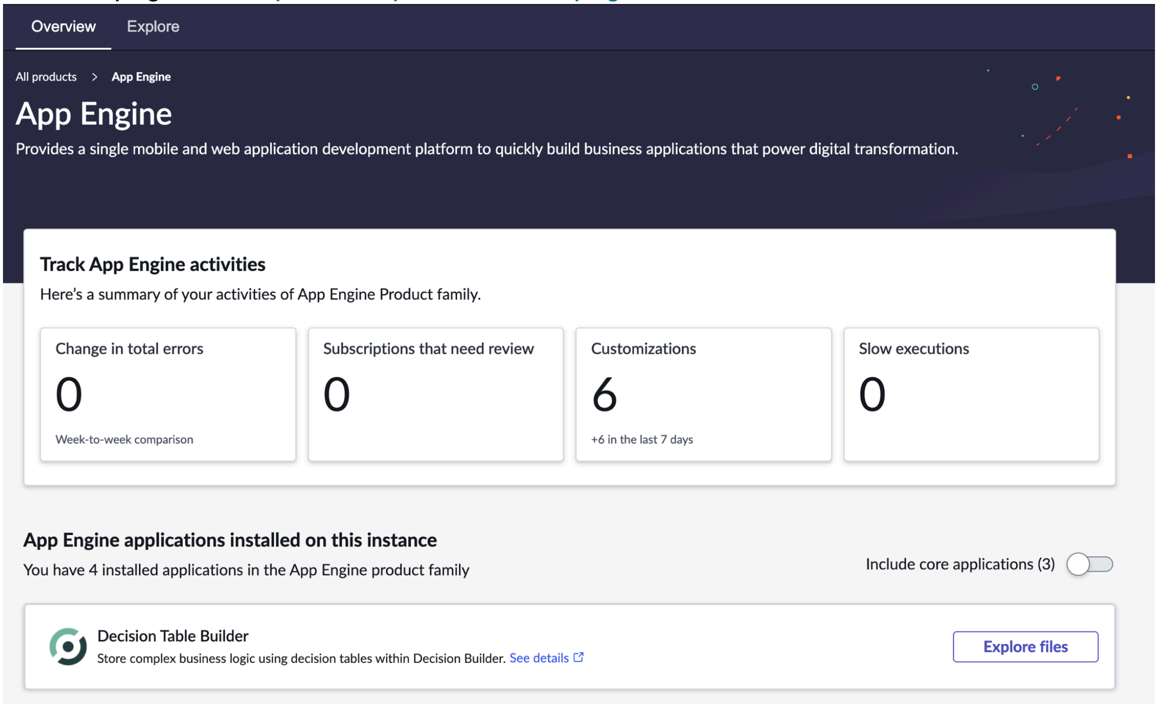

Explore the product home page

Track delta changes, over-allocated subscriptions, customizations, and slow execution on the product home page by selecting a product family

Before you begin

Role required: admin

Procedure

1. Access Configuration Hub in one of the following ways.

2. Select one of the product families as per your requirement.

The product home page shows 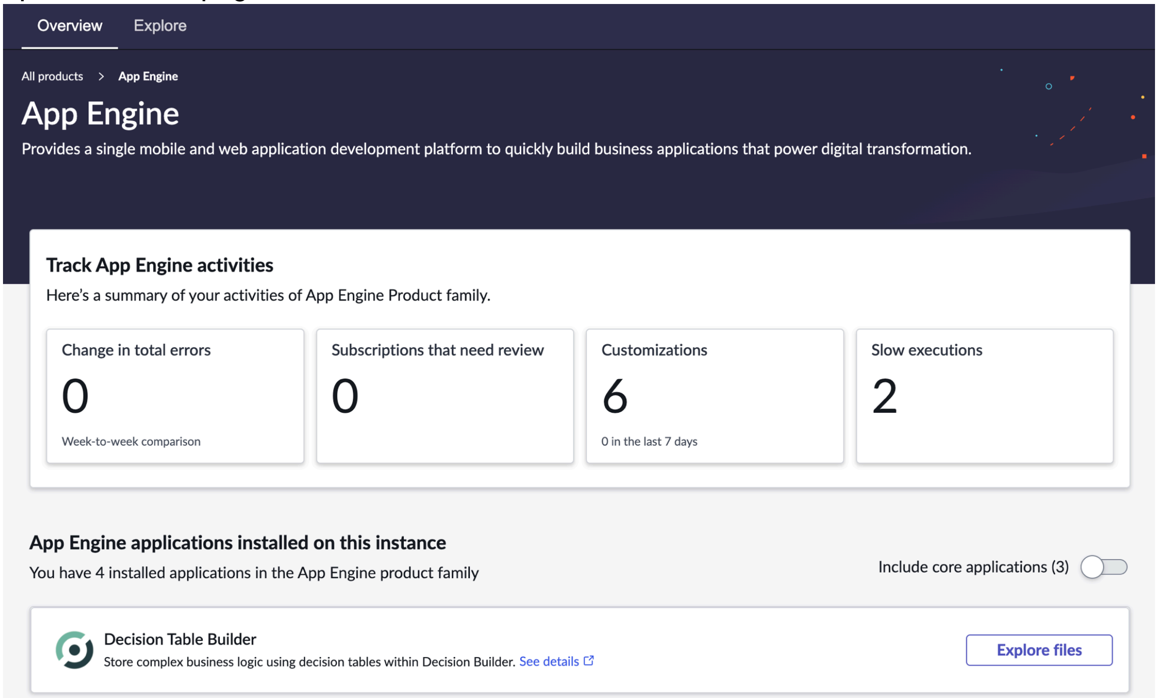

up.

3. Track the following activities in the Overview tab.
   - Change in total errors: Weekly delta change in errors
   - Subscriptions that need review: Number of licenses that have been over-allocated
   - Customizations: Total modified files related to the installed applications within the selected product family.

     Note:
     The number mentioned under the total customizations is the number of files that have been modified recently within last 7 days.

   - Slow execution: The number of slow transactions, events and scripts on the instance.

   By default, the core applications are hidden. Core applications are installed by default as you upgrade your instance. You can't uninstall the core applications. But you can use the toggle to show all the applications including the core apps.

   Note:
   It is recommended not to modify the core applications and its functionality so that during an upgrade you don't have too many skipped files.

   Select **See details** to read more about the store applications that have been installed. Select **Explore files** to list the records in the selected application of the product.
4. Drill down to a particular application by selecting the **Explore** tab.

   Note:
   The first time the **Explore** option is loaded, it takes in all the applications displayed in the **Overview** section.

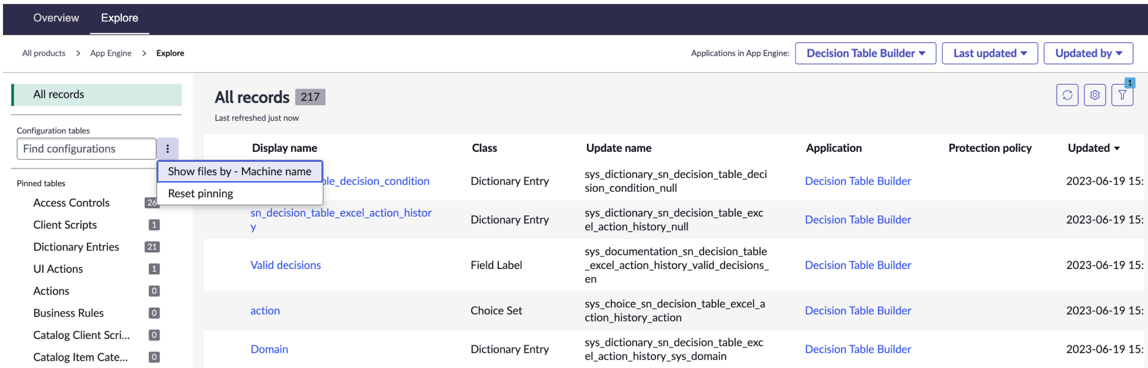

You can do the following in the Explore tab:

- Search for table names, sys_id, or relevant records in the Configuration tables search bar.

  Note:
  You can also display the files as per machine names.

- Select a particular table from the left navigation bar and all the records in that table show up.

  Note:
  The number mentioned next to the table name is the number of records in that table. The records also satisfy the Applications in filter. You can also save your table preference by pinning the required tables.

- You can make edits in the records directly within Configuration Hub. You don't have to open the records in a conventional way to make edits in them.

Configure data synchronization for applications

As of Washington DC, certain ServiceNow applications have the ability to provide data visibility across instances within a customer's account.

Data visibility is protected by a trust configuration per instance, per application. You can configure the data sharing for production and non-production instances for the applicable applications by navigating to **All** > **Multi-Instance Management** > **Trust Configuration**.

Note:

By default, non-production instances allow application data sharing with production only. Production instances do not allow application data sharing with any other instance. However, these can be updated by the admin for each instance based on business needs.

If you demote a production instance to non-production, or promote a non-production instance to production, all previous customizations to data sharing settings are reset to the default configuration.

The **Grant access** column indicates the permission granted by the instance that you're currently logged in to for the instance named in the same row's instance column. The **Is granting access** column displays whether or not the instance mentioned in the Instance column has granted access to the instance you're currently logged in to; it can't be edited.

The following are some examples for data sharing restrictions between instances.

Example: Logged in to Prod1

In the following example, you're logged in to Prod1. Prod1 has granted access to the instance Prod2 for the application Subscription Management, as indicated by the `True` value in the **Grant access** column.

Prod2 hasn't granted access to Prod1, as indicated by the `False` value in the **Is granting access** column.

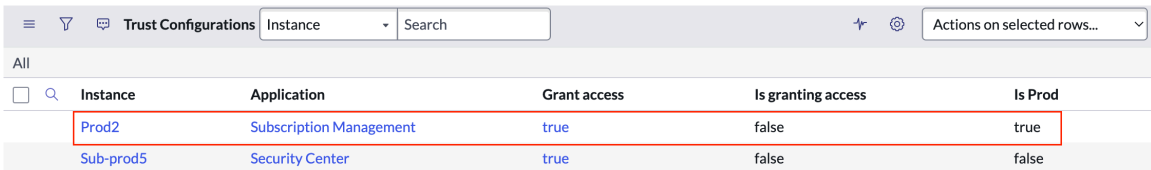

Instance	Application	Grant access	Is granting access	Is Prod
Prod2	Subscription Management	true	false	true
Sub-prod5	Security Center	true	false	false

To revoke access for the Subscription Management app from Prod1 to Prod2, update the value in the **Grant access** column to `False` while logged in to Prod1.

Example: Logged in to Prod2

In the following example, you're logged in to Prod2. Prod1 has granted access to the instance Prod2 for the application Subscription Management, as indicated by the `True` value in the **Is granting access** column.

Prod2 hasn't granted access to Prod1, as indicated by the `False` value in the **Grant access** column.

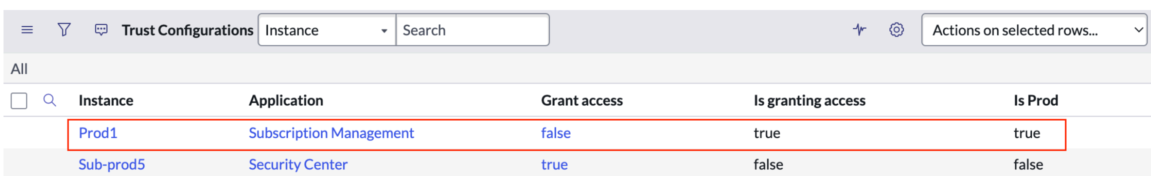

Instance	Application	Grant access	Is granting access	Is Prod
Prod1	Subscription Management	false	true	true
Sub-prod5	Security Center	true	false	false

To grant access from Prod2 to Prod1 for the Subscription Management application, update the value in the **Grant access** column to `True` while logged in to Prod2.

Example: Logged in to Sub-prod2

In the following example, you're logged in to Sub-prod2. Prod1 hasn't granted access to the instance Sub-prod2 for the application Subscription Management, as indicated by the `False` value in the **Is granting access** column.

Sub-prod2 has granted access to Prod1, as indicated by the `True` value in the **Grant access** column.

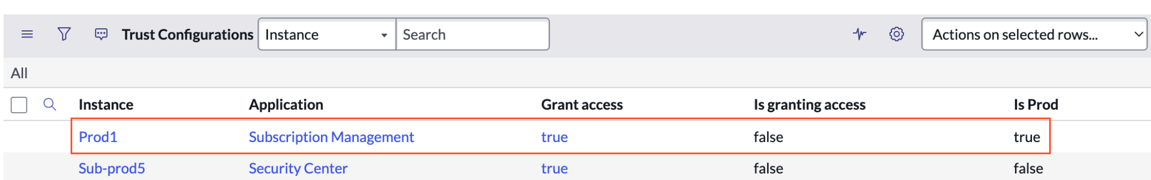

Instance	Application	Grant access	Is granting access	Is Prod
Prod1	Subscription Management	true	false	true
Sub-prod5	Security Center	true	false	false

To revoke access from Sub-prod2 to Prod1 for the Subscription Management application, update the value in the **Grant access** column to `True` while logged in to Sub-prod2.

Example: Logged in to Sub-prod3

In the following example, you're logged in to Sub-prod3. Sub-prod4 has granted access to the instance Sub-prod3 for the application Subscription Management, as indicated by the `True` value in the **Is granting access** column. Sub-prod3 has also granted access to Sub-prod4, as indicated by the `True` value in the **Grant access** column.

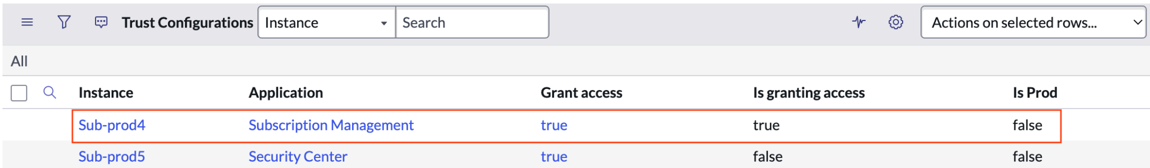

Instance	Application	Grant access	Is granting access	Is Prod
Sub-prod4	Subscription Management	true	true	false
Sub-prod5	Security Center	true	false	false

Admin Center roles

Admin Center is installed with these roles.

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

System Administrator [admin]

The admin role allows you to access admin home features, adoption blueprints and configuration hub. You can create your own personalized admin home page on your instance with the information that is most important to you.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevate to a privileged role.

No.

Special considerations

None.

Subscription Management

Subscription Management enables you to proactively manage your subscriptions and monitor subscription usage on your instances.

https://player.vimeo.com/video/1018055648?h=5c896ae137&badge=0&autopause=0&player_id=0&app_id=58479

Watch this short video to see an overview of Subscription Management.

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 Learn about Subscription Management benefits and features.	**Configure** Configure Subscription Management settings.	**Manage per-user subscriptions** Allocate per-user subscriptions using recommended groups.
Map custom tables and applications Stay in compliance by mapping custom tables and applications to subscriptions.	**Monitor capacity subscriptions** Learn how capacity subscriptions are calculated on your instance.	**Reference** Get details about types of subscriptions, meters, and domain separation.

Helpful resources

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

ServiceNow Community

ServiceNow Community

ServiceNow Store

http://servicenow.com

Support

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

Exploring Subscription Management

Learn how Subscription Management helps you track and manage your product subscriptions and their applications.

Subscription Management Overview

Subscription Management gives you an accurate, consolidated view of your entitlements where you can proactively manage your subscriptions.

- Gain visibility and control of your product subscriptions and their applications.
- Access a real-time view of your entitlement data.
- Maximize subscription usage and make informed decisions about purchases by viewing insights.
- Stay in compliance and avoid unexpected costs by tracking subscription usage and allocations over time.

Subscription Management users

Users

User	Description
Instance admin	The instance admin allocates per-user subscriptions by adding groups to a product subscription. The instance admin maps custom tables and custom applications to subscriptions with available entitlements to maximize subscription usage and stay in compliance.

Subscription Management workflow

This infographic shows a sample end-to-end workflow of how different users in an organization work together to purchase, allocate, and manage subscriptions.

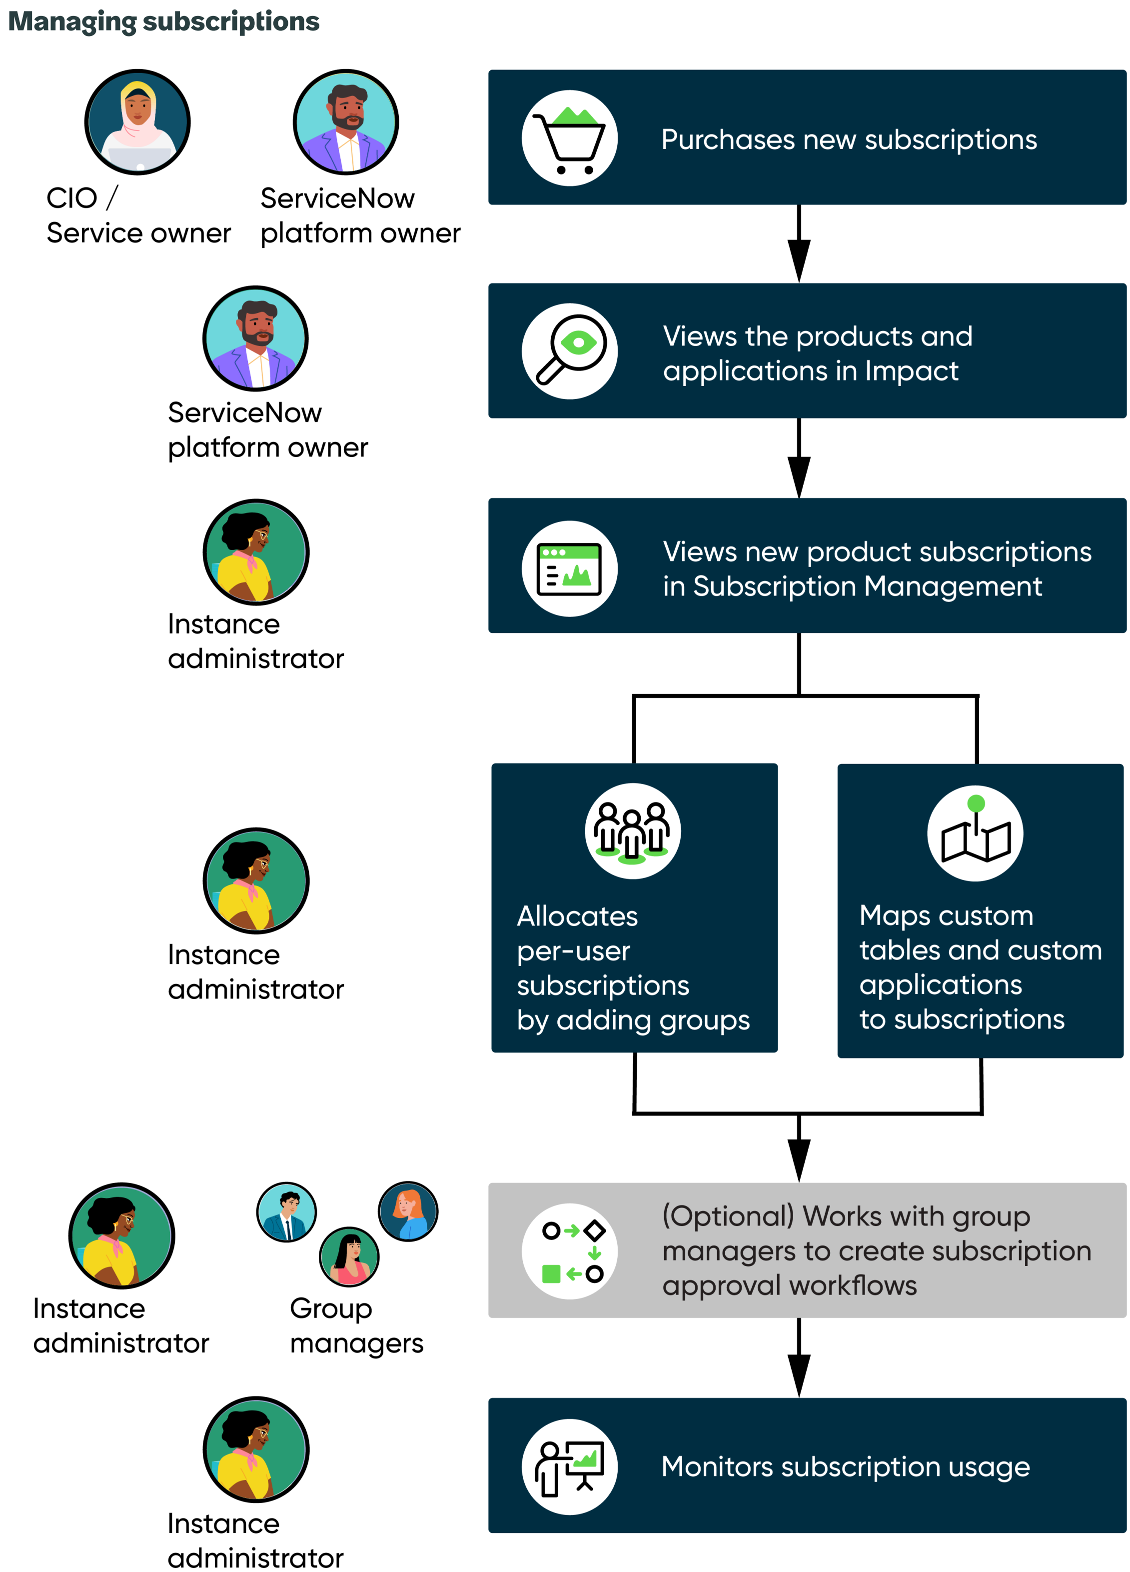

In this workflow:

1. The chief information officer or service owner works with the ServiceNow platform owner to purchase new subscriptions.
2. After the purchase is complete, the platform owner views the products and applications in Impact.
3. The instance admin views the new product subscriptions in Subscription Management on the production instance. For more information, see Product subscriptions overview.

4. The instance admin performs the following tasks in Subscription Management:
    - Allocates per-user subscriptions by adding groups to a product subscription, which grants entitlement to all group members. For more information, see Managing per-user subscriptions in Subscription Management.
    - Maps custom tables and custom applications to subscriptions with available entitlements to maximize subscription usage and stay in compliance. For more information, see Managing custom tables and applications in Subscription Management.
5. If the organization requires security approvals before adding users to groups or financial approvals before subscriptions are allocated to users, the instance admin works with group managers to create subscription approval workflows that determine who gets a subscription when new users are added to groups.
6. The instance admin monitors subscription usage to stay in compliance and make informed decisions on future purchases, consistently monitoring the following data:
    - Over-allocated subscriptions
    - Unmapped custom applications and tables
    - Usage patterns over time
    - Application popularity
    - Custom table usage

Subscription Management benefits

Benefit	Feature	User
Fix over-allocated subscriptions	Insights	Instance admin
View unallocated subscriptions	Insights	Instance admin
Allocate available subscriptions using recommended groups	Allocate subscriptions	Instance admin
Map global custom tables to a product subscription	Issues	Instance admin
Map custom applications to a product subscription	Issues	Instance admin
View a summary of subscription allocations	Subscription details	Instance admin
View subscription allocation details over time	Subscription details	Instance admin
View custom table usage over time	Subscription details	Instance admin

What to explore next

To learn more about configuring and using Subscription Management, see:

- Configuring Subscription Management
- Managing per-user subscriptions in Subscription Management
- Managing custom tables and applications in Subscription Management
- Monitoring capacity subscriptions
- Subscription Management reference

Product subscriptions overview

View a summary of your product subscriptions in Subscription Management.

Key benefits

- Determine whether you have any over-allocated subscriptions.
- Check for unallocated subscriptions that you can allocate to maximize usage and give users entitlement to ServiceNow products.
- Determine whether any custom tables in the global scope or any custom applications aren't mapped to subscriptions.
- Stay in compliance by monitoring the **Status** column in the list of product subscriptions.

Subscriptions overview

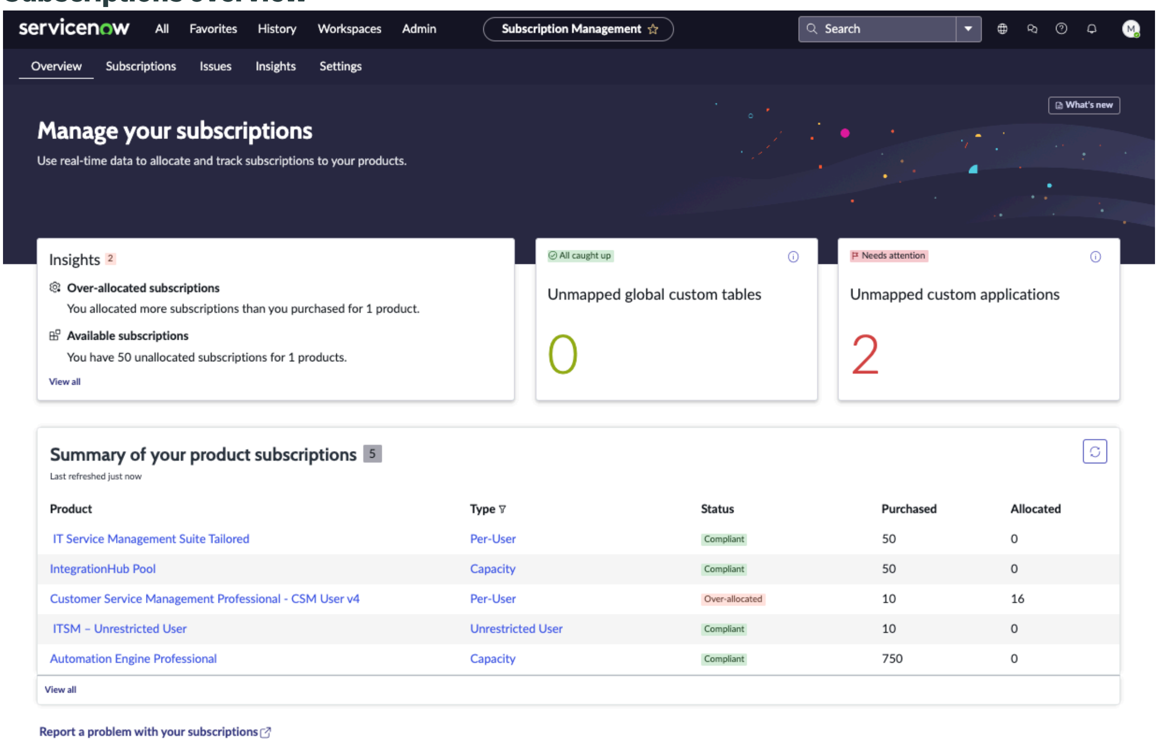

Required ServiceNow AI Platform roles

The usage_admin or admin roles are required to view and edit Subscription Management.

Accessing product subscription overview information

View subscription overview information on the **Overview** tab in Subscription Management by navigating to **Admin** > **Subscription Management** or **All** > **Subscription Management** > **Subscription Management**.

Use cases

View a summary of your instance-level entitlements in the **Instance entitlements** tab.

- Identify when subscriptions are over-allocated and view insights into actions you can take to maximize subscription usage.
- Determine whether any custom tables and custom applications on your instance aren't mapped to a valid product subscription. You must map custom tables and custom applications that you create to a product subscription.
- View product subscription details by selecting the product name.
- Identify which type of subscriptions you have by viewing the values in the **Type** column. The subscription type determines the properties and attributes of the product subscriptions that your company receives and whether to allocate subscriptions manually.
    - Product subscription types include per-user, capacity, unrestricted user, unlimited, and display only.
    - Per-user subscriptions require manual allocation by an administrator. Capacity, unrestricted user, and unlimited subscriptions are allocated on your instance automatically. Display-only subscriptions aren't measured on your instance.
    - Product subscriptions provide entitlements for different applications and can include different add-ons.

  For more information, see Types of subscriptions in Subscription Management.
- Determine whether your subscription allocations are in compliance by viewing the values in the **Status** column. Adjust allocations if subscriptions are over-allocated.

  Compliant

  The number of allocated subscriptions is below the number of purchased subscriptions.

  Near capacity

  The number of allocated subscriptions exceeds the threshold for your instance. You define the threshold on the **Settings** tab.

  Even

  The number of allocated subscriptions equals the number of purchased subscriptions.

  Over-allocated

  The number of allocated subscriptions exceeds the number of purchased subscriptions.

  Account-level only

  The subscription allocation status isn't calculated. Only applies to Creator Plus products.

  You can customize the near capacity threshold for your instance in the **Settings** tab. For more information, see Configure the subscription allocation status in Subscription Management.
- View a summary of the Creator Plus subscriptions on the instance that you're currently accessing.
    - View the total number of Creator Plus subscriptions that your company has purchased in the **Purchased** column.
    - Monitor the total number of Creator Plus allocations on your current instance in the **Allocated** column.

View a summary of your account-level entitlements in the **Account entitlements** tab.

- View an account-level summary of Now Assist usage.
  - View the aggregated total of Now Assist usage entitlements from all the subscriptions that your account has purchased in the **Purchased** column.
  - View the total number of Now Assist usage entitlements consumed across all your instances in the **Allocated** column.
  - View additional details about Now Assist usage by selecting the product name.
- View an account-level summary of Now Assist creators.
  - View the total number of Creator Plus subscriptions that your account has purchased in the **Purchased** column.
  - Monitor unique Creator Plus allocations across all your instances by checking the **Allocated** column.

    Note:
    Users allocated to a Creator Plus subscription on multiple instances are automatically de-duplicated in the count on the **Account entitlements** tab. When a user is allocated to a Creator Plus subscription on more than one instance, that user counts toward the allocation total only once.

  - View additional details about Now Assist creators by selecting the product name.

Reports

Title	Type	Source Table	Description
Insights	List	Subscription Insights definition [sn_sub_man_st_subscription_insights]	List of over-allocated subscriptions and available subscriptions
Unmapped global custom tables	Single Score	Custom Table Inventory [ua_custom_table_inventory]	Count of custom tables that aren't mapped to a subscription
Unmapped custom applications	Single Score	Custom Table Inventory [ua_custom_table_inventory]	Count of custom applications that aren't mapped to a subscription

Viewing product subscription details in Subscription Management

View a summary of your subscription allocations, allocation totals over time, and track your custom table entitlements in Subscription Management.

Key benefits

- See how a product subscription is measured by viewing the subscription's meter type.
- Match a product subscription in Subscription Management to your purchase order.
- View the groups and users that you've added to a subscription.
- View your entitlements for custom tables, portal visits, or transactions.
- View the custom tables and applications that you've mapped to a product subscription.

Subscription details

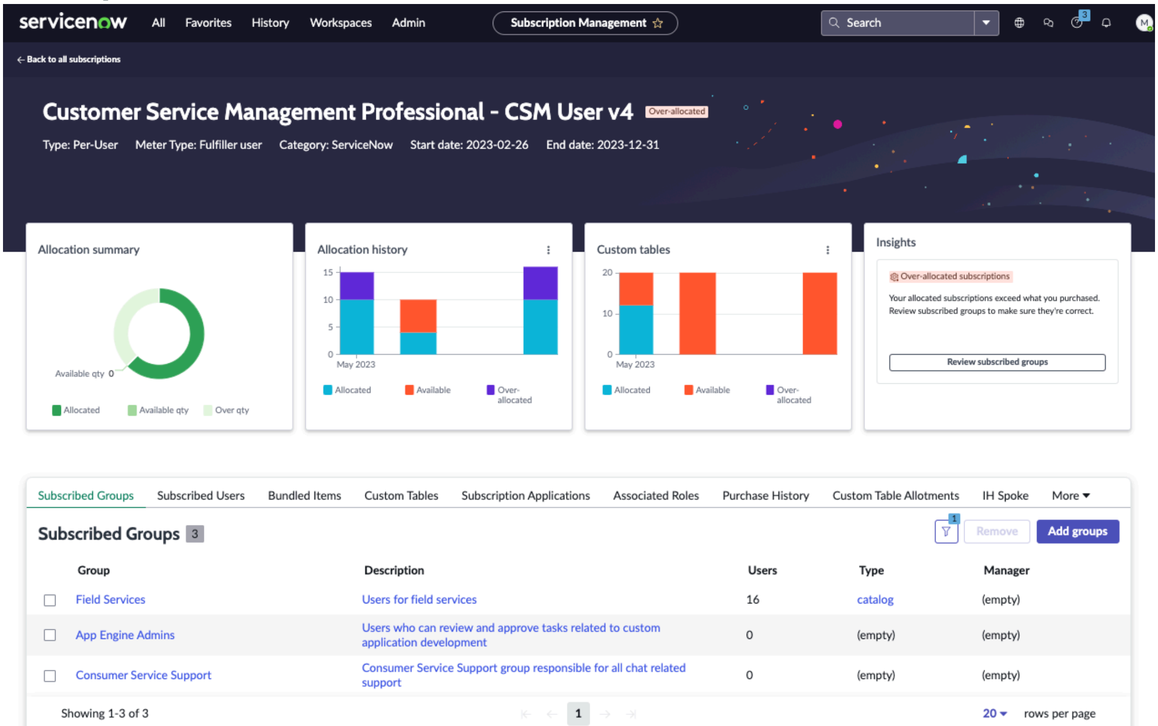

Required ServiceNow AI Platform roles

The usage_admin role is required to view and edit Subscription Management.

Accessing subscription details

View subscription details on the details page by navigating to **Admin** > **Subscription Management** or **All** > **Subscription Management** > **Subscription Management**, and selecting a product or a subscription from the list.

Use cases

- See how a subscription is measured and the type of entitlement it provides by checking the **Meter Type** field.
  - Entitlements for users with fulfiller-type access are tracked through a per-user subscription with a Fulfiller User meter type.
  - Entitlements for users who have fewer access rights than fulfiller users are tracked through a per-user subscription with a Business Stakeholder User meter type.
  - Entitlements for users with the now.assist.creator role are tracked through a per-user subscription with a Creator User meter type.
  - The creation and retention of platform objects such as custom tables is tracked through a capacity subscription with a Capacity meter type.

  For more information, see Types of subscriptions in Subscription Management.
- See how a capacity subscription total is calculated by selecting the Capacity Definition record and viewing the definition details. For more information, see Monitor capacity subscription calculations in Subscription Management.

- Stay in compliance by monitoring your product subscription add-on usage on the **Bundled Items** tab. Product subscription add-ons can include bundled custom tables, portal visits, and transaction entitlements.
- Plan for future purchases and keep your subscriptions current by monitoring the start and end dates for each of your product subscriptions in the **Purchase History** tab. Note that multiple entries for the same product might appear if modifications are made to a product.
- View the applications that are included with a product subscription on the **Subscription applications** tab.
  - View application entitlement details, including the application's related plugins and associated roles by selecting an application.
  - Determine whether an application is available to install from the current product subscription by checking the **Installation status** column.
- Maximize subscription usage and grant entitlement to the right users by allocating per-user subscriptions. You allocate per-user subscriptions by adding one or more groups to a product subscription. Subscription Management helps you with the allocation process by recommending groups based on their assigned roles. For more information, see Allocate subscriptions in Subscription Management.

Consolidated subscriptions

When you purchase multiple qualifying subscriptions for certain types of products, such as Integration Hub, those individual subscriptions combine into a single subscription pool. Any usage against that subscription deducts from the common pool.

- Capacity subscriptions with a common capacity definition are consolidated into one capacity subscription. For example, if you purchased two different Integration Hub subscriptions that are measured using the same capacity definition, they're consolidated into one Integration Hub Pool capacity subscription. Other types of capacity subscriptions with matching definitions are consolidated as well, but aren't labeled as a pool. For any consolidated capacity subscription, you can view the combined products in the **Subscription line items** tab on the details page.
- Capacity subscriptions can also consolidate entitlements for secondary line items such as custom tables, transactions, or portal visits from other product subscriptions. For example, if you have two Integration Hub transaction subscriptions combined in a pool, transaction entitlements that are included as a secondary line item from another capacity subscription are added to the pool if the capacity definition matches. In this scenario, three separate entries appear on the **Subscription line items** tab, one for each product that contributes to the pool.
- Custom table entitlements that come with per-user subscriptions can also be consolidated under one product subscription. When custom table entitlements are consolidated under one subscription, the total number of custom tables appears in the **Purchased** column on the **Bundled Items** tab. The **Bundled Items** tab only appears on one of the product subscriptions that contributed to the pool of custom tables.
- Unrestricted User subscriptions aren't consolidated into a single product subscription.

Reports

Title	Type	Source table	Description
Allocation summary	Donut	Subscription license detail definition [sn_sub_man_st_subscription_license_detail_metric] table	View a real-time summary of your subscription allocation. If an

Title	Type	Source table	Description
			allocation is nearing capacity, consider reducing the number of allocated subscriptions or purchasing more subscriptions.
Account-level allocations	Donut	Account level entitlements [sn_sub_man_st_account_level_entitlement] table	View a summary of Creator Plus allocations across all of your instances. If the account-level allocation is nearing capacity, consider reducing the number of allocated subscriptions or purchasing more subscriptions.
Allocation history	Bar	Subscription license detail definition [sn_sub_man_st_subscription_license_detail_metric] table	View subscription allocation details over time. Prepare for future purchases or plan for changes according to the allocation trends.
Custom tables	Bar	Custom Table Inventory [ua_custom_table_inventory] table	View custom table usage over time in subscriptions that include bundled custom tables. Anticipate when to increase custom table entitlements by monitoring usage month-to-month.

Viewing instance-level entitlements in Subscription Management

View a complete list of the product subscriptions purchased for your current instance.

Key benefits

- View details for product subscriptions, including purchase information and usage history by selecting the product name.
- Plan for future purchases and keep subscriptions current by monitoring the start and end dates for each of your product subscriptions.
- Stay in compliance by monitoring the subscription allocation values in the **Status** column.

Instance-level entitlements

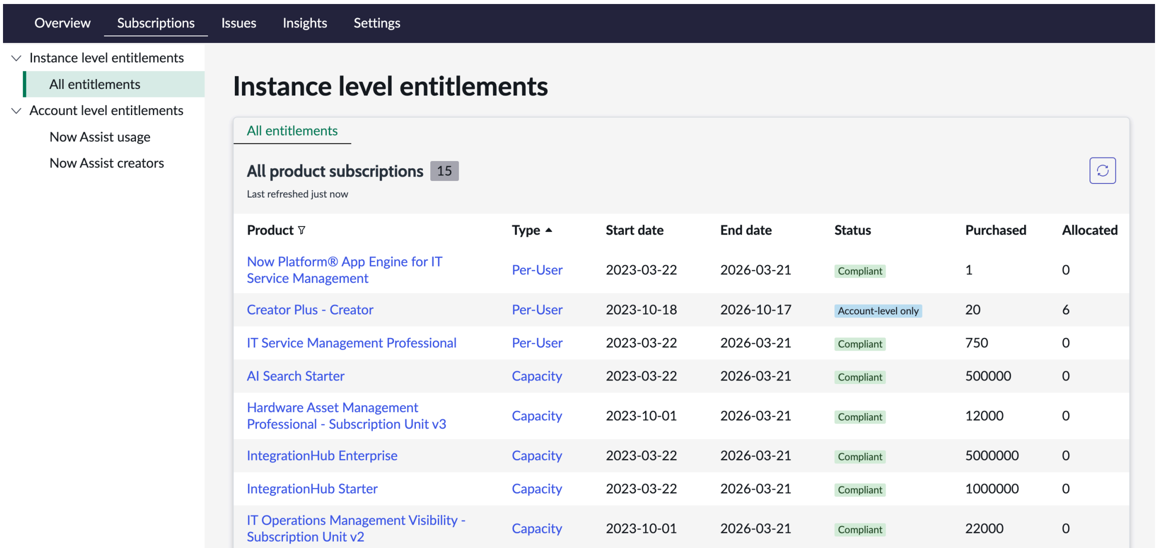

Accessing your instance-level entitlements

Access your instance-level entitlements by navigating to **Admin** > **Subscription Management** > **Subscriptions** or **All** > **Subscription Management** > **All Subscriptions**.

Subscription types

Identify which type of subscriptions you have and whether to allocate subscriptions manually by viewing the values in the **Type** column.

Per-User

Provides entitlements to users according to their assigned roles and associated access rights. You allocate per-user subscriptions manually by adding groups with metered roles or access rights to a product subscription. Subscription Management helps you with the allocation process by recommending groups based on their role assignments.

Unrestricted User

Provides entitlements for an organization's number of active users regardless of their role assignments. An active user is any user whose record in the Users [sys_user] table has a value in the **User ID** field and has the **Active** field set to true.

Capacity

Automatically tracks measured units, and counts of transactions, records created, or resources managed such as devices, software, or nodes.

Unlimited

Provides entitlements for active users, with no limit to the number of users that can be allocated. An active user is any user whose record in the Users [sys_user] table has a value in **User ID** field and has the **Active** field set to true.

Display Only

Displays information. This type of subscription doesn't support allocation.

Subscription status

Determine whether your subscription allocations are in compliance by viewing the values in the **Status** column.

Compliant

The number of allocated subscriptions is below the number of purchased subscriptions.

Near capacity

The number of allocated subscriptions exceeds the threshold for your instance. You define the threshold on the **Settings** tab.

Even

The number of allocated subscriptions equals the number of purchased subscriptions.

Over-allocated

The number of allocated subscriptions exceeds the number of purchased subscriptions.

Account-level only

The subscription allocation status isn't calculated. Only applies to Creator Plus products.

Missing subscriptions

Purchased subscriptions might not appear on the **Subscriptions** tab for one or more of the following reasons:

- Subscription data arrives on production instances only.
- Self-hosted instances don't receive subscription information. Contact your account manager for assistance.
- Subscription data might not yet have arrived. The data is downloaded daily.

Viewing account-level entitlements in Subscription Management

View an account-level summary of entitlements and consumption in Subscription Management.

Key benefits

- View and track Now Assist usage across all your instances.
- View and track Now Assist Creator allocations across all your instances.

Accessing your account-level entitlements

Access your account-level entitlements by navigating to **Admin** > **Subscription Management** > **Subscriptions** or **All** > **Subscription Management** > **All Subscriptions** > **Account level entitlements**.

Viewing production and non-production instances

To view the entitlement data for all your instances, access your account-level entitlements on the **Subscriptions** tab in Subscription Management on a production instance. On a non-production instance, only entitlement data and usage for the instance that you're currently accessing is displayed.

Note:
On a production instance, you can view entitlement data for all the instances owned by your account. However, if data sharing is restricted on an instance, that instance's entitlement data is hidden. You can disable the restriction by updating the trust configuration on that instance. See Share subscription data from another instance.

Monitoring Now Assist creators in Subscription Management

View a summary of your Creator Plus entitlements and track allocations for distinct users across all of your instances in Subscription Management.

Key benefits

- View the total number of Creator Plus entitlements that you purchased across all your instances.
- View the total number of Creator Plus entitlements allocated to unique users across all your instances.
- Track the availability of the remaining Creator Plus entitlements that you can allocate to users.

Now Assist Creator Plus entitlements

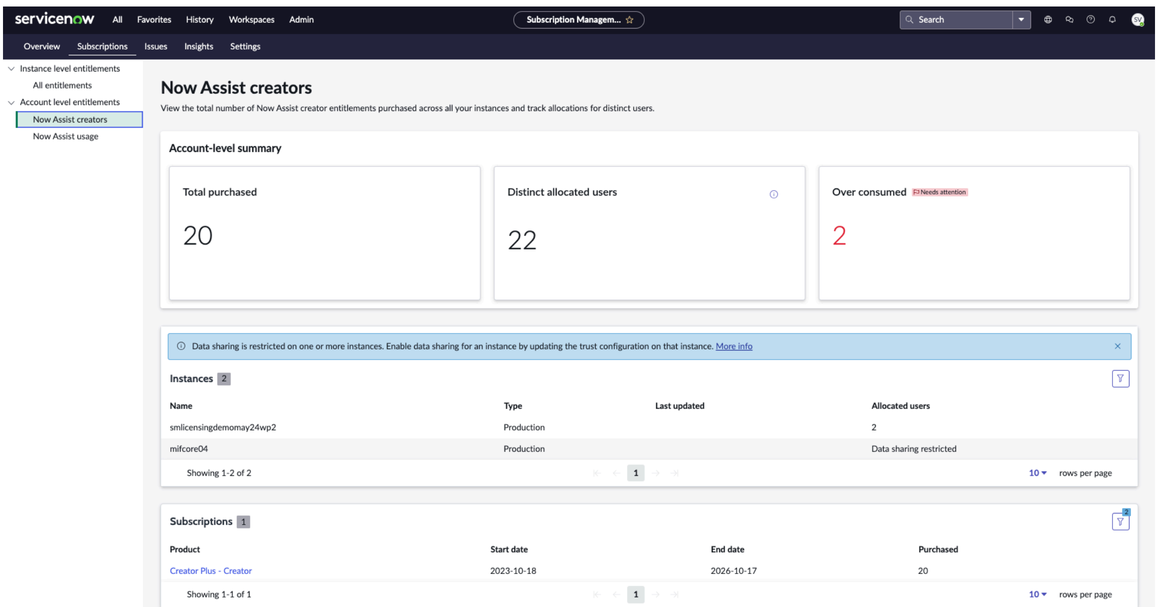

Required ServiceNow AI Platform roles

The usage_admin role is required to view and edit Subscription Management.

Accessing Creator Plus entitlements

Access details about Creator Plus entitlements by navigating to **Admin** > **Subscription Management** > **Subscriptions** > **Now Assist creators**.

Use cases

- Compare the total number of Creator Plus entitlements that your company has purchased to the total number of Creator Plus entitlements allocated across all your instances.
- Monitor the number of Creator Plus entitlements that are allocated to unique users across all your instances in the **Distinct allocated users** card.

Note:
Users allocated to a Creator Plus subscription on multiple instances are automatically de-duplicated in the total count. When a user is allocated to a Creator Plus subscription on more than one instance, that user counts toward the allocation total only once.

- View Creator Plus allocations by instance in the **Allocated users** column in the **Instances** list.

  In some scenarios, data isn't displayed in the **Instances** list.

  - Only instances on Washington DC Patch 2 and higher are listed in the **Instances** list.
  - If data sharing is restricted on an instance, usage data for that instance is hidden. You can disable the restriction by updating the trust configuration on that instance. See Share subscription data from another instance.
- View Creator Plus product subscription details by selecting a product in the **Subscriptions** list. On the details page, allocate available Creator Plus subscriptions by adding one or more groups with the Creator role.

Monitoring Now Assist usage in Subscription Management

View a summary of your Now Assist entitlements and track Now Assist usage across all of your instances in Subscription Management.

Key benefits

- View an account-level summary of your Now Assist entitlements.
- Track Now Assist usage across your production and non-production instances.
- View a summary of Now Assist usage by skill.
- Plan for future purchases by analyzing Now Assist usage over time.

Note:
Now Assist usage data is refreshed daily in Subscription Management. The data reflects a 24-hour delay.

Now Assist usage

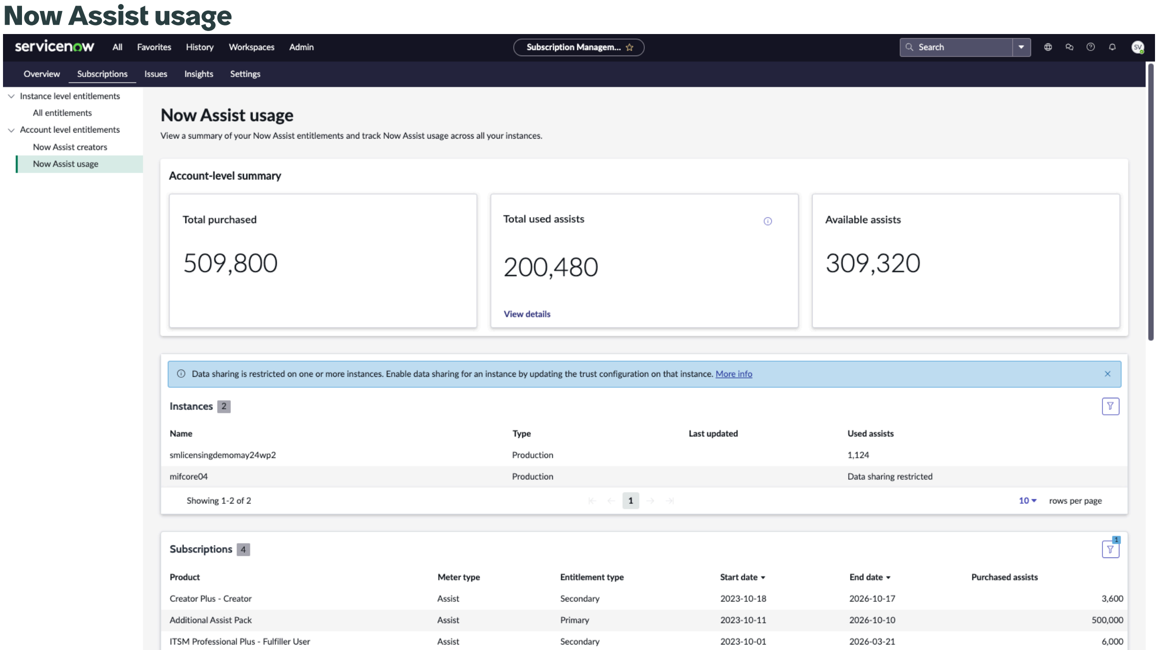

Required ServiceNow AI Platform roles

The usage_admin role is required to view and edit Subscription Management.

Accessing Now Assist usage details

Access details about Now Assist usage by navigating to **Admin** > **Subscription Management** > **Subscriptions** > **Now Assist usage**.

Use cases

- Compare the total number of Now Assist entitlements that your company has purchased to the total number of Now Assist entitlements consumed across all your instances. Monitor the remaining number of available number of Now Assist entitlements for your account.
- Visualize Now Assist consumption over time by selecting **View details** in the Total used assists card. Monitor compliance for your account over the last 365 days or determine when Now Assist usage exceeded the total purchased count.
- View the number of assists consumed by each of your instances in the **Used assists** column in the **Instances** list.

  In some scenarios, data isn't displayed in the **Instances** list.

  - Only instances on Washington DC Patch 2 and higher are listed in the **Instances** list.
  - If data sharing is restricted on an instance, usage data for that instance is hidden. You can disable the restriction by updating the trust configuration on that instance. See Share subscription data from another instance.
- View details for each of your products that include assists in the **Subscriptions** list.
  - Determine whether assists are the primary or secondary line item in a purchase according to the value in the **Entitlement type** column.
  - Verify the start date and end date for each product.
  - View the total number of entitlements provided by each product in the **Purchased assists** column.
- Determine cumulative Now Assist usage by skill on your current instance. Data in the table is updated nightly. If your current instance doesn’t have any active subscriptions with assists, the list of skills is empty.
  - View the number of times a skill has been used in the **Number of actions** column.
  - View the relative cost for each assist in the **Assist ratio** column.
  - Determine which skills are consuming the most Now Assist entitlements by checking the **Total assists** column. Total assists are calculated by multiplying values in the **Number of actions** column by values in the **Assist ratio** column.
  - Analyze Now Assist usage over time on your current instance by viewing the Utilization history chart. Monitor monthly or daily usage using the filter. Look for trends related to Now Assist skills. If usage continues to grow every month, consider purchasing more entitlements to stay in compliance.

Addressing issues in Subscription Management

Stay in compliance by mapping custom tables and custom applications to subscriptions in Subscription Management.

Key benefits

- Identify custom tables and custom applications that aren't mapped to subscriptions.
- Maintain accurate entitlements by mapping custom tables and custom applications to a recommended product subscription or a product subscription of your choice.
- Plan for future subscription needs by assessing your custom table entitlements.

Subscription issues

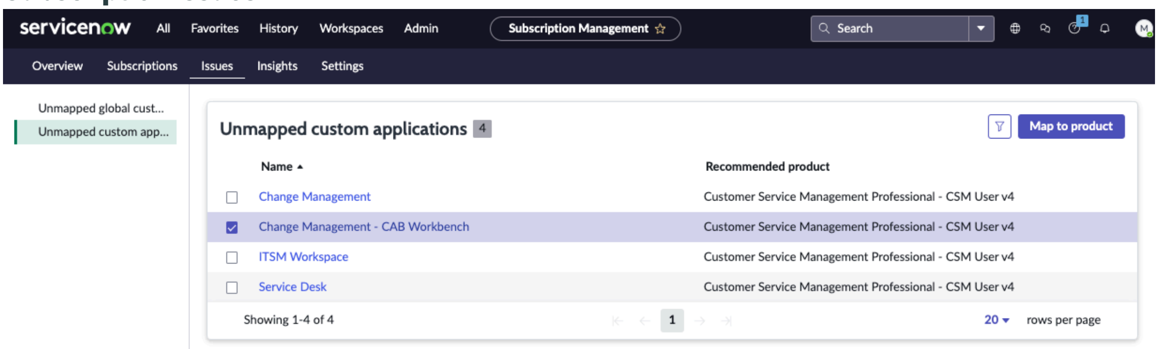

Accessing subscription issues

Access information on subscription issues on the **Issues** tab by navigating to **Admin** > **Subscription Management** > **Issues** or **All** > **Subscription Management** > **Issues**.

Unmapped global custom tables

Determine whether you have custom tables in the global scope that aren't mapped to a valid product subscription by accessing the **Unmapped global custom tables** tab. Keep custom table entitlements current by mapping the tables to a product subscription. For more information, see Map custom tables to a product subscription in Subscription Management.

Custom tables that aren’t mapped to a subscription appear on the **Issues** tab the day after they’re created.

Unmapped custom applications

Determine whether you have custom tables in a scoped application that aren't currently mapped to a valid product subscription by accessing the **Unmapped custom applications** tab. Ensure that entitlements for these tables are current by mapping the application to a product subscription.

You map tables in a scoped application at the application-level, which eases administration because you don't need to map tables individually after the application is mapped. For more information, see Map a custom application to a product subscription in Subscription Management.

Custom applications that aren’t mapped to a subscription appear on the **Issues** tab the day after they’re created.

Discovering insights in Subscription Management

Identify when subscriptions are over-allocated and view actions you can take to maximize subscription usage in Subscription Management.

Key benefits

- View product subscriptions that are over-allocated and require adjustments or more entitlements.
- View product subscriptions with available subscriptions that you can allocate by adding groups.
- Maximize subscription usage by allocating your available subscriptions.

Subscription insights

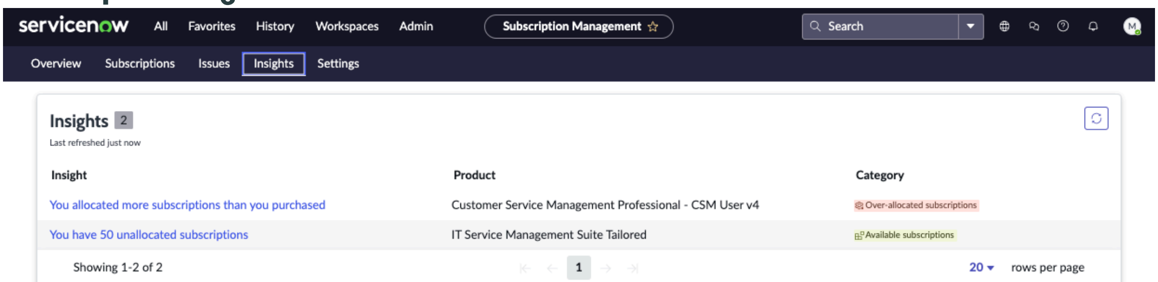

Accessing subscription insights

Access subscription insight information on the **Insights** tab by navigating to **Admin** > **Subscription Management** > **Insights** or **All** > **Subscription Management** > **Insights**.

Over-allocated subscriptions

Identify over-allocated per-user subscriptions on your instance by accessing the **Insights** tab. Stay in compliance by performing one or more of the following actions:

- Remove a group from a subscription
- Remove a role from a group
- Remove users from a group
- Purchase more subscriptions

For more information, see Remove a group from a product subscription in Subscription Management.

Available subscriptions

Identify available subscriptions on your instance by accessing the **Insights** tab. Maximize subscription usage and grant entitlements to the right users by allocating available subscriptions.

You allocate per-user subscriptions by adding one or more groups to a product subscription. Subscription Management helps you with the allocation process by recommending groups based on their assigned roles. For more information, see Allocate subscriptions in Subscription Management.

Configuring Subscription Management

Set up and configure Subscription Management.

Configuration overview

Subscription Management is free and active by default for all production instances and non-production instances. You do not have to activate a plugin.

1. Create a Subscription Management administrator

   Assign the usage_admin role to at least one user who will manage subscriptions for your organization.

2. Share subscription data from another instance

   View subscription data from other production instances in areas of Subscription Management that provide account-level summaries of subscription usage.

3. Configure the subscription allocation status in Subscription Management

   Configure the allocation status threshold using a percentage value that represents what near capacity means for your instance.

Create a Subscription Management administrator

Create a Subscription Management administrator by assigning the usage_admin role to an administrator.

Before you begin

Role required: admin

About this task

Usage admins review and manage how purchased subscriptions are used on your production instance. They add groups to appropriate subscriptions, monitor usage of applications, and update subscriptions as needed.

Assign the usage_admin role to users that manage subscriptions for your organization. Assign the usage_admin role to at least one person in your organization. Only admins and usage_admins can use Subscription Management.

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Select an existing user.
3. Update the user's role assignment by selecting **Edit** in the Roles related list.
4. Add the usage_admin role.
5. Select **Update**.

Share subscription data from another instance

View subscription data from another instance by enabling data sharing on that instance.

Before you begin

Role required: admin

About this task

By default, data sharing is restricted between production instances, which means you can't see subscription data for other production instances in Subscription Management. To view subscription data from another production instance, modify the trust configuration settings on that instance.

For more information about modifying the trust configuration for an instance, see Configure data synchronization for applications.

Note:
Subscription data from a non-production instance is only viewable from a production instance.

In this scenario, the instance for viewing subscription data across all your instances is considered the primary production instance.

Procedure

1. Log in to the instance that has subscription data that you want to view on your primary production instance.
2. Navigate to **All** > **Multi-Instance Management** > **Trust Configuration**.
3. In the list of instances, find the primary production instance.
4. Share data from your current instance by changing the value in the **Grant access** column to true.

Result

Subscription data is shared with your primary production instance.

What to do next

Log in to your primary production instance to view the subscription data that is now shared with that instance. Repeat these steps for each additional production instance that has subscription data that you want to view from the primary production instance.

Configure the subscription allocation status in Subscription Management

Identify when you're close to allocating all of your subscriptions by configuring an allocation status threshold if you prefer a threshold other than the default value of 90 percent in Subscription Management.

Before you begin

Role required: usage_admin or admin

Procedure

1. Navigate to the **Settings** tab in one of the following ways.
   - Navigate to **Admin** > **Subscription Management** > **Settings**.
   - Navigate to **All** > **Subscription Management** > **Settings**.
2. In the **Allocation Status** tab, enter a percentage value that represents near capacity for your instance.
   If you want the **Status** column to display Near capacity before reaching the 90 percent threshold, enter a lower percentage value.
3. Select **Save**.

Result

When the number of allocated subscriptions exceeds the percentage you configured, the value Near capacity appears in the **Status** column throughout Subscription Management.

Managing per-user subscriptions in Subscription Management

Manage your per-user subscription allocations using Subscription Management.

Overview of managing per-user subscriptions

To successfully manage your per-user subscriptions, follow the process in this infographic.

servicenow

Managing per-user subscriptions

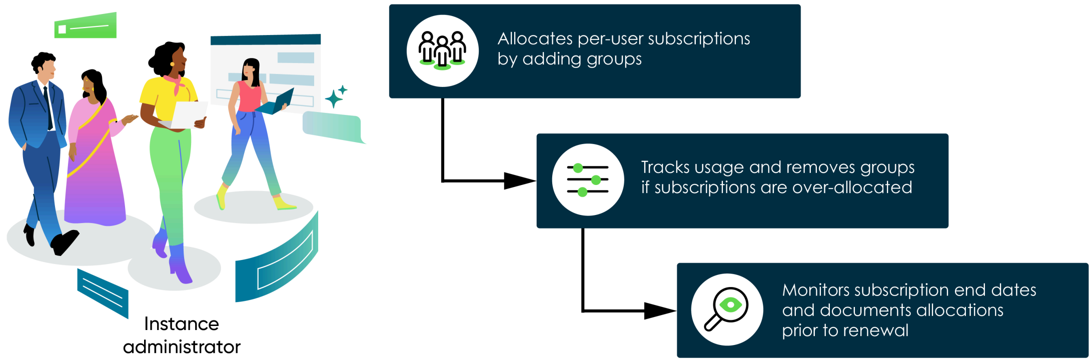

- Allocate subscriptions in Subscription Management

  Subscriptions are managed in Subscription Management using groups. Subscriptions are measured and audited in Subscription Management using roles. You allocate subscriptions by adding one or more groups with measured roles to a product subscription.

- Remove a group from a product subscription in Subscription Management

  You can potentially over-allocate subscriptions, meaning more users possess an entitlement than the number of purchased subscriptions. Subscription Management flags over-allocated subscriptions for you. To stay in compliance, remove users or groups from the subscription or purchase more subscriptions.

- Maintain per-user subscription allocations in Subscription Management at renewal

  When a per-user subscription renews, a new product SKU might be added to your account. Maintain your current allocations during the renewal process by documenting your allocations before the renewal date.

Recommended groups

When allocating per-user subscriptions, you can add groups based on the roles they have or you can add groups recommended by Subscription Management based on their roles. For example:

- If the Service Desk group has a fulfiller-type role in the ITSM application, you can add the Service Desk group to the ITSM product subscription to allocate the fulfiller subscriptions included with ITSM to all of the users in the group.
- If you have available ITSM subscriptions, you can add groups that are recommended by Subscription Management based on the types of roles they have. If you don't want to add the recommended groups, Subscription Management identifies other groups with relevant roles that you can add instead.

Groups that you add to a product subscription are considered subscribed groups. Consider assigning a group manager who is responsible for managing group membership and role assignments.

Measured roles

Roles that require a subscription are considered measured roles. Measured roles grant users access to an application and are categorized by the level of access they provide.

When possible, simplify role administration by assigning measured roles to groups instead of assigning roles directly to individual users.

Allocate subscriptions in Subscription Management

Give users entitlement to ServiceNow products by allocating subscriptions in Subscription Management.

Before you begin

Role required: admin

Procedure

1. Navigate to Subscription Management in one of the following ways.
   - Navigate to **Admin** > **Subscription Management**.
   - Navigate to **All** > **Subscription Management** > **Subscription Management**.
2. Select a product with subscriptions that you want to allocate.
3. Allocate subscriptions by either adding recommended groups based on role type or adding other groups with relevant roles.

Result

Subscriptions are allocated to all group members that grant each member entitlement to use the product.

Remove a group from a product subscription in Subscription Management

Fix an over-allocated subscription and free up entitlements by removing a group from a product subscription in Subscription Management.

Before you begin

Role required: usage_admin or admin

Procedure

1. Navigate to Subscription Management in one of the following ways.
   - Navigate to **Admin** > **Subscription Management**.
   - Navigate to **All** > **Subscription Management** > **Subscription Management**.
2. Identify an over-allocated subscription by checking either the **Status** column on the **Overview** tab or the list of over-allocated subscriptions on the **Insights** tab.
3. Select the over-allocated product subscription.
4. On the subscription details page, select the **Subscribed Groups** tab.
5. Select the check box next to one or more groups that you want to remove.
6. Select **Remove**.
7. In the dialog box, confirm that you want to remove entitlements for all users in the selected groups by selecting **Remove**.
8. **Optional:** Verify that the subscription is no longer over-allocated by checking the value in the **Status** column on the **Subscriptions** tab.

What to do next

The group is removed from the subscription. If the group still has a measured-role that requires a subscription, add the group to a different subscription or remove the measured-role from the group.

Maintain per-user subscription allocations in Subscription Management at renewal

Ensure subscription allocations remain consistent throughout the renewal process.

Before you begin

Role required: admin

About this task

When a per-user subscription renews, a new product SKU might be added to your account. This SKU is treated as a new product subscription in Subscription Management, requiring you to re-allocate subscriptions by adding groups with measured roles to the subscription.

Procedure

1. Before the renewal date, monitor the subscription end dates on the details page and document the current group allocations.
   See Viewing product subscription details in Subscription Management.
2. After the renewal date, add the documented groups to the renewed subscription in Subscription Management.
   See Allocate subscriptions in Subscription Management.

What to do next

Continue to monitor subscription renewal dates to ensure that your allocations are accurate throughout each renewal.

Managing custom tables and applications in Subscription Management

Keep your custom table entitlements updated by mapping all of your custom applications and tables to valid subscriptions in Subscription Management.

Overview of managing custom tables and applications

To successfully manage your custom tables and applications, follow the process in this infographic.

Managing custom tables and applications

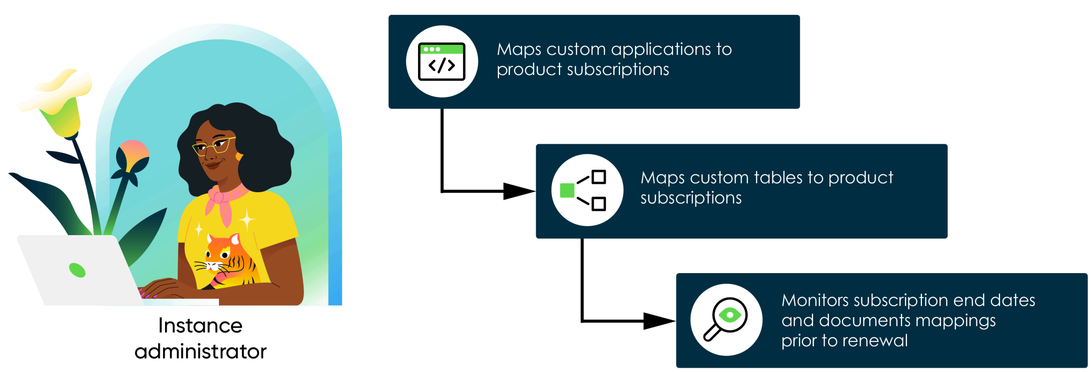

- Map a custom application to a product subscription in Subscription Management

  Map a custom application to a product subscription so that when developers add more tables to the application, they're automatically mapped to the subscription. Track your custom table usage to ensure that your organization doesn't run out of custom table entitlements.

- Map custom tables to a product subscription in Subscription Management

Map custom applications and custom tables on your production instance to a valid subscription. On non-production instances, mapping is recommended but not required.

- Maintain custom table and application mappings at renewal in Subscription Management

  When a subscription renews, a new product SKU might be added to your account. Maintain your current custom table mappings during the renewal process by documenting your mappings before the renewal date.

Custom tables

A custom table is defined as a table that you create on the ServiceNow AI Platform that isn't provided by ServiceNow. Some tables are exempted from being classified as a custom table.

> **Note:**
> For more information on custom tables on the ServiceNow AI Platform, see Table administration. For a detailed description of custom table use rights and exempt tables, see the Custom Table Guide on http://www.servicenow.com/upgrade-schedules.html

Subscriptions come with custom table entitlements that dictate the number of custom tables that you can map to a subscription. The number of custom table entitlements depends on the subscription.

Custom applications

A custom application is defined as an application that you create on the ServiceNow AI Platform that isn't provided by ServiceNow. If your organization purchased App Engine, you can use Subscription Management to manage your custom application subscriptions.

Recommended subscriptions

When you map a custom table or a custom application to a subscription, Subscription Management provides a recommended subscription based on the following criteria:

The table belongs to an app scope or package

The custom table or tables in the application belong to an app scope or package associated with the recommended subscription.

The table extends another table

The custom table or tables in the application extend another table mapped to the recommended subscription.

The table has a relationship to another table

The custom table or tables in the application are related to a different custom table or a ServiceNow out-of-the-box table associated with the recommended subscription.

The table doesn't meet any of the above criteria

The custom table or tables in the application don't meet any of the above criteria. If you have an App Engine subscription, Subscription Management recommends mapping the table or application to an App Engine subscription.

If you don't want to map to the recommended subscription, you can map to a subscription of your choice. For details, see KB1644296.

In some cases, groups aren't recommended because the measured role is inherited and not directly assigned. To work around this problem, assign the measured role directly to the group.

Once the group is assigned the measured role, the group appears in the list of recommend groups in Subscription Management.

App Engine usage

App Engine subscriptions entitle you to develop applications and deploy custom tables in production instances. Based on your App Engine version, you monitor the subscription usage of your App Engine applications differently.

- You can view and manage App Engine V2 subscriptions in Subscription Management just like any other product subscription. You can map your applications and custom tables to an App Engine V2 subscription and allocate per-user subscriptions by adding groups with measured roles to an App Engine V2 subscription.
- For details on auditing App Engine V1 usage, see KB0999383.

Map a custom application to a product subscription in Subscription Management

Maintain accurate entitlement for custom tables in a scoped application by mapping the application to a product subscription in Subscription Management. To stay in compliance, you must map custom tables that belong to custom applications to a product subscription.

Before you begin

Role required: usage_admin or admin

About this task

As your developers create and deploy new custom applications and tables on your instances, use Subscription Management to map them to subscriptions. After you map a custom application to a product subscription, additional tables that developers add to the application are automatically mapped to the subscription.

Procedure

1. Navigate to the **Issues** tab in Subscription Management in one of the following ways.
   - Navigate to **Admin** > **Subscription Management** > **Issues**.
   - Navigate to **All** > **Subscription Management** > **Issues**.
2. In the **Unmapped custom applications** tab, look for custom applications that aren't mapped to a subscription.
3. Update your entitlements by mapping one or more custom applications to a recommended product or a product of your choice.
   Subscription Management recommends product subscriptions with available custom table entitlements in the **Recommended Product** column.

Result

One or more custom applications are mapped to a product subscription and your custom table entitlement count is updated.

Map custom tables to a product subscription in Subscription Management

Maintain accurate entitlement for custom tables in the global scope and stay in compliance by mapping the tables to a product subscription in Subscription Management. Mapping your custom tables keeps your custom table allotment updated and helps you avoid running out of custom table entitlements.

Before you begin

Role required: usage_admin or admin

Procedure

1. Navigate to the **Issues** tab in Subscription Management in one of the following ways.
   - Navigate to **Admin** > **Subscription Management** > **Issues**.
   - Navigate to **All** > **Subscription Management** > **Issues**.
2. In the **Unmapped global custom tables** tab, look for custom tables that aren't mapped to a subscription.
3. Update your entitlements by mapping one or more custom tables to a recommended product or a product of your choice.
   Subscription Management recommends product subscriptions with available custom table entitlements in the **Recommended Product** column.

Result

One or more custom tables are mapped to a product subscription and your custom table entitlement count is updated.

Maintain custom table and application mappings at renewal in Subscription Management

Ensure custom table and application mappings remain consistent throughout the renewal process.

Before you begin

Role required: admin

About this task

When a subscription renews, a new product SKU might be added to your account. This SKU is treated as a new product subscription in Subscription Management, requiring you to re-map custom applications and custom tables to the subscription to stay in compliance.

Procedure

1. Before the renewal date, monitor the subscription end dates and document the current list of custom application and custom table mappings.
   See Viewing product subscription details in Subscription Management.
2. After the renewal date, map the documented tables and applications to the renewed subscription in Subscription Management.
   See Map a custom application to a product subscription in Subscription Management and Map custom tables to a product subscription in Subscription Management.

What to do next

Continue to monitor subscription renewal dates to ensure that your mappings are accurate throughout each renewal.

Grandfathered and exempted tables

Grandfathered and exempted custom tables are treated differently than other custom tables on your instance. Please refer to your organization's ServiceNow® contract for additional information.

Grandfathered tables

Depending on your organization's legacy subscription entitlements and contractual migration, grandfathered tables may be allotted to your production instance. To preserve an organization's previous custom table entitlements through the license migration, a grandfathered table entitlement may be provided to the production instance. This one-time process ensures that

custom tables that currently exist in a production instance don't count against your new custom table entitlements.

Grandfathered custom tables may not be transferred, reused, or otherwise classified as another custom table type. If you delete a grandfathered custom table, you can't classify any other custom table as a grandfathered custom table to take its place. Once a custom table has been associated with a grandfathered table subscription, it remains in the Custom Table Inventory even if the table is deleted from the instance.

To view your grandfathered tables, enter `ua_custom_table_inventory.list` in the filter navigator, and then add a filter with the condition **[Allotment type] [is] [Grandfather]**.

If you have a combination of custom tables and grandfathered tables in one application, the grandfathered tables must be mapped to a grandfathered table subscription and the remaining custom tables to one other subscription. Open a case with Customer Service and Support for help with mapping a grandfathered table to a grandfather table subscription.

Example: Grandfathered tables on an instance

Scenario: Your organization purchases three new subscriptions with entitlements for a total of 55 App Engine Starter custom tables (formerly Bundled Custom Tables). To ensure your organization doesn't lose the previous custom table entitlements, a grandfather subscription with an entitlement for 25 grandfathered tables is allotted to your instance. This subscription requires the admin to identify and map the existing 25 tables. These subscriptions result in a total custom table entitlement of 80. Once a custom table has been assigned to the grandfather subscription there is no option to revert this.

Multiple subscriptions with grandfathered and custom table entitlements

Subscription	Number of custom table entitlements	Number of mappable tables
Grandfathered subscription	25	25
New Subscription 1	25	25
New Subscription 2	25	25
New Subscription 3	5	5

Exempted tables

Exempted tables are custom tables that do not count towards your subscription entitlements. These tables consist mostly of custom extensions of certain system-related tables, such as sys_portal_page and sys_user_preference, that you extend for operational improvements in your production instance.

To view a list of exempted tables, enter `ua_exempted_table_inventory.list` in the filter navigator.

The Custom Table Guide legal schedule provides the complete definition of which tables are considered exempt.

Note:
Legal schedules are enforced based on when the contract is executed.

To find past versions of the Custom Table Guide, see the Legal Schedules Custom Table Guide.

Monitoring capacity subscriptions

Manage and track capacity subscriptions using Subscription Management.

Overview of monitoring capacity subscriptions

Learn how capacity totals are calculated

Capacity subscription usage is calculated using a table query or scripted definition. You can view a Capacity Definition record to view the details of the table query or scripted definition. You can also execute the definition on demand to view the latest counts.

After you determine how a capacity total for a subscription is calculated, you can adjust the number of allocated resources or increase your capacity subscription to keep your subscription in compliance. For example, you can find out which tables are queried for active users or how usage is defined for devices or software when the system allocates CIs to a capacity subscription.

Monitor capacity subscription calculations in Subscription Management

View how a capacity total is calculated to keep your subscription in compliance.

Before you begin
Role required: admin or usage_admin

Procedure

1. Navigate to Subscription Management in one of the following ways.
   - Navigate to **Admin** > **Subscription Management** > **Subscriptions**.
   - Navigate to **All** > **Subscription Management** > **All Subscriptions**.
2. In the list of all your product subscriptions, select a capacity subscription.
3. On the subscription details page, preview the Capacity Definition record by selecting the record ID next to the **Capacity Definition** field.
4. Depending on the definition type, review the following details to view how usage is calculated.
5. **Optional:** Select **Open Record** to view the full Capacity Definition record details.
6. **Optional:** Select **Execute** to run the query or script immediately and view the latest results.

   Note:
   Executing a script or table query in a Capacity Definition record does not count toward the execution total that is monitored by the system.

   The total time required to execute and the latest count are displayed. The count from prior executions is displayed in the Previous Executions related list.
7. If the capacity subscription is out of compliance, adjust the number of allocated resources or increase the capacity subscription.

Subscription Management reference

Reference topics provide information about subscription types, meters, and domain separation support.

Types of subscriptions in Subscription Management

Subscriptions to ServiceNow applications come in different types. The type of subscription determines the allocation of users, access to applications, and custom application and table entitlements.

Each type of product subscription is measured according to a meter. For details on the different types of meters and what they measure see KB0727967.

Subscription type	Associated meter types	Usage allocation method	Description
Capacity	• Devices • Subscription Units • Transactions	No manual steps are required. Usage is automatically calculated and shown in your instance.	Automatically tracks measured units, and counts of transactions, records created, or resources managed such as devices, software, or nodes.
Display Only	None	Allocation isn't measured	Displays information. This type of subscription doesn't support allocation.
Per-User	• Rights-based user: measure of Fulfiller Users and Business Stakeholder Users • Application user: measure of users with any level of access to a product; includes Creator Users	Manually allocated	Provides entitlements to users according to their assigned roles and associated access rights. You allocate per-user subscriptions manually by adding groups with measured roles or access rights to a product subscription. Subscription Management helps you with the allocation process by recommending groups based on their role assignments.
Unlimited	Unlimited: No measurement limit	Automatically allocated by your instance	Provides entitlements to users, with no limit to the number of users that can be allocated.
Unrestricted User	Unrestricted user: Instance-wide measure of all active users	Automatically allocated by your instance	Provides entitlements for an organization's number of active users regardless of their role assignments. An active user is any user whose record in the Users [sys_user] table has a value in the **User ID** field and has the **Active** field set to true.

Subscription form

The Subscription form contains important information about your subscription.

You can view the subscription form by entering `license_details.list` in the filter navigator to open the Subscriptions [license_details] table. Select a subscription name to open the Subscription form. Some fields do not appear for ServiceNow Store apps.

Subscription form

Field	Value
Name	Name of the subscription.
Category	Source of the purchased subscription: ServiceNow platform or ServiceNow Store.
Type	Type of subscription: • Per-User • Unrestricted User • Capacity • Unlimited • Display To learn more about subscription types, see Types of subscriptions in Subscription Management.
Start date/End date	Start and end of the subscription period as agreed in your service contract. This information appears in the General section of the service contract for the subscription.
Capacity Definition	The ID of the Capacity Definition. Learn how the total is calculated by selecting the Capacity Definition record ID on the product subscription details page. See Monitor capacity subscription calculations in Subscription Management.
Display only	Value of **true** indicates that user allocation or monitoring is not currently supported for the subscription.
Product Code	The product SKU shown on your subscription order form.

Related links

The related links on the subscription form may vary depending on the subscription.

Link	Description
Manage Your Tables	See custom tables mapped to the subscription.

Link	Description
View Subscription History	View subscription status, dates, and trend information.
View User Subscription History	Lists user allocation history for the subscription. The list shows a user's current subscription status and how they were allocated.

User sets and groups in Subscription Management

Subscription Management uses groups (previously user sets) to provide entitlement to sets of users.

User set and groups overview

Prior to Quebec, Subscription Management used user sets to group users together. Usage_admins grouped users by conditions, such as what city or department they are in. User sets could then be assigned to subscriptions to give every user in the user set a license to the subscription.

In Quebec and later releases, all user sets are converted to groups. A group is a set of users who share a common purpose. Similar to user sets, groups can be added to subscriptions. Adding a group to a subscription gives every user in the group entitlement.

User set conversion

On upgrade to Quebec or later releases, all current user sets convert to groups. The new groups contain the same users and properties as your user sets. The conversion is a one-time occurrence, and users are not automatically added or removed from the new group after the conversion.

Group management

Admins and user_admins create groups and add users in the **User Administration** application.

To learn more about creating and editing groups, see Creating groups.

Components installed with Subscription Management

Several components are installed with activation of the Subscription Management plugin (com.glide.subscription_management), including user roles and tables.

Roles installed

Role title [name]	Description	Contains roles
Usage administrator [usage_admin]	Administers Subscription Management. Allocates subscriptions, maps custom tables and custom applications to subscriptions.	sn_sub_man.admin

Tables installed

Table	Description
Subscribed Groups [sn_sub_man_st_subscribed_groups]	[Remote table] Groups that you add to a product subscription
Subscription Insights definition [sn_sub_man_st_subscription_insights]	[Remote table] Records for available subscriptions and over-allocated subscriptions
Subscription license detail definition [sn_sub_man_st_subscription_license_detail_metric]	[Remote table] Subscription allocation metrics and history
Unallocated Group Recommendations [sn_sub_man_st_unallocated_group_recommendation]	[Remote table] Product recommendations for unallocated groups

Subscription Management roles

Subscription Management is installed with these roles.

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

Usage administrator [usage_admin]

Administers Subscription Management. Allocates subscriptions, maps custom tables and custom applications to subscriptions.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevate to a privileged role.

No

Special considerations

None

Usage administrator [usage_admin]

Administers Subscription Management. Allocates subscriptions, maps custom tables and custom applications to subscriptions.

Domain separation and Subscription Management

Domain separation is unsupported for Subscription Management. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

Related topics

Domain separation for service providers

Now Support administration

The Now Support portal is where you can access self-help, get technical support, and manage your ServiceNow instances.

Who uses the Now Support portal

If you are responsible for managing, monitoring, or maintaining your company's ServiceNow instance(s), then you will likely use the Now Support portal. ServiceNow Customer Administrators are the gatekeepers and manage their companies' Now Support users. Administrators can then add or remove other users, such as developers, architects, and partners. What you can do within the portal will vary based on your role.

> **Note:**
> For information about the types of notifications different Now Support users may receive, see KB0547262: Company Key Contacts and Notification List Overview on Now Support.

View Now Support users

As a Now Support admin, you can view active and inactive users and their roles.

Before you begin

Role required: admin

Procedure

1. Navigate to **Now Support**
2. Select the drop-down profile menu.
3. Select **Manage users and accounts**.
4. Select **View all Users** or search for a specific user.

Create a user in Now Support

You can create users in Now Support and assign them access roles.

Before you begin
Role required: admin

Procedure

1. Navigate to **Now Support**
2. Select the drop-down profile menu.
3. Select **Manage users and accounts**.
4. Select **View all Users**.
5. Select **Add new user**.
6. Fill in the required fields.
7. Select the appropriate Now Support access roles.
8. Select **Submit**.

Manage user access in Now Support

You can manage users in Now Support and change their access roles.

Before you begin
Role required: admin

Procedure

1. Navigate to **Now Support**
2. Select the drop-down profile menu.
3. Select **Manage users and accounts**.
4. Select **View all Users** or search for a specific user.
   You can also select **Locked users** or **Inactive users** to find the user you are looking for.
5. After you select the user you want to change access for, select **Edit Role(s)** from the **Actions** pull-down field.
6. Add or remove Now Support access roles.
7. Select **Confirm**.

Configure ServiceNow AI Platform Core Features

The ServiceNow AI Platform provides for a multitude of customization options to your applications. Customize your UI, handle user and data administration, and localize your instance for time zones, currencies, and more.

User administration

Set up users, groups, and roles.

ServiceNow AI Platform tables and data

Manage data with external sources, and configure data tables.

Forms, fields, and lists

Control the action and appearance of forms, fields, and lists.

Search administration

Enable users to find the information through Zing or AI search engines and specialized search experiences.

Notifications

Manage system email, notifications, and responses to inbound messages.

ServiceNow AI Platform translation and localization

Match languages and currency everywhere on the globe.

Time configuration

Create calendars, schedule events, and display time info to users around the world.

Currency administration

Manage the way financial data is displayed.

User administration

Create account access types and related roles to work within your instance. User records give you control over individual single account types and the records or actions that they can access. Group records expand control further, to assign access to teams or departments, as well as to subgroups or by skill associations. Customize instances even more with Roles, and assign access to various apps or modules. Troubleshoot and test it all with Tools that provide auditing, impersonation, and more.

See User administration.

ServiceNow AI Platform tables and data

Manage and use your business information effectively on the ServiceNow AI Platform. Take advantage of Data Management to store, transfer, import, and archive data. Work with multiple records at once using Table Administration, and connect with third-party or external instances to handle data with Financial Services Remote Tables.

See ServiceNow AI Platform tables and data.

ServiceNow AI Platform forms, fields, and lists

Take control of every aspect of data collection and usage in your instance. Forms on the ServiceNow AI Platform are fully customizable, with options for personalization, templates, attachments, UI, and more. Create and modify Fields, used in forms and displayed in lists or tables. Define user access to all your data with Lists, including custom controls and functions.

See ServiceNow AI Platform forms, fields, and lists.

Search administration

Help your users find information with ease using ServiceNow search administration products. AI Search and the Zing search engine provide intelligent search results and proper indexing. Focus your data requests with Table searches, and go wider with Global searches, all with configurable search systems across the ServiceNow AI Platform. Improve search assistance with search suggestions, Contextual Search based on intent, and search data analysis.

See Search administration.

Time configuration

Keep your business on-time with the ServiceNow AI Platform. Handle schedule and event info collection with Date/Time fields, and manage personnel, users, and data across Time Zones. Business and fiscal calendars and event system schedulers are all available and tailored to best help you and your users. Time displays, counts, and time-based data analysis all help you manage your work and track historical info.

See Time configuration.

Currency administration

Monitor currency values with ServiceNow currency administration systems. Configure Standard Currency fields for financial and non-financial records and applications, or use expanded Foreign Exchange (FX) Currency fields for added configuration options.

See Currency administration.

Service Level Management

Ensure your services to your customers are properly managed and documented. The full series of ServiceNow modules for Service Level Management lets you define service contracts and service level agreements (SLA) as needed. Repair, monitor, and validate SLA records for time and duration accuracy, get analyses and reports on service, and separate record domains for information display and control.

See Service Level Management.

ServiceNow AI Platform translation and localization

Communicate effectively with users and customers around the world. The Localization Framework provides info to users in a variety of languages and currencies across a multitude of apps—all within the same instance. Dynamic Translation helps accommodate users even more fully, by actively translating text in forms and other applications.

See ServiceNow AI Platform translation and localization.

ServiceNow AI Platform integrations

Bring ServiceNow and external applications together with integrations on the ServiceNow AI Platform. Make onboarding and service configuration easy by adding to your instance. Integrate other instances, third-party apps, and data sources, and make helpful features such as Single Sign-on available throughout your enterprise.

See ServiceNow AI Platform integrations.

Notifications

Use Notifications to automate status updates by email, system, or user (through individual providers). Set up Email administration and your instance email client to work with notifications, and proper sending and receiving of email. System notifications enable real-time updates through push notifications, SMS, chat messages, and more.

See Notifications.

State Management

Monitor workflow in your instance with State Management. Define state models and transitions to determine your business's service processes, as well as UI actions possible based on state.

See State Management.

User administration

Manage the individuals who can access your instance by defining them as users in the system. Create user groups and assign users to them. Use roles to specify what different users and user groups can see and do.

Get started

Explore	Create users, companies, and departments	Create groups
Learn about User administration.	Create records for users or enable self-registration. Optionally create companies and departments to further categorize users.	Create and manage groups that you can grant roles to and assign users to for efficiency.
Manage roles	**Monitor user activity**	**Monitor instance usage**
Manage roles that govern what users and groups with that role can do.	Monitor user activity and troubleshoot issues by impersonating users and managing sessions.	Track application usage on your instance.

Troubleshoot and get help

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

Exploring user administration

Learn more about user administration for your instance.

Overview

Creating users, groups, and roles provide a flexible and scalable way to manage access to features on the ServiceNow AI Platform. By creating user accounts, assigning users to groups, and defining roles and permissions, administrators can ensure that users have the appropriate level of access to applications and data. This enables organizations to control access to sensitive data, maintain conformance with regulatory requirements, and improve overall security. Additionally, users, groups, and roles can be easily managed and modified over time as organizational needs change.

Work flow

Subscriptions, users, groups, and roles work together to help you define who can access features on your instance.

1. Subscription Management

   Understand your subscriptions. Subscription management enables you to manage your subscriptions proactively and monitor subscription usage on your instances.

   Subscriptions may include per-user subscriptions. For more information, see Managing per-user subscriptions in Subscription Management.

2. Creating users, companies, and departments

   Create an account record for the individuals who have access to your instance. Each user account has a unique login ID, password, and set of permissions (roles) that define what they can do and access within the platform.

3. Creating groups

   Define groups that have similar roles or permissions. Groups enable you to apply permissions (roles) to multiple users at the same time. When a user is a member of a group, that user has the same permissions that have been defined for the group.

   You can view group members by navigating to **All** > **User Administration** > **Groups**. Select a group name and view the members in the Group Members related list.

4. Managing roles

   Roles describe the types of activities that a user can perform on the instance. Each role has a set of permissions that can govern what the users and groups can do, such as read, write, create, or delete records. Roles can be assigned to users and groups. Users can have multiple roles.

   For a complete list of the roles included with the ServiceNow AI Platform, see Base system roles.

   Role records are stored in the Roles [sys_user_role] table.

5. Monitoring instance usage

   Users with the admin or usage_admin role can view the **Application usage overview** and ServiceNow Store **usage overview** dashboards to track instance usage.

6. Monitoring user activity

   Users with the admin role can impersonate users, manage user sessions, and leverage non-interactive sessions.

User administration workflow

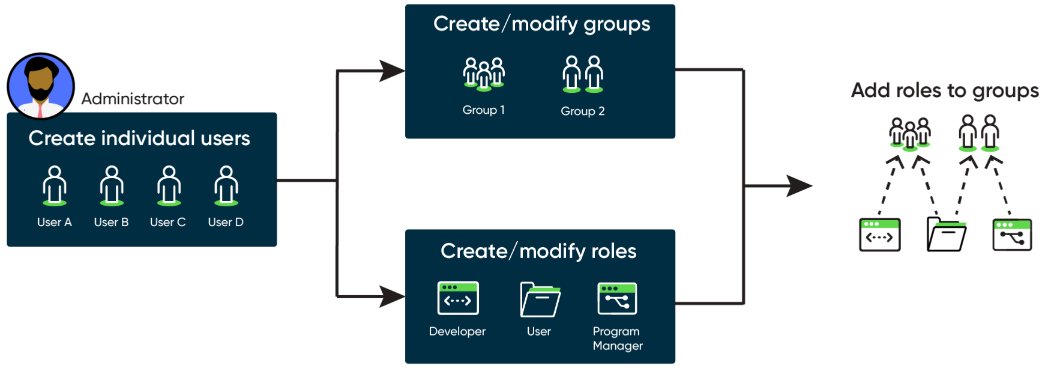

User preferences

Users can configure many UI features. Some examples include the number of rows per page in a list or whether the response time displays at the bottom of a list or form. Administrators can modify or delete these preferences as needed. For more information, see User preferences.

User groups

A group is a set of users who share a common purpose. Groups may perform tasks such as approving change requests, resolving incidents, receiving email notifications, or performing work order tasks. Any business rules, assignment rules, system roles, or attributes that refer to the group apply to all group members automatically. Users with the user_admin role can create and edit groups.

When possible, simplify user administration by assigning roles to groups. Create groups that contain all the roles necessary for specific personas, and then assign users to those groups.

Note:
You can view group members by navigating to **All** > **User Administration** > **Groups**. Select a group name and view the members in the **Group Members** related list.

Group records are stored in the Groups [sys_user_group] table.

User roles

Roles control access to features and capabilities in applications and modules. The admin role provides access to all features and capabilities.

After access has been granted to a role, all the groups or users assigned to the role are granted the access. Roles can contain other roles, and any access granted to a role is granted to any role that contains it.

For a complete list of the roles included with the ServiceNow platform, see Base system roles.

Note:

When possible, simplify user administration by assigning roles to groups. Create groups that contain all the roles necessary for specific personas, and then assign users to those groups.

Role records are stored in the Roles [sys_user_role] table.

The User record

Learn about user records and their use within the ServiceNow AI Platform.

User account records

User records establish a relationship between an individual and your ServiceNow instance. User records consist of a user name, a password, and information relating to the individual, such as contact information, location, and job title.

User record

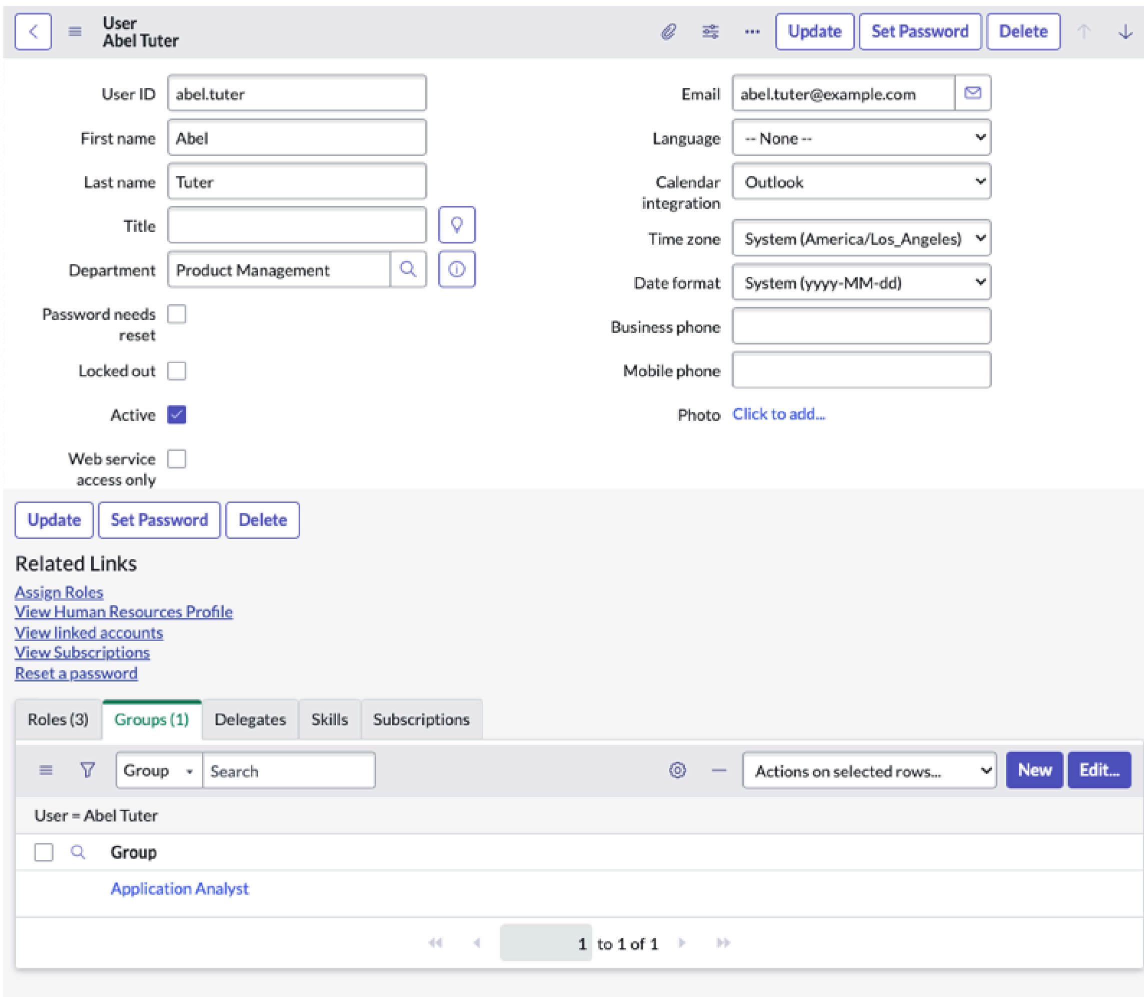

User records are stored in the Users [sys_user] table.

Related records

User records are associated with records on several other tables to control permissions, preferences, and other features.

Roles

Roles control access to features and capabilities in applications and modules. For more information on roles, see Managing roles.

Note:

When possible, simplify user administration by assigning roles to groups. Create groups that contain all the roles necessary for specific personas, and then assign users to those groups.

Groups

A group is a set of users who share a common purpose. Users assigned to groups are automatically assigned to all roles associated with that group. For more details, see Creating groups and Managing roles.

Delegates

In addition to role and group assignments, users can be assigned as delegates, giving them permission to act with the same permissions as a delegator user. See Delegating roles for more information on delegation.

Skills

Use skill management to associate users with their areas of training and expertise. For more information on skill management, see Skill Management.

Subscriptions

Administrators use subscriptions to control which users have access to purchased subscriptions on their instances. Details on subscription management can be found at Subscription Management.

User preferences

User accounts are also connected with user preferences. Users can save personalized preferences to configure many UI features, as well as preferences regarding the notifications they receive. Details on administering user preferences are found at User preferences.

System and guest users

Some automated processes use the system or guest user to apply and track changes to records. As a result, some records may show that they were last updated by system or guest.

For example, when a user logs in for the first time in a day, some fields on that user's record are updated by the system user, such as **Last login** and **Last login time**. If a user has a failed login attempt or is locked out, some fields on that user's record are updated by the guest user, such as **Failed Login Attempts** or **Locked Out**.

If a record was last updated by the system or by guest users, identify the fields that were updated by enabling auditing for the table and viewing the audit history set. For more information, see Configuring auditing for a table and Knowing about History sets.

Creating users, companies, and departments

Create user records for the individuals who access your instance. Users can be assigned to groups with defined roles to determine what records and actions they can access.

Creating users

Users are typically added through Lightweight Directory Access Protocol (LDAP) directory integrations. Admins can also manually add users to the instance, enable self-registration for new users, and impersonate users to ensure that they have the proper access privileges.

For more information on LDAP integrations, see LDAP integration.

Create a user

You can add a user to your instance to enable them to log in and use designated application features.

Before you begin

Role required: user_admin

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Select **New**.
3. On the form, fill in the fields.

Field	Description
User ID	Create a unique identifier for this user's ServiceNow login user name. Typical examples of user IDs are `cwitherspoon` and `charlie.witherspoon`. You can't create a user whose User ID duplicates an existing user. If you do import duplicates from an update set, the more recently created name takes the duplicate User ID.
Given name	Enter the user's given (often their first) name.
Family name	Enter the user's family name (often their last name). **Note:** You can clear the **First Name** field, or the **Last Name** field in an existing user record, but you can't clear both at the same time.
Title	Enter a title or job description, or select one from the list.
Department	Select the user's department from the list.
Password	Assign a password to the user. To set up the user's password, fill in the fields on the form and **save** the record. Then, select **Set Password**. This password can be permanent or temporary.
Password needs reset	Select this check box to require the user to change the password during the first login.
Locked out	Select this check box to lock the user out of the instance and terminate all their active sessions. The system helps prevent users with the admin role from locking themselves out.
Active	Select this check box to make this user active. Only the administrator sees an inactive user in: ◦ Lists of users ◦ The selection list on reference fields (magnifying glass icon) ◦ The auto-complete list that appears when you type into a reference field
Web service access only	Select this check box to designate this user as a non-interactive user. This field is available with Non-Interactive Sessions.
Internal Integration User	Select this check box to Mark service accounts as internal integration users.
Email	Enter the user's email address. To enter a non-standard email address that doesn't pass field validation, you must deactivate the validation script first:

Field	Description
	a. Navigate to **System Definition** > **Validation Scripts**. **b.** Select the **email** record. **c.** Clear the **Active** check box and save the change. **d.** Complete the user profile, including the email address, and update or submit the record. **e.** Reactivate the email validation script.
Language	Select your preferred language from the list.
Calendar integration	Select **Outlook** to have this user receive meeting notifications via email directly to the calendar. Otherwise, select **None**.
Time zone	Select the user's time zone.
Date format	Select the user's preferred format for dates.
Business phone	Enter the user's business phone number.
Mobile phone	Enter the user's mobile phone number.
Photo	Attach a photo of the user, if appropriate.

The minimum fields required to create a user record are: **User ID** and either **First name** or **Last name**.

You can select the **Personalize Form** icon to remove fields.

4. **Optional:** Add **Roles**, **Groups**, **Delegates**, **Skills**, and **Subscriptions** to the user.
   For more information, see Managing roles, Creating groups, and Delegating roles.
5. Select **Submit**.
   The new user record appears in the list.

User self-registration

The User Registration Request [com.snc.user_registration] plugin provides the ability for unregistered users to request access to a ServiceNow instance. An administrator can activate the plugin.

A user can request an account by navigating to the instance. If the plugin is installed, the following section is added to the welcome screen.

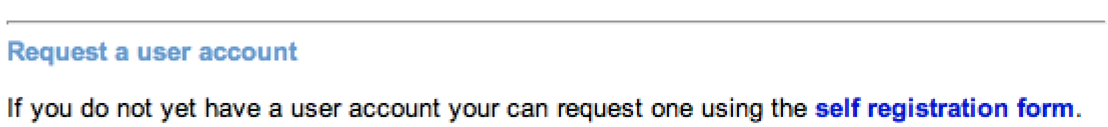

The user can complete and submit the self-registration form, and see a confirmation that it was submitted. The user receives an email when the account is registered.

Note:
If the email address entered in the self-registration form is already in the system, the request is not submitted.

Approve a self-registered user account

When a user submits a self-registration form, an admin can review and approve it.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Pending User Registrations** and open the request.
2. Use the **Create User** and **Reject** related links on the registration request form to approve or deny the request.
   - If **Create User** is selected, a new user is created using the email address as the **User ID**.
   - If **Reject** is selected, the request is marked **Rejected**.

   If the request was accepted, the user is sent an email notification with the login information. If it was denied, the user is sent an email with the rejection information.
3. **Optional:** To view past registration requests, remove the **[State] [is] [Pending]** condition from the filter in the list view.

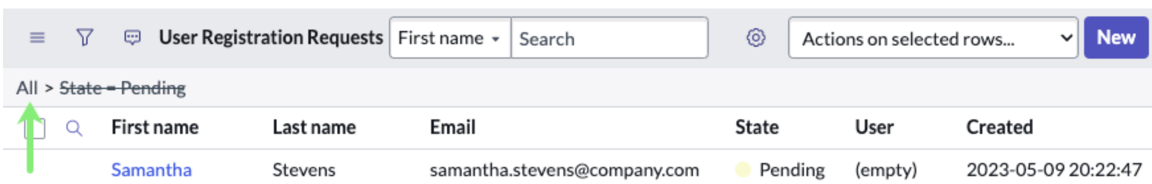

Enable automatic self-registered user account approval

Admins can enable users to self-register. Enable automatic approval of such accounts to streamline user registration.

Before you begin

Role required: admin

Procedure

1. Navigate to **System Properties** > **System**.
2. Select **Enable auto processing of user registration request**
3. Select **save**

   If enabled, registration requests do not require approval. Instead, the business rule *`Auto-Process User Registration`* creates the user record from the information provided.

What to do next

Request activation of the Explicit Roles (`com.glide.explicit_roles`) plugin. The Explicit Roles plugin creates two roles to differentiate between internal and external users.

Add a new company

You can add companies that represent vendors, manufacturers, or customers with whom you do business. These companies provide a way to categorize users, groups, and assets.

Before you begin

Role required: user_admin or admin

Procedure

1. Navigate to **All** > **User Administration** > **Companies**.
2. Click **New**.
3. Complete the form with the company details.

Company fields

Field	Description
Name	Name of the company.
Phone	Company phone number.
Fax phone	Company fax number.
Manufacturer	Whether the company is a manufacturer.
Vendor	Whether the company is a vendor.
Stock symbol	Three or four letter stock symbol for the company.
Stock price	Current price at which company stock is sold.
Street	Mailing street address of the company.
City	City in which the company is located.
State	State or province in which the company is located.
Zip/Postal code	Zip or postal code for the company.
Notes	Any information about the company that would be helpful for others to know.
Fields that can be added by Personalize a form:	
Latitude	The latitude of the company, if applicable. This field is populated by a business rule called `get_lat_long`. Deactivate this business rule to prevent the system from overwriting any values populated in the field manually. Latitude is expressed as a floating point data type. Latitude from upgraded versions of ServiceNow expressed in any format other than floating point appears in a column called **Old Latitude**. The system attempts to convert all latitude values from previous versions to the floating point notation, where possible. **Note:** This field doesn't display by default. You can add it by personalizing the form.
Longitude	The longitude of the company, if applicable. This field is populated by a business rule called `get_lat_long`. Deactivate this business rule to prevent the system from overwriting any values populated in the field manually. Longitude is expressed as a floating point data type. Longitude from upgraded versions of ServiceNow expressed in any format other than floating point appears in a column called **Old Longitude**. The system attempts to convert all longitude values from previous versions to the floating point notation, where possible.

Field	Description
	Note: This field doesn't display by default. You can add it by personalizing the form.

Note:
The IT Finance application adds a Finance view to the Company form. The Finance view adds a chart that shows expenses that were allocated to the company.

What to do next
Normalize company data to create consistency when referring to a company name, such as a vendor or manufacturer. For more information see, Normalization data services.

Normalization data services

The Normalization Data Services plugin helps maintain consistency for table fields that refer to a company name.

Tables related to configuration items and assets usually contain columns referring to a company name, such as a vendor or manufacturer. Often, these tables refer to the same company by several different names. This situation can happen for many reasons. For example, vendors sometimes use different company names depending on the device. This inconsistency creates problems, especially for reports that rely on these company names.

The Normalization Data Services plugin helps alleviate this problem.

When you enable and configure the Normalization Data Services plugin, the system downloads a list of standard company names that ServiceNow has compiled. It also downloads a list of common variants of that name. Anytime the plugin finds a company-name field with one of those variants, it substitutes the standard name in its place.

Note:
The Normalization Data Services plugin refers to a standard name as a normalized name and to a variant name as a discovered name.

While the Normalization Data Services plugin provides services similar to the field normalization feature, this plugin doesn't require you to manually set up the mappings from one name to another. If you like, you can add your own mappings to the set. However, you automatically start with extensive lists of common company names (normalized names) and their variants (discovered names).

The Normalization Data Services plugin stores data in two tables. The Normalized Company Names table contains the list of normalized company names. The Normalized Mappings table contains the mappings between each discovered name and its normalized name.

Important:
If you edit a field whose value is a normalized name, you change the normalized name for ALL discovered names that map to it. These types of updates occur regardless of the table in which you edit the field.

The Normalization Data Services plugin adds a unique index or hash on the [core_company] table. You can store only one company record for a company name. The unique hash is the same for two companies with the same name. The uniqueness is required for many features to work correctly. For example, **Discovery**, which creates customer management database (CMDB) models and CIs. On a domain separated instance, the unique index is expanded to

hash,sys_scope. This enables each domain to have its own version of a company with the same name.

Implementing normalization data services using guided setup

You can implement Normalization Data Services using a guided setup.

Normalization Data Services guided setup provides a sequence of tasks that help you configure the Normalization Data Services plugin on your instance. To open Normalization Data Services guided setup, navigate to **Normalization Data Services** > **Guided Setup**. For more information about using the guided setup interface, see Using guided setup.

Note:
Because the Normalization Data Services feature requires an internet connection to download mapping information, this feature isn't supported for on-premise instances.

Activating the plugin

In the first task of the guided setup, you activate the Normalization Data Services plugin (com.glide.data_services_canonicalization.client). Activating the plugin installs all the necessary schema to support normalizing references to Company records in the customer management database (CMDB) and Asset Management.

Downloading normalized data

In this task, you download the repository of all enterprise vendor names and the mappings to their normalized counterparts. After the download finishes, you can view the downloaded data in the Normalized Company Names [cds_client_name] table. You can verify mappings in the Normalized Mappings [cds_client_mapping] table. For example, you can verify that "Dell", "DELL", "Dell (UK)" are mapped to "Dell Inc.".

Updating reference qualifiers

This task modifies the reference qualifiers in reference fields linked to the Companies [core_company] table, helping to ensure that users can only select from a list of normalized company names. For example, when creating a Computer (cmdb_ci_computer) record, users can only select "Dell Inc." and not "Dell" or "Dell Incorporated".

Some reference qualifiers may not get updated. To normalize these reference qualifiers manually, see KB0784201.

Activating normalization properties

In this task, select the properties you want to activate. See Normalization properties.

Enable the **Normalize existing canonical core_company records** property if the Companies [core_company] table has two or more records that were both set to Normalized ("canonical=true") before any mapping. This option ensures only one [core_company] record is set to Normalized. When you normalize the CMDB tables later in guided setup, the normalization job reassigns the proper normalized value.

For example, prior to mapping, both "DELL Inc" and "Dell Inc." are set to Normalized ("canonical=true"). You define a new mapping for "DELL Inc" to "Dell Inc.". If this property is true, guided setup denormalizes "DELL Inc", and updates the manufacturer field in the CMDB records to "Dell Inc." when the CMDB tables are normalized.

Normalizing data

To finish the guided setup, complete the remaining tasks by normalizing data in the following tables:

- Configuration Items (CMDB)
- Configuration Items (CMDB) Model
- Software Asset Management tables

Normalized company names table

The Normalization Data Services plugin stores the normalized company names in the Normalized Company Names table.

Normalized company names fields

Field	Description
Name	The normalized name for this company.
Table	The table in which this name is stored.
Description	[optional] A description with further information about this company or record.
Field	The field in which this name is stored.
Customer override	**True**, if the customer has an override in place for this name. Otherwise, **false**.

Normalized company names related lists

List	Description
Normalized mappings	A list of all the discovered names that map to this normalized name.

See KB0819618 for more information.

Normalized Mappings table

The Normalized Mappings table lists all the discovered names and the normalized name to which each maps.

Normalized mappings fields

Field	Description
Discovered name	A variant of a normalized name in the database.
Normalized name	The normalized name to which the discovered name maps.
Table	The table in which this name is stored.
Field	The field in which this name is stored.

Normalized company names related lists

List	Description
Related mappings	A list of all the discovered names that map to the same normalized name as the selected record.

Actions

List	Description
Promote discovered name	Replaces the normalized name with the discovered name for the selected record and for all the records in its related mappings list.

Normalization properties

On the normalization properties form, you can see and change the configurable properties for the Normalization Data Services plugin.

Normalization Properties

Field	Description
Reference qualifiers on all tables that reference the company will be updated to use the Normalized field.	If **Yes**, any reference field for the company – for any table across the platform – uses the Normalized name. This service only works if you have run the Update Reference Qualifiers task in the Guided Setup.
Enable a business rule that automatically normalizes manufacturer names for configuration items.	If **Yes**, the system normalizes company names when you add or update configuration items by a mechanism other than Discovery (such as by manual import sets).
Enable Discovery to use the normalization service for the manufacturer name.	To have Discovery use the normalized company name for the manufacturer name, select **Yes**.
This property is to enable or disable the Normalization process.	To enable the Normalization Data Service process, select **Yes**. To disable, select **No**. For details, see Normalization data services.
Normalize existing canonical core_company records.	Select **Yes** to update an existing normalized (canonical=true) Companies record to not-normalized (canonical=false) if a mapping is present. For details, see KB0957144.

Changing normalized company names

You can change a normalized company name several different ways. In all cases, that change affects all normalized fields referring to that company.

You have several options for changing a normalized company name:

- Edit the Normalized Name field in the Normalized Mappings table. This method is preferred.
- Edit the Normalized Company name table.
- Edit the Company Name field on any table that refers to the Company [core_company] table.

> ⚠ **Warning:**
> If you edit a field whose value is a normalized name, you change the normalized name for ALL discovered names that map to it.

Change a normalized company name

You can change normalized company names by editing records in the Normalized Company Name table.

Before you begin
Role required: admin

About this task
You can add and edit records in the Normalized Company Names table.

Procedure

1. Navigate to **All** > **User Administration** > **Normalization Data Services** > **Normalized Mappings**.
2. Find the record with the name that you want to replace and edit the Normalized name field.
   The system changes the Normalized Company name for every discovered name that maps to that normalized name.

Add a department

Departments provide another way to categorize users, groups, and assets. You can add departments and assign them to users.

Before you begin
Role required: user_admin or admin

About this task
An administrator may need to configure the form to show all the fields listed in the steps. For more information see Personalize a form.

Procedure

1. Navigate to **All** > **User Administration** > **Departments** > **Create a new record**.
2. Enter or modify the department name, ID, and description.
3. Select the company that the department is associated with.
4. **Optional:** Add a department head, primary contact, or both from your list of users.
5. Select **Submit**.

Creating groups

Admins can add users to one or more groups.

Create a user group

Create groups and assign roles to them. Users assigned to the group inherit the roles.

Before you begin
Role required: user_admin

Procedure

1. Navigate to **All** > **User Administration** > **Groups**.
2. Select **New**.
3. On the form, fill in the fields.

   Some fields appear after personalizing the form. Personalize a form.

Field	Description
Name	Name of the group.
Manager	Group manager or lead.
Type	Category for this group. For example, a group designated as type **catalog** is a service catalog group and can also be accessed under the **Service Catalog** > **Catalog Policy** > **Fulfillment Groups** module. See Configure assignment group types for more information. You may need to Personalize a form to add the **Type** field. Activating the Work Management plugin adds the **Type** field automatically. **Note:** ITIL is added for groups with an empty group type. Also, the default reference qualifier for tasks enables these groups to assign tasks and other task types to the group.
Group email	Group email distribution list or the email address of the point of contact, such as the group manager.
Parent	Other group of which this group is a member. If a group has a parent, the child group inherits the roles of the parent group. The members of the child group aren't members of the parent group. For example, if an incident is assigned to the parent group and you select the **Assigned to** lookup icon, only the members in the parent group are available. The members of the child group aren't available.
Active	Check box that indicates whether the group is active or inactive. Inactive groups still appear in any reference field that already references the group, but aren't visible by non-admin users in: ◦ lists of groups ◦ the reference lookup list for reference fields ◦ the auto-complete list of groups displayed when you type into a reference field
Exclude manager	Check box that controls whether the group manager receives email notifications.
Include members	Check box that controls whether the group members receive individual emails when someone sends an email to the **Group Email** address. The only exception to this functionality is for approval notifications, whereby all members of a group receive an approval notification, regardless of the **Include members** selection. See Receive notifications for more information.
Description	Helpful information about the group.

4. **Optional:** Remove or hide a group.

a. **Note:**
Only users in the hidden group are able to see the hidden group when selecting a group in a reference field.

Create a true/false field labeled `Hidden` on the Group form.
For more information on creating fields, see Add and customize a field in a table
The system creates a field called **u_hidden** on the **Users [sys_user_group]** table and enables use of the **Hidden** check box to designate a hidden group.

b. Create a new before query business rule on the [sys_user_group] (Group) table with the following script:

```
if (!gs.hasRole("admin") && !gs.hasRole("groups_admin") &&
 gs.getSession().isInteractive()) {
  var qc = current.addQuery("u_hidden", "!=",
 "true"); //cannot see hidden groups...
  qc.addOrCondition("sys_id",
 "javascript:getMyGroups()"); //...unless in the hidden group
}
```

Note:
For more information, see Before Query business rules.

Add a user to a group

Add a user to a group so that the user inherits all the roles assigned to the group.

Before you begin

Role required: user_admin

About this task

If you're a non-admin user, you can't add a user to a group that contains the admin role. Likewise, if you don't have a security_admin role, you can't add a user to a group that contains the security_admin role.

Procedure

1. Navigate to **All** > **User Administration** > **Groups**.
2. Select a group **Name**.
3. In the **Group Members** related list, select **Edit**.
4. Select one or more names in the **Collection** list.
5. Select **Add** and **Save**.
6. **Optional:** Remove a user from a group when they change roles.

   a. Navigate to **All** > **User Administration** > **Groups**.

   b. Select a group **Name**.

   c. In the **Group Members** related list, select the check box next to each group member name you want to remove.

d. From the **Actions on selected rows** menu, select **Delete**.

> **Note:**
> Before selecting **Delete**, first make sure you have properly selected the rows containing the users you want to remove from the group.

Configure assignment group types

Use the **Type** field to define categories of groups. Once defined, you can use these categories to filter assignment groups based on the group type using a reference qualifier.

For example, when selecting an assignment group from the Incident form, **Type** can be used to filter groups based on whether they're typically involved in the Incident management process. Groups such as Network or Help Desk are displayed as they're typically involved. Groups such as HR or New York are omitted.

The following items are provided in the base system.

- The types **catalog**, **itil**, and **survey**.
- The reference qualifier on [task.assignment_group] filters on **[Type] [equals] [itil]**.
- A reference qualifier named `GetGroupFilter` is available to filter for group types using Create a dynamic filter option.

> **Note:**
> Dictionary overrides enable administrators to filter a group type on an extended table using a simple reference qualifier override.

Add a group type

You can add additional group types to filter assignment groups for tasks.

Before you begin

Role required: admin

About this task

You may need to configure the form to display the **Type** field. For more information see: Personalize a form

Procedure

1. Navigate to **All** > **User Administration** > **Groups**.
2. Select a group record.
3. Select the lock icon beside **Type**.
4. Select the lookup icon beside the selection field.
   The Group Types dialog opens.
5. Complete the following steps.

   a. Select **New**.

   b. Enter the group type name and description.

Example
For example, to define a type for a group as **incident** and **problem**, enter: **incident,problem**.

Example
Select **Submit**.

The Group form reopens with the new type listed.

6. **Optional:** Add additional group types if needed.
7. Select **Update**.

Assign a group type

You can assign group types to filter assignment groups for tasks.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Groups** and select the desired group.
2. Select the lock icon beside **Type**.
3. Select the lookup icon beside the selection field and select one or more group types.

   Note:
   Because the default behavior of `task.assignment_group` is to filter out groups with group types defined, adding a type to a group filters it out of the **Assignment Group** field on tasks. To change the behavior, set up the reference qualifier.

4. Select **Update**.

Managing roles

Admins can create and configure roles that grant specific permissions, which govern what users and groups with that role can do.

All ServiceNow roles

Use this filterable table to review a complete list of roles used in ServiceNow instances. The table includes important details for all base system roles as well as any roles added by installed plugins.

Product

Product	Roles
AWA	AWA manager [awa_manager] AWA integration user [awa_integration_user] AWA agent [awa_agent] AWA administrator [awa_admin]
Document Intelligence	DocIntel Admin [sn_docintel.admin] DocIntel Extraction Agent [sn_docintel.extraction_agent]

Product (continued)

Product	Roles
	DocIntel Viewer [sn_docintel.viewer] DocIntel Manager [sn_docintel.manager] DocIntel Creation Agent [sn_docintel.creation_agent]
Document Services	Administrator [platform_document_management_admin]
Hermes Messaging Service	administrator [hermes_admin] Kafka namespace administrator [kafka_namespace_admin] Kafka administrator [kafka_admin]
IDR	administrator [idr_admin] rest [idr_rest] read [idr_read]
Knowledge Management	KCS contributor [kcs_contributor] Knowledge administrator [knowledge_admin] Knowledge manager [knowledge_manager] KCS publisher [kcs_publisher] KCS candidate [kcs_candidate] Knowledge domain expert [knowledge_domain_expert] Knowledge coach [knowledge_coach] Knowledge group member [knowledge_group_member] Knowledge group manager [knowledge_group_manager] Administrator [admin] Knowledge [knowledge]
MetricBase	administrator [clotho_admin] REST put [clotho_rest_put] REST accumulate [clotho_rest_accumulate] plus writer [mbplus_writer] plus reader [mbplus_reader]

Product (continued)

Product	Roles
Natural Language Understanding	feedback admin [sn_nlu_workbench.nlu_feedback_admin] editor [nlu_editor] administrator [nlu_admin] user [nlu_user]
Predictive Intelligence	ML Labeler [ml_labeler] ML Report User [ml_report_user] ML Admin [ml_admin]
Reporting	Group report user [report_group] Report scheduler [report_scheduler] Report description administrator [report_description_admin] Report administrator [report_admin] Global report user [report_global] Report user [report_user]
Service Catalog	Administrator [admin] Catalog builder editor [catalog_builder_editor] Catalog editor [catalog_editor] Catalog manager [catalog_manager] Catalog administrator [catalog_admin]
Service Portal	administrator [sp_admin] Announcements administrator [announcement_admin] Portal analytics administrator [portal_analytics_admin] Portal analytics viewer [portal_analytics_viewer]
Task Intelligence	Task Intelligence Admin [tia_admin] Task Intelligence User [tia_user] Task Intelligence Analyst [tia_analyst]
Virtual Agent	Chat analytics administrator [chat_analytics_admin]

Product (continued)

Product	Roles
	UI builder administrator [ui_builder_admin] Virtual Agent administrator [virtual_agent_admin] Chat analytics viewer [chat_analytics_viewer] UI builder administrator [ui_builder_admin] Virtual Agent administrator [virtual_agent_admin] Chat analytics administrator [chat_analytics_admin] Chat analytics viewer [chat_analytics_viewer] Issue auto resolution administrator [iar_admin] Automated testing framework designer [atf_test_designer] Administrator [nlu_admin] admin [virtual_agent_admin] [action_designer] External application installation administrator [external_app_install_admin] Self-service manager [sn_ssa_core.self_service_manager] Localization editor [localization_editor] Localization fulfiller [localization_fulfiller] Image administrator [image_admin] Automated Customer Experiences (ACE) user [sn_ace.ace_user] Activity subscriptions administrator [actsub_admin] Conversational Interfaces guided setup administrator [sn_ci_gs_content.ci_guided_setup_admin] Virtual Agent branding [va_branding] Localization requestor [localization_requestor] Virtual Agent branding [va_branding]

Product (continued)

Product	Roles
	Search application administrator [search_application_admin] Chat analytics administrator [chat_analytics_admin] administrator [evam_admin] administrator [evam_admin] Localization requestor [localization_requestor] Activity subscriptions administrator [actsub_admin] [action_designer] Chat survey administrator [chat_survey_admin] Flow designer [flow_designer]

Base system roles

Administrators can assign one or more base system user roles to grant access to base system platform features and applications.

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

Important:
None.

Administrator [admin]

The administrator role. This role has access to all system features, functions, and data because administrators can override access control list (ACL) rules and pass all role checks. Avoid assigning this role to your users when more targeted roles are available.

Contains Roles

List of roles contained within the role.

- ais_admin
- announcement_admin
- catalog
- catalog_admin
- catalog_builder_editor
- catalog_lookup_admin
- catalog_template_editor
- chat_admin
- evam_admin

- image_admin
- import_admin
- import_scheduler
- import_set_loader
- import_transformer
- live_feed_admin
- ml_admin
- ml_labeler
- nlu_admin
- nlu_editor
- nlu_user
- pa_data_collector
- pa_viewer
- personalize_dictionary
- platform_ml_create
- platform_ml_read
- platform_ml_write
- search_application_admin
- search_relevancy_model_admin
- sn_ace.ace_user
- sn_employee.admin
- sn_hr_sp.admin
- sn_hr_sp.esc_admin
- sn_nlu_workbench.nlu_feedback_admin
- sn_templated_snip.template_snippet_admin
- sn_templated_snip.template_snippet_reader
- sn_templated_snip.template_snippet_writer
- sp_admin
- taxonomy_admin
- user_criteria_admin

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

Grant this privilege carefully. If you have sensitive information, such as HR records, that you must protect, create a custom admin role for that area. Train any users authorized to see those records to act as the administrator. Also note the Special Administrative Roles.

Note:
Users with roles related to the Key Management Framework can only be modified by admins with the kmf_admin role. For details on KMF roles, see Roles installed with Key Management Framework.

Agent administrator [agent_admin]

Agent administrators can download and administer the built-in system agent. They can manage MID Server-related scripts.

Contains Roles

List of roles contained within the role.

- agent_security_admin
- view_changer

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

AI search administrator [ais_admin]

AI search administrators can query, create, update, and delete indexing and search settings and log messages through the AI Search application.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Application client company installer [app_client_company_installer]

Users assigned the app_client_company_installer role can install applications containing the same company as the currently logged in instance. Assigning this role enables first-time installation of applications for the company associated with the current instance. Users with this role can't install an application for another company.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Application client user [app_client_user]

Application client users can install applications containing the same company as the currently logged in instance.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Approval administrator [approval_admin]

Approval administrators can view or modify approval requests not directly assigned to them. Use the approver_user role to enable approvers to view or modify only requests directly assigned to them.

Fulfillers may approve within the product to which they are subscribed (ITSM Fulfiller approving within ITSM). This approval may be in the platform or via email. No additional entitlement is required.

Fulfillers may not approve beyond the product to which they are subscribed (ITSM Fulfiller approving within Procurement, GRC, etc.). This approval would need an additional approval entitlement for the user.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Approver users [approver_user]

Approver users can modify requests for approval routed to them. They also have all the capabilities of requesters.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

There's a fee associated with this role. Don't assign it to users without confirming your organization has the appropriate entitlement.

Asset user [asset]

Asset users can manage hardware and software assets.

Contains Roles

List of roles contained within the role.

- inventory_user
- cmdb_query_builder
- canvas_user
- financial_mgmt_user
- cmdb_read
- contract_manager
- category_manager

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Assignment rule administrator [assignment_rule_admin]

Assignment rule administrators can manage assignment rules.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Business process administrator [business_process_admin]

Business process admins can create, read, update, and delete all records and their relationships in the business process.

In the context of Governance, Risk, and Compliance (GRC), users with the sn_grc.admin role who manage GRC applications and their setup automatically gain access to this role. This access enables the GRC administrators to administer a business process and its records similar to other GRC tables.

Contains Roles

List of roles contained within the role: business_process_manager.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

This role is assigned to users who are administrators and have thorough information and training on business processes. None.

Business process manager [business_process_manager]

Business process managers can create, read, and update any business process and manage the relationship of business processes with other records. This role is assigned to business process managers who are usually specialists and manage multiple business processes in the organization. Assign this role to users who generally work with other employees and are experts around business processes.

In the context of GRC, users with the sn_grc.manager role automatically inherit this role that enables them to manage the business processes for the entire organization.

Contains Roles

List of roles contained within the role - business_process_user.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Business process user [business_process_user]

Business process users can update the business processes that a user owns and can also read any business process. This role must be assigned to the respective process owners. This role can also be provided to users who are required to view the business processes in the organization and understand them better.

In the context of GRC, users with the sn_risk.user role are automatically assigned this role. This role enables users to manage the business processes that they own as well as read all business processes.

Contains Roles

List of roles contained within the role- cmdb_read.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Catalog administrator [catalog_admin]

Catalog administrators can manage the Service Catalog application, including catalog categories and items.

Contains Roles

List of roles contained within the role.

- user_criteria_admin
- catalog_builder_editor
- catalog_template_editor
- catalog
- catalog_lookup_admin

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Catalog editor [catalog_editor]

Catalog editors can create, modify, and publish items within categories that they're assigned to.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Catalog item designer [catalog_item_designer]

Catalog item designers can view the status of their category requests. This role is granted automatically to users when they make a request for an item designer category.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Catalog manager [catalog_manager]

Catalog managers can view and assign catalog editors to their categories. They can also create, modify, and publish items within their categories.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Catalog user [catalog]

Catalog users have read and some write access to all Service Catalog Requests, Tasks, and Items.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

- Catalog Request Approvers > $1000
- Catalog Request Approvers for Sales
- Field Services

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Category manager [category_manager]

Category managers can create, edit, and delete model categories.

Contains Roles

List of roles contained within the role - model_manager.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

CMDB administrator [sn_cmdb_admin]

The configuration management data base (CMDB) administrator is a key role required for interacting with the CMDB Workspace. CMDB administrators can access all CMDB data, tools, and UIs within the CMDB Workspace. Users with this role can set policies that an editor can't, such as class manager and app service requirements.

As you drill down in the CMDB Workspace, there are some dashboards and list views that require specific roles in addition to the CMDB Admin, CMDB Editor, or CMDB User roles.

Contains Roles

List of roles contained within the role.

- canvas_admin
- cmdb_ms_admin
- data_manager_admin
- sn_cmdb_editor

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

CMDB de-duplication administrator [cmdb_dedup_admin]

CMDB de-duplication admins can review and remediate CMDB de-duplication tasks.

Contains Roles

List of roles contained within the role - cmdb_read.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

CMDB editor [sn_cmdb_editor]

A key role required for interacting with CMDB Workspace. CMDB editors can create, edit, and delete CMDB records but can't change policies such as data manager, class manager within CMDB Workspace.

As you drill down in the CMDB Workspace, there are some dashboards and list views that require specific roles in addition to the key CMDB Admin, CMDB Editor, or CMDB User roles.

Contains Roles

List of roles contained within the role.

- cmdb_ms_editor
- sn_cmdb_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

CMDB multi-source administrator [cmdb_ms_admin]

The CMDB multi-source administrator can create and run a query and can modify CMDB 360 properties. Contains the cmdb_ms_write role.

Contains Roles

List of roles contained within the role - cmdb_ms_editor.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

CMDB multi-source editor [cmdb_ms_editor]

CMDB multi-source editors can create and query, read, and write CMDB records, but can't perform recomputing actions.

Contains Roles

List of roles contained within the role - cmdb_ms_user.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

CMDB multi-source user [cmdb_ms_user]

CMDB multi-source users have read and execute access to the multi-source queries.

Contains Roles

List of roles contained within the role - cmdb_read.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

CMDB reader [cmdb_read]

CMDB reader users can read data from the CMDB hierarchy.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

CMDB user [sn_cmdb_user]

A key role required for interacting with CMDB Workspace. CMDB users have read-only access to CMDB data and basic UI within CMDB Workspace.

As you drill down in the CMDB Workspace, there are some dashboards and list views that require specific roles in addition to the key CMDB Admin, CMDB Editor, or CMDB User roles.

Contains Roles

List of roles contained within the role.

- app_service_user
- canvas_user
- cmdb_ms_user
- cmdb_query_builder
- data_manager_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Contract manager [contract_manager]

Contract managers can create, edit, and delete contracts through the Contract Management application.

Contains Roles

List of roles contained within the role.

- canvas_user
- financial_mgmt_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

CreateNow unlimited [unlimited_createnow]

Role for CreateNow unlimited licensed users.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Data classification administrator [data_classification_admin]

Data classification administrators manage all aspects of the Data Classification application, data classification code setup, and assignment.

Contains Roles

List of roles contained within the role - data_classification_auditor.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Data classification auditor [data_classification_auditor]

Data classification auditors audit Data Classification code assignments.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Enterprise CMDB administrator [ecmdb_admin]

Enterprise CMDB administrators can perform administrative tasks and access tables and records in Enterprise CMDB.

Contains Roles

List of roles contained within the role - cmdb_read.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Filter administrator [filter_admin]

Filter administrators can create, edit, and delete filter [sys_filter] records.

Contains Roles

List of roles contained within the role.

- filter_global
- filter_group

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Filter group user [filter_group]

Filter group users can create filters that belong to groups of which the user is a member.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Gauge maker [gauge_maker]

Gauge makers can create gauges from reports. Starting with Helsinki, reports are no longer made into gauges.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Global filter user [filter_global]

Global filter users can create global filter [sys_filter] records.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Global template editor [template_editor_global]

Users with the template_editor_global role can create templates for global use.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Group template editor [template_editor_group]

Users with the template_editor_group role can create templates for groups. (Users must also be assigned the template_read_global role)

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Guided tour administrator [guided_tour_admin]

Guided tour administrators can create, modify, and delete guided tour [sys_embedded_tour_guide] records.

Contains Roles

List of roles contained within the role - sn_tourbuilder.tour_admin.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Image administrator [image_admin]

Image administrators can create, modify, and delete image [db_image] records.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Impersonator [impersonator]

Impersonators can impersonate users.

> **Warning:**
> This role doesn't enable the impersonation of admin users.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

For details on impersonation, see Base system roles.

Import administrator [import_admin]

Import administrators can manage all aspects of import set [sys_import_set] records and imports.

Contains Roles

List of roles contained within the role.

- import_set_loader
- import_transformer
- import_scheduler

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Import scheduler [import_scheduler]

Import schedulers can schedule imports.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

> **Warning:**
> Grant this role carefully. The import_scheduler can execute scripts with administrator level privileges.

Import set loader [import_set_loader]

Import set loader users can load import set [sys_import_set] records.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Import transformer [import_transformer]

Import transformer users can manage import set transform map [sys_transform_map] records and run transforms.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Inventory administrator [inventory_admin]

Inventory administrators administer stockrooms, stock models, stock rules.

Contains Roles

List of roles contained within the role.

- inventory_user
- canvas_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Inventory user [inventory_user]

Inventory users have access to stock information.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

ITIL

Information Technology Infrastructure Library (ITIL) users can open, update, close incidents, problems, changes, and configuration management items. This role is the base system technician role. Users with the ITIL role can have tasks assigned to them.

Contains Roles

List of roles contained within the role.

- agent_workspace_user
- app_service_user
- cmdb_query_builder
- cmdb_read
- dependency_views
- email_composer
- snc_platform_rest_api_access
- sn_change_write
- sn_cmdb_editor
- sn_comm_management.comm_plan_viewer
- sn_gd_guidance.guidance_user
- sn_incident_write
- sn_nb_action.next_best_action_user
- sn_problem_write
- sn_request_write
- sn_sow.sow_list
- sn_sow.sow_user
- sn_sttrm_condition_read
- survey_reader
- template_editor
- tracked_file_reader
- view_changer
- viz_creator

Groups

List of groups this role is assigned to by default.

- Field Services
- Catalog Request Approvers > $1000
- Catalog Request Approvers for Sales

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

ITIL administrator [itil_admin]

ITIL administrators can delete incidents, problems, changes, and other related records. This role is intended for team leads.

Contains Roles

List of roles contained within the role.

- sn_cmdb_admin
- assessment_admin
- sn_bm_client.benchmark_data_viewer
- cmdb_read

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

Avoid granting admin roles when more specialized roles are available.

Knowledge [knowledge]

Knowledge users can write, edit, and review knowledge management articles.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Knowledge administrator [knowledge_admin]

Knowledge administrators can manage knowledge bases.

Contains Roles

List of roles contained within the role - knowledge.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

List updater [list_updater]

List updater users can select the **Update Entire List** and **Update Selected** menu options on a list.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Maintenance

This role is reserved for ServiceNow use.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

This role can't be assigned or impersonated, and is reserved for ServiceNow use.

MID server [mid_server]

MID server users can access to the tables that MID servers ordinarily use. This role should be granted to your MID servers.

Contains Roles

List of roles contained within the role - soap.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

This role should be assigned to user accounts created for MID servers to interact with your instance. For details, see Create the MID Server user and grant the role.

Model manager [model_manager]

Model managers can create, modify, and delete base model [cmdb_model] records.

Contains Roles

List of roles contained within the role - catalog_editor.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Personalize [personalize]

Users with the personalize role can personalize forms, lists, rules, controls, and scripts.

Contains Roles

List of roles contained within the role.

- personalize_control
- personalize_rules
- personalize_dictionary
- personalize_choices
- personalize_styles
- personalize_responses
- personalize_list
- personalize_form

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize choices [personalize_choices]

Users assigned to the personalize_choices role can personalize the choices for a list field.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize control [personalize_control]

Personalize control users can personalize controls on lists, such as filters, links, and buttons.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize dictionary [personalize_dictionary]

Users with the personalize_dictionary role can personalize dictionary entries and labels.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize form [personalize_form]

Users with the personalize_form role can personalize forms.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize list [personalize_list]

Users with the personalize_list role can personalize lists.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize responses [personalize_responses]

Users with the personalize_form role can personalize predefined responses for suggestion fields, such as the additional comments field.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize rules [personalize_rules]

Personalize rules users can personalize business rules and scripts. This role contains additional roles for granting selective, administrative access to rules and scripts.

Contains Roles

List of roles contained within the role.

- ui_action_admin
- business_rule_admin
- client_script_admin
- ui_policy_admin

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

Avoid granting this role to users who don't need access to all the roles contained in this role.

Personalize styles [personalize_styles]

Users with the personalize_styles role can personalize field styles.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Personalize UI [personalize_ui]

Users with the personalize_ui role can personalize forms and lists.

Contains Roles

List of roles contained within the role.

- personalize_form
- personalize_list

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Platform Rest API access [snc_platform_rest_api_access]

Allows access to Platform Rest APIs. This role is contained with in the ITIL [itil] role.

- Table API
- Import Set API
- Aggregate API
- Attachment API

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Public [public]

No login is required to access features or functions with the public role.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Release administrator [release_admin]

Release administrators can edit the release history for a release.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Report administrator [report_admin]

Report administrators can manage, share, publish, and schedule all reports. Users assigned this role can access the **Reports** > **Administration** module and manage all report-related objects. The report_admin role inherits all other report roles.

Contains Roles

List of roles contained within the role.

- gauge_maker
- report_alias_admin
- report_global
- report_group
- report_publisher
- report_scheduler
- viz_admin

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Report alias administrator [report_alias_admin]

Report alias administrators can maintain field and value aliases.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Report global [report_global]

Report global users can manage reports that are shared with everyone (listed in Global).

Contains Roles

List of roles contained within the role - report_user.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Report group [report_group]

Report group users can manage and share reports that are shared with them (listed in Group).

Contains Roles

List of roles contained within the role - report_user.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Report publisher [report_publisher]

Report publisher users can publish reports any that they can manage. Publishing a report creates a public link to that report. Users with this role must also have another role that grants permission to create, edit, and share reports.

Contains Roles

List of roles contained within the role - report_user.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Report scheduler [report_scheduler]

Report scheduler users can schedule emailing of all reports that they can see, including reports they can't manage. Users with this role must also have another role that grants permission to create, edit, and share reports.

Contains Roles

List of roles contained within the role - report_user.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Report user [report_user]

Report users can create and view reports that have been shared with them. Users with this role can't share, edit, or delete reports that have been shared with them.

Contains Roles

List of roles contained within the role - viz_creator.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Script fix administrator [script_fix_admin]

Script fix administrators can manage fix scripts.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Search application administrator [search_application_admin]

Search application administrators can insert, update, and delete search user experience-related configuration tables:

- sys_search_context_config
- sys_search_source
- m2m_search_context_config_search_source
- sys_search_facet
- sys_search_filter

Search application admin is granted the ais_admin role to enable AI search configuration.

Contains Roles

List of roles contained within the role.

- ais_admin
- personalize_dictionary

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

SOAP [soap]

users with the soap role can query, create, update, and delete records on all tables, as well as execute scripts.

Contains Roles

List of roles contained within the role.

- soap_create
- soap_delete
- soap_ecc
- soap_query
- soap_script
- soap_update

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP create [soap_create]

Users with the soap_create role can create records in all tables and columns.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP delete [soap_delete]

Users with the soap_delete role can delete records in all tables and columns.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP ECC [soap_ecc]

Users with the soap_ecc role can query, create, and update on the external communication channel (ECC) Queue table only.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP query [soap_query]

Users with the soap_query role can query records on all tables and columns.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP query update [soap_query_update]

Users with the soap_query_update role can query and update all tables and columns.

Contains Roles

List of roles contained within the role.

- soap_query
- soap_update

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP script [soap_script]

Users with the soap_script role can execute business rule endpoint functions via **script.do**.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

SOAP update [soap_update]

Users with the soap_update role can update records on all tables and columns.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Survey administrator [survey_admin]

Survey administrators can see all surveys, definitions, questions, and instances created by them and others. Users with this role can use all modules in the **Survey** application menu.

Contains Roles

List of roles contained within the role.

- assessment_admin
- sn_bm_client.benchmark_data_viewer
- sn_publications_recipients_list_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Survey creator [survey_creator]

Survey creators can manage survey definitions, questions, and instances created by them.

Contains Roles

List of roles contained within the role.

- sn_bm_client.benchmark_data_viewer
- sn_publications_recipients_list_user

Groups

List of groups this role is assigned to by default - Survey Creators.

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Survey reader [survey_reader]

Survey readers can view surveys and related information, such as survey responses, survey groups, scorecards, and reports. Users with this role can't change or modify surveys or survey responses.

Contains Roles

List of roles contained within the role - sn_bm_client.benchmark_data_viewer.

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Task editor [task_editor]

Task editors can edit protected task fields.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Template editor [template_editor]

Template editors can create templates for personal use, and modify or delete personal templates. This role is included in the itil role in the base system.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Template scheduler [template_scheduler]

Template schedulers can schedule template-based record creation.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Text search administrator [text_search_admin]

Text search administrators can customize Global Text Search groups and tables.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Timecard administrator [timecard_admin]

Timecard administrators can access all timecard records.

Contains Roles

List of roles contained within the role.

- timecard_approver
- timecard_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Timecard approver [timecard_approver]

Timecard approvers approve or reject time cards for users.

Contains Roles

List of roles contained within the role.

- pa_viewer
- timecard_user

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Timecard user [timecard_user]

Timecard users can create time cards themselves, and view their own time cards.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

User [user]

The user role has no functionality and doesn't grant access to any assets on your instance. Users with this role are counted as licensed fulfillers.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

User administrator [user_admin]

User administrators can administer users, groups, locations, skills, and companies.

Contains Roles

List of roles contained within the role.

- fsm_skill_admin
- skill_admin
- territory_admin

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

View changer [view_changer]

View changers can switch active views.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Workflow administrator [workflow_admin]

Workflow administrators can create, edit, publish, or delete graphical workflows.

Contains Roles

List of roles contained within the role.

- activity_creator
- itom_admin
- workflow_creator
- workflow_publisher

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren’t assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Workflow creator [workflow_creator]

Workflow creators can create graphical workflows.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Workflow publisher [workflow_publisher]

Workflow creators can publish graphical workflows.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Workflow report viewer [workflow_report_viewer]

Workflow report viewers can access the workflow scratchpad for reports.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None

Zing text search administrator [ts_admin]

Users with the ts_admin role can administer the Zing text indexing and search engine.

Contains Roles

List of roles contained within the role.

None

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevated privilege roles.

No

Special considerations

None.

Special administrative roles

Certain roles grant specific administrative rights without the full privileges of the admin role. For example, an administrator can grant a user the right to change UI policy but not client scripts.

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

These roles don't change the behavior of the admin role, which grants full administrative privileges.

> **Important:**
> You can't rename roles of any kind in the ServiceNow AI Platform

Assignment rule administrator [assignment_rule_admin]

Create, modify, and delete assignment rules.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Business rule administrator [business_rule_admin]

Create, modify, and delete Classic Business rules.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Client script administrator [client_script_admin]

Create, modify, and delete Client scripts.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Form administrator [form_admin]

Create, modify, and delete forms, and form sections and section elements.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Report administrator [report_admin]

Manage, share, publish, and schedule all reports. Users assigned this role can access the **Reports** > **Administration** module and manage all report-related objects. The report_admin role inherits all other report roles.

Contains Roles

List of roles contained within the role.

- gauge_maker
- report_alias_admin
- report_global
- report_group
- report_publisher
- report_scheduler
- viz_admin

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Script include administrator [script_include_admin]

Create, modify, and delete script includes.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

UI action administrator [ui_action_admin]

Create, modify, and delete UI actions.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

UI macro administrator [ui_macro_admin]

Create, modify, and delete UI macros.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

UI page administrator [ui_page_admin]

Create, modify, and delete UI pages.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

UI policy administrator [ui_policy_admin]

Create, modify, and delete UI policies.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

UI script administrator [ui_script_admin]

Create, modify, and delete UI scripts.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Read-only role

The read-only role (snc_read_only) restricts a user or a group of users to read-only access on the tables to which the user already has access.

This role is designed to complement other roles a user possesses. Its purpose is to restrict actions like the insert, update, or delete operations on the tables accessible through their existing roles.

After you assign this role to a user, they can no longer create, update, or delete records on ANY tables.

Note:
Assign this role only to users. Don't assign this role to other resources in the system, including applications, access control levels (ACLs), and so on.

The snc_read_only role can be assigned to any user to limit access to data without having to create ACLs for system tables, custom tables, and fields. This practice is useful for performing internal or external audits without enabling a user to have insert or update access to data.

Users with the snc_read_only role have the following restrictions regardless of other roles and privileges that they have.

- Can't insert, update, or delete records from the UI or when using the GlideRecord API.
- Can't activate or upgrade plugins.
- Can't directly run SQL.
- Can't upload XML files.
- Can only run background scripts when on an instance in the public sandbox environment.

Note:
These role restrictions are in place even if impersonating another user with write access such as an admin.

Activate the read-only role

If it isn't already active, an administrator can activate the Read-Only User Role (com.snc.read_only.role) plugin.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Read-Only User Role (com.snc.read_only.role) plugin using the filter criteria and search bar.

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

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

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

Read-only role properties

These system properties control the snc_read_only role. The following default values are used for the properties.

Name	Description
glide.security.snc_read_only_role.tables.exempt_create	Specifies which tables are exempt from the read-only role enforcement and enable the creation of new records. • Type: string • Default value: sys_user_session, sysevent, syslog, syslog_transaction, sys_user_preference, sys_ui_list, sys_ui_list_element, sys_db_cache, user_multifactor_auth • Location: System Properties [sys_properties] table
glide.security.snc_read_only_role.tables.exempt_write	Specifies which tables are exempt from the read-only role enforcement and enable the updating of existing records. • Type: string • Default value: sys_user_session, sysevent, syslog, syslog_transaction, sys_user_preference, sys_ui_list, sys_ui_list_element, sys_db_cache, user_multifactor_auth • Location: System Properties [sys_properties] table
glide.security.snc_read_only_role.tables.exempt_delete	Specifies which tables are exempt from the read-only role enforcement and enable the deletion of existing records.

(continued)

Name	Description
	• Type: string • Default value: sys_user_preference, sys_ui_list, sys_ui_list_element, sys_db_cache, user_multifactor_auth • Location: System Properties [sys_properties] table

After you configure these properties, assign the read-only role as needed. When you log in, you're restricted from creating, updating, or deleting records on ANY tables unless you modified these properties.

Note:
Test the read-only role by assigning it to a user and then impersonating that user.

Application specific roles

Applications you install on your instance may include additional roles. Follow the links in this section to see roles installed along with applications.

Workflow	Product	Roles Documentation Links
Customer Service Management	Customer Access Management	Roles installed with customer access management
Customer Service Management	Customer Service Management	Roles installed with Customer Service Management
Customer Service Management	Major Issue Management	Components installed with Major Issue Management
Customer Service Management	OpenFrame	Components installed with OpenFrame
Customer Service Management	Public Sector Digital Services	
Customer Service Management	Special Handling Notes	Components installed with Special Handling Notes
Customer Service Management	Targeted Communications	Components installed with Targeted Communications
Employee Service Management	Case and Knowledge Management	Components installed with Case and Knowledge Management

servicenow

Workflow	Product	Roles Documentation Links
Employee Service Management	HR Service Delivery	Setting up your Alumni Service Center
Employee Service Management	HR Service Delivery	Installed with HR Service Management Legacy
Employee Service Management	HR Service Delivery	Reference for Document Templates
Employee Service Management	HR Service Delivery	vReference for Listening Posts
Employee Service Management	Journey Designer	Reference information for Journey designer
Employee Service Management	Journey Designer	Components installed with Journey Accelerator
Employee Service Management	Manager Hub	Components installed with Manager Hub
Employee Service Management	Workplace Case Management	Components installed with Workplace Case Management
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Calendar Synchronization
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Case Management
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Core
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Move Management
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Reservations for Microsoft Outlook Add-in
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Reservation Management

Workflow	Product	Roles Documentation Links
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Space Management
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Space Mapping
Employee Service Management	Workplace Service Delivery	Components installed with Workplace Visitor Management
Employee Service Management	Vendor Management Workspace	Installed with Vendor Management Workspace
Environmental, Social, and Governance Management	Environmental, Social, and Governance Management	Components installed with ESG Management
Governance, Risk, and Compliance	Audit Management	Components installed with Audit Management
Governance, Risk, and Compliance	GRC: Metrics	Components installed with the GRC: Metrics application
Governance, Risk, and Compliance	Operational Resilience	
Governance, Risk, and Compliance	Privacy Management	Roles installed with Privacy Management
Governance, Risk, and Compliance	Risk Management	Components installed with Risk Management
Healthcare and Life Sciences	Healthcare and Life Sciences Service Management Core	Components installed with Healthcare and Life Sciences Service Management Core
Healthcare and Life Sciences	Patient Support Services	Components installed with Patient Support Services
Healthcare and Life Sciences	Pre-Visit Management	Components installed with Pre-Visit Management
IT Asset Management	Model Management	Installed with Model Management

Workflow	Product	Roles Documentation Links
IT Asset Management	Product Catalog	Roles installed with Product Catalog
IT Operations Management	Automation Discovery	Components installed with Automation Discovery
IT Operations Management	Event Management	Components installed with Event Management
ServiceNow AI Platform Capabilities	Advanced Work Assignment	Components installed with Advanced Work Assignment
ServiceNow AI Platform Capabilities	AI Search	Components installed with AI Search
ServiceNow AI Platform Capabilities	Application Portfolio Management	Installed with Application Portfolio Management
ServiceNow AI Platform Capabilities	Automation Center	Components installed with Automation Center
ServiceNow AI Platform Capabilities	Client Software Distribution 2.0	Components installed with
ServiceNow AI Platform Capabilities	Conversational Analytics	Conversational Analytics dashboard roles
ServiceNow AI Platform Capabilities	Document Intelligence	Document Intelligence roles
ServiceNow AI Platform Capabilities	Interaction Management	Components installed with Interaction Management
ServiceNow AI Platform Capabilities	Live Feed	User roles installed with Live Feed
ServiceNow AI Platform Capabilities	Notify	Roles installed with Notify
ServiceNow AI Platform Capabilities	Orchestration (Legacy)	Installed with Orchestration ROI

Workflow	Product	Roles Documentation Links
ServiceNow AI Platform Capabilities	Build Workflows	Installed with catalog item designer
ServiceNow AI Platform Capabilities	Build Workflows	
ServiceNow AI Platform Capabilities	Build Workflows	User access to Flow Designer
ServiceNow AI Platform Capabilities	Build Workflows	Installed with catalog item designer
ServiceNow AI Platform Capabilities	Service Bridge	Personas and roles for Service Exchange (legacy)
ServiceNow AI Platform Capabilities	Service Bridge	Personas and roles for Service Exchange for Providers (legacy)
ServiceNow AI Platform Capabilities	State Model	Installed with State Model
Platform Analytics	Automation Discovery	Components installed with Automation Discovery
Platform Analytics	Performance Analytics	Performance Analytics roles
Platform Analytics	Process Mining	Components installed with Process Mining
Platform Analytics	Reporting	Reporting roles
Platform Analytics	User Experience Analytics	Components installed with Usage Insights
IT Service Management	Asset Management	Installed with Asset Management
IT Service Management	Change Management	Components installed with ITSM Roles - Change Management
IT Service Management	Continual Improvement Management	Continual Improvement Management roles

Workflow	Product	Roles Documentation Links
IT Service Management	Contract Management	User roles installed with Contract Management
IT Service Management	Digital Portfolio Management	Roles for Digital Portfolio Management
IT Service Management	Employee Center	Employee Center users and roles
IT Service Management	Expense Line	Components installed with Expense Line
IT Service Management	Facilities Move Management	Roles installed with Facilities Move Management
IT Service Management	Facilities Service Management	Roles installed with Facilities Service Management
IT Service Management	Incident Management	Components installed with ITSM Roles — Incident Management
IT Service Management	ITSM Predictive Intelligence Workbench	Components installed with ITSM Predictive Intelligence Workbench
IT Service Management	Knowledge Management	Knowledge Management roles
IT Service Management	Performance Analytics	Performance Analytics roles
IT Service Management	Planned Maintenance	Installed with SM Planned Maintenance
IT Service Management	Problem Management	Components installed with Problem Management
IT Service Management	Procurement	User roles installed with Procurement
IT Service Management	Request Management	Request ITSM Roles - Request Management
IT Service Management	Service Builder	Roles installed with Service Builder

Workflow	Product	Roles Documentation Links
IT Service Management	Service Operations Workspace	Components installed with Service Operations Workspace ITSM Applications
IT Service Management	Service Operations Workspace	Components installed with Agent Client Collector for Investigation
IT Service Management	Service Operations Workspace	Components installed with Microsoft Endpoint Configuration Manager for Investigation
IT Service Management	Service Operations Workspace	Components installed with Remedial Actions Framework
IT Service Management	Service Operations Workspace	Components installed with Metrics and CI Actions Framework
IT Service Management	Service Level Management	Installed with Service Level Management
IT Service Management	Service Management Core	Installed with Service Management Core
IT Service Management	Service Portfolio Management	
IT Service Management	Vendor Management Workspace	Installed with Vendor Management Workspace
IT Service Management	Walk-up Experience	Installed with Walk-up Experience
IT Service Management	Workforce Optimization	Components installed with Process Mining
Security Operations	Security Incident Response	Components installed with Security Incident Response
Security Operations	Threat Intelligence Sharing	
Security Operations	Vulnerability Response	Components installed with Vulnerability Response

Workflow	Product	Roles Documentation Links
Strategic Portfolio Management	Data Separation	Components installed with Data Separation
Strategic Portfolio Management	Enterprise Release Management	
Strategic Portfolio Management	Field Service Management	Components installed with Field Service Management
Strategic Portfolio Management	Field Service Management	Field Service with Service Locations Support components
Strategic Portfolio Management	Financial Management	User roles installed with Financial Management
Strategic Portfolio Management	Investment Funding	Components installed with Investment Funding
Strategic Portfolio Management	Project Portfolio Management	Components installed with Project Portfolio Management (PPM) Standard
Strategic Portfolio Management	Scrum Programs	Components installed with Scrum Programs
Strategic Portfolio Management	Teamspaces	Installed with teamspaces

System roles

Administrators can control access to features and capabilities on a ServiceNow instance by assigning roles to users.

Your ServiceNow includes roles to grant access to the platform features and applications included a base system instance. Applications you install on your instance may include additional roles to control access to those installed features. For more information about roles, see Exploring user administration.

Important:
Base system and roles installed with applications can be deactivated by administrators, but can't modify or rename these roles.

Base system roles

Base system roles are present in all ServiceNow instances and don't require the installation of additional plugins.

Role	Description
admin	The system administrator role. This role has access to all system features, functions, and data because administrators can override access control list (ACL) rules and pass all role checks. Avoid assigning this role to your users when more targeted roles are available. **Warning:** Grant this privilege carefully. If you have sensitive information, such as HR records, that you need to protect, you must create a custom admin role for that area. You must also train any users authorized to see those records to act as the administrator. Also note the Special Administrative Roles.
agent_admin	Agent administrators can download and administer the system's built-in agent. They can manage MID Server-related scripts.
ais_admin	AI search administrators can query, create, update, and delete indexing and search settings and log messages through the AI Search application.
approval_admin	Approval administrators can view or modify approval requests not directly assigned to them. Use the approver_user role to enable approvers to only view or modify requests directly assigned to them. Use of this role requires a Fulfiller license. Use of the approver_user role requires an Approver license.
approver_user	Approver users can modify requests for approval routed to them. They also have all capabilities of requesters. **Note:** There's a fee associated with this role. Don't assign it to users without confirming your organization has the appropriate entitlement.
assignment_rule_admin	Assignment rule administrators can manage assignment rules.
asset	Asset users can manage hardware and software assets.
business_process_admin	Business process admins can create, read, update, and delete all records and their relationships in the business process. In the context of Governance, Risk, and Compliance (GRC), users with the sn_grc.admin role who manage GRC applications and their setup automatically gain access to this role. This access enables the GRC administrators to administer a business process and its records similar to other GRC tables. **Important:** This role is assigned to users who are administrators and have thorough information and training on business processes.
business_process_manager	Business process managers can create, read, and update any business process and manage the relationship of business processes with other records. This role is assigned to business process managers

servicenow

Role	Description
	who are usually specialists and manage multiple business processes in the organization. These users generally work with other employees and are experts around business processes. In the context of GRC, users with the sn_grc.manager role automatically inherit this role that enables them to manage the business processes for the entire organization.
business_process_user	Business process users can update the business processes that a user owns and can also read any business process. This role must be assigned to the respective process owners who manage the business process that they own. This role can also be provided to users who are required to view the business processes in the organization and understand them better. In the context of GRC, users with the sn_risk.user role are automatically assigned this role as this role enables them to manage the business processes they own as well as read all business processes.
catalog	Catalog users can access service catalog requests.
catalog_admin	Catalog administrators can manage the Service Catalog application, including catalog categories and items.
catalog_editor	Catalog editors can create, modify, and publish items within categories that they’re assigned to.
catalog_item_designer	Catalog item designers can view the status of their category requests. This role is granted automatically to users when they make a request for an item designer category.
catalog_manager	Catalog managers can view and assign catalog editors to their categories. Can also create, modify, and publish items within their categories.
category_manager	Category managers can create, edit, and delete model categories.
cmdb_dedup_admin	CMDB de-duplication admins can review and remediate CMDB de-duplication tasks.
cmdb_ms_user	CMDB multiscource readers can access and run a multi-source CMDB query, but can't create a query. This role contains Contains cmdb_read role.
cmdb_ms_editor	Can create and run a query, has full read and write access, but can't do Recompute. Contains cmdb_ms_read role.
cmdb_ms_admin	Can create and run a query, and can modify CMDB 360 properties. Contains cmdb_ms_write role.
cmdb_read	Can read any CMDB table. Contained in admin and itil.
communication_manager	Manages communication for major incidents and is responsible for communicating with all stakeholders.
contract_manager	Can create, edit, and delete contracts through the Contract Management application.

servicenow

Role	Description
data_classification_admin	Administers all aspects of the Data Classification application, data classification code setup and assignment,
data_classification_auditor	Audits Data Classification code assignments.
ecmdb_admin	Can administer the CMDB.
filter_admin	Can manage filters.
filter_global	Can create global filters.
filter_group	Can create filters that belong to groups of which the user is a member.
gauge_maker	Can create gauges from reports. Starting with Helsinki, reports are no longer made into gauges.
guided_tour_admin	Can manage and administer Guided Tour functionality.
image_admin	Can manage image files on the Images [db_image] table.
impersonator	Can impersonate users. This role doesn't allow impersonation of admin users.
import_admin	Can manage all aspects of import sets and imports.
import_scheduler	Can schedule imports. **Warning:** Grant this role carefully. The import_scheduler role is equivalent to giving the user the admin role, because the import_scheduler has the ability to execute scripts with administrator level privileges.
import_set_loader	Can load import sets.
import_transformer	Can manage import set transform maps and run transforms.
incident_manager	Manages Incident properties and Major Incident trigger rules.
inventory_admin	Can create and delete stock information. Only users with the inventory_admin role can edit stock rules, stockrooms, and stockroom types.
inventory_user	Has access to stock information. Can create and manage transfer orders.
itil	Can perform standard actions for an ITIL helpdesk technician. Can open, update, close incidents, problems, changes, configuration management items. By default, only users with the itil role can have tasks assigned to them.
itil_admin	Possesses more privileges than the itil role and is intended for team leads. This role has the ability to delete incidents, problems, changes, and other related entities when both the itil and itil_admin roles are assigned. In addition, the itil_admin role grants full control of the CMDB. The itil_admin role includes all of the permissions granted to the sn_cmdb_admin role, which provides full access to CMDB data, tools, and UIs.
knowledge	Can create, edit, and review knowledge base articles.

Role	Description
knowledge_admin	Can manage the knowledge base.
list_updater	Can use Update Entire List and Update Selected menu options on lists.
maint	Reserved for ServiceNow use.
mid_server	Role that any MID server user should be granted. This role gives the MID server access to the tables it ordinarily uses.
model_manager	Can create CMDB models. Model manager can control the base models and any model extensions that aren't software or consumables. Consumable models are controlled by the asset manager role (asset). Software models are control by the software asset manager role (SAM).
major_incident_manager	Initiates the major incident process by assessing and approving major incident candidates or creating a major incident. Maintains the ownership and accountability for the life cycle of the incident. Identifies the users and groups to be involved in the resolution activities and sets up communication channels.
nobody	The nobody role means that no user has access - not even admin or maint. Use the nobody role carefully. The nobody role takes precedence over the admin override option on ACLs, so even admins can't have access. See Create an ACL rule. Don't assign it to specific users. You can use this role in ACLs that control access to resources, such as UI pages, processors, script includes, and records. **Warning:** Applying the nobody role may be irreversible if applied to some important system functions.
personalize	Can configure forms, lists, rules, controls, scripts.
personalize_choices	Can configure choices and predefined responses for non-journal fields designated as choice or suggestion fields.
personalize_control	Can configure controls on lists, such as filters, links, and buttons.
personalize_dictionary	Can configure dictionary entries and labels.
personalize_form	Can configure forms.
personalize_list	Can configure lists and list calculations.
personalize_responses	Can configure predefined responses for journal fields designated as suggestion fields.
personalize_rules	Can configure business rules and scripts. This role contains the following specialized roles for granting selective, administrative access to rules and scripts: • business_rule_admin • client_script_admin • ui_policy_admin • ui_action_admin

Role	Description
personalize_styles	Can configure field styles.
personalize_ui	Can configure forms and lists.
public	No login is required to access features or functions with the public role.
release_admin	Can edit Release history for a release.
report_admin	Can manage reports.
report_global	Can create global reports.
report_group	Can create reports and share reports with groups that the user is a member of. Users with this role can edit reports shared by other users in the group.
report_publisher	Can make reports available on a public page.
report_scheduler	Can schedule a report to be emailed.
script_fix_admin	Can create and manage fix scripts but can’t run fix scripts.
search_application_admin	Can query, create, update, and delete records on search UX-related tables. Contains the ais_admin role.
sn_appclient.app_client_company_installer	Can install applications containing the same company as the currently logged in instance. User role that enables first-time installation of applications for the company associated with the current instance. A user with this role can’t install an application for another company.
sn_appclient.app_client_user	Can install applications containing the same company as the currently logged in instance.
sn_cmdb_admin	Provides full access to CMDB data, tools, and UIs. A CMDB Admin, for example, sets policies in the CI Class Manager and application service requirements. CMDB Admin provides the highest level of access to the CMDB.
sn_cmdb_editor	Provides access to CMDB records. A CMDB Editor can't change policies such as in the CMDB Data Manager or in the CI Class Manager.
sn_cmdb_user	Provides read-only access to CMDB data and to basic UIs such as CMDB reports and dashboards.
soap	Can query, create, update, and delete records on all tables, as well as execute scripts.
soap_create	Can create records on all tables and columns.
soap_delete	Can delete records on all tables and columns.
soap_ecc	Can query, create, and update on the ECC Queue table only.
soap_query	Can query records on all tables and columns.
soap_query_update	Can query and update records on all tables and columns.
soap_script	Can execute business rule endpoint function via script.do.
soap_update	Can update records on all tables and columns.

Role	Description
survey_admin	Can see all Surveys, their definitions, questions, instances created by them and others. Survey administrators can use all modules in the Survey application menu.
survey_reader	Users with survey reader role can view surveys and related information, such as survey responses, survey groups, scorecards, and reports. Survey_reader can't change or modify a survey or survey responses.
task_editor	Can edit protected task fields.
template_editor	Can create templates for personal use, and modify or delete personal templates. Included in the itil role in the base system.
template_editor_global	Can create templates for global use.
template_editor_group	Can create templates for groups.
template_scheduler	Can schedule template-based record creation.
text_search_admin	Can customize Global Text Search groups and tables.
timecard_admin	Can approve, modify, and delete the time cards of other users.
ts_admin	Can administer Zing text indexing and search engine.
unlimited_createnow	Role for CreateNow unlimited licensed users.
upgrade_app	Can upgrade installed applications containing the same company as the currently logged in instance. Can't perform first-time installations of applications published to the Application Client page.
user	Available for customer use, has no function in the base system.
user_admin	Can administer users, groups, locations, and companies.
view_changer	Can switch active views.
workflow_admin	Can create, edit, publish, or delete graphical workflows.
workflow_creator	Can create new graphical workflows.
workflow_publisher	Can publish graphical workflows.

Application-specific roles

Applications you install on your instance may include additional roles. Follow the links in this section to see roles documentation on roles installed along with applications.

Product	Application
Platform Capabilities	Advanced Work Assignment

Create a role

Create a role to control access to features and capabilities in applications and modules. The new role doesn't have access to any application or module until you add other roles to it, or add it to the appropriate applications and modules.

Before you begin

Role required: admin

About this task

After access has been granted to a role, all groups or users assigned to the role are granted the access. Roles can contain other roles, and any access granted to a role is granted to any role that contains it.

For a complete list of the roles included with the base instance, see Base system roles

Note:

You can't rename roles of any kind in the ServiceNow AI Platform. If you manually create a role, you can't rename it after you save it.

Move this action to the background so you aren't waiting when adding a number of group members by adding the system property `glide.ui.schedule_slushbucket_save_for_group_roles`. Set the system property to **true**. The system user is used to create records or update existing ones since the action is running in the background.

Procedure

1. Navigate to **All** > **User Administration** > **Roles** and create a record.
2. Complete the form.

Field	Description
Name	Enter a name for the role.
Requires Subscription	Indicates whether users with this role require a subscription to stay in compliance. **Yes** Measured-role that requires a subscription. You can allocate subscriptions to users with this role by adding one or more groups to a product subscription. To learn more, see Managing per-user subscriptions in Subscription Management. **No** Not a measured-role. **Unspecified** Neither
Application	Select the application that contains this record.
Elevated privilege	A role that requires elevated privilege can't be assigned to a user at login by the system. Instead, a user must manually elevate privileges to receive the elevated role. Select this option to mark this role as required to elevate it to high security. See Elevated privilege roles for more information.
Description	Select the roles to delegate to the group member.

Grant a role access to applications and modules

Roles control access to features and capabilities in applications and modules. You add a role to an application or module to enable the role to grant access to the application or module for all users with the role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Applications** or **System Definition** > **Modules**.
2. Select the appropriate application or module to open it in form view.
3. Select the lock to open the **Roles** field.
4. Add the desired roles to the application or module.
5. Select the lock to close the **Roles** field, and then save your changes.

Assign a role to a group

You can assign a role to a group to grant access to applications and modules to group members.

Before you begin

Role required: user_admin or admin

About this task

When you assign roles to groups rather than to individual users, members of the group inherit the role. When a user switches groups, the new group role is assigned automatically. For information about the Service Mapping roles, see Control user access to application services.

Procedure

1. Navigate to **All** > **User Administration** > **Groups**.
2. Select the group to assign a role.
3. In the **Roles** related list, select **Edit**.
4. Add the desired roles to the group.
5. Select **Save**.

Assign a role to a user

A user inherits roles from all groups to which they belong. You can also assign roles directly to a user. Whenever a user is assigned a new role, it only takes effect after logging in with a new session.

Before you begin

Role required: user_admin or admin
When possible, simplify user administration by assigning roles to groups. Create groups that contain all the roles necessary for specific personas, and then assign users to those groups.

About this task

To grant the admin role to a user, you must also have the admin role. To grant the security_admin role to a user, you must also have the security_admin role. You must elevate to the security_admin role before granting the security_admin role to other users. See Elevate to a privileged role.

> **Note:**
>
> The System administrator(admin) role provides access to all system features, functions, and data, regardless of security constraints. Avoid assigning this role to your users when more targeted roles are available.

You can't delete roles that are assigned to the group from the user record. You must remove the user from the group record.

Procedure

1. Navigate to **All** > **User Administration** > **Users** and then open a user record.
2. In the **Roles** related list, select **Edit**.
3. In the **Collection** list, select the desired roles, and then select **Add**.
4. Select **Save**.

What to do next

Note:
If the user is logged in when you update their roles and they're unable to access records enabled by the new role, they may need to log out and back in again.

Add a role to an existing role

When you add a new role to an existing role for a user, the user inherits the access that is granted by the new role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Roles** and open the existing role.
2. Click **Edit** in the Contains Roles related list.
3. Use the slushbucket to add one or more roles to the existing role.
4. Click **Save**.

   The users with the existing role inherit the access that is granted by the new role.

Create a group role

Create a group role to control access to features and capabilities in applications for all members in a group.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **User Administration** > **Group Roles**.
2. Select **New**.
3. Fill in the fields on the form and then select **Submit**.

Field	Description
Group	Select a group.
Role	Select the role to apply to the group.

Field	Description
Inherits	Select this option to have all members of the group inherit the role. This option is selected by default.

Note:
To move this action to the background so you aren't waiting when adding a number of group members, add the system property *`glide.ui.schedule_slushbucket_save_for_group_roles`* and set it to true.

Audit user roles

Changes to user roles are automatically tracked in the Audit Roles [sys_audit_role] table.

Before you begin

Role required: admin

Note:
If the Prevent duplicate entries with Contextual Security: Role Management V2 plugin is installed, you must Enable role auditing with Contextual Security: Role Management V2.

Procedure

Navigate to **All**, and then enter `sys_audit_role.list` in the filter.
The Audit Roles [sys_audit_role] table displays changes to user roles.

Field	Description
Changed by	The user who made the change.
Count after change	Direct role added as a result of the change (if any), plus the number of inherited roles added.
Granted by group	If the role was inherited, the group that the role was inherited from.
Operation	The type of change. Values include: • Added • Removed
Role	The affected role.
User	The affected user.

Delegating roles

Administrators can grant users the ability to assign roles within groups. However, these users can only assign roles that they already have.

Role delegators are users responsible for assigning roles to other users within a group. This role can be assigned manually by granting the role_delegator role to the user, or through the **Group Manager Change** business rule.

The **Group Manager Change** business rule, which is disabled by default, automatically grants the role_delegator role to a user who is designated manager of a group in the **Manager** field on the **Group** form. The role is removed when the user is no longer the manager of the group.

Activate the business rule to take advantage of it. For more information on **Business rules**, see Classic Business rules.

Reviewing Roles and Role Changes

To view roles and role changes:

- Individual users: Navigate to **User administration** > **Users** select a user, and review the **Roles** related list.
- Role delegators: Review who can give roles in which groups by navigating to **User administration** > **Role delegators**.
- For all role changes: To see a history of role changes, navigate to **System security** > **Reports** > **Role audit**. See Audit user roles for more information.

Designate role delegators

Designate role delegators to assign roles to users who are in a particular group.

Before you begin

Role required: admin

About this task

Delegators can assign roles that they inherit from a group. They can also assign roles that an administrator assigned to them.

Procedure

1. Navigate to **All** > **User Administration** > **Designate Role Delegator**.
2. Select the group that includes the user who you want to be the role delegator.
3. Select the user.
4. Select **Submit**.
   A change request for the role delegator request is created and automatically approved.

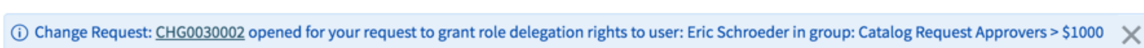

Assign roles as a role delegator

If you're a role delegator, you can delegate roles that are assigned to you for groups that you manage.

Before you begin

Role required: role_delegator

Procedure

1. Navigate to **User Administration** > **Delegate Roles in Group**.
2. Fill out the form.

Field	Description
Group	Select the group in which to delegate a role or roles to a member. Any group can be selected, including groups that the role_delegator doesn’t belong to or groups that the role_delegator doesn’t manage.
User	Select the group member to delegate a role or roles.
Roles to delegate	Select the roles to delegate to the group member. The roles available for delegating are only the roles that the role_delegator has.

(Optional) To remove a delegated role from a user, open the delegation record and remove the unwanted role or roles.

Prevent a role from being delegated

You can prevent roles from being delegated to users.

Before you begin

Role required: admin

About this task

By default, the following roles can’t be delegated.

- admin
- public
- role_delegator

Note:
A user with the role_delegator role can’t delegate this role to other group members.

Procedure

1. Navigate to **All** > **User Administration** > **Roles**.
2. Open the role.
3. Configure the form to add the **Grantable** or **Can delegate** fields.
4. Clear the check box for one or both of these fields.
5. Select **Update**.

Grant a time-limited user role

Learn how to assign a role to a user temporarily. Use this feature if you have a user who needs to perform a one-time action that is normally outside their roles.

Before you begin

Role required: admin

About this task

Warning:
Usage of Time-limited User Roles may consume access and use rights procured by customer and could result in additional subscription fees.

Procedure

1. Navigate to **All** > **User Administration** > **Time-Limited User Roles**.
2. Select **New**.
3. Fill in the fields on the form.

   All fields except Comments are mandatory.

Fields	Description
Active	Determines if the time limited role assignment is active.
Role	The role assigned to the user, only admin, impersonate and snc_readonly is allowed
User	The user to be assigned the time limited role
Start Time	The start time and date for the time limited role
End Time	The end time and date for the time limited role **Important:** The end date should not be more than 5 days after the start date.
Comments (Optional)	Additional comments or information for the limited role assignment

4. Select **Submit**.

Result

The specified user now has the role. They must complete their otherwise restricted task between the start time and the end time.

Note:
Time-limited roles are not displayed to the assigned users.

Monitoring instance usage

Track application usage on your instance through the Application Usage Overview and the ServiceNow Store Usage overview dashboards.

The usage analytics process collects data on all your instances and regularly updates the reports in the Usage Overview and ServiceNow Store Usage Overview dashboards. Application usage data is collected whenever an application is opened, and counts on tables are collected once a day. Data is collected on:

- The number of active users in the system
- The hardware CIs discovered (for instances that use Discovery)
- The number of cloud management service catalog items available to users in instances that use Cloud Management

Users with the admin or usage_admin role can view the Application Usage Overview and ServiceNow Store Usage Overview dashboards.

Application usage overview dashboard

The Application Usage Overview dashboard displays reports on the use of ServiceNow applications on your instance.

Navigate to **Self-Service** > **Dashboards**. From the Dashboards list, select **Usage Overview**.

To manage dashboards, create tabs, reset filters, or create a favorite, select the menu icon (). To see dashboards that you own or are shared with you, select the dashboards overview icon (). To add a widget, select the add widgets icon (). To change the dashboard layout or background color, select the configuration icon (). To share the dashboard with another admin, select the sharing icon ().

Application Usage Overview reports

Chart or report	Description
Usage of ServiceNow Applications	Shows the number of views or operations performed for each of the listed applications, grouped by month.
Users of ServiceNow Applications	Shows the number of active users who have used the instance, grouped by month. An active user is any user who could have accessed the applications, not only users who actually did access the applications.
Usage of Custom Applications	Shows the number of views or operations performed for each of the listed custom applications, grouped by month.
Users of Custom Applications	Shows the number of active users who have used custom applications, grouped by month. An active user is any user who could have accessed the custom applications, not only users who actually did access the applications.
`Google` **Maps Usage By Month**	Shows the number of views or operations performed for the Google maps application.
Additional Metrics	Shows any of the following items, depending on your licensing agreement: • Primary database size of this instance in megabytes • ServiceNow® Performance Analytics of all indices • NotifyNow conference calls • ServiceNow® Cloud Provisioning plugins active • ServiceNow® Discovery plugin active • ServiceNow® Password Reset plugin active • ServiceNow® Cloud Management — VMware instance • Vendor custom applications

Application Usage Overview reports (continued)

Chart or report	Description
	• Custom applications • ServiceNow® Event Management plugin active • Active benchmark indicators • Customer updated PA indicator • Active users • ServiceNow® Cloud Management — EC2 instances • ServiceNow® Orchestration core plugins active • Custom tables extending base system tables other than Task • Virtual node count • Active users for subscriptions • NotifyNow Messages • ServiceNow® Edge Encryption plugin active • Number of Google map page hits during the current month - as of yesterday • Physical node count

Application Usage Overview dashboard

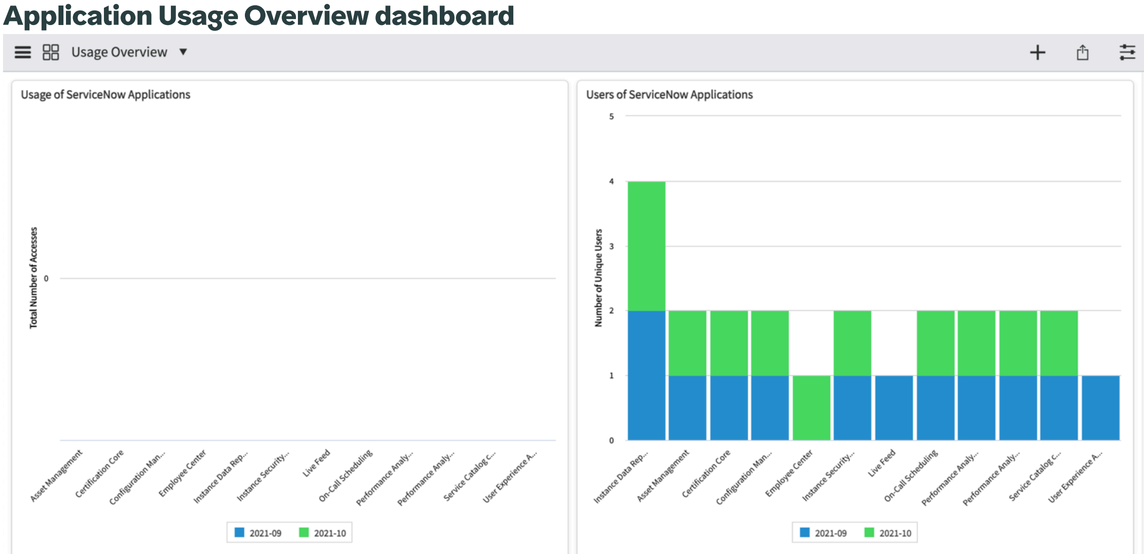

ServiceNow Store Usage overview dashboard

The ServiceNow Store Usage Overview dashboard displays reports on the use of ServiceNow Store applications on your instance.

Navigate to **Self-Service** > **Dashboards**. From the Dashboards list, select **ServiceNow Store Usage Overview**.

To manage dashboards, create tabs, reset filters, or create a favorite, select the menu icon ( ). To see dashboards that you own or are shared with you, select the dashboards overview icon ( ). To add a widget, select the add widgets icon ( ). To change the dashboard layout

or background color, select the configuration icon ( ). To share the dashboard with another admin, select the sharing icon ( ).

ServiceNow Store Usage Overview reports

Chart or report	Description
Usage of ServiceNow Store Applications	Shows the number of views or operations performed for each of the listed applications, grouped by month.
Users of ServiceNow Store Applications	Shows the number of active users who have used the instance, grouped by month. An active user is any user who could have accessed the applications, not only users who actually did access the applications.

ServiceNow Store Usage Overview dashboard

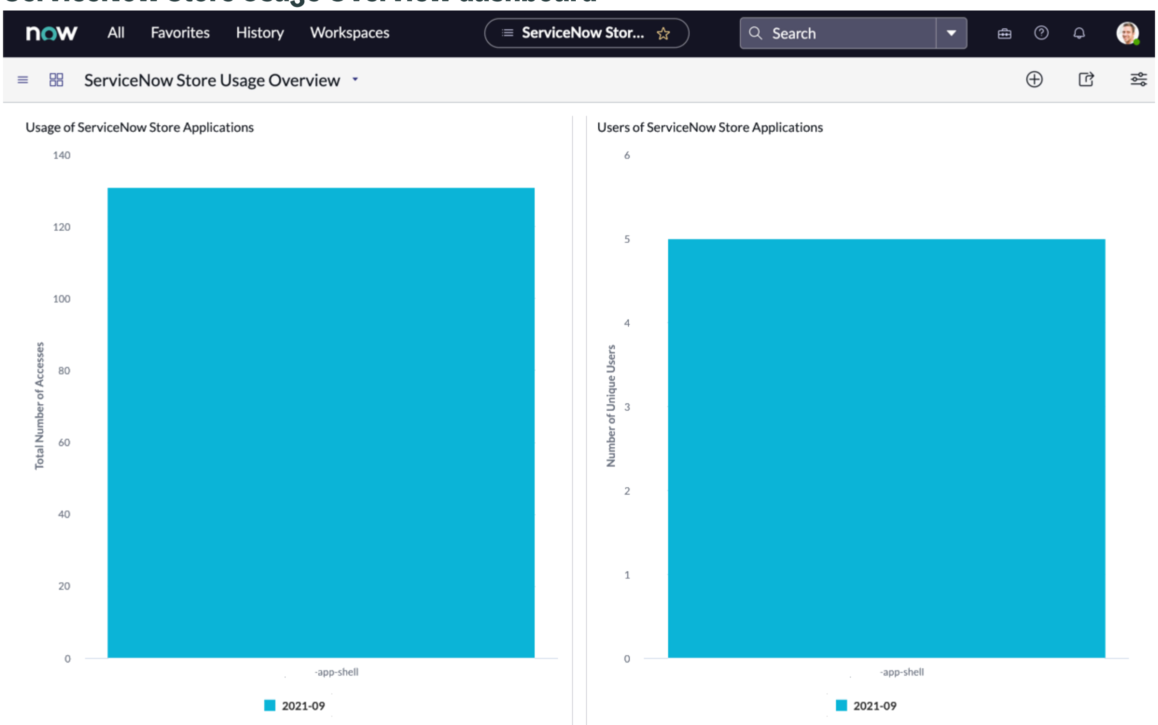

Monitoring user activity

Learn about the tools available to administrators for monitoring users on their instance.

User administration tools

In addition to storing information on a user, user accounts are also associated to groups and roles. Your instance uses groups and roles to determine which records and features an individual may access.

Impersonation

Administrators can select user records for impersonation. Use this feature to experience the instance as another user, with that user's preferences and permissions. User impersonation can be a valuable tool for testing and troubleshooting. For more information on impersonation, see Impersonating users.

Investigating user account activity

At any time there is a need to review specific user behavior, below are the recommended steps on how to review the transaction logs and event logs:

- Locate the IP address of successful/failed login for a particular ServiceNow user for their instance
- Modify the time frame of the search
- Limiting the scope of the search by user name
- View successful/failed login attempts

For more information about investigating user activity, see the How to investigate user account activity [KB0564981] article in the Now Support Knowledge Base.

User sessions

With user session management, and administrator can view and terminate individual user sessions, lock out users from the instance, and make users inactive. For details on managing user sessions, see Managing user sessions.

User records are also associated with transaction logs. Administrators can use these logs to track all browser activity for an instance.

Impersonating users

Administrators are able to impersonate other authenticated users, a feature primarily used for testing.

This function enables the administrator to access the system exactly as the impersonated user, including identical menus and modules. All actions performed by the administrator during impersonation are recorded as if they were executed by the impersonated user.

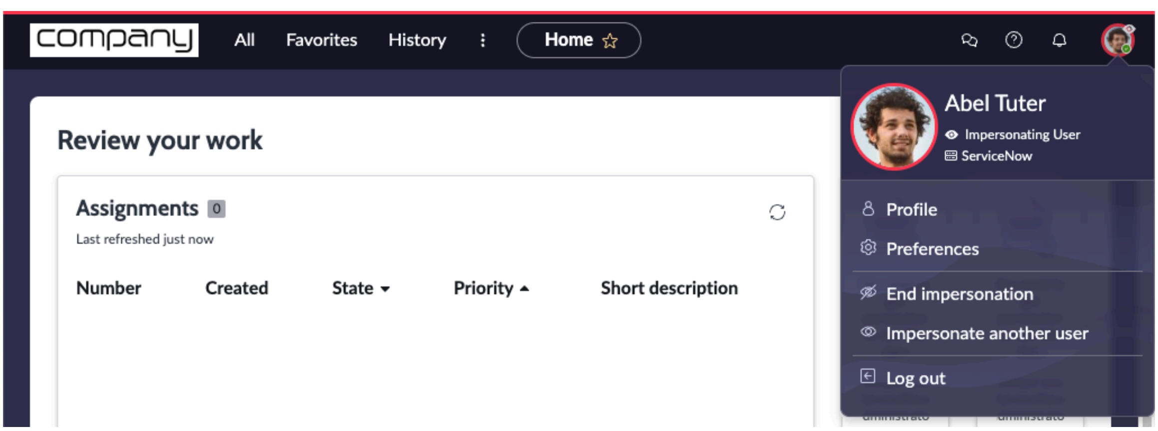

Impersonation limitations

When you impersonate a user, all scope-protected roles and encryption module roles are supported if the **Impersonation** option is configured in the module access policy. See Create a module access policy for details.

Impersonating a user enables access to scope-protected and encryption roles, as defined in the access policy. However, if impersonating a user with an admin role, access to certain features and modules is limited unless the impersonator already possesses those roles.

Impersonating a user with an application-specific admin role, like Human Resources admin or Security Incident Response, limits access to certain features such as security incidents and

profile information, unless these roles are already assigned to the impersonating admin. This restriction extends to certain modules and applications in the navigation bar, and admins can't change the password of users with application admin roles.

The following actions or conditions cause a user impersonation to end:

- The user impersonates a different user
- The user session ends, for example after a user logs out of their instance

Note:
When an administrator starts impersonating a user, the 'Impersonate Begin' event is logged in the system log. Similarly, the 'Impersonate End' event is recorded when impersonation concludes under one of the two conditions listed above.

Impersonation requirements

The user account to be impersonated must have a user ID. You can find this ID in the User [sys_user] record for the account. If this value is missing, the message `The user you selected could not be impersonated` appears.

You need several different accounts to test the system.

- An admin account to do work
- An information technology infrastructure library (ITIL), or similar, account to test as a technician
- An ESS account to test as an end user

More logins may be required to adequately test the system.

Note:
If you try to impersonate a user who is either locked out or inactive, the system will automatically log you out if you initiate an action or select a link. Remember that all changes made during impersonation only apply to that session. To help ensure accuracy, log out and then log back in after completing the impersonation.

Mobile impersonation

Mobile impersonation is available on ServiceNow mobile apps. For information on mobile impersonations, see Mobile impersonation.

Manage the visibility of the impersonation feature

Before users can impersonate another user, an administrator must make the feature visible.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System properties** > **UI properties**.
2. Select a check box under **Enable impersonation button in the banner line** to display or conceal the button.
3. Select **save** at the top of the page.
   Admins can still impersonate users via the **impersonate_dialog** UI page.

Impersonate a user

You can select a user or enter a different user name to perform an impersonation.

Before you begin

Role required: impersonator

Procedure

1. In the banner frame, select your user profile picture to open the user menu.
2. Select **Impersonate User**.
    - In Next Experience:

      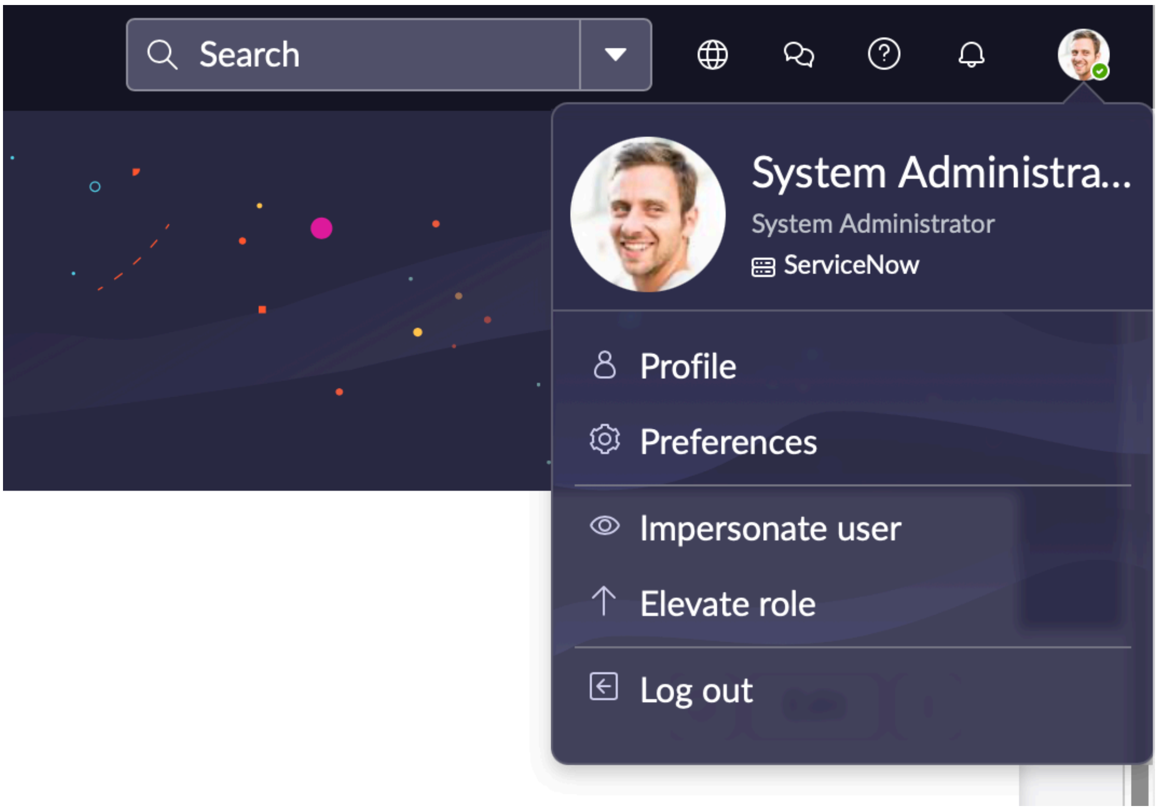

    - In Core UI:

      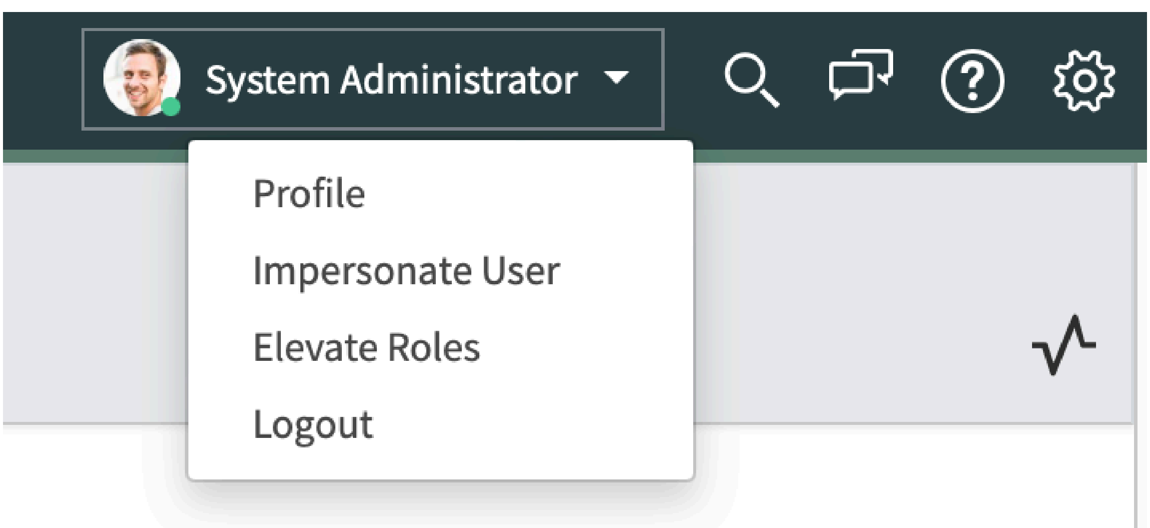

   The Impersonate user dialog box displays.
3. Select a user from the **Recent Impersonations** list or enter a different user's name in the user selection field.

- In Next Experience:

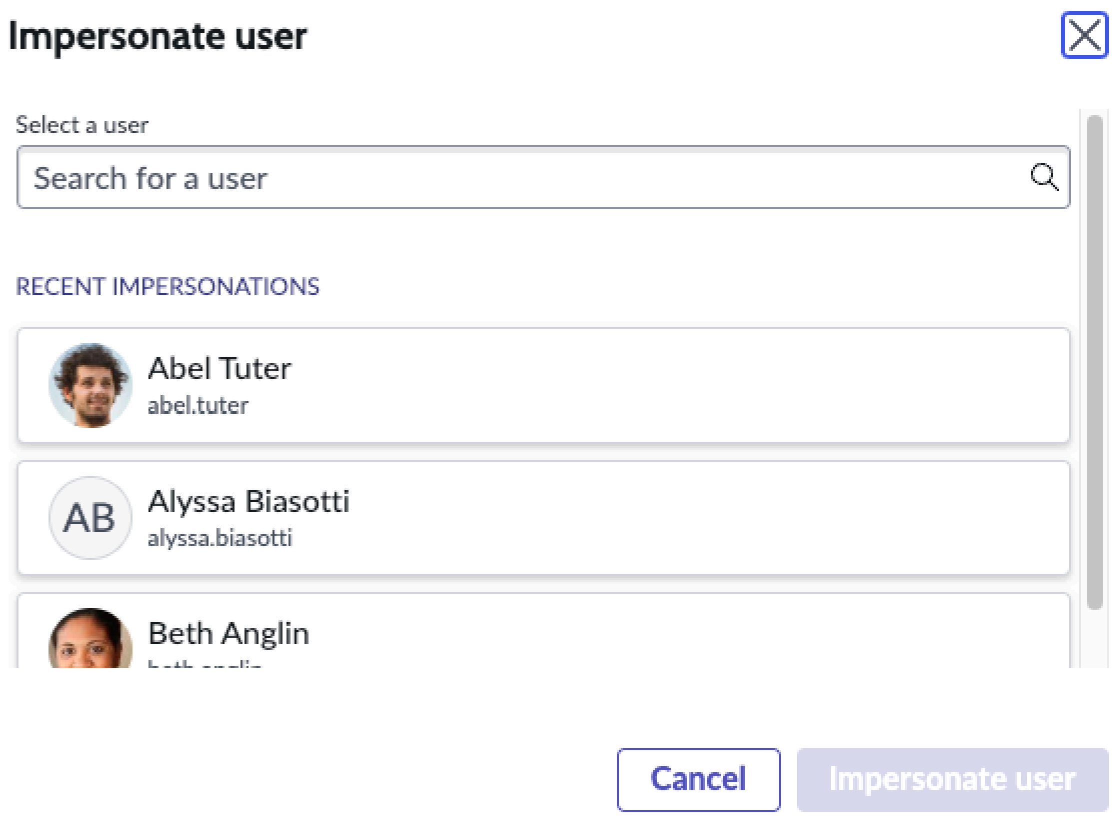

- In Core UI:

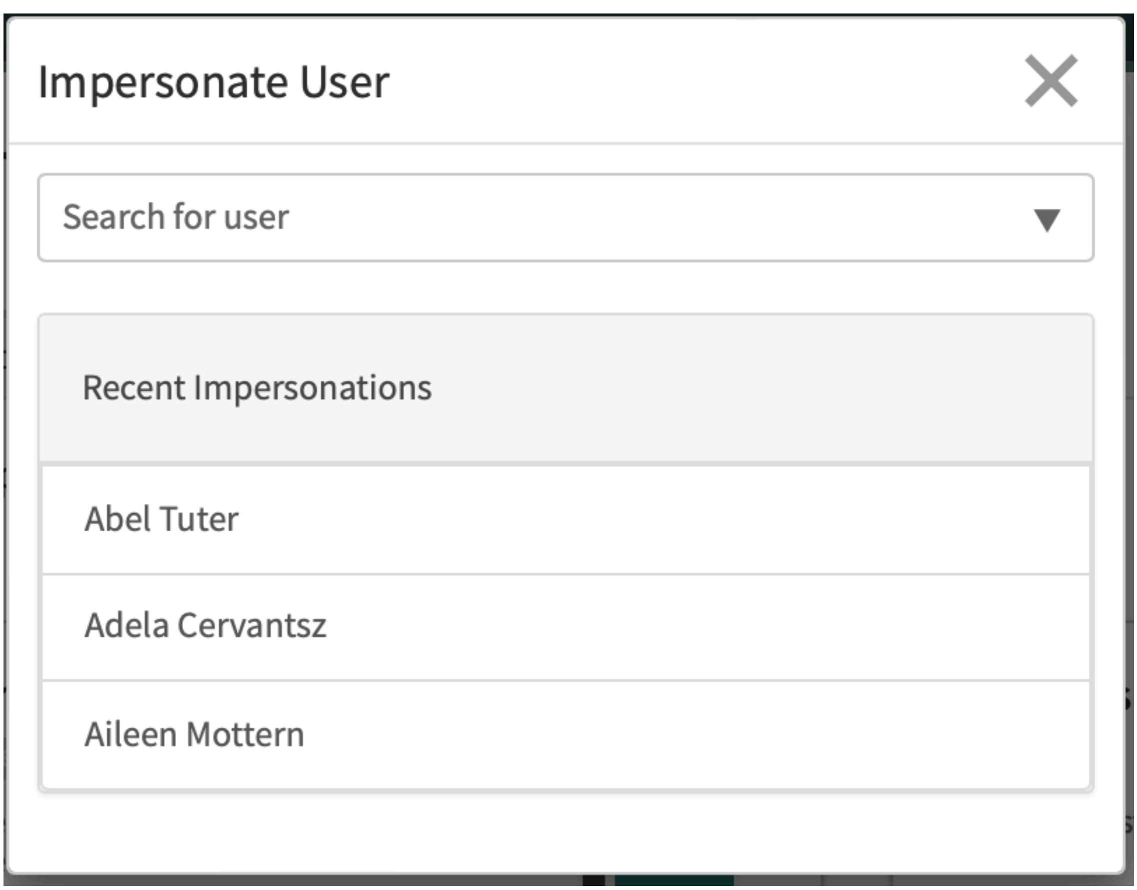

4. Select **Impersonate User**.
5. **Optional:** To impersonate another user while impersonating a user, open the user menu and select **Impersonate another user**.
6. To return to your original login, open the user menu and select **End Impersonation**.

Note:
In some cases, impersonating a user might cause an issue that makes it difficult to switch back. If you're presented with a broken page while impersonating a user in a test environment, you may need to force a logout. To do so, navigate to http://<instance name>.service-now.com/logout.do and log back in.

Impersonation logs

Impersonations are logged in the system log.

Log impersonations for either interactive (UI) or non-interactive sessions.

Impersonation logging for interactive sessions

Interactive sessions are performed through the user interface (UI). Enable or disable impersonation logging for interactive sessions using the *glide.sys.log_impersonation* property.

If you enable impersonation logging for interactive sessions by setting *glide.sys.log_impersonation* to **true**, all interactive sessions are recorded in the log.

Property	Description
glide.sys.log_impersonation	Enables (**true**) or disables (**false**) impersonation logging for interactive sessions.

Impersonation logging for non-interactive sessions

Non-interactive sessions are performed by applications and scripts, not through the UI.

Impersonation logging of non-interactive sessions is turned off by default. If you enable impersonation logging by setting the *glide.sys.log_impersonation.non_interactive* property to **true**, impersonations of non-interactive sessions are recorded in the impersonate log.

Even with *glide.sys.log_impersonation.non_interactive* set to **true**, the system doesn't log certain common impersonation tasks performed on behalf of the default users (*system*, *soap.guest*, and *guest*) because the application impersonates those default users to perform a variety of tasks.

Use the *glide.sys.log_impersonation.non_interactive.exclusion* property to exclude impersonations by other users in addition to the default users.

Property	Description
glide.sys.log_impersonation.non_interactive	Enables (**true**) or disables (**false**) impersonation logging for

Property	Description
	non-interactive sessions.
glide.sys.log_impersonation.non_interactive.exclusion	Excludes impersonation logging of non-interactive sessions for specified users. Enter user names as a comma-separated list. Default users (**system**, **soap.guest**, and **guest**) don’t need to be included in the list.

Managing user sessions

The ServiceNow AI Platform provides the ability to view and terminate individual user sessions, lock out users from the system, and make users inactive.

Regardless of the number of windows you have open in a browser, it’s considered to be one session. However, if you have two separate browsers open (such as Google Chrome and Safari), it’s considered to be two separate sessions.

- Terminating a user session effectively logs that user out of the next transaction, which is usually the next browser select. Use the terminate sessions feature when you want to perform system maintenance.
- Locking a user out of the system means they can no longer log in or generate any actions from any email messages that they send to the instance. Locking users out also terminates their sessions.
- Making a user inactive means they don't show up in any fields that reference active users on the **User** table.

Configure a maximum active time for user sessions

Secure your instance by enforcing a maximum time for sessions regardless of user activity.

Before you begin

Role required: admin

About this task

By default, sessions expire only after a period of inactivity. Enforcing a maximum active session time ends sessions regardless of whether a user has been active recently, including whether they recently selected to extend a session. The active session timeout should be greater than the value configured for the inactive session timeout. For example, if sessions are configured to time out after 30 minutes of inactivity, the active session timeout should be greater than 30 minutes.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Filter the System Properties [sys_properties] list for the following properties and then select a property to open its record.
   - *`glide.ui.active.session.life_span`*: Sets the maximum session time for authenticated user sessions regardless of user activity. Authenticated users are logged out of the instance after the time specified and must enter their credentials again to access the instance.
   - *`glide.guest.active.session.life_span`*: Sets the maximum session time for guest user sessions regardless of user activity. This setting helps secure an instance using applications that involve guest user sessions, such as Agent Chat.
3. In the **Value** field, enter the number of minutes before the session times out regardless of user activity.
   The value should be greater than the value of the corresponding properties for an inactive session timeout: *`glide.ui.session_timeout`* for authenticated users or *`glide.guest.session_timeout`* for guest users. By default, the inactive session timeout for both authenticated and guest users is 30 minutes.
4. Select **Update**.

Related topics

Configure a maximum active time for integration sessions

Modify user session timeout after inactivity

Specify when to time out user sessions after a period of inactivity.

Before you begin

Role required: admin

About this task

By default, after 30 minutes of inactivity in the application, the platform logs the user out automatically, unless the **Remember Me** check box in the login screen is selected. Making the interval longer can lead to the unnecessary maintenance of inactive sessions in memory. Adjust this timeout setting to no more than a few hours, although up to 24 hours is workable.

Note:

- Ajax calls to the server keep the session alive (such as Labels and Refreshing dashboards).
- Polling keeps the session alive when the chat desktop is open (requires the Chat plugin).

Procedure

1. Clear the **Remember Me** check box in the login screen.
2. In the navigation filter, enter `sys_properties.list`.
3. From the System Property [sys_properties] list, search for *`glide.ui.session_timeout`*.
4. If *`glide.ui.session_timeout`* doesn't exist, select **New** to add a new property using the following values:
   - **Name**: *`glide.ui.session_timeout`*
   - **Description**: Type a brief description. In this case, enter something like `Override the default session timeout (30). This value is in minutes.`

- **Type**: Select the appropriate data type. In this case, select **Integer**.
- **Value**: Change the default value from 30 minutes to a value of your choice.

Note:
The session timeout can also be set through installation exit customizations.

What to do next

Administrators may also want to add the following properties to the System Properties table.

- `glide.security.csrf.handle.ajax.timeout`: Handles errors for timed out Ajax requests when set to **true**.
- `glide.security.auto.resubmit.ajax`: Automatically resubmits timed-out Ajax requests when set to **true** and the Log in to an instance check box is selected or Change the default value of the Remember me check box. A pop-up appears to users asking them to continue.
- `glide.ui.auto_req.extend.session`: When set to **true**, the system automatically extends a user's session by the value they select for the homepage refresh time. If there's no homepage refresh time, the standard timeout value applies. Tablet and mobile devices don't support this property. When set to **false**, user sessions time out when the **Remember me** check box is clear. The timeout is based on whether there's a homepage refresh time. When there's no homepage refresh time, the standard timeout value applies. When there's a homepage refresh time, the user session times out after the timeout value plus one interval of the homepage refresh time. For example, if a user selects a refresh interval of five minutes, then that session expires after the timeout value plus five minutes.

Note:
Users who select the **Remember me** check box are unaffected by session timeout properties.

Administrators can also add the following properties to configure additional timeout settings for user sessions. These additional settings help to conserve system resources:

- `glide.session.unauthorized.timeout.enabled`: Enables an alternate session timeout for unauthenticated, guest sessions. Guest sessions are created for HTTP requests to the instance that don't contain authentication information. By default this property is set to **true**.
- `glide.unauthorized.session_timeout`: Specifies the time, in minutes, after an authenticated user logs out of an instance before the session ends. Set the property to a value greater than **0** and less than the value in the `glide.ui.session_timeout` property.

Related topics

Modify integration session timeout after inactivity

Prompting users to extend a session

Configure how much time users have to extend a session before it times out due to inactivity.

Before you begin

Role required: admin

About this task

By default, users are prompted to extend their session two minutes before it expires with an "Extend your session" dialog box. This procedure explains how to adjust the timing of when users are presented with this prompt.

Procedure

1. Navigate to **All**, and then enter `sys_properties.list` in the navigation filter.
2. Open the *`glide.ui.session_timeleft`* property.
3. In the **Value** field, enter the number of minutes before the session timeout to prompt users to extend their session.
   Setting the value to zero turns off prompting users to extend their sessions.
4. Select **Update**.

Lock out a user

Lock out a user when you don't want the user to access the instance.

Before you begin

Role required: user_admin or admin

Procedure

1. Navigate to **All** > **User Administration** > **Users** and select the user from the list.
2. Select the **Locked Out** check box, and update the record.

Mark a user inactive

You can mark a user inactive so the user doesn't show up in any fields that reference active users on the **User** table.

Before you begin

Role required: admin

About this task

Note:
By default, when you mark a user inactive, you lock out that user. The **Lock Out Inactive Users** business rule governs this behavior.

If you clear the **Active** checkbox, the user is locked out and cannot access the instance. This functionality is controlled through a Glide property *`glide.authenticate.only.allow.active.user.login`* which will only allow the active users to access the instance.

Procedure

1. Navigate to **All** > **User Administration** > **Users** and select the user from the list.
2. Clear the **Active** check box, and update the record.

Terminate a user session

You can terminate a user session, for example, if you're going to perform system maintenance and users are still logged in.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Logged in users**.
   You can only see users who are logged in to the same application node as you. If the Active field on a user record value is false, the user is logged in but not currently running a transaction. Most users appear as inactive at any given time.
2. Select the session that you want to end.
3. Select **Lock Out Session**.
   The session is terminated, and the user is redirected to the login page at the next attempted transaction. The user isn't locked out. Multiple user sessions may be associated with one user. Terminating a user session only affects the specific session.

   Note:
   Mobile user sessions can't be terminated using this process.

Non-interactive sessions

The Non-Interactive Sessions plugin creates a distinction between interactive and non-interactive users.

Interactive users

New users added to the instance automatically become interactive users. Interactive users can perform the following actions:

- Use their user name and password to log in to the UI or a service portal.
- Connect to an instance from a URL that calls a UI page, form, or list, for example, https://<instance name>.service-now.com/incident.do.
- Connect with single sign-on, for example, digest authentication or a Security assertion markup language (SAML).
- Use their credentials to authorize SOAP connections if allowed by strict security.
- Use their credentials for other API connections such as WSDL, JSON, XML, or XSD without restriction.

Non-interactive users

Non-interactive users can only use their credentials to authorize API connections such as JSON, SOAP, and WSDL. They can't log in to the ServiceNow UI. The strict security high security setting determines if non-interactive users are subject to Contextual Security Manager requirements.

Distinguishing between interactive and non-interactive users increases instance security by helping to ensure that users abide by the principle of least privilege.

Installed with Non-Interactive Sessions

Note:
Non-Interactive Sessions is enabled for all new instances since the Calgary release. If you don't see it in the list of plugins, request it using the Activate Plugin service catalog item in Now Support.

- Adds a column titled **Web Service Access Only** [web_service_access_only] to the **User** [sys_user] table.
- Changes all existing users to be interactive users (web_service_access_only=false).
- Updates the User form to display the **Web Service Access Only** [web_service_access_only] field by default.

Create a non-interactive user for web services

Non-interactive users can only connect to a ServiceNow instance from an API protocol. Use this feature to set up user accounts for web service authentication purposes.

Before you begin

Role required: user_admin or admin

About this task

Non-interactive users can't log in to an instance or a service portal or connect through single-sign-on. They can be used as a MID Server user if they're flagged as an Internal Integration User.

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Search for the user to be updated.

   Example

   For example, SOAP user.
3. Select the **Web Service Access Only** check box.
4. Select **Update**.

   Note:

   The ServiceNow platform uses any user name and password credentials supplied with a request even if the High Security Settings don't require authorization for a given API protocol. For example, if a SOAP request supplies a user name and password, the instance verifies those credentials even if SOAP requests don't require authorization. To avoid verifying user credentials, the request must not include them.

Make a non-interactive user record interactive

Manually switch a non-interactive user to an interactive user.

Before you begin

Role required: user_admin or admin

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Search for the user you want to update.
   For example, **System Administrator**.
3. Clear the **Web Service Access Only** check box.
4. Select **Update**.

Update web service user accounts for strict security

If your instance requires strict security, add the soap role to any user accounts used for web services.

Before you begin

Role required: user_admin or admin

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Select a web service user from the list.
3. From the **Roles** related list, select **Edit**.
4. Add soap to the **Roles List**.
5. Select **Save**.
6. Select **Update**.

Require authentication

You can specify whether non-interactive sessions require authentication from the **High Security Settings** module.

Before you begin

Role required: admin with elevated privileges

About this task

A non-interactive session bypasses the UI to connect to the instance at an API level. Typically, non-interactive sessions use set protocols such as JSON, SOAP, XSD, or WSDL. By default, all non-interactive sessions require authentication.

Procedure

1. Log in with an administrator user with the security_admin role.
2. Elevate your privileges to use security_admin.
3. Navigate to **All** > **System Security** > **High Security Settings**.
4. Select the matching **Requires authorization...** option for the protocol you want to set.
   For example, **Requires authorization for incoming SOAP requests**.
5. Select the check box to require authentication for the non-interactive session method.

   Clear the check box to enable the non-interactive session method to connect without providing any credentials.

   Note:

   Activating the Non-Interactive Sessions plugin might restrict existing users who manage SOAP and WSDL-based integrations from logging in, unless they already possess the SOAP role.

ServiceNow AI Platform tables and data

ServiceNow provides options for managing your data on the ServiceNow AI Platform.

Table administration	Data management	Export data	Roll back and delete recovery
Learn about working with complex, large amounts of data, by using tables and records on the ServiceNow AI Platform.	Explore options for storing and deleting data with ServiceNow.	Learn about exporting data from the ServiceNow AI Platform.	Learn how to roll back actions and recover deleted records.

Table administration

Handle data and process management for all your ServiceNow applications. Administrative roles on the ServiceNow AI Platform give you access to tables and records to manage incidents, problems, database changes, and more. Build graphical representations of tables and their relationships (such as parent or child tables) using Schema maps, and define elements in your database with Data Dictionary tables. Work with complex relationships between multiple sets of records using Many-to-Many Task Relations.

Data Management

Manage the growth of data on your instance with our set of Data Management tools. Keep performance smooth by archiving records, delete older records using table cleaner, and manage large tables with database rotations.

Exporting data

Export individual or multiple records from the ServiceNow AI Platform in a variety of formats.

Roll back and delete recovery

Roll back a software upgrade or plugin activation. Recover deleted records and related changes.

Table administration

A table is a collection of records in the database. Each record corresponds to a row in a table, and each field on a record corresponds to a column on that table.

This video covers the ServiceNow® infrastructure and introduces the concepts of tables, records and fields; table classes other table relationships; and schema maps.

Watch this short video to see an overview of table administration.

Get started

Exploring ServiceNow AI Platform® tables Learn the basics of tables and table administration.	Managing tables and indexes Manage tables and indexes, and create relational connections between records.
Working with the Task table Use the Task table to track task-specific functions, and connect with Incident and Problem tables.	Creating database views for reporting Create database views for reporting purposes.

Helpful resources

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

ServiceNow Community

ServiceNow Community

Support

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

Exploring ServiceNow AI Platform® tables

Applications use tables and records to manage data and processes, such as Incident, Problem, and CMDB. Tables can extend other tables, creating parent tables and child tables.

Table administration overview

The ServiceNow AI Platform® uses a table-based data structure to store and organize information. Pre-built tables are included for common IT service management (ITSM) processes, with their data fully importable and exportable in CSV, XML, or other formats.

Table administration users

Users

User	Description
Administrator	Admins can create custom tables for specific business requirements, define field properties, create relationships between tables, and extend existing tables without modifying the originals. Administrators also define access controls to allow only approved user access to records, business rules for various tasks, and table cleanup policies for unused and outdated records.

Table administration workflow

This infographic depicts how an administrator can extend both core tables and custom tables.

Extending core tables and custom Tables

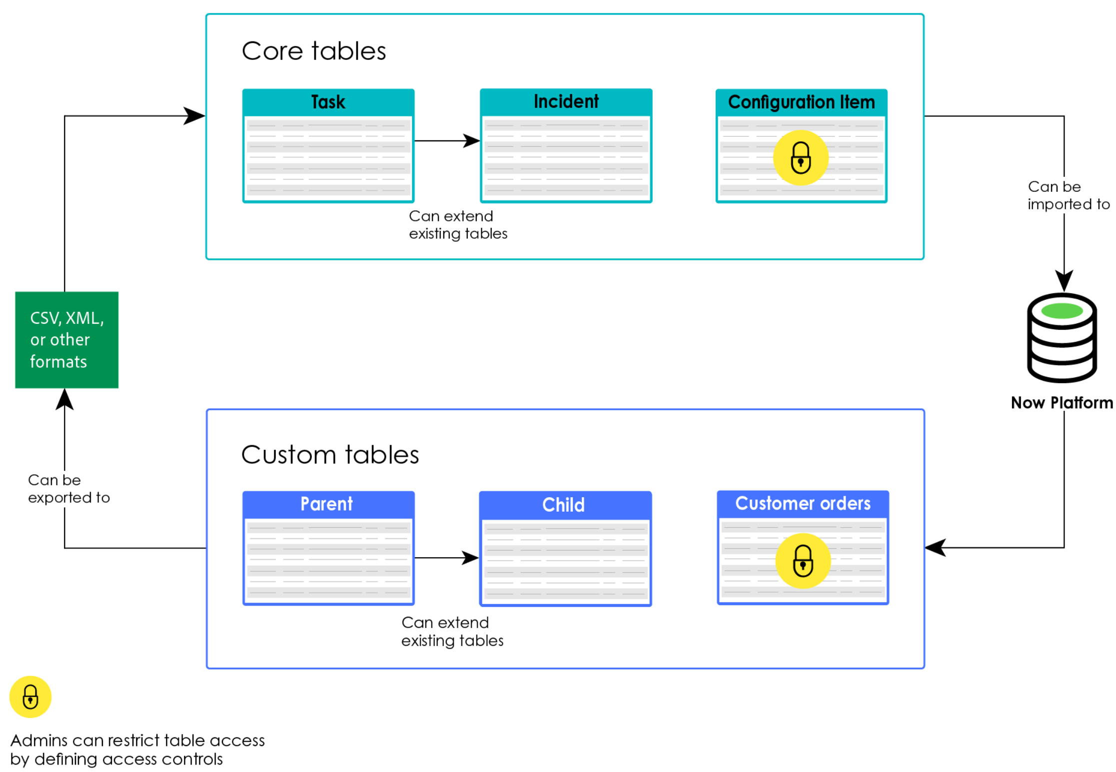

1. An administrator views table relationships and learns how table extension works on the ServiceNow AI Platform®.
2. The administrator creates a custom parent table and extends it by creating a custom child table.
3. The administrator defines access controls to restrict access to each table.
4. The administrator exports and imports data to core and custom tables as needed.

Table administration benefits

Benefit	Feature	Users
Define models and entity relationships across multiple tables.	Data dictionary tables	Admin
Access visual references for tables and their contents.	Viewing table references and extensions	Admin
Store records for your instance.	Create a custom table	Admin
Remove unneeded tables from your instance.	Deleting custom tables	Admin
Quickly find data stored in a table.	Create a table index	Admin
Connect a list to another list of multiple entries instead of a single field.	Create a many-to-many table relationship	Admin
Organize and store task-related data.	Working with the Task table	Admin
Connect tables for reporting purposes.	Creating database views for reporting	Admin

What to explore next

To learn more about using table administration features, see:

- Managing tables and indexes
- Working with the Task table
- Creating database views for reporting

Data dictionary tables

Access details related to tables, columns, and field labels in your instance.

Tables

All of the information on an instance is stored in tables, which consist of a series of records. The record in turn holds a series of fields that hold the individual bits of data and can be viewed either as a list or a form.

The Tables [sys_db_object] table contains a record for each table in the database.

Access the Tables [sys_db_object] table by navigating to **System Definition** > **Tables**. From the Tables list, an administrator can create a custom table, or select an existing table to view table details. In the table form, an administrator can access table administration options.

- View, add, or modify columns with a searchable and sortable embedded list, define the auto-number format, make the table extendable by other tables, and create modules for the table.
- Launch a schema map for a table by clicking the **Show Schema Map** related link.
- Open the dictionary entries for the table by right-clicking the form header and selecting **Show Dictionary Record**.

- Navigate directly to the default list or form view for the table by clicking the **Show List** or **Show Form** related link.
- Delete all records from a table by clicking the **Delete All Records**.

The following image shows a list of the tables that extend the Application File table.

Tables

Tables | Name | Search | Actions on selected rows... | New

All > Extends table = Application File

Label	Name	Extends table	Extensible	Updated
Agent Assist Recommendation	agent_assist_recommendation	Application File	false	2022-04-25 04:58:49
Record Producer Configuration	aisa_rp_config	Application File	false	2022-04-25 05:05:44
Search Actions	aisa_ui_action	Application File	false	2022-04-25 05:05:43
AI Search ACL Overrides	ais_acl_overrides	Application File	false	2022-04-25 03:48:21
AI Search Child Table	ais_child_table	Application File	false	2022-04-25 03:48:16

Dictionary Entries

The Dictionary Entries [sys_dictionary] table, also called the **System Dictionary**, defines every table and field in the system. It contains information about data type, character limit, default value, dependency, and other attributes of a field.

Access the system dictionary in one of these ways:

- To see the system dictionary list view, navigate to **System Definition** > **Dictionary**.
- To view particular dictionary definition, right-click the list header, form header, or field label, and select **Configure Dictionary**.

The following image shows a filtered list of dictionary entries for the Incident table and the Task table, which it extends.

Dictionary Inc Task

Dictionary Entries | Table | Search | Actions on selected rows... | New

All > Table = incident .or. Table = task

Table	Column name	Type	Reference	Default value	Display	Text index	Audit	Updated
incident	resolved_at	Date/Time	(empty)		false	false	false	2022-04-25 04:29:00
incident	resolved_by	Reference	User		false	false	false	2022-04-25 04:29:00
incident	rfc	Reference	Change Request		false	false	false	2022-04-25 04:16:53
incident	severity	Integer	(empty)	3	false	false	false	2022-04-25 04:16:53
incident	subcategory	String	(empty)		false	false	false	2022-04-25 04:16:53
incident	sys_id	Sys ID (GUID)	(empty)		false	false	false	2022-04-25 04:16:53
task		Collection	(empty)		false	true	false	2022-04-25 05:43:21
task	active	True/False	(empty)	true	false	false	false	2022-04-25 03:45:17

Field Labels

The Field Labels [sys_documentation] table, also called the **Language File**, contains information about the labels and hints for each table and column in the system.

Access the language file in one of these ways:

- To see the list view, navigate to **System Definition** > **Language File**.
- To see the field label for a particular field, right-click the field label on the form.

The following image shows the language file filtered to display only labels on the Incident table.

Field Labels

Field Labels | Table | Search | Actions on selected rows... | New

All > Table = incident

Table	Element	Language	Label	Plural	Hint	Updated
incident	problem_id	en	Problem	Problems	Related problem, if one exists	2022-04-25 04:16:53
incident	notify	en	Notify	Notifies		2022-04-25 04:16:53
incident	business_impact	en	Business impact	Business impacts		2022-04-25 05:00:23
incident	severity	en	Severity	Severities		2022-04-25 04:16:53
incident	reopened_by	en	Last reopened by	Last reopened by		2022-04-25 04:16:53

System dictionary

View a list of all tables in columns in your instance from the system dictionary.

Each row in the system dictionary represents either a table or a column in one of the tables. The system dictionary provides options for administrators to modify tables and fields, which in turn define lists and forms.

Use caution when changing system dictionary records because changes can have a high impact on functionality. In particular, changes to dictionary entries for system tables, which are tables that begin with sys_, can create system-wide issues such as the inability to use update sets.

Dictionary changes are difficult to reverse. Also, dictionary changes automatically apply to all extended tables unless a dictionary override is defined. Be sure that changes are well-tested before applying them to a production instance.

Creation options

When you create a field from the system dictionary, it is automatically added at the end of the first section of the default form view.

In most cases, use the following interfaces rather than creating entries directly on the system dictionary:

- To create tables and fields, use the Tables module.
- To create fields, configure the table form.

Table relationships

You can create relationships between tables by extending tables, referencing records in another table, creating many-to-many relationships, and joining tables in a database view.

Tables can be related to each other in several ways.

Extensions

A table can extend another table. The table doing the extending (child class) includes all of the fields of the other table (parent class) and adds its own fields. For instance, the Incident [incident] table has all of the Task [task] table fields (because an incident is a special form of task) and has its own incident-specific tasks. See Table extension and classes.

One-to-many

Within a table, a field can hold a reference to a record on another table. There are three types of one-to-many relationship fields.

Reference fields

Allow a user to select a record on a table defined by the reference field. For instance, the **Caller** field on the Incident table allows the user to select any record on the User table.

Glide lists

Allow a user to select multiple records on a table defined by the glide list. For instance, the **Watch list** field on the Incident [incident] table allows the user to select records on the User [sys_user] table.

Document ID fields

Allow a user to select a record on any table in the instance. These fields are much less common, but one example is the **Document** field on the Translated Text [sys_translated_text] table.

Many-to-many

Two tables can have a bi-directional relationship, so that the related records are visible from both tables in a related list.

Database views

Two tables can be joined virtually in a database view to enable reporting on data that might be stored over more than one table.

Table extension and classes

Enable one or more child tables to share fields and records with a parent table. Administrators and application developers can only extend tables during table creation.

Administrators and application developers typically extend tables to create a set of related records that share information. For example, in the base system, the Task and the Configuration Item tables have multiple extensions:

Sample extended tables in the base system

Original table	Related tables extended from original table
Task [task]	• Incident [incident] • Problem [problem] • Change Request [change_request]
Configuration Item [cmdb_ci]	• Application [cmdb_ci_appl] • Computer [cmdb_ci_computer] • Database [cmdb_ci_database]

A table that extends another table is called a child class, and the table it extends is the parent class. A table can be both a parent and child class both extending and providing extensions for other tables. A parent class that is not an extension of another table is called a base class.

Administrators can use these tools to see the relationships between classes.

- Schema map
- System dictionary
- Tables module

Extending a table:

- Links the new table to the extending table.
- Creates system fields in the new table.
- Creates one or more database tables to store the parent and child classes. The number of tables the system creates depends upon the extension model selected during table creation.

Extension models

The ServiceNow AI Platform offers these extension models.

- Table per class
- Table per hierarchy
- Table per partition

The extension model determines these attributes.

- The number of database tables created
- The derivation of fields from parent classes
- The replication of records from child classes

Table per class

Tables created

Creates a separate database table for the parent class and each child class.

Fields derived from parent class

Child classes derive fields from parent classes.

Dictionary records created for parent class

A parent class has a Dictionary record for the collection and for each field that can be derived from it. For example, the Contract [ast_contract] table has 59 Dictionary records, which define the table and its fields.

Dictionary records created for child classes

Each child class only has Dictionary entries for fields unique to the class.

Records replicated

The parent class replicates each record stored in its child classes. Each child class only stores records unique to the class. Replicated records have the same Sys ID value in each table. The system replicates any change you make to a child record to the matching record in the parent table. For example, Contract [ast_contract] table replicates records from the Lease [ast_lease] and Warranty [ast_warranty] tables.

Table per hierarchy

Tables created

Creates one database table for the parent class, which stores all records for the parent and child classes. Child classes do not have separate database tables.

Fields derived from parent class

Child classes derive fields from parent classes. For example, the Incident table derives fields from the Task table.

Dictionary records created for parent class

A parent class has a Dictionary record for the collection and for each field that can be derived from it. For example, the Task table is a parent class that has 66 Dictionary records, which define the table and its fields.

The Dictionary entry for the parent class contains a sys_class_name column whose value indicates which child class each record belongs to. For example, Incident records have a sys_class_name value of incident, and change records have a sys_class_name value of change.

Dictionary records created for child classes

Each child class only has Dictionary entries for fields unique to the class. For example, the Incident table only has 22 Dictionary records, which are not already defined in the Task table.

Records replicated

Record replication is not needed, because the parent class stores all records that belong to the hierarchy. For example, the Task table contains all records from its child classes such as the Change, Incident, and Problem tables.

Table per partition

Tables created

Creates one database table for the parent class, which stores all records for the parent and child classes. Child classes do not have separate database tables. As the database table reaches a storage limit, the system dynamically adds storage tables (partitions) to store additional records.

Fields derived from parent class

Child classes do not derive fields from parent classes. Instead each child class has its own list of fields. For example, the Base Configuration Item [cmdb], Configuration Item [cmdb_ci], and Hardware [cmdb_ci_hardware] tables all have their own field definitions.

Dictionary records created for parent class

A parent class has a Dictionary record for the collection and for each field relevant to it. For example, the Base Configuration Item [cmdb] table is a parent class that has 48 Dictionary records.

The system replicates changes made to parent class Dictionary entries to child class Dictionary entries. For example, when you change the name column in the parent class Base Configuration Item [cmdb] table, the system replicates it to child class Dictionary entries such as the Configuration Item [cmdb_ci] and Hardware [cmdb_ci_hardware] tables.

The Dictionary entry for the parent class contains columns for sys_class_name and sys_class_path whose values indicate which child class each record belongs to. For example, Hardware records have a sys_class_name value of cmdb_ci_hardware, and computer records have a sys_class_name value of cmdb_ci_computer.

When the database table reaches a storage limit, the system updates the Dictionary entry for the parent class to include columns for sys_storage_alias and storage_table_name. These storage column Dictionary entries allow administrators to manage the parent class and its storage tables as a single logical unit.

Dictionary records created for child classes

Each child class has a Dictionary record for the collection and for each field relevant to it. For example, the Hardware table has 73 Dictionary records with some records duplicating columns in the parent class.

Records replicated

Record replication is not needed, because the parent class stores all records that belong to the hierarchy. For example, the Base Configuration Item [cmdb] table contains all records from its child classes such as the Application [cmdb_ci_appl], Computer [cmdb_ci_computer], and Hardware [cmdb_ci_hardware] tables.

Storage aliases

Learn about the role storage aliases play in data manipulation and field creation in the ServiceNow AI Platform.

Understanding storage aliases is important for effective data management and schema customization in the ServiceNow AI Platform, particularly when dealing with complex table hierarchies like those in the Task [task] table.

By default, administrators have access to the Storage Column Aliases [sys_storage_alias] table within an instance. However, transactional processes against this table can't be performed by an administrator from the user interface.

Table hierarchy and models

Understanding storage aliasing requires knowledge of table hierarchies within the Task [task] table, which involves two models: Table per hierarchy and Table per class.

Table per hierarchy

This model uses a single physical table, the Task [task] table, which features a flattened hierarchy where all columns within the task hierarchy exist only on the task table. For example, fields related to change requests aren't on a separate Change Request [change_request] table but are integrated into the Task [task] table. You can verify whether a table is Table per hierarchy by checking the extension model field in the Tables [sys_db_object] table. If the table's parent is a direct child of the Task [task] table, the table uses Table per hierarchy.

An extended table inherits the hierarchy of its parent. For example, the IMAC [change_request_imac] table is a child table of the Change Request [change_request] table, which extends the Task [task] table. Because the Change Request [change_request] table is Table per hierarchy, the IMAC [change_request_imac] table model is also Table per hierarchy. Legacy tables such as the Incident [incident] table, Change Request [change_request] table, and Problem [problem] table are all part of the flattened Task [task] table hierarchy.

Table per class

This model applies to tables that physically exist in the database. It's used for new tables directly extending from the Task [task] table when the task row count exceeds 1 million rows. Unlike Table per hierarchy, Table per class doesn't use glomming because it doesn't operate within a flattened hierarchy.

Storage alias definition

A storage alias is created for every field on a table within an instance. Familiarize yourself with the key fields on the Storage Column Aliases [sys_storage_alias] table.

Element Name

Indicates how the field appears to users, which is important for user interface interactions and scripting.

Storage Alias

Indicates the exact physical column where the data for a field is stored. The storage alias value is used in combination with the table_name value to identify what data to manipulate. The storage alias value is the actual physical column on the Task [task] table.

Storage Table Name

Specifies the physical table that houses the element. For Table per hierarchy tables, the value is always task. For Table per class tables, the value is the name of the physical table.

Table

Specifies the logical class that each element links to on the physical Task [task] table. The Table element holds the table_name value which is the class discriminator.

Storage column aliases

Storage Column Aliases | Element name | Search | Actions on selected rows... | New

All > Storage table name = task > Storage alias starts with a_ref_2

Element name ▲	Offrow storage	Storage alias	Storage table name	Table
Search	Search	a_ref_2	=task	Search
cab_delegate	false	a_ref_2	task	Change Request [change_request]
category	false	a_ref_2	task	Standard Change Proposal [std_change_proposal]
chat_channel	false	a_ref_2	task	Chat Queue Entry [chat_queue_entry]
feedback	false	a_ref_2	task	Knowledge Feedback Task [kb_feedback_task]
guidance_step	false	a_ref_2	task	Guidance Task [help_guidance_task]
policy_id	false	a_ref_2	task	CMDB Data Management Task [cmdb_data_management_task]
request	false	a_ref_2	task	Requested Item [sc_req_item]
requested_for	false	a_ref_2	task	Asset Task [asset_task]
request_item	false	a_ref_2	task	Catalog Task [sc_task]
rfc	false	a_ref_2	task	Incident [incident]

In this example, the storage alias for the cab_delegate element is a_ref_2 and the physical storage table where the data is stored is task. The example depicts 10 logical elements in different logical classes on the Task [table] that all link to the same alias a_ref_2 on the physical Task [task] table.

The sibling elements are glommed, which means they share one physical column on the Task [task] table. You can query data from the logical element cab_delegate using a query like:

```
SELECT a_ref_2 from task WHERE sys_class_name='change_request'
 AND a_ref_2 IS NOT NULL
```

The query specifies data in the physical column a_ref_2. The class discriminator change_request is used in combination with the storage alias a_ref_2 to query the logical element cab_delegate from the logical class change_request on the physical Task [task] table.

The naming convention for fields created in the actual physical tables can vary depending on the type of field. In this example, a_ref_2 is an alias on the Task [task] table that holds values for reference fields.

Functionality and usage

Storage aliases serve multiple purposes.

Mapping

Aliases map logical elements (Table per hierarchy) or physical elements (Table per class) to the actual physical columns in the backend database. The logical element can be glommed to a physical column that is larger than the logical element's max_length.

Glomming

Aliases allow multiple sibling elements to share a single physical column in a Table per hierarchy model.

Label mapping

Aliases associate sys_documentation (label) records with their respective elements, enhancing visibility in forms, reports, and list views.

Restrictions and rules

- Two logical elements within the same class can't share a physical column. For example, two string fields created on the Incident [incident] table can't map to the same physical column in the database.
- A parent element and its child can't share a physical column. For example, a field created on the Incident [incident] table can't be mapped to a physical column when that physical column is already mapped to a field on the Task [task] table.
- Only sibling elements can share a physical column. For example, a reference field on the Change Request [change_request] table and a reference field on the Incident [incident] table and can both map to the same physical column.
- Fields created directly on the Task [task] table (where sys_class_name is 'task') can't be glommed.

Managing tables and indexes

Administrators can modify the database structure using table administration tools.

Custom tables and fields

Administrators can add custom tables, define fields, and manage field labels. Administrators can also delete custom tables that aren't needed or delete all the records in a table without deleting the table itself.

Dictionary entries

The system dictionary is a table, called Dictionary Entry `[sys_dictionary]`, that contains details for each table and defines every column within each table in an instance. Administrators can modify tables and fields in the system dictionary, which subsequently define the lists and forms that users access throughout the instance.

Dictionary attributes

Dictionary attributes define global behavior for fields. An administrator applies attributes directly to the dictionary entry to control various aspects of how the table or field behaves across the entire system.

Dictionary overrides

Dictionary overrides enable administrators to override field settings on extended tables. Dictionary overrides enable administrators to customize the behavior of a field for a specific table that inherits from another table. Overrides are used to change the inherited dictionary attributes on child tables without affecting the parent table or other child tables.

Table flattening

With table flattening, the Task [task] table and Base Configuration Item [cmdb] table have a hierarchy of related tables where one or more child tables extend from a parent table.

The extension model does not affect how the tables appear or function in the instance. All table functionality, including database views, remains unchanged by the extension model.

Table indexes

An index puts unordered tables into order, and efficiently speeds up queries to columns in your tables, by creating pointers to where information is stored in your database. Administrators can add indexes to access table data more easily.

Many-to-many table relationships

Many-to-many (M2M) table relationships enable administrators to associate records between two tables in a bidirectional and flexible manner. This relationship type is used to create links between records where each record from one table can relate to multiple records in another table, and vice versa.

Create a custom table

Manage and store application data by creating a custom table, adding columns, and defining access controls.

Before you begin

Check your custom table entitlements before creating custom tables.

Role required: admin

About this task

Note:
To minimize the existence of unneeded and extraneous tables in an instance, don't overdo table creation. Create only those tables you need to minimize administration tasks and shorten upgrade times. If possible, create tables when building applications with tools. To learn more, see Building applications.

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Select **New**.
3. On the form, fill in the fields.

Control	Description
Label	Enter a unique label for the table (such as **Laptops** or **Thin Clients**). The label appears on list and form views for the table. Updating the **Label** field also updates the label record in the language file for the current language. See Field Labels in Data dictionary tables.
Name	Edit the table name, which is automatically populated based on the table label and a prefix as follows:

Control	Description
	◦ For a table in a scoped application, the name is prefixed with a namespace identifier to indicate that it is part of an application. ◦ For a table in the global application, the name is prefixed with the string `u_`. ◦ For a remote table in a scoped application, the name is prefixed with a namespace identifier and the string `st_` to indicate that it is remote and part of an application. ◦ For a remote table in the global application, the name is prefixed with the string `u_st_`. You cannot modify the prefix; however, you can modify the rest of the table name. The name can contain only lowercase, alphanumeric ASCII characters and underscores (_).
Extends Table	Select the table to extend. Extending a base table incorporates all the fields of the original table and creates system fields for the new table. If they are in the same scope or if they can be configured from other scopes, you can extend tables that are marked as extensible. **Important:** Don't extend the [sys_choice] table. Extending the [sys_choice] table is not supported. The Extends Table option is not available if you select the **Scriptable Table** check box to designate that this table is scriptable.
Application	[Read only] Displays the application associated with this table. If you are working on an application or are creating a table from an application record, the field defaults to that application. Otherwise, the field defaults to **Global**. Any records that are created from the table record, such as modules and security rules, are assigned to this application by default.
Remote Table	Option to designate this table as a remote table. A remote table is a table to which you can associate a script definition for retrieval of data from an external data source. Unlike a standard internal table, the data does not come from records in the current instance. If you select this check box, you cannot access the **Extends Table** option because scriptable tables cannot be extended from other tables. **Note:** To learn more about scriptable tables, see: ◦ Retrieving and caching external data using remote tables ◦ Create a remote table ◦ Create a script definition for a remote table
Create module	Select the check box and then complete the **Add module to menu** field to create a list module in the application menu. This option is available only when creating a table.

Control	Description
Add module to menu	Select an existing menu or select **Create new** and enter a new menu name. This option is available only when the **Create module** check box is selected.

4. Add columns to the table using the **Table Columns** embedded list in the **Columns** section.

Column	Description
Column label	Define a unique label for the column. The label appears on list headers and form fields for the column. When you update the **Column label** field, the system also updates the label in the language file for the current language.
Type	[Mandatory] Define the field type for the column. See Field administration and Field types. To preserve existing data, only change fields between the same basic type (for example, **Choice** and **String**). A warning appears if a change to a custom field results in data loss. You cannot change a base system field, that results in data loss.
Reference	Make the field into a Reference field type by entering the referenced table name. **Note:** Dynamic reference creation is enabled for this field. So, if you enter a table name that does not match an existing table, a new table is created when you save changes to the current table record. If the current table has a module in the application navigator, then a module for the newly created table is automatically created in the same application menu.
Max length	[String fields only] Limit the length of the field. A length of under 254 appears as a single-line text field. Anything 255 characters or over appears as a multi-line text box.

Column	Description
	Note: ◦ You can only change this value if the **Type** of the field is **String**. Changes for any other type of field are ignored. ◦ On an Oracle instance, you cannot increase the maximum length of a string field to greater than 4000 because it requires the CLOB datatype in Oracle. To increase beyond this size, log an incident with ServiceNow Technical Support to request the change. ◦ To prevent data from being lost, only decrease the length of a string field when you are developing a new application and not when a field contains data. A warning appears if a change to a custom field results in data loss. For a base system field, you cannot make a change that results in data loss.
Default value	Specify the default value of the field for any new record. Ensure that this value uses the correct field type. For example, an integer field can use a default value of 2 but cannot use a default value of *two*. These values can be overridden with dictionary overrides.
Display	Indicate whether this field is the Display values (appears on records that reference this table). **Note:** This option does not control whether this field is displayed on lists or forms.

5. Define additional table options in the **Controls** section.

Control	Description
Extensible	Select the check box to enable other tables to extend this table. Clear the check box to prevent the creation of additional child tables; existing child tables remain unchanged.
Live feed	Select the check box to enable record feeds for the table. This option adds the **Show Live Feed** icon ( ) in the form header.

Control	Description
Auto-number	Select the check box, and then define the number format to add an auto-numbered field to the table. The check box is available only when a number format does not exist for the table. Otherwise, you can edit the existing number format.
Create access controls	Select the check box and then complete the **User role** field to create basic security rules for the table.
User role	Enter a new name or select an existing user role that is required to access this table. This option is available only when the **Create access controls** check box is selected.

6. Define the scope protection for the table in the **Application Access** section.
   For more information, see Application access settings.

Control	Description
Accessible from	Specify which application scopes can access the table: **All application scopes** Can be accessed from any application scope. **This application scope only** Can be accessed only from the current application scope.
Can read	Select the check box to enable script objects from other application scopes to read records stored in this table. This option offers runtime protection. For example, a script in another application can query data on this table. First select read access to grant any other API record operation.
Can create	Select the check box to enable script objects from other application scopes to create records in this table. This option offers runtime protection. For example, a script in another application can insert a new record in this table. This option is available only when the **Can read** check box is selected. Clear the check box to prevent script objects from other application scopes from creating records in this table.
Can update	Select the check box to enable script objects from other application scopes to modify

Control	Description
	records stored in this table. This option offers runtime protection. For example, a script in another application can modify a field value on this table. This option is available only when the **Can read** check box is selected. Clear the check box to prevent script objects from other application scopes from modifying data stored in this table.
Can delete	Select the check box to enable script objects from other application scopes to delete records from this table. This option offers runtime protection. For example, a script in another application can remove a record from this table. This option is available only when the **Can read** check box is selected. Clear the check box to prevent script objects from other application scopes from deleting records from this table.
Allow access to this table via web services	Select the check box to enable users to make inbound Web services web service queries to this table. This option offers both design-time and runtime protection. The user performing the query must have the correct permissions to access this table, even when this check box is selected. Clear the check box to prevent users from making web service queries to this table.
Allow configuration	Select the check box to enable applications from other application scopes to create configuration records for this table that change its functionality. For example, an application designer can select this table from the **Tables** list on business rules, client scripts, or UI actions. This option offers design-time protection. Clear the check box to prevent application designers from selecting this table when creating configuration records.

7. Click **Submit**.

What to do next

To change the field labels in your table or the label of the table itself, perform the following task:

1. Navigate to the Field Label [sys_documentation] table, and filter the table to show the table that has the fields you want to change. To find the label that represents the table itself, filter the records where **Element** is empty.
2. Open the field record that you want to change, and make your updates. For a description of each field, see Field Label table.

Global default fields

When you create a new custom table, several fields appear in the **Table Columns** embedded list. For all tables, required system fields are added automatically. You cannot delete or modify these fields.

For tables that extend another table, fields on the parent table also appear on the **Table Columns** embedded list for the current table. If you modify these fields, remember that all changes to fields on the parent table also affect all child tables, not just the current table.

These required system fields are added to all tables:

Required system fields

Field	Type	Description
Class [sys_class_name]	System Class Name	If the table is extensible, a string field that indicates which child table contains the record.
Created [sys_created_on]	Date/Time	A time-stamp field that indicates when a record was created.
Created by [sys_created_by]	String	A string field that indicates the user who created the record.
Sys_id [sys_id]	Sys ID	The Unique record identifier (sys_id) for the record.
Updates [sys_mod_count]	Integer	A numeric field that counts the number of updates for this record since record creation.
Updated by [sys_updated_by]	String	A string field that indicates the user who most recently updated the record.
Updated [sys_updated_on]	Date/Time	A time-stamp field that indicates the date and time of the most recent update.

Deleting custom tables

Administrators can delete custom tables that are no longer needed. For example, delete a table from an application that is under development because the business requirements change.

A table is custom if an administrator created it and it is not part of a system upgrade or plugin activation. Custom table names always begins with `u_`, or `x_` for scoped tables.

> **Warning:**
> You cannot delete base system tables. If you inadvertently delete such a table, it is automatically recreated when you upgrade an instance.
>
> You cannot delete a table that has associated tables that extend from it.
>
> You must first delete all records in the table before you delete the table itself. If you do not first delete the records, errors can result on the parent table if it references the records. Deleting the records removes any references from the parent table. To learn more, see delete the records in the table.

Delete a custom table

If you no longer need a custom table, you can delete it after you delete all the records in the table.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Open the table to delete.
3. Select **Delete All Records**.

   Deleting all records before deleting the table ensures that the business logic is properly executed (for example, reference cascade rules or other delete business rules). If you do not delete all records from the table first, then you must manually fix any other records or tables that the table deletion impacts. Cascade deletions are also taken care of in this step.
4. Select **Delete**.
   The **Delete** button is only available for custom tables beginning with u_ (global) or x_ (scoped).
5. In the confirmation dialog box, enter `delete` and click **OK**.

   The table and all items that reference the table are deleted, including:

   - Choice list items
   - Forms, form sections, lists, and related lists
   - Reports and Performance Analytics widgets
   - Reference fields that reference the table
   - Access controls

Related topics

Deleting custom tables

Delete all records from a table

You may decide to delete all the records on a table without deleting the table itself. For example, the administrator may want to delete all incidents on a test instance without deleting the incident table itself.

Before you begin

Role required: admin

About this task

When you prepare to delete many records from a table, consider the following guidelines to minimize impact on performance.

- Limit the number of records to be deleted in a single delete action to prevent the table from being locked. Use the *`setLimit()`* method described at setLimit.
- Minimize triggering an excessive number of business rules as a result of this deletion. Use the *`setWorkflow(Boolean e)`* method described at setWorkflow.

Note:
Deleting all records for a table also deletes records from tables that extend the table.

There are several methods for deleting table records. Depending on the number of records to be deleted, choose the method that would be most efficient in your environment. Use these methods with caution. Before performing this procedure, be sure that you do not need any of the records.

Procedure

- From the **Tables and Columns** module, complete the following steps.

  1. Navigate to **System Definition** > **Tables & Columns**.

  2. Select the table for which to delete records.

     Note:
     If the table you want to delete records from belongs to a scoped application, select that application in the scope picker. Alternatively, configure scope-to-scope access. See Restricted caller access privilege settings.

     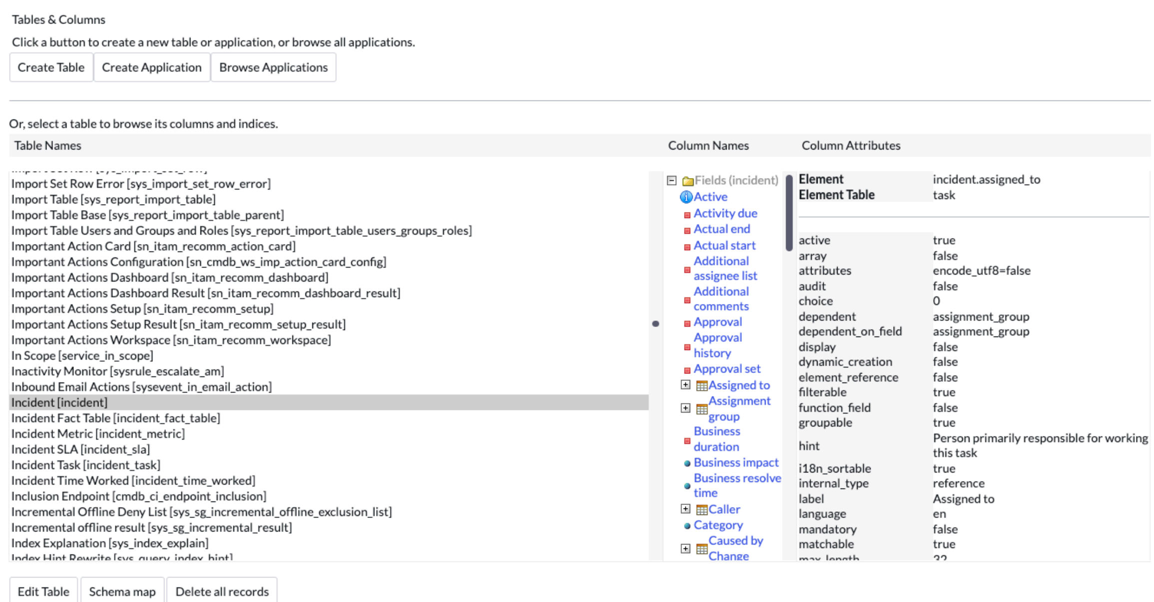

     Note:
     Some system tables do not allow this method of deleting records, for example, you cannot delete all user records [sys_user]. The list of tables does not include system tables that you cannot delete records from.

  3. Click **Delete all records**.

  4. In the confirmation dialog box, enter `delete` and click **OK**.

- From a list view, complete the following steps.

  1. In List v2, navigate to the list view that displays the table records, click the context menu, and then navigate to **Show**.
  2. In List v3, click **<Number> rows per page** in the footer.
  3. Set the number of rows per page to display the max number.
  4. Select all the rows on the page.

5. Click **Actions on selected rows**, and then click **Delete**.
6. Continue deleting all rows on a page until all records are deleted.

Result

The selected table is empty of records. The table still exists, and any references to the table on other tables (such as business rules or reference fields) are preserved.

Modify dictionary entries

Modify dictionary entries by configuring a field on a form or from the Dictionary module.

Before you begin

Role required: admin

Procedure

1. Do one of the following steps:
   - Navigate to a field on a form, right-click the field, and select **Configure Dictionary** or **Show <field name>**. The system dictionary entry for the field opens.
   - Navigate to **System Definition** > **Dictionary**, and click an entry for a field or table. Entries for tables have **Type** set to **Collection**.
2. Update the dictionary entry fields.
3. Click **Update**.

Dictionary entry form

The Dictionary Entry form provides an Advanced view and additional fields. You might need to configure the form to see all fields.

Dictionary Entry form

Field	Description
Table	Defines the table in which the element is created. **Note:** This list shows only the tables that meet the scope protections for adding fields.
Type	Defines the field type of the column or, if the **Type** is **Collection**, indicates that the dictionary entry represents the table. For more information, see Field administration. You can change the type of a field. To preserve existing data, only change between logical types that map to the same physical type on the database. For example, Choice and String.
Active	Enables or disables the field. When this check box is cleared, the field is hidden in lists and forms in the Classic Environment. However, inactive fields are still available for queries. Additionally, disabling a field by clearing this option doesn't remove it from a configured list. Inactive fields display in admin-configured lists until an admin updates the list layout to remove them. See below for additional behavior notes.

Dictionary Entry form (continued)

Field	Description
	• Forms: Deactivated fields are removed from forms. • Lists: Inactive fields continue to display in admin-configured lists until an admin updates the list layout to remove them. If an inactive field is removed from an admin-configured default view, any personal lists will reset to follow the default view, thus removing the inactive field from the personal list. • Reports: Existing reports that include inactive fields will continue to work, because the data still exists in the records. However, deactivated fields are not available when creating new reports or list layouts if they are not present in admin-configured lists. • Condition Builder: Deactivated fields are not available in the Condition Builder, unless a you are using an existing URL with filters containing the inactive field. **Show Matching**, **Filter Out**, and **Group By** can also be used on an inactive field if the column is present in the list view. • ServiceNow API: The ServiceNow API does not check whether a column is active or inactive. The old values for the field remain in the records and cannot be undone. • Scripts: Inactive fields may still be referenced in scripts. (Server side scripts that use GlideRecord can read and write to inactive fields).
Function field	If selected, creates a field that displays the results of a database function, such as a mathematical operation, field length computation, or day of the week calculation. Once the new function record is saved, you cannot clear the check box to make the field a regular field.
Read only	Determines whether the users can change the field value in the user interface. When this checkbox is selected, users can't change the value. The system calculates and displays the data for the field. **Note:** You can override this option for extended tables.
Audit	Enables or disables auditing for a table. Turning on Auditing (History) for a Table. **Note:** This option only applies to tables.
Text index	Determines whether searches index the text in a table. **Note:** This option only applies to tables. To exclude fields from indexing, see Remove an index for a specific field.
Column label	Defines a unique label for the column. The label appears on list headers and form fields for the column.

Dictionary Entry form (continued)

Field	Description
	• Updating the **Column label** field also updates the label in the language file (for the current language). • When you create a new column, the column name is populated automatically based on the label, which is prefixed with u_ to indicate that it is custom. For example, if you enter **Activity Description** as the column label, the column name defaults to **u_activity_description**.
Column name	Defines the field name of the column. When you create a new field, this name is populated automatically based on the label and a prefix as follows: • For a field on a table in a different scope, the name is prefixed with the scope to indicate that it is custom and not part of that application. • For a field on a table in the same scoped application, the name does not have a prefix, which indicates that it is part of the application. • For a field in a global application, the name is prefixed with u_ to indicate that it is custom. You cannot modify the prefix; however, you can modify the rest of the name. The name can contain only lowercase, alphanumeric ASCII characters and underscores (_). You cannot change the name of an existing dictionary record. **Note:** You cannot enter consecutive (double) underscores in a column name. For example if you enter `my__sample_name` as the column name, it is automatically changed to `my_sample_name` when you save the dictionary entry.]
Max length	Provides a logical limit for the size of string fields to determine how the system displays them in the user interface and how to map them to physical database data types. String fields with a length under 255 characters appear as a single-line text fields. String fields with a length over 254 characters appear as a multi-line text box. The system maps the field length to the closest physical data type available on the database. Sometimes, this results in more available length than originally specified. For example, entering a length of 50 maps to the closest physical data type of VARCHAR(100), which provides up to a 100 character limit or double the requested field length. Likewise, entering a length of 1000 maps to the closest physical data type of MEDIUMTEXT, which provides up to a 4000 character limit or four times the requested field length.

Dictionary Entry form (continued)

Field	Description
	Note: • You can only change this value for a **String** field. Changes for any other type of field are ignored. • Users on an Oracle instance cannot increase the maximum length of a string field to anything greater than 4000 through the application UI as this requires the CLOB datatype in Oracle. To increase beyond this size, log an incident with technical support to request the change. • To prevent data from being lost, only decrease the length of a string field when you are developing a new application and not when a field contains data. A warning appears if a change to a custom field results in data loss. For a base system field, you cannot make a change that results in data loss.
Mandatory	Determines whether this field must contain a value to save a record. For more information, see Make a field mandatory. **Note:** You can override this option for extended tables.
Display	Indicates that this field is the Display values for reference fields. Set this to true for the one field whose value you want to use as the text displayed in links to this table on lists and forms. By default, the **Number** field is the display value for all task tables. **Note:** • This option does not control whether a list or form displays this field as part of the layout. Instead, see List configuration and Personalizing Forms. • The display value becomes part of the form title when viewing an individual record from a table. • You can set a different display value on an extended table than the display value on a parent table by using a dictionary override.
Function definition	Defines the function that the field performs. Starts with `glidefunction:`, followed by the operation to be performed (for example, `concat`), followed by function parameters. Constants must be enclosed in single quotes. For example, the following function definition creates a field that shows the short description, followed by a space, followed by the caller name: `glidefunction:concat(short_description, ' ', caller_id.name)` **Note:** Function parameters support dot-walking. See Dot-walking to data in related tables.

Dictionary Entry form (continued)

Field	Description
Attributes [Advanced view]	Alters the behavior of a field or functionality that depends on the field. For more information, see Dictionary Attributes. Attributes can be overridden for extended tables with dictionary overrides. You can also configure attributes for this dictionary entry through the **Attributes** related list.
Default Value	
Use dynamic default [Advanced view]	Allows you to specify a default value that is generated dynamically based on a dynamic filter.
Dynamic filter value [Advanced view]	Specifies the Create a dynamic filter option that determines the default value if the **Use dynamic default** option is selected.
Default value	Specifies the default value of the field for any new record. Ensure that this value uses the correct field type. For example, an integer field uses a default value of 2 but cannot use a default value of two. These values can be overridden with dictionary overrides.
Reference Specification	
Reference	Makes the field into a Reference field type. If you enter a name that does not match an existing table, a new table is created when you save your changes to the dictionary record. If the current table has a module in the application navigator, a module for the new table is automatically created in the same application menu.
Use reference qualifier [Advanced view]	Specifies the type of qualifier to use: • Simple: A set of choice lists where you can specify a reference qualifier condition. • Dynamic: A Create a dynamic filter option that you can use to build the qualifier. • Advanced: A static encoded query string or JavaScript code that you can use to build the qualifier.
Reference qual condition	Specifies a condition when the reference qualifier runs if the **Simple** qualifier type is selected.
Dynamic ref qual [Advanced view]	Specifies the Create a dynamic filter option that determines the reference qualifier when the **Dynamic** qualifier type is selected.
Reference qual [Advanced view]	Filters the records available for a reference field if the **Advanced** qualifier type is selected. Reference qualifiers can be overridden with dictionary overrides.

Dictionary Entry form (continued)

Field	Description
Reference key [Advanced view]	Identifies a field other than sys_id to use as the unique identifier for reference fields.
Reference cascade rule [Advanced view]	Defines what happens to a record if the record it references is deleted. Select one of the following options: • Clear: Clears the references (default). • Delete: Deletes all referencing records. • Restrict: Prevents record deletion if there is a referencing record. • None: Does not change referencing records.
Reference floats [Advanced view]	Enables the **Edit** button on related lists for one-to-many relationships.
Dynamic creation [Advanced view]	For reference fields, determines whether entering a value that does not match an existing record creates a record on the referenced table. If selected, use the **Dynamic creation script** field to define how to create the record.
Dynamic creation script [Advanced view]	When the **Dynamic creation field** is selected, allows you to enter a script for creating a record on the referenced table.
Dependent Field	
Dependent on field [Advanced view]	Specifies a field on which the current field depends. For more information, see Make a field dependent **Note:** You can override this value for extended tables. .
Choice List Specification	
Choice	Allows users to see a list of suggested values in one of the following ways: • List menu without -- None -- • List menu with -- None -- • Suggestion field type If a choice is used, either define a Choice list field type or use the fields **Choice** table and **Choice** field to copy choices from another field elsewhere in the dictionary.
Choice table	Populates the field choices with the same values as another choice field. If the **Choice** field is set to anything besides **None**, select a table to draw choice values from. The field **Choice field** must also be populated.

Dictionary Entry form (continued)

Field	Description
[Advanced view]	For example, if **Choice table** is set to the Incident [incident] table, this field has the same choice list as one of the choice fields on Incident. **Choice field** (see below) determines which field.
Choice field [Advanced view]	Populates the field choices with the same values as another choice field. If the **Choice** field is set to anything besides **None**, select a field from the table you selected for **Choice table**. For example, if the **Choice table** field is set to the Incident [incident] table, and **Choice Field** is set to **Priority**, this field has the same choices as the **Priority** field on Incident, even if those choices change. **Note:** This field must be a choice field.
Calculated Value	
Calculated [Advanced view]	Determines whether the value of the field is calculated from other values. If selected, use the **Calculation** field to define how the calculation is performed. The **Calculation Type** field allows you to select script or formula based calculation for the column value. When filtering, sorting or grouping by a calculated field, the sort order is based on the stored field value from the last time the field was updated, not the last time the field appeared. **Note:** In relation to business rules, calculated fields are populated first before any business rule, even a before business rule, is run. Calculated fields are then populated again if necessary after any before business rules run.
Calculation Type [Advanced view]	Allows you to select script based calculation or formula based calculation for a column. For more information about predefined functions and example formulas, see Formulas for column values in Table Builder.
Calculation [Advanced view]	Allows you to enter a script or formula based on the selected Calculation Type. When the **Calculation Type** is **Formula**, you can enter a predefined function or formula for calculating the value of the field. You can use a predefined formula or combine other formulas to create a formula according to your requirements. When the **Calculation Type** is **Script**, you can enter a script for calculating the value of the field. These can be overridden for extended tables with dictionary overrides. You can use the current object in this script. As with access control rules, the script can: • Evaluate to true or false. • Return an answer variable set to true or false. • Set a field value directly, such as: `current.display_name="name"`. **Note:** Fields display as read-only when calculated scripts are applied. For example, by adding a script to the calculated field on a dictionary record that is set to Active, the Active field is made read-only on the form. Inline editing of the Active field from the list view is also disabled.

Dictionary Entry form (continued)

Field	Description
Additional fields	
Class	Identifies the table that the current table extends. Tables that do not extend other tables specify their own name in this field. For more information, see .
Defaultsort	Obsolete.
Size class	Determines whether the platform handles this table as a large table by reducing the amount of memory stored for each row during queries. There is a scheduled job which runs and sets the value of this field.
Spell check	Enables or disables spell check on the field.
Unique	Requires the field value to be unique. **Warning:** Making a field unique when the corresponding table already has different values for that field causes data loss. Before you enforce uniqueness on a field, verify that no records in the table for the field have values, or that they all have the same value.
UI action	
Default view or Advanced view	Changes the form view to the default or advanced view. The fields change based on the view. If you are using the default view, you must write a script to accomplish the same tasks that advanced view fields provided.
Related Lists	
Access Controls	Provides access to the access controls that permit or limit access to the data in the table.
Choices	Provides access to the options in the choice list field you are editing.
Dictionary overrides	Provides access to the dictionary overrides for this field.
Attributes	Provides access to the dictionary attributes for this entry.
Labels	Provides access to the labels used for the table or field you are editing.

Altering tables and fields using dictionary attributes

Dictionary attributes alter the behavior of the table or field that the dictionary record describes. Administrators can add or modify dictionary attributes.

Adding an attribute

To add or remove an attribute to a table or field, open a dictionary record, select the **Advanced** link, and modify the **Attributes** field. Alternatively, in a dictionary record, select **New** in the Attributes related list. For details on modifying dictionary entries, see Modify dictionary entries.

Attributes are comma-separated. If attributes exist on a dictionary record, add a comma with no spaces before adding an attribute.

For an attribute that accepts true/false values:

- To specify a value of true, you can either enter `attribute` or `attribute=true`.
- To specify a value of false, you can either ensure that the attribute does not appear or enter `attribute=false`. To maintain the values during upgrades, do not remove an attribute that is on a table by default.

Maintaining attribute values for upgrades

If you remove an attribute that is part of the base system, it is automatically restored during an upgrade. To prevent upgrades from changing the behavior of your system, leave the attribute on the table or field, but set its value as desired.

For example, if a field has the attribute `knowledge_search=true` by default, do not remove the attribute to set it to false; rather set it to `knowledge_search=false`.

Available dictionary attributes

Name	Value	Target Element	Description
allow_null	true/false	field_name field	If present or true, enables you to enter `None` as the field
allow_public	true/false	table_name field	If true, the table name field displays tables from all scopes instead of only the current scope.
allow_references	true/false	field_name field	If true, a tree is displayed to select from that includes reference fields so you can dot-walk.
allowHugeAlter	true/false	any table	If true, enables you to add a column to a table with more than 100 million rows.
approval_user	name of field containing the users for the approval type this field represents	integer field	Specifies the table fields that are used to perform lookups using a matcher. You specify approvals as table fields that have an attribute of *`approval_user=<field_name>`*, where <field_name> indicates the field in the table that contains the users for this approval type. • Fields with this attribute contain an integer value that indicates the sequence for the approvals. • All approval fields with the same sequence number indicate that multiple approvals are required before continuing. • Approvals are requested in the order of the sequence numbers. For example, all approvals with sequence number 100 must

Name	Value	Target Element	Description
			be approved before approvals with sequence number 200 are requested.
attachment_index	true/false	any table	If true, attachments on the table are indexed for search purposes. To learn more, see Index attachments on a table.
barcode	true/false	string field	Enables a string field to access a camera on a mobile device to scan and process a bar code.
base_table	name of base table type	table_name field	Specifies the table_name field that enables you to select any table that is derived from the table specified by this attribute. By default, the base table itself is also included in the list (but see `skip_root` to turn off this behavior).
calendar_elements	list of field names, separated by semicolons (";")	any calendar event table	Specifies the list of fields used when constructing the description of a calendar event. If not specified, the display name plus short description are used. The `calendar_elements` attribute does not support derived (dot-walked) fields.
collection_interval	interval specified as "HH:MM:SS" (like "01:02:30" for one hour, two minutes, and 30 seconds)	collection field	Specifies the interval of metrics collection.
close_states	inactive state integer values	task state field	Used by the `TaskStateUtil` API to identify the list of inactive state values delimited by semicolons (;)
critical	true/false	any field in the apm_application table	Defines fields that have critical information about an application, which enables tracking the entry of critical information.
current_location	true/false	string field	Enables a string field to access the GPS location of a mobile device.
disable_execute_now	true/false	any table derived from sys_auto	If present or true, disables the usual **Execute Now** button. The ServiceNow AI Platform applications that use schedules, such as Discovery, use this attribute to substitute their own more appropriate action.

Name	Value	Target Element	Description
default_rows	integer value	multitext fields	Sets the default number of rows in a multitext field.
default_close_state	state integer value	task state field	Used by the `TaskStateUtil` API to identify the default close state value for a task table.
default_work_state	state integer value	task state field	Used by the `TaskStateUtil` API to identify the default working state value for a task table.
detail_row	name of field to display in detail row	any table	Displays the value of the specified field as a Lists for each record in the list view. UI15 is required to use this attribute. **Note:** If different `detail_row` attributes are defined for a parent table and a child table, the system uses the child table attribute.
display_image	true/false	table with fields of the user_image type	If present or true, displays records of the user_image type in an image form. If the record is in the [sys_user] table, and the image is not available, the user initials are displayed.
document_viewer_audit	view, download, or all	any table	When the Document Viewer is enabled, this attribute keeps a record of the user actions that are performed on attachments. The values are as follows: • **view**: Enables an audit when users view an attachment • **download**: Supplies audit information for only downloaded attachments • **all**: Shows audit information when attachments are viewed and downloaded
email_client	true/false	any table	If present or true, causes an icon (an envelope) to appear in the more options menu in the form header. If you click it, a pop-up email client appears.
exclude_auto_recovery	true/false	any table	Disables automatic recovery of draft records for this table and its extensions.
exclude_from_rollback	true/false	any table	Excludes this table (and any of its descendant tables) from being

Name	Value	Target Element	Description
			recorded for rollback. For example, use this attribute to preserve records that are created from an automated test. For tables excluded by default, see Tables excluded from rollback after running an automated test.
extensions_only	true/false	any table	Specifies that the table should only have records in tables that extend it. For example, the Task table has this attribute because you would create incident, problem, change records, and not task records.
field_list_selector	true/false	any glide_list	Enables the user to select a field from the dependent table (or current if the dependent is not specified). Some workflow activities use this attribute.
field_decorations	UI Macro name list, separated by semicolons (";")	most fields (except list and multiline text fields)	Similar to `ref_contributions`, causes the named UI macro to be invoked when the field is rendered on a form.
format	format name	any numeric field	Specifies a named format to use instead of the standard numeric formatting. Options are as follows: • **glide_duration**: Formats a time specified in milliseconds as **ddd hh:mm:ss**. • **none**: Disables automatic number formatting (for example, changes 2,500–2500). **Note:** This attribute applies only to display values and not user input values.
glide.db.oracle.ps.query	true/false	any table	If present and false, prevents the use of Oracle prepared queries on the table.
global_visibility	true/false	any table with a sys_domain column	If present or true, makes this table visible globally even if there are domain restrictions (that is, the sys_domain field has a value).
hasLabels	true/false	any table	If present or true, marks this table as being the target of a label at some point. You can manually set this attribute, but it is normally set automatically whenever a label is generated. When true, the label engine runs on any change to the

Name	Value	Target Element	Description
			table, and updates the labels as needed.
hasListeners	true/false	any table	If present or true, marks this table as available for listeners to get events (insert, update, delete).
hasWorkflow	true/false	any table	Instructs the workflow engine to listen for changes to the table and fires events to a workflow when a record associated with a particular workflow has changed.
html_sanitize	true/false	any HTML or translated HTML field	If present or true, enables HTML sanitization for the selected HTML or Translated HTML field. By default, performs HTML Sanitization on all HTML and Translated HTML fields. Set this attribute to false to disable sanitization.
html_sanitize_config	name of custom sanitizer configuration file	any HTML or translated HTML field	If present, performs sanitization using the specified configuration file instead of the default platform sanitizer. For example, the Embedded Help [sys_embedded_help_content] table uses a specified file (EmbeddedHelpSanitizerConfig) for the help **Content** field.
icons	name of JavaScript class	any workflow field	Specifies a JavaScript class that produces workflow icons.
ignore_filter_on_new	true/false	any field	If true, prevents the fields on a form from populating with values from a filtered list.
image	relative path of image file	any table	Specifies an image file to use when a module or BSM map uses this table. This specification overrides the icons that would otherwise be used for the table.
include_container_types	true/false	any internal_type field	Causes the field to render with container (split) types as well other types.
ip_data_control	• none • canonical • expanded • canonicalize_when_possible (default value)	field type on any table	Checks whether the data should be normalized through the `ip_data_control` attribute. • If the attribute is none, no conversion happens and data is inserted as-is. • If the attribute is canonical, values are inserted in their smaller canonical form.

Name	Value	Target Element	Description
			• If the attribute is expanded, values are inserted in the larger expanded form to facilitate comparisons. • If the attribute is canonicalize_when_possible, values are inserted in canonical form when it is a valid IP or as-is when it is not.
isOrder	true/false	any field	Sets the field as the default field that is used for the sort order in lists. Overridden by the **ORDERBY** URL parameter, or the user's sort preferences.
iterativeDelete	true/false	any table	If present or true, forces all row deletes to execute iteratively. Otherwise, you can perform some deletes using a more efficient bulk method.
json_view	true/false	any table	if a JSON object is associated with the field, places an icon next to a field that, when you click it, displays the data interchange information.
knowledge_custom	name of JavaScript function	any field	Specifies a JavaScript function to implement a custom knowledge search (see *knowledge_search*).
knowledge_search	true/false	string fields	If present or true, causes a knowledge search icon (a small book) to appear next to the field. Clicking this icon launches a pop-up window for searching the knowledge base, unless a custom knowledge search function has been specified. To learn more, see *knowledge_custom*.
largeTable	true/false	any table	If present or true, marks this table as "large" to preventing table locking with specific MySQL database operations (adding/removing a column/index, compacting a table). Without this attribute (or the *smallTable* attribute), the *glide.db.large.threshold* property, or the default value of **5,000**, determines whether a table is large.
listen	true/false	any field	If present or true, calls a JavaScript function named *<tableName>_<fieldName>Listen*,

Name	Value	Target Element	Description
			or *globalListen* if that function does not exist. It calls the function with arguments (*tableName*, *fieldName*, *oldValue*, *newValue*).
live_feed	true/false	any field	If present or true, creates a toggle option on the activity formatter header for incidents, tasks, and problems. The toggle provides a choice between the **Live Feed** for that record (also known as a document feed) or the activity formatter fields already in use. To learn more, see Activity formatter for more details.
long_label	true/false	any reference field	Long or short labels refer to the label that appears for reference fields on a form. For example, if the reference field contains the caller's email address, the long label would be Caller Email, while the short label would be Email. • Usually the placement of the reference field on the form makes it clear what the field represents. • It uses the global property (*glide.short.labels*) to specify the type of labels that appear for all reference fields on any form. • You can override the global property for any reference field by setting the *short_label=true* or *long_label=true* attribute in the dictionary for the reference field.
maintain_order	true/false	any glide_list	If present or true, the up/down arrow order buttons appear to the right of the list of selected items. **Note:** You need to use the **Add multiple** button to see the order buttons.
mode_toggler	true/false	any composite_name field	If present or true, causes a name mode toggle icon (a small right-pointing triangle) to appear to the right of the label. Clicking this icon causes the field's rendering to change from a text field that accepts

Name	Value	Target Element	Description
			<tablename>.<fieldname> to a pair of reference choice boxes (one for the table, the other for the field). The latter is the default.
model_class	binary Java class name	any field of type glide_var	Specifies a model variable within the Java code. The model must have a class that implements the IVariablesModel interface.
model_field	see description	any field of type glide_var	Identifies a reference field in the record that has the model defined for it. For example, a workflow activity is associated with an activity definition. The activity definition has a related list of questions that make up the model for that activity definition. By using the activity_definition as the model_field for the activity, it builds the workflow activity model by reading the questions that are defined for the referenced activity definition.
nibble_size	positive integer	any table affected by the table cleaner.	Specifies the maximum number of records that the table cleaner can delete in a single operation. The default value for this attribute is **250**.
nibble_sleep	true/false	any table affected by the table cleaner.	If false, causes the table cleaner to perform cleanup operations without a pause between each operation.
no_add_me	true/false	any table	If present or true, hides the display of the Add Me icon that appears next to certain fields, such as the **Watch List** field.
no_attachment	true/false	any table	If present or true, prevents the attachment icon (a paperclip) from appearing on the form header.
no_attachments	true/false	any table	If present or true, attachments are not checked for and deleted when a record from this table is deleted. Meant for high-activity tables that never have attachments.
no_audit	true/false	any table	If present, this field is not audited, even if the table is being audited. The true or false value is ignored.
no_audit_delete	true/false	any table	If present or true, a sys_audit_delete record is never created when you delete a record from this table. Meant for high-activity tables

Name	Value	Target Element	Description
			that never need sys_audit_delete information.
no_auto_map	true/false	any table	If true, this field is not mapped during an import set. Primarily used for LDAP imports.
no_email	true/false	any glide_list field referencing sys_user	If present or true, the email box is removed from the glide_list field like the **Watch list** field.
no_multiple	true/false	any glide_list field	Hides the select multiple icons.
no_optimize	true/false	any table affected by the table cleaner.	If present or true, prevents the MySQL table compaction operation from running on the specified table. The table compaction operation typically runs after the table cleaner deletes at least 50 percent of the data in the specified table.
no_separation	true/false	any table	If present or true, marks this table as not participating in domain separation. To learn more, see Domain separation for service providers.
no_text_index	true/false	any field on a text indexed table, or any child table of a text indexed table	If a table is text indexed, the `no_text_index` field attribute prevents this field or child table from being included in the text index.
no_truncate	true/false	any string field	In a list view, shows the entire text value of the multitext value in a list, without truncating it. Without this attribute, the string is truncated based on the UI property **Number of characters displayed in list cells** which is **40** by default.
no_update	true/false	table	True for tables in which records are inserted or deleted but not updated. Prevents the system from creating sys_mod_count, sys_updated_by, sys_updated_on fields in the table when created. Does not stop updates of the table. This attribute setting is used to save space on high volume system tables, such as syslog and sys_audit.
no_view	true/false	any glide_list field	Hides the view-selected item icon.
omit_sys_original	true/false	price and currency type fields	When true, avoids retrieving the original value of a field during calculations. Omitting the

Name	Value	Target Element	Description
			sys_original value from being referenced during calculations can reduce performance issues when dealing with currencies or prices.
onlineAlter	true/false	any table	Tables with the `onlineAlter` attribute perform MySQL database operations using online schema changes. • Online schema changes provide a lock-free table upgrade when adding, modifying, or removing columns and when adding or dropping indexes. Without online schema changes, these changes to the database lock write access during execution. • Online schema changes use additional system resources. Oracle databases do not lock tables by default and do not use online schema changes.
order	numeric value	model variable fields	Used internally only (for model variables).
pop-up_processor	binary Java class name	any field or table	Specifies a custom pop-up processor for processing the field (or all fields in a table).
readable	true/false	any conditions field	When true, causes rendering of the conditions field in any list view as a human-readable condition (instead of the encoded query that is stored in the database).
ref_ac_columns	list of field names separated by semi-colons	any reference field with an auto completer (see ref_auto_completer)	Specifies the columns whose display values should appear in an auto completion list, in addition to the name. To learn more, see the cmdb_ci field (**Configuration Item**) on the Incident form for a working example.
ref_ac_columns_search	true/false	any reference field with an auto completer (see ref_auto_completer)	Causes auto-complete to work with all fields that are specified in the `ref_ac_columns` attribute. This attribute overrides the default behavior, which searches only the display value column. To learn more, see Configure auto-complete to match text from any reference field.
ref_ac_display_value	true/false	any reference field with an auto	Causes the reference field to hide the display value column so that the auto-complete function

Name	Value	Target Element	Description
		completer (see ref_auto_completer)	only matches the text from the columns that are listed in the `ref_ac_columns` attribute. This feature requires the use of the `AJAXTableCompleter` class and the `ref_ac_columns`, `ref_ac_columns_search`, and `ref_ac_display_value` attributes. To learn more, see Remove the display value column. **Note:** The ref_ac_display_value function does not work with Catalog Item variables.
ref_ac_order_by	field name	any reference field with an auto completer (see ref_auto_completer)	Specifies the column that is used to order the auto completion list.
ref_auto_completer	JavaScript class name	any reference field (can be applied to a table to affect all reference fields on the table.)	Specifies the name of a JavaScript class (client side) that creates the list auto completion choices. Valid class values include: • `AJAXReferenceCompleter`: Matching auto-complete choices appear as a list. Only the Display values column of the reference table appears. If there is no other auto-completion class specified, reference fields automatically use this class. • `AJAXTableCompleter`: Matching auto-complete choices appear as rows in a table. The display value column of the reference table appears, with any columns listed in the `ref_ac_columns` attribute. • `AJAXReferenceChoice`: Displays matching auto-complete choices as a list. Only the display value column of the reference table, and up to 25 matching choices, appear. If there are more than 25 auto-complete choices, the choices for the `AJAXTableCompleter` class appear for the reference field instead.

Name	Value	Target Element	Description
			For more information, see Auto-complete for reference fields.
ref_contributions	UI Macro name list, separated by semicolons (";")	any reference field	When the field renders on a form, it invokes the named UI macro.
ref_decoration_disabled	true/false	any reference field	When set to `true`, disables the display of the reference icon on a selected field.
ref_list_label	label text	any table	Specifies the title to use in a list banner.
ref_qual_elements	field name list, separated by semicolons (";")	any reference field with a reference_qual field	Specifies a list of fields to send back to the server to get an updated reference.
ref_sequence	list of fields in referenced table, separated by top hats ("^")	any reference field	Specifies the fields in the referenced table that are used to order the list. This attribute works like an ORDER BY clause in SQL, with each element in ascending order.
reference_types	list of valid reference types that are clickable separated by semicolons (";")	field_name field	Limits the reference fields that appear in the tree to the specified types.
remoteDependent	name of database and table (like "model.matcher")	any script field	Defines the remote (such as in another database) table that the script depends on.
repeat_type_field	field name	a repeat count field for schedule rotation	Specifies the field that contains the repeat type (daily, weekly, monthly, or yearly).
ro_collapsible	true/false	any multiline field	If present or true, causes an icon (either a "+" or a "-") to appear next to the field label, allowing the field itself to be expanded or collapsed.
saver_exempt	• exempt_always= field will be skipped from writing to xml for all install contexts. • exempt_never= field will not be skipped from writing to xml (Default). • exempt_vcs_only= field will be	Fields of metadata records	When present, determines if the field is written to the xml representation of the record. An absence of the attribute is equivalent to exempt_never. When exempt_always, the field is skipped from writing to xml when exporting (publish, commit to source control, update set). When exempt_vcs_only, the field is skipped from writing to xml only in the source control context.

Name	Value	Target Element	Description
	skipped from writing to xml only in source control context.		
scale	integer	decimal field	Sets the number of decimal places to use on a Decimal field type. The default is 2. This attribute is applied to the max length attribute setting for the field. The decimal places for a field are a combination of the max length attribute setting and the scale attribute. For example, if you want four decimal places for the field, you need a max length attribute setting of 4, plus the default of 15 or 19. You would set the scale attribute to 4 to indicate the number of decimal places. Six decimal places require a max length setting of 21, a scale attribute setting of 6, and so on. **Note:** Increase the max length attribute to a value greater than 15 to increase this attribute.
script	a function that returns the contents of the field	any slushbucket field	Enables you to write a script to define what is loaded into the slush bucket field.
short_label	true/false	any reference field	Long or short labels refer to the label that appears for reference fields on a form. For example, if the reference field contains the caller's email address, the long label would be Caller Email, while the short label would be Email. • Usually the placement of the reference field on the form makes it clear what the field represents. • The global property (*glide.short.labels*) specifies the type of labels that appear for all reference fields on any form. • You can override the global property for any reference field by setting the *short_label=true*

Name	Value	Target Element	Description
			or *long_label=true* attribute for the reference field in the dictionary
show_all_tables	true/false	document ID fields	Enables users to select documents from system tables. For example, sys_script or sys_user. By default, users cannot select records from system tables.
show_condition_count	true/false	condition fields	Enables or disables the condition count widget to preview how many records are a set of conditions. To learn more, see Add the condition count to a condition field.
skip_root	true/false	table_name field	If present or true, removes the base table from the list. To learn more, see base_table.
sla_basis	list of table names separated by semicolons (";")	any field of date type (glide_date_time, glide_date, due_date, date, or datetime)	Defines the tables for which this field determines the start (open) time of an SLA.
sla_closure	list of table names separated by semicolons (";")	any field of date type (glide_date_time, glide_date, due_date, date, or datetime)	Defines the tables for which this field determines the ending (close) time of an SLA.
slushbucket_ref_no_expand	true/false	any reference field	If present or true, prevents users from expanding the field from a form or list slush bucket.
smallTable	true/false	any table	If present or true, marks this table as Small (that is, not large) for querying purposes. Without this attribute (or the *largeTable* attribute), the *glide.db.large.threshold* property, or a default value of **5,000** determines whether the table is large.
start_locked	true/false	any glide_list field	Determines whether the field is locked or unlocked by default. Set the value to false to unlock the field by default.
staticDependent	name of table	any script field	Specifies the table that the script depends on.
strip_html_in_pdf	true/false	any field	Attempts to remove HTML tags from a field when that field is exported to a PDF. This attribute is most useful on HTML fields.

Name	Value	Target Element	Description
synch_attachments	true/false	any table	Similar to *update_synch* but writes the file attachments of the record to update sets. To learn more, see Index attachments on a table.
table	name of table	field_name field	Displays the specified table fields.
tableChoicesScript	name of script include	table_name field	Specifies the name of a script include whose `process()` method returns an array of table names from which to select.
target_field	percent_complete_target decimal field name	any percent complete field any decimal field	Works in combination with the target_threshold_colors attribute to define thresholds for percent complete fields. See the description for target_threshold colors. Example: `target_field=percent_complete_target,target_threshold_colors=0:tomato;50:khaki;90:lightgreen` Specify an additional target_field attribute to compare the actual completion percentage of a task with a target percentage in a different decimal field that specifies where the task should be at this point. Use the column name of the decimal field. Example: `target_field=percent_complete_target,target_threshold_colors=0:tomato;50:khaki;90:lightgreen,target_field=u_expected_completion` If you do not specify an additional target_field, a target value of 100 is used, allowing you to use the color thresholds with a single field value. See Target threshold colors attribute.
target_form	name of form	any table	Specifies the alternative form to be used when this table is referenced through a pop-up window on a reference field.
target_threshold_colors	threshold limit:color	any percent complete field with the	Works in combination with the target_field attribute to define thresholds for percent complete

Name	Value	Target Element	Description
		target_field attribute	fields. See the description for target_field.
text_index_filter_junk	true/false	any table	Sets the value to false to Enable or disable the Zing junk filter for the table. • By default, Zing does not index or search for two-digit numbers and single character words (unless they are Chinese or Japanese characters). • Regenerate the index after disabling the junk filter. This attribute results in a larger table index. • For optimal performance, do not apply it unless it is required.
text_search_only	true/false	table_name field	Limits the tables listed to those tables that are searchable by text.
text_index_translations	true/false	any table	If present or true, forces indexes to be recalculated when translated strings are added. You must have a sys-admin role to modify this attribute. This one is automatically set for indexed fields that are translated, and to fields that have a translation and are being indexed. The `glide.i18n.force_index` `system` property, which defaults to true, overrides this attribute.
time_zone_field	name of field containing the time zone	any schedule date/time field	Specifies the field in the parent record that contains the reference time zone for this field.
timeDimension	true/false	any field of date type (glide_date_time, glide_date, due_date, date, or datetime) in a table subclassed from the task table	If present or true, enables the production of time dimension data for use by OLAP (to produce reports based on quarters, weeks, or other time periods). **Note:** OLAP functionality has been deprecated.
tree_picker	true/false	reference field to a hierarchical table	Displays the hierarchy of reference values in a tree display (such as locations).

Name	Value	Target Element	Description
			Note: The tree display hierarchy does not appear in reference fields when you are using mobile applications.
ts_weight	integer value	any field	Controls the relative importance of a match in the field for a text search. To learn more, see Set the relative weight of a field.
types	list of valid element types separated by semicolons (";")	field_name field	Limits the display of fields to the specified types.
update_exempt	true/false	any field from a table that has update_synch set to true	If present or true, you can change this field without skipping updates to the rest of the record. During software upgrades, the value of this field is preserved, while the rest of the record receives upgrades. By default, the Active field on a tracked table is treated as update_exempt even if the attribute is not present.
update_synch	true/false	any table	Indicates that changes in the table are tracked in update sets. Administrators cannot modify this attribute. To migrate data, use an Importing from another ServiceNow instance.
use_document_viewer	true/false	any table	If present or true, enables users to open supported attachments in a document viewer within the platform, rather than downloading the files directly to their own file system.
use_workflow	true/false	any table that has delivery plans or uses workflow	If present or true, causes a workflow to be used instead of delivery plans.
user_preference	true/false	any field	If present or true, causes any user preferences to be used instead of the normal default value.

Modify the Glide durations format

To convert fields that are displayed in milliseconds (such as 'Resolution Time' on the Incident table) to a duration format (Months/Days/Hours/Seconds), populate the attribute field on the dictionary with: `format=glide_duration`.

Now the value is displayed as a duration in lists and can be utilized in the Configure Calculations functions. This may be useful to illustrate ITIL KPIs such as Mean Time To Repair.

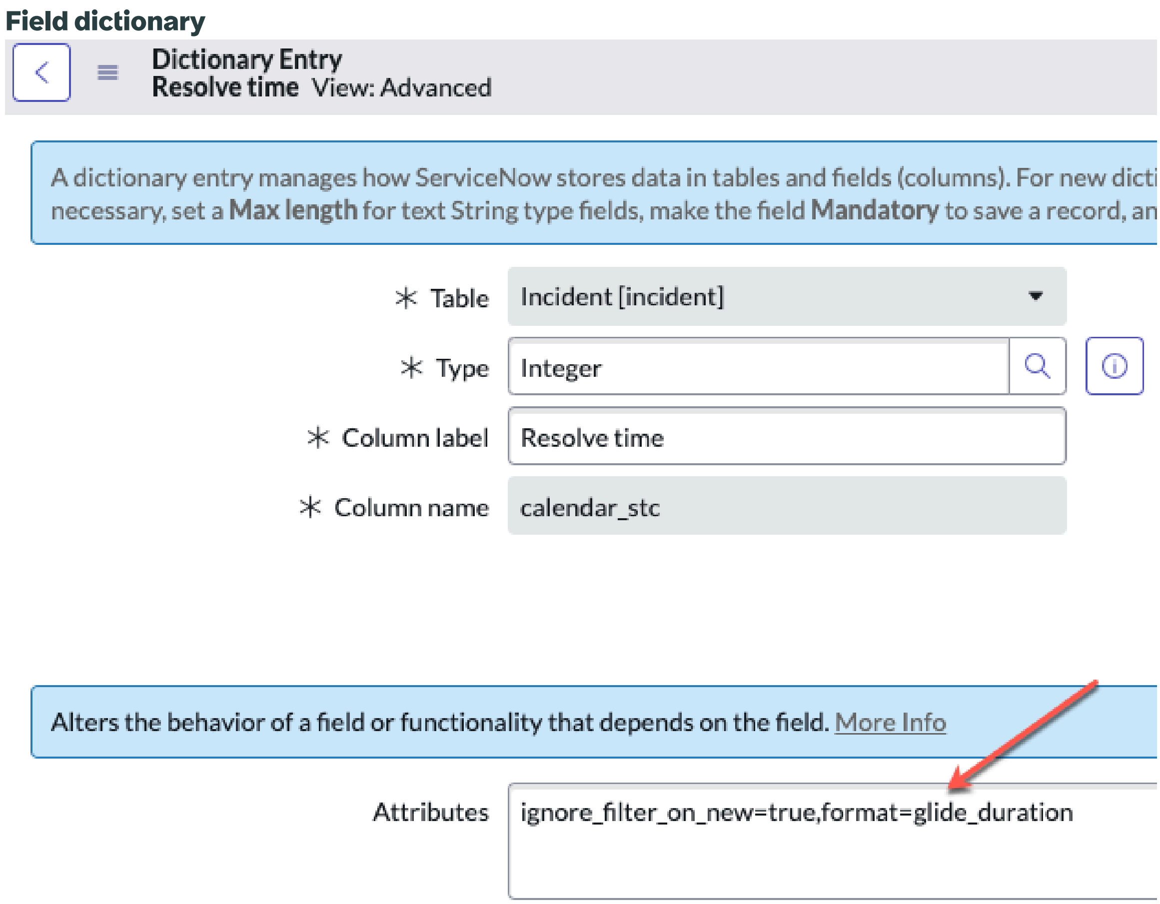

Results in list:

Results in list

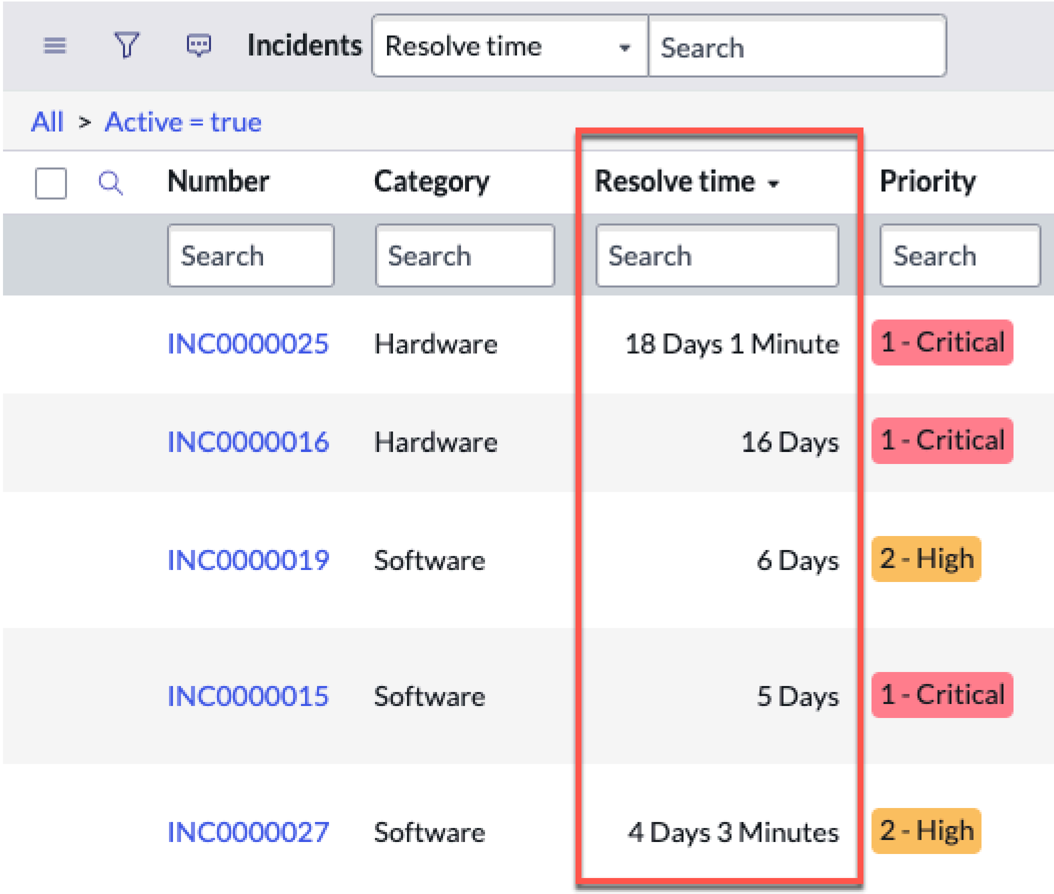

Dictionary overrides

Dictionary overrides provide the ability to define a field on an extended table differently from the field on the parent table.

For example, for a field on the Task [task] table, a dictionary override can change the default value on the Incident [incident] table without affecting the default value on Task [task] or on Change [change].

Administrators can override these aspects of a field:

- Reference qualifiers
- Dictionary attributes
- Default values
- Calculations
- Field dependencies
- Default column display values
- Mandatory and read-only status

Define a dictionary override

Use a dictionary override to allow a field in a child table to have a different value or behavior than the same field in a parent table. For example, a dictionary override changes the default value of the priority field from 4 in the parent table to 5 in the Incident table.

Before you begin

Role required: admin.

About this task

Dictionary overrides are only available for tables that support table extension.

Dictionary overrides are inherited by extended tables, so it may be necessary to define additional overrides to change the value of fields on extended tables. For example, if you define the default value of cmdb_ci.install_status to be 3, and create an override for the same field on cmdb_ci_hardware to be 5, all tables extended from the Hardware table will also default to 5. So if you want the default cmdb_ci_computer.install_status to be 3, you will need another override.

You can add a dictionary override in a scoped app, but you must add the override while working in the scoped app context and you can only select child tables from within that scope. Open the scoped app, edit the field's dictionary entry, add a new dictionary override, select the child table, and then provide the override details.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the record for the field.
3. In the **Dictionary Overrides** related list, click **New**.
4. Fill in the fields on the form, as appropriate (see table).

Field	Description
Application	The application scope to which the record belongs.
Base table	The parent table containing the field to override.
Table	Select the extended table to which the dictionary override applies. **Note:** The list shows only tables and database views that are in the same scope as the dictionary override.
Override reference qualifier	Select the check box to display the **Reference qualifier** field, which overrides the Reference qualifiers for the field on the extended table.
Override dependent	Select the check box to display the **Dependent** field, which overrides the field on which the current field depends.
Override attributes	Select the check box to display the **Attributes** field, which overrides the dictionary attributes for the field on the extended table.

Field	Description
	Note: Any attributes defined on the base table are ignored. If there are attributes on the base table that should still apply to the extended table, make sure to include them in this field.
Override default value	Select the check box to display the **Default Value** field, which overrides the default value for the field on the extended table.
Override calculation	Select the check box to display the **Calculation** field, which overrides the calculation of the value for the field on the extended table.
Override mandatory	Select the check box to display the **Mandatory** field, which overrides whether the field on the extended table must contain a value to save a record.
Override read only	Select the check box to display the **Read only** field, which overrides whether a user can change the field value on the extended table.
Override display value	Select the check box to use this field as the display value on the extended table. For example, the Story [`rm_story`] table uses the short description as the display value in reference fields instead of the number, as defined in the Task [`task`] table. **Note:** Only fields that exist at the task table can be used as display value overrides.

5. Click **Submit**.

Table flattening

Table flattening stores a hierarchy of related tables as one table in a relational database.

Extension models

The system offers these extension models to store a table hierarchy on a relational database.

Available extension models

Extension model	Flattens tables?
Table per class	No
Table per hierarchy	Yes
Table per partition	Yes

Table per class

The Table per class extension model stores each table of the hierarchy in its own physical table on the relational database. Each physical table uses the table prefix of the source table each stores a different class of records. An example of the Table per class extension model is the Asset [alm_asset] table, and its child tables: Hardware [alm_hardware], Consumable [alm_consumable], Facility [alm_facility], and Software License [alm_license]. The parent table of the hierarchy, Asset, stores a copy of every record in its descendant tables.

To find records in the Table per class extension model, the system queries records from multiple tables and joins the results. For example, when searching for hardware in a related facility, the system must join results from the Hardware, Facility, and Asset tables.

Table joins cause a performance bottleneck on relational databases. The more classes a query includes, the worse the query performance. Therefore any query for records from the top of the table hierarchy has the worst performance because it requires joining all descendant tables.

The system uses the Table per class extension model by default when creating tables. Most system tables also use the Table per class extension model as there is no performance benefit from flattening them.

Table per hierarchy

The Table per hierarchy extension model stores an entire table hierarchy in a single flat physical table on the relational database. The physical table is named after the parent table of the hierarchy, such as Task. The physical table contains all records of the table hierarchy and assigns a class name column value to each descendant table of the hierarchy. The system uses the name of the source table as the class name value. For example, Task records can have class names such as Change, Incident, or Problem.

To find records in a table hierarchy, the system queries the physical table and uses the class name column to constrain the results. Since such queries do not require joining results from multiple tables, the system provides better search performance.

The system uses the Table per hierarchy extension model for the Task table hierarchy on MySQL databases. Other tables use the Table per class extension model because there is no performance benefit to flattening them. To use Table per hierarchy on an Oracle database, contact Technical Support.

Table per partition

The Table per partition extension model stores an entire table hierarchy in a single flat logical table on the relational database. Each logical table can have multiple physical storage tables called partitions supporting it. Each partition optimizes the database resources available to a physical table such as the column count, index count, and row size. The system adds a partition whenever the logical table needs additional relational database resources.

Each logical table is named after the parent table of the hierarchy, and each supporting physical partition consists of the logical name plus a partition name. For example, the Base Configuration Item [cmdb] table starts as a logical table with no partitions. Suppose your hardware configuration items consume enough database resources that the system creates a partition called **cmdb$par1** to store them. Later, computer configuration items could consume enough database resources to warrant the system creating a second partition called **cmdb$par2** to store these records.

Within each logical table, the system assigns a class name column value to each descendant table of the hierarchy. For example, within the Base Configuration Item logical table there are records with class names for Application, Computer, and IP Router. The system also assigns a two-digit class path value to each descendant table of the hierarchy. The class path is based on the table location in the hierarchy. For example, the parent class Hardware might have a class path such as **/!!/!D** and the child class Computer might have a class path such as **/!!/!D/!!**.

To find records in the Table per partition extension model, the system queries the logical table and its partitions and uses the class path column to constrain the results. Since these queries do not require joining results from multiple tables, the system provides better search performance. In addition, the class path reduces the total number of records to search, which further improves search performance.

The system uses the Table per partition extension model for the Base Configuration Item [cmdb] table hierarchy on MySQL databases. To use Table per partition on an Oracle database, contact Technical Support.

View a table hierarchy and the extension model

The extension model used by a table is not immediately apparent. While a hierarchy can use a single physical table, the platform displays tables as if each logical table has a unique physical table. For example, when specifying a table for a workflow, you can select Change [change_request] or Incident [incident] even though the parent table, Task [task], uses a single physical table.

Before you begin

Role required: admin

About this task

Administrators can view the status of flattened table hierarchies, but cannot flatten additional hierarchies. You can configure the form to add the **Extension model** field if necessary.

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Select a table record.
3. Review the **Extension model** field value, which indicates whether the table hierarchy uses multiple unique tables or a single flat table.
   - **None**: Indicates that the table uses the table per class model. Defines a unique physical table per logical table in a hierarchy. For example, there are separate physical tables for the Asset [alm_asset] table, and its child tables: Hardware [alm_hardware], Consumable [alm_consumable], Facility [alm_facility], and Software License [alm_license].

Note:
If a table has an empty Extension Model field, it is not necessarily a Table Per Class table. If the table extends another table, it actually inherits the model from any parent tables it is extending, and only when none of the parents have a defined extension model is the table's extension model Table Per Class. This can be seen on many tables that extend the Task table. Incident is a logical table that extends Task and is physically located on Task, but by default has a blank Extension Model value.

When a table doesn't extend anything at all, and has a blank Extension Model field, it is a Table Per Class table.

- **Table per hierarchy**: Defines a single physical table per logical table hierarchy. For example, there is a single Task [task] table on the MySQL database that contains all the Task, Incident, Problem and Change records. This single physical table is represented as separate logical tables.

Viewing table references and extensions

View table relationships in a visual manner using the schema map.

The schema map displays the details of tables and their relationships visually, allowing administrators to view and easily access different parts of the database schema.

The schema map can also be printed directly from a browser.

Table relationships depicted in the schema map

Relationship type	Color used
Referenced by	Red
Referencing	Orange
Extended by	Green
Extending	Blue

By default, each type of relationship is displayed. However, you can view or hide each type as needed.

Generate a schema map

Generate a schema map to view different parts of the database schema.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Tables & Columns**.
2. In the **Table Names** pane, select a table.
3. Click **Schema map**.

   The schema map for the selected table opens in a separate tab or window.

   If required, you can print this map from the browser.

Analyzing table relationships

The schema map shows the selected table in yellow, typically centered, and all tables related to that table, typically shown at the sides.

Example schema map

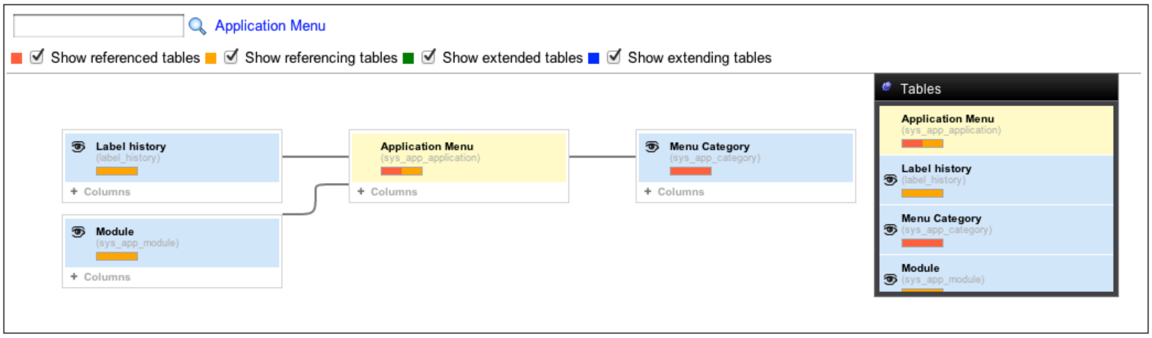

From this map:

- The check boxes at the top allow you to control which relationships to display. Select or clear a relationship type to display or hide tables with that relationship to the selected table.
- Each related table has a colored bar indicating the relationship to the selected table.
- You can point to the connector lines to display the details of a relationship between the two tables.

Note:
Since relationships are shown as single lines for simplicity, the diagrams rendered are not entity relationship diagrams.

Using the Table Selector:

To view a schema map as a list, point to the table selector in the right

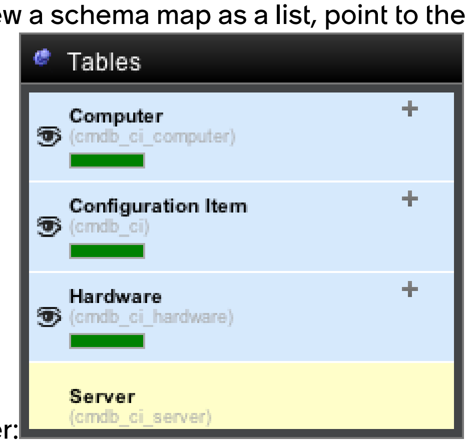

corner:

You can:

- Click a table in the list to scroll the schema map to that table.
- Click the eye icon beside a listed table to hide or show that table in the schema map.
- Click the pin icon in the selector to keep the list open.

Using Related Tables:

Right-click a table node header to display a context menu with these functions:

- **Focus on this table**: Make the selected table the new focus table and redraw the schema map based on the new selection.

  The new focus table is added as a breadcrumb at the top, allowing you to return to the previous table at any time.
- **Go to list**: Display the list of records for the table.
- **Go to dictionary**: Display the system dictionary, filtered for the selected table.

To hide a related table from view, click the eye icon in the node header (the node can be made visible again with the table selector).

For tables that are part of their own derivation hierarchy, click the expand button (**+**) in the node header to add their derivation hierarchy to the schema map.

Viewing More Information

Click the expand button (**+**) beside **Columns** to expand the table fields.

Columns -	
Business resolve time:	integer
Caller:	reference to User
Category:	string
Caused by Change:	reference to Change Request
Change Request:	reference to Change Request
Child Incidents:	integer
Close code:	string
Incident state:	integer
Notify:	integer
Parent Incident:	reference to Incident
Problem:	reference to Problem
Reopen count:	integer
Resolve time:	integer
Severity:	integer
Subcategory:	string

The reference fields show a red notation of the table they refer to.

If any tables extend from a table, their columns are displayed in reverse derivation order. For example:

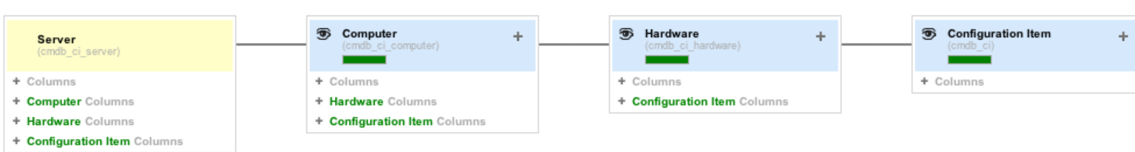

Here, the Server `[cmdb_ci_server]` table extends from Computer `[cmdb_ci_computer]`, Hardware `[cmdb_ci_hardware]`, and Configuration Item `[cmdb_ci]`, and displays the columns from those tables.

Similarly, the Computer table displays the columns from the Hardware and Configuration Item tables.

Create a table index

Build indexes to access the data held in your tables more easily.

Before you begin

Role required: admin

About this task

Note:
Constructing an effective index requires specialized knowledge in database architecture. If you don't have this expertise, consult with a database administrator.

Creating an index on a large table can significantly affect system performance. A large table in this case is generally considered to be any table around 1 TB in size. Therefore, consider creating indexes on larger tables only during off-peak hours. For example, avoid adding an index to a CMDB table while a discovery job is running, as this can lead to increased load and potential delays.

Procedure

1. Access the Index creator in one of the following ways.
2. Select the fields you want included in the index.
   The order in which you select the fields affects how the index works. If you do not have expertise in database design, you should consult someone who does.
3. To create a unique index, select the **Unique Index** check box.
4. Select **Create Index**.
   The **Table Name** field is there for your reference only. Overriding the default has no effect.

Drop a custom index

Remove a custom index by dropping it from a table.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Select the table that contains the custom index you want to drop.
3. On the **Table** form, select the Database Indexes related list.
4. In the list, select the custom index that you want to drop.
5. On the Database Indexes form, select **Drop**.

Result

The custom index is dropped from the table.

Create a many-to-many table relationship

Create a bi-directional relationship between two tables, so that the related records are visible from both tables in a related list.

Before you begin

Role required: admin.

About this task

Many-to-many relationships allow a list to point to a list of entries, rather than to single field.

If a knowledge base article points to a list of related configuration items, it uses a related list. Not all lists can be used as related lists, however. For a list to be related to another list, a many-to-many relationship that refers to both tables must exist.

The Many to Many Definitions `[sys_m2m]` table allows administrators to create custom many-to-many relationships.

Many-to-Many tables aren't considered custom tables and don't count toward your custom table allotment.

Procedure

1. In the navigation filter, enter `sys_m2m.list`.
2. Click **New**.
3. In the **From table** field, specify a parent table.
4. In the **To table** field, specify a child table.
   The Many-to-Many form automatically populates the other fields with suggested values.
5. **Optional:** Edit other field values, if appropriate.

   > **Note:**
   > Many-to-Many table names cannot exceed 30 characters.

Related topics

Reference default many-to-many relationships

Reference default many-to-many relationships

Some many-to-many relationships are defined by default.

To reference many-to-many relationships that are available in the base system, administrators can enter `sys_collection.list` in the navigation filter.

> **Note:**
> Only use this table to view many-to-many relationships in the base system. To create a new relationship, always use the Many-to-Many Definitions table.

Working with the Task table

The Tasks [task] table is one of the core tables provided with the base system.

The Tasks [task] table provides a series of standard fields used on each of the tables that extend it, such as the Incident [incident] and Problem [problem] tables. In addition, any table that extends the Tasks [task] table can take advantage of task-specific functionality for driving tasks.

The Planned Task plugin provides the Planned Task [planned_task] table, which extends the Tasks [task] table to provide more fields for tasks to measure duration and effort.

Creating tasks

The Tasks [task] table provides means for creating and completing tasks, as well as tracking their workflow in your instance.

Tasks aren't created directly on the Tasks [task] table. Instead, tasks are created on task child tables. To create a record in a task child table, access the Tasks [task] table by entering `task.list` in the navigation filter, and then select **New**. In the Task Interceptor, select the desired task child table.

Tools for completing tasks

These tools can be run on any table which extends Task.

- Approvals

  Approvals can be generated to a list of Approvers, either manually or automatically, according to Approval Rules. Approvals can be incorporated into workflows or can stand alone. For more information, see Process approvals.

  Approvals can be used on tables that do not extend Task.

- Assignments

  Assignment rules can automatically assign tasks to users or groups, ensuring that the most appropriate team members handle the tasks. For more information, see Defining Assignment Rules.

- Service levels

  Service level agreements can track the amount of time that a task has been open, to ensure that tasks are completed within an allotted time.

- Inactivity monitors

  Inactivity monitors ensure that tasks do not fall by the wayside by notifying users when tasks have been untouched for a predefined period of time. For more information, see Set an inactivity monitor.

- Flow Designer

  A process owner can design a flow for tasks that meet certain conditions. After a task is created that meets the conditions, the flow applies an automated process to the task. The process is defined in Flow Designer. For more information, see Flow Designer.

Modifying the Task table

Modifications made to the Tasks [task} table are applied to all child tables.

Be sure that changes are appropriate for all the child tables. Adding fields is a low-impact change, because you can hide fields on tables that don't need them. However, if the fields are being used across tables, deleting fields may cause unwanted data loss.

Note:
When adding choice list entries to a choice list on the Tasks [task] table, make sure that the entry value is unique.

You can use dictionary overrides to change some parts of a field definition in a way that does not to not apply to all child tables.

Tasks workflow

An administrator can specify a specific workflow process to apply to tasks that meet certain conditions.

Once a task is created that meets the conditions, the workflow applies the process to the task, asking for approvals, notifying users, generating other tasks, running scripts, and so on.

Task table structure

The Task table structure provides a framework to organize and store task-related data, and includes extensive customization options.

Important Task table fields

The Task table is a base class that provides fields for the core ITSM applications such as Incident, Problem, and Change Management. All applications that extend the Task table share these fields in common.

Label	Name	Type	Description
Active	active	boolean	Specifies whether work is still being done on a task or whether the work for the task is complete. By default, only application-specific business rules set the value of this field. For example, the `incident autoclose` business rule closes resolved incidents that have not been updated for one day.
Additional comments	comments	journal_input	Displays and allows the entry of comments about the task record. Each comment is inserted into the **Activity** field. For more information, see Journal field type.
Approval History	approval_history	journal	Displays the history of approvals for the record. For more information, see Tools for completing tasks.
Assigned To	assigned_to	reference	Specifies the user assigned to complete the task. By default, this field uses a Reference qualifiers to only display users with the itil role. Some applications override the parent reference qualifier to display relevant users. For example, the Service Order and Project Task tables filter users based on their skills.
Created	sys_created_on	glide_date_time	Displays the date and time when the task record was created.
Description	description	string	Displays and allows the entry of a multi-line description of the work to be done.
Escalation	escalation	integer	Indicates how long the task has been open. Escalations are dynamically populated using service level agreements, which specify how long a task remains in each escalation state. Escalation states go from **Normal** to **Moderate** to **High**, and finally to **Overdue**. Record lists color code each task by escalation state.
Number	number	string	Displays an identifying number for each task record. This field is the display value for

Label	Name	Type	Description
			the Task table. The system generates this number when the task is created. To manage number generation, see Record numbering.
Opened	opened_at	glide_date_time	Displays the date and time when a human opened the task record for the first time.
Priority	priority	integer	Specifies how high a priority the task should be for the assignee. By default, the *calculatePriority*business rule calculates this value based on the **Impact** and **Urgency** values. Record lists color code each task by degree of priority.
Short Description	short_description	String	Displays and allows the entry of a short description of the task, which is a human-readable title for the record.
State	state	Integer	Displays a choice list for status of the task: • Pending • Open • Work in Progress • Closed Complete • Closed Incomplete • Closed Skipped Applications typically use Dictionary overrides to display application-specific states.
Sys ID	sys_id	GUID	Displays the Unique record identifier (sys_id).
Task Type	sys_class_name	sys_class_name	Specifies the type of task, which corresponds to the child class. The system populates this value when a record is created on a child table.
Time Worked	time_worked	timer	Display a timer which measures how long a record is open in the form view.
Watch list	watch_list	glide_list	Specifies users who receive Email and SMS notifications when the record is updated. By default, only the Incident, Change, and Service Catalog applications notify users listed in this field. For other Task-based applications, you must create custom email notifications sent to the users listed in this field.
Work notes	work_notes	journal_input	Displays and allows the entry of comments viewable only by ITIL users. Each comment

Label	Name	Type	Description
			is inserted into the **Activity** field. For more information, see Journal field type.
Work notes list	work_notes_list	glide_list	Specifies users who receive Email and SMS notifications when work notes are added to the record. By default, only the Change, Problem, and Service Catalog applications notify users listed in this field. For other Task-based applications, you must create custom email notifications sent to the users listed in this field.

Journal fields

Journal fields work together to create a log of changes and comments as tasks are worked on.

Note:
Journal fields work on audited tables only.

Fields of the journal_input type are multi-line text boxes which, upon save, add the comments into the Activity field with a notation.

Fields that accept input into the journal

Field	Description
Additional comments	Any user can update.
Work notes	ITIL users can update.

These comments, as well as any changes to the record or email notifications sent out because of the record, are displayed in the activity formatter, which can be added to the form like a field.

Activity formatter for journal fields

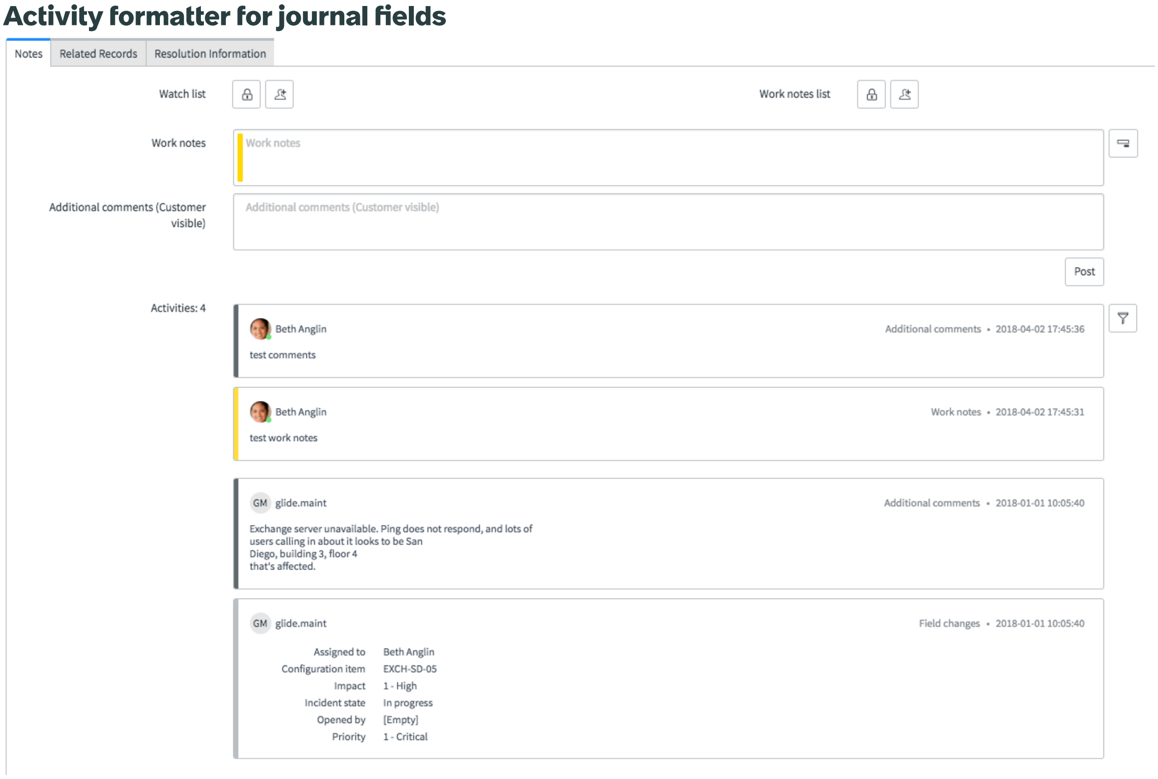

Reminder table

The Reminder [reminder] table provides a way to auto-generate reminders for a task.

Any table that extends the Task table, such as the Incident [incident] table, can use the Reminder [reminder] table.

You can add Reminders as a related list to the Incident form by opening the form context menu, navigating to **Configure** > **Related Lists**, and then adding **Reminders->Task**.

By default, only an administrator can create or modify a Reminder record. To enable non-administrators to create reminders, you must add create and read ACL rules to the Reminder [reminder] table and specify the user's role.

Note:
You can only create reminders for active task records.

To access the Reminder table, type `reminder.do` in the filter navigator.

Reminder table

Field	Description
Task	Refers to the parent task record.
User	Specifies the user who has logged in to the system.
Remind me	Specifies the time before which the reminder must be sent. Available values are: **15 minutes**, **30 minutes**, **1 hour**, and **2 hours**. You can add values to this field as required. **Note:** Internally in the system, the value is always converted to minutes and then stored in the **Remind me** field.
Before	Refers to a date field on the parent task record. **Note:** The **Activity due** and **SLA due** fields are legacy fields with an **UNKNOWN** value that you can't use. You can only use the **Due date** and **Follow up** fields. Note the following details about using the **Before** field. • The parent task record must have a value in the field you specify. For example, if you select **Due date**, the parent task record must have a value in **Due date** before you can save the reminder. • You can use a date field other than **Due date** or **Follow up** by performing the following steps: 1. Select either **Due date** or **Follow up**in the **Before** field. 2. Save the reminder. 3. Find the reminder in a list view using one of the following methods:

Reminder table (continued)

Field	Description
	▪ In the parent task record, configure the form, and then add Reminders as a related list. ▪ Navigate to `reminder.do` and search for the reminder. **4.** In the list view, change the **Before** field value to a different date field from the parent task. Use the field name and make sure the field has a value. For example, enter `work_start` to use the **Actual start** field instead. Note that you can't make additional changes on the reminder form without changing the **Before** field value back to either **Due date** or **Follow up**.
Using	Specifies the method by which the reminder must be sent. Available options are: **Send an Email** and **Outlook Calendar**.
Subject	Specifies the subject or the reason of the reminder. Maximum number of characters allowed on this field is 100.
Notes	Provides a space where you can enter your comments regarding the reminder. Maximum number of characters allowed on this field is 8000.

Defining assignment rules

The instance can automatically assign a task to a user or group based on pre-defined conditions by using data lookup rules and assignment rules.

Assignment rules module

The Assignment rules module allows you to automatically set a value in the assigned_to and assignment_group fields when a set of conditions occurs.

An assignment rule must also meet these additional criteria to run:

- The task record has been created or updated. Assignment rules do not apply to unsaved changes on a form.
- The task record must be unassigned. The record cannot have an existing value for either the *assigned_to* or *assignment_group* fields. Assignment rules cannot overwrite existing assignments (including assignments set by a default value or a previously run assignment rule).
- The assignment rule is the first rule that matches the table and conditions. If more than one assignment rule matches the conditions, only the rule with the lowest order value runs.

Data lookup rules

Data lookup rules offer a generic way to change any field value, not just assignment fields.

Data lookup and record matching support offer the following improvements over the *Assignment* module:

- Ability to change any field value not just an assignment field
- More options to define when a rule runs:

    - On form change (Allows assignment rules to apply to unsaved changes on a form)
    - On record insert
    - On record update
- Option to replace existing values (including default values)

Note:
You can define data lookup and *Assignment* rules at the same time. The system ignores any duplicate rules after an incident has been assigned unless you are using a data lookup definition option to replace existing values.

Related topics

Precedence between data lookup, assignment, and business rules

Precedence between data lookup, assignment, and business rules

Scripts, assignment rules, business rules, workflows, escalations, and engines all take effect in relation to a database operation, such as insert or update. In many cases, the order of these events is important.

Note:
Client-based code that executes in the browser, using Ajax or running as JavaScript, will always execute before the form submission to the server.

The order of execution is as follows:

1. Before business rules: Scripts configured to execute before the database operation with an order less than 1000.
2. Before engines. The following are not executed in any specific order:
    - Approval engine (for task and sys_approval_approver tables)
    - Assignment rules engine (for task tables)
    - Escalation engine
    - Data policy engine
    - Field normalization engine
    - Role engine - keeps role changes in sync with sys_user_has_role table (for sys_user, sys_user_group, sys_user_grmember, and sys_user_role tables)
    - Execution plan engine (for task tables)
    - Update version engine - creates version entry when sys_update_xml entry is written (for sys_update_xml table)
    - Data lookup engine inserts or updates
    - Workflow engine (for default workflows)
3. Before business rules: Scripts configured to execute before the database operation with an order greater than or equal to 1000.
4. The data base operation (insert, update, delete).
5. After business rules: Scripts configured to execute after the database operation with an order less than 1000.
6. After engines. The following are not executed in any specific order:
    - Label engine
    - Listener engine

- Table notifications engine
- Role engine - keeps role changes in sync with sys_user_has_role table (for sys_user, sys_user_group, sys_user_grmember and sys_user_role tables)
- Text indexing engine
- Update sync engine
- Workflow engine (for deferred workflows)
- Trigger engine (for all Flow Designer flows)

7. Email notifications. The following are executed based on the weight of the notification record:
   - Notifications sent on an insert, update, or delete
   - Event-based notifications
8. After business rules (Only active records). Scripts configured to execute after the database operation with an order greater than or equal to 1000.

Note:
Like After business rules, Async business rules execute their logic after a database operation occurs. Unlike After business rules, Async business rules execute asynchronously, running in the background simultaneously with other processes. Async business rules run after the user submits the form and after the scheduler runs the scheduled job created from the business rule. The system creates a scheduled job from the business rule after the user submits the form but before any action is taken on the record in the database.

Related topics

Business rules

Workflow assignments

An alternative to creating data lookup or assignment rules is to create one or more workflow tasks that assign a task record as part of a workflow.

Consider using Task workflow activities for assignment if your process includes multiple steps or conditions such as requiring a particular group approve a request.

When using a workflow to manage task assignments, add a brief Timer workflow activity to the start of the workflow. Without this timer activity, the workflow runs before the parent record, the current record, is inserted into the database. After the timer activity completes, the workflow resumes using the parent record information from the database instead of the original current. Pausing a workflow in this way does not change a default workflow to a deferred workflow. For more information on how the workflow engine interacts with the database, see Workflow engine operation order.

Baseline assignment rules example

A baseline instance contains certain assignment rules.

Baseline instance rules

Name	Table	Conditions	Assigned to user, Assigned to group, or Script
Networking	incident	Category is Network	User: System Administrator

Baseline instance rules (continued)

Name	Table	Conditions	Assigned to user, Assigned to group, or Script
			Group: Network
Database or Software	incident	Category is one of Request, Inquiry / Help, Software, Hardware	User: System Administrator Group: Software
SC Item fulfillment - Field Services	Ticket	Parent.Task type is Request item	Group: Field Services
Release Planning	release_phase	Name is Plan	Script: `current.release.product.service.assigned_to;`
IT Hardware	sc_req_item	Approval is Approved and Item.Category is Request Computers and Hardware	User: System Administrator Group: Hardware
Service Desk	incident	Active is true	Group: Service Desk

Create an assignment rule

Automatically assign a record according to one or more conditions in an assignment rule. Assignment rules are designed to run at the time you open a record.

Before you begin

Role required: assignment_rule_admin or admin

Procedure

1. Navigate to **All** > **System Policy** > **Rules** > **Assignment**.
2. Select **New**.
3. On the form, fill in the fields.

Assignment rule form

Field	Description
Name	The descriptive name for the assignment rule.
Active	Option to activate the assignment rule.
Applies to	
Table	The table with the records that the assignment rule applies to. The list shows only tables and database views that are in the same scope as the assignment rule. If you select a custom table that extends the task table, and for the assignment rule to work properly, you must clear the instance cache by navigating to https://<instance_name>.service-now.com/cache.do.

Field	Description
	Important: Clearing the system cache can affect overall performance, and degrade system response times. Do not run cache flushes during business hours, and do not trigger cache flushes automatically.
Conditions	The conditions under which the assignment rule applies.
Assign to	
User	The user the event is assigned to.
Group	The group the event is assigned to.
Script	
Script	A script to specify advanced assignment rule functionality. The current.variable_pool set of variables is available. **Note:** Make sure the input in the script is correct, and that the input type matches the field type in the Assignment Rule script. For example, if the assignment rule script sets the value of an Integer field, and the value in the script is set to String, the assignment rule may yield unexpected results.
Optional fields	
Match conditions	**Any** If any of the conditions are met, assignment rule applies. **All** If all the conditions are met, assignment rule applies.
Execution Order	The order in which the assignment rule is processed. If assignment rules conflict, a rule with a lower-order value takes precedence over a rule with a higher value. If the order values are set to the same number, the assignment rule with the first matching condition takes precedence over the others. Only the first assignment rule with a matching condition runs against a record.

Example: Condition editor

In this example, the condition statement automatically assigns any incident opened in the Network category to the system administrator in the Network assignment group.

Applies To

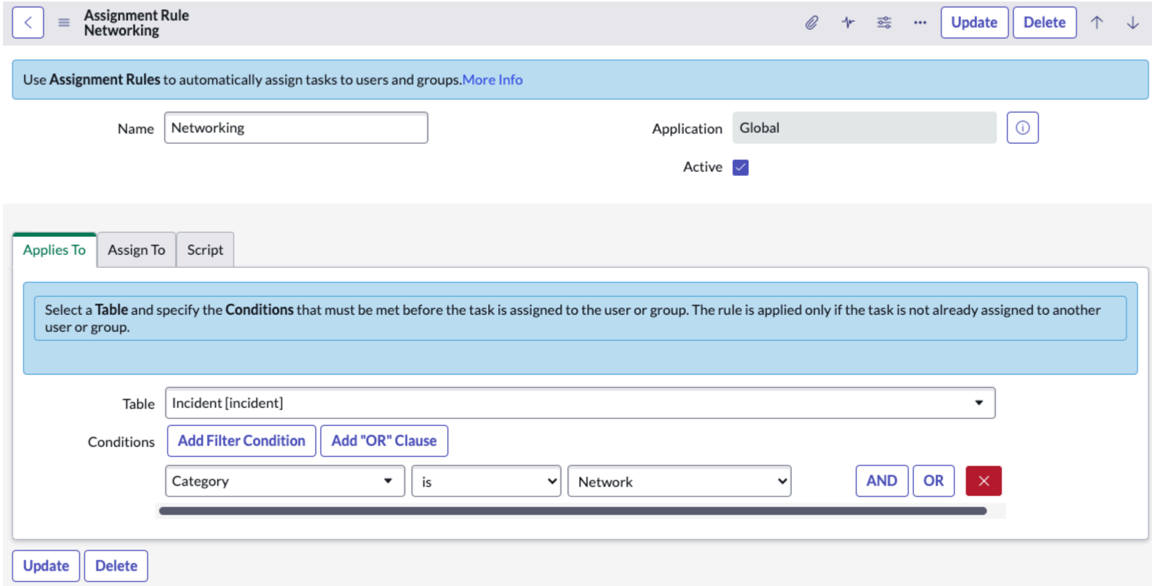

Assign To

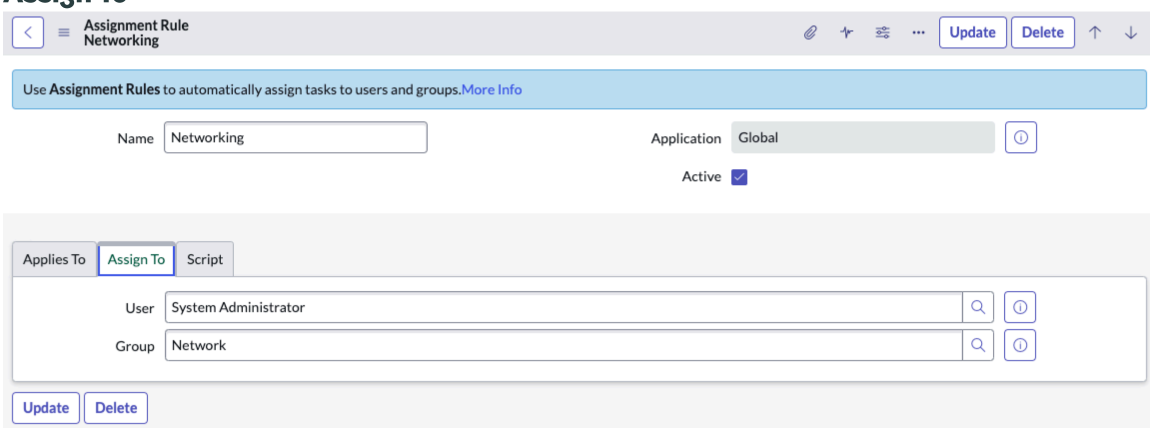

Create an assignment data lookup rule

Automatically assign a record using Data Lookup and Record Matching.

Before you begin

Role required: assignment_rule_admin or admin

Procedure

1. Navigate to **All** > **System Policy** > **Rules** > **Assignment Lookup Rules**.
2. Select **New**.
3. On the form, fill in the fields.

Assignment Data Lookup fields

Field	Description
Category	The category the data lookup matches against.
Subcategory	The subcategory the data lookup matches against.
Configuration item	The configuration item the data lookup matches against.

Field	Description
Location	The location the data lookup matches against.
Assignment group	The assignment group to assign the incident to.
Assigned to	The user to assign the incident to.
Active	Option to run the rule or deactivate the rule.
Order	Order in which the rule runs compared to other rules on the same table. The Data Lookup Plugin runs the rule with the lowest order and matching values.

4. Select **Submit**.

Result

The rule assigns incidents to the values in the *Assignment Group* and *Assigned To* fields based on the values selected in the **Category**, **Subcategory**, **Configuration Item**, or **Location** fields. A valid assignment lookup rule requires at least one matcher field and one setter field.

Example: Assignment Data Lookup

In this example, the Assignment Data Lookup rule automatically assigns any incident with the Category of Inquiry/Help and Subcategory of Email to Fred Luddy.

Assignment data lookup

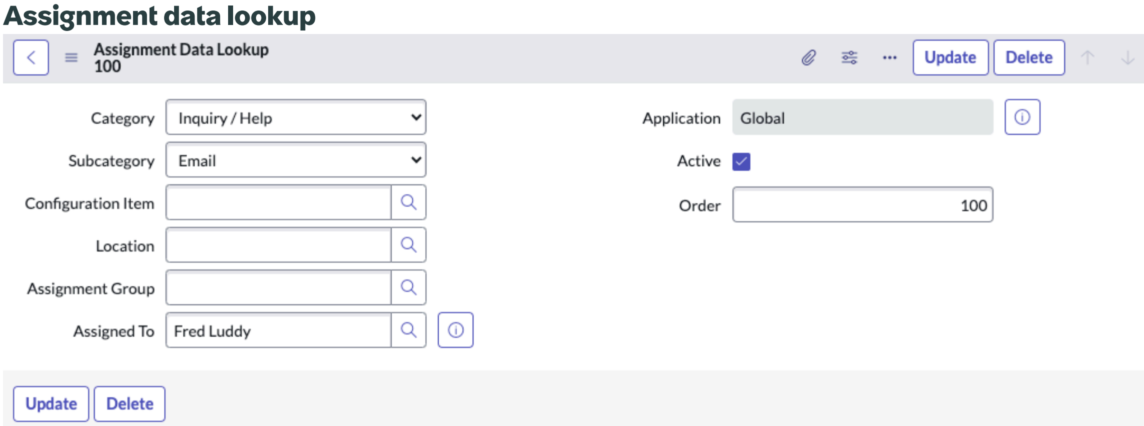

Extending the Task table with Planned tasks

The Planned Task plugin provides a Planned Task `[planned_task]` table that extends the Task `[task]` table.

The Planned Task plugin cannot be activated independently. It gets activated when activating the Project Management plugin (through PPM Standard plugin).

Planned tasks provide additional fields for tasks pertaining to time and effort as part of a planned, multi-stage process.

Note:
If the Planned Task_v2 (`com.snc.planned_taskv2`) plugin is active when you upgrade to the Geneva or later releases, the system adds the Task column to the Planned Task [planned_task] table. The Task column is of type Composite Field, and stores the Short Description and Number of the task. During the upgrade, the system updates all records in the Planned Task table.

Create a planned task

Planned Tasks are created on planned task child tables.

Before you begin

Role required:

About this task

Selecting the **New** button on the Planned Task record list launches the Planned Task Interceptor, which prompts the user to select a child table to create the planned task on:

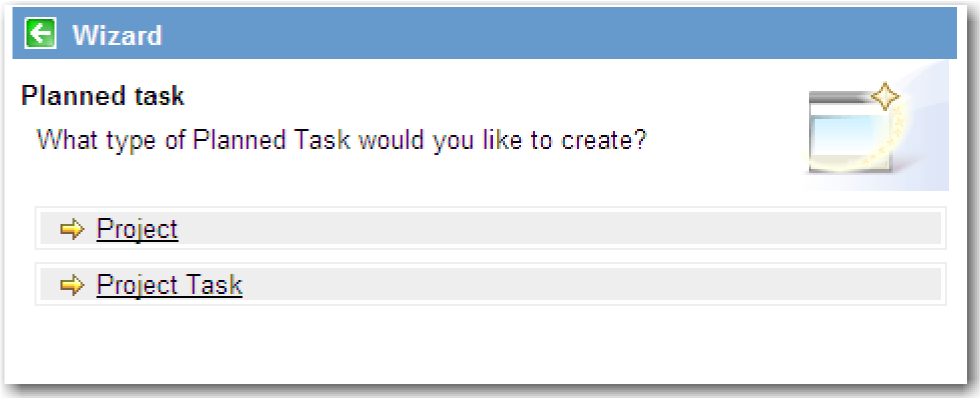

To modify the planned task interceptor:

Procedure

1. Navigate to **All** > **System Definition** > **Interceptors**.
   This module may need to be activated.
2. Select the **Planned Task** Interceptor.
3. Modify the **Answers** related list as desired.
   This related list specifies what choices are presented, and where the user is redirected to once they select the choice.

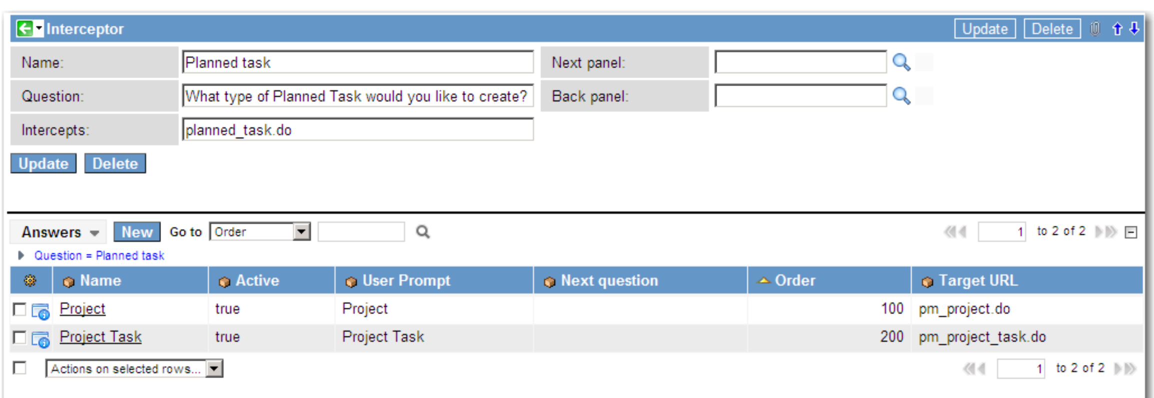

Create a baseline

A Planned Task Baseline is a record of the start and end times of the planned task at a particular moment in time.

To create a baseline, navigate to the form of the top planned task and select the **Create a Baseline** related link:

Create baseline

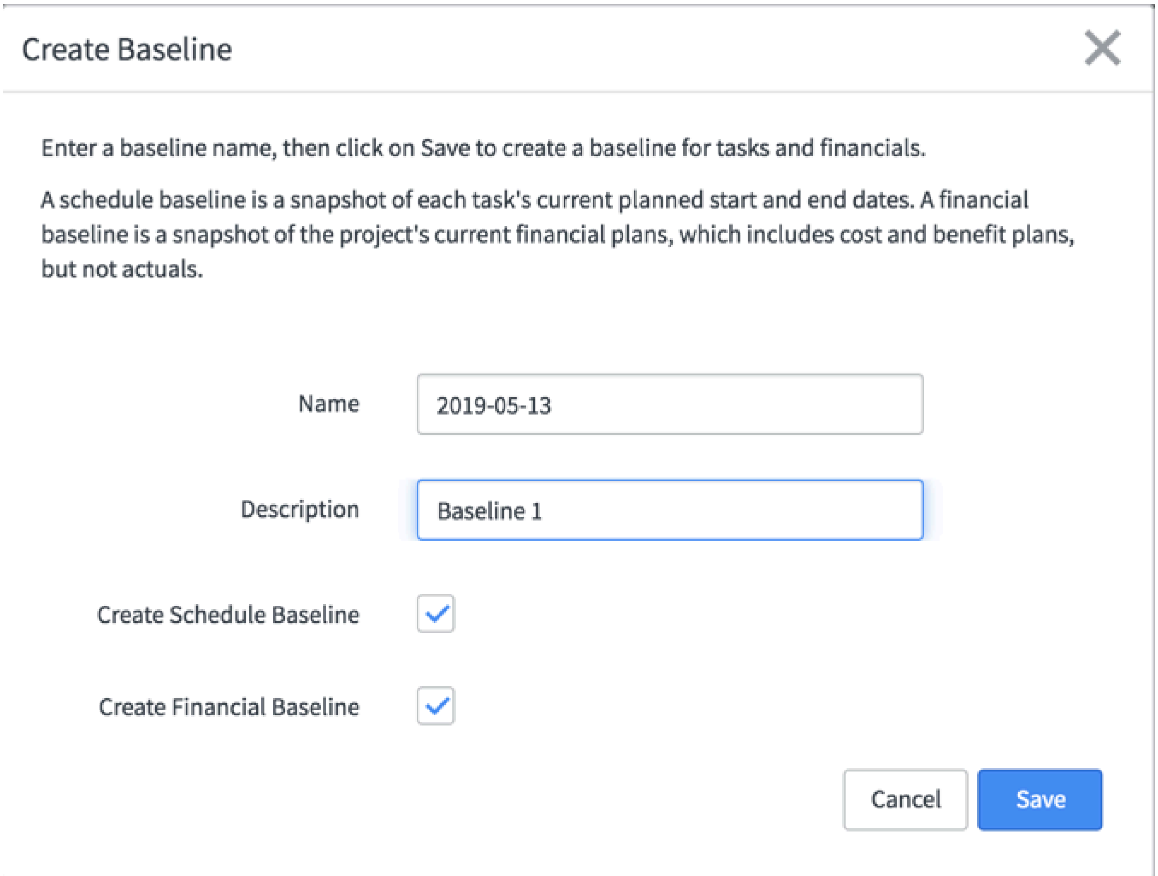

To view the baseline, configure the related lists to add a related list of baselines:

Baselines tab

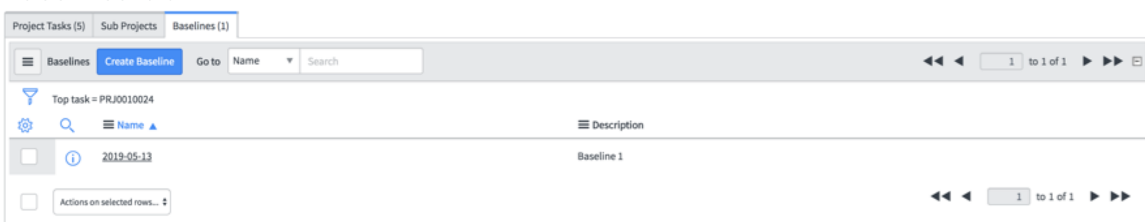

The baseline can be viewed on a Gantt Chart using the related link.

Measure time and effort

The Planned Task [planned_task] table provides standard fields for tracking duration and effort.

Duration measures the time from start to end date. Effort measures the hours of work exerted on the project.

Duration

- Planned duration: The projected length of time for the planned task.
- Actual duration: The actual length of time so far for the planned task.
- Remaining duration: The Planned duration minus the Actual duration, which represents the projected length of time left.

Effort

- Planned effort: The projected amount of time that will be spent on the planned task.
- Actual effort: The actual amount of time that has already been spent on the planned task.
- Remaining effort: The Planned effort minus the Actual effort, which represents the amount of project work left.
- Percent complete: For the parent task, the percent complete field is not editable and is rolled up from the percent completion of child tasks using the following formula:

  Percent complete of parent task = sum (planned duration of child tasks * percent complete of child tasks) / sum (planned duration of child tasks)

  For a child task, the percent complete is entered manually.

Date Calculation: All planned dates and actual dates of the child tasks are rolled up to the parent task automatically. Planned end date for a child task is calculated based on the planned start date and planned duration. If actual start is present, then planned end date is calculated from actual start date and planned duration.

Dependencies: All type of task dependencies, excluding external dependencies, are supported by the planned tasks. For more information about various dependency types, see Project task relationships and dependencies.

Important planned task table fields

The Planned Task table has these fields.

Important Planned Task Table Fields

Label	Name	Type	Description
Actual cost	`work_cost`	currency	The actual cost of the planned task, to be compared with the Estimated cost.
Actual duration	`work_duration`	glide_duration	The actual length of time (from start time to end time) of work on the planned task, to be compared with the Planned duration.
Actual effort	`work_effort`	glide_duration	The actual time spent working, to be compared to the Planned effort.
Critical Path	`critical_path`	boolean	
Estimated cost	cost	currency	An estimation of the cost of the planned task, to be compared with the actual cost.
HTML Description	`html_description`	html	A description field that accepts HTML mark-up.
Percent Complete	`percent_complete`	decimal	A percentage of the completed effort. Generated using the Planned effort and Actual effort fields.
Planned duration	duration	glide_duration	The estimated length of time (from start time to end time) of the planned task.

Important Planned Task Table Fields (continued)

Label	Name	Type	Description
Planned effort	effort	glide_duration	The estimated amount of time spent working on the planned task.
Planned end date	`end_date`	glide_date_time	The estimated date and time for the planned task to end.
Planned start date	`start_date`	glide_date_time	The estimated date and time for the planned task to start.
Remaining duration	`remaining_duration`	glide_duration	The difference in planned and actual duration, representing the time left for the planned task.
Remaining effort	`remaining_effort`	glide_duration	The difference in planned and actual effort, representing the amount of work time left for the planned task.
Rollup	rollup	boolean	Read-only field managed by the system that identifies the task as having child tasks. A rollup task has several its fields calculated from the children so those fields are read-only.
Time constraint	`time_constraint`	string	A description of time constraints that apply to the planned task.
Top Task	`top_task`	reference (planned_task)	When different planned tasks are stacked in a hierarchy, this field populates with the highest-level parent task. For example, if Project A has a child Project B, and Project B has a child Project C, then the Top Task for Project C is Project A. The Top Task field for Project A is blank.
Actual start date	`work_start`	glide_date_time	The actual date when the work starts on the planned task.
Actual end date	`work_end`	glide_date_time	The actual date when the work ends on the planned task.

Planned task scripts

Several business rules and one script include determines the dynamic calculation of crucial Planned Task fields.

Planned Task Scripts

Business Rule	Description
Set actual work start value	Sets the Actual Start Date of the planned task when State is set to the default work state.
Set close data on inactive	Sets the close data of the planned task when task becomes inactive.

Planned Task Scripts (continued)

Business Rule	Description
Recalculate	Recalculates the planned task schedule fields when one of the schedule fields changes.
Planned task global events	Raises global event for every insert or update of planned task record. For more information, see Global events.
Update parent actual effort	Rolls up actual effort from child task to parent task.
Update parent effort	Rolls up planned effort from child task to parent task.
Update parent percent complete	Rolls up percent complete from child task to parent task.
Validate duration	Validates the value in the planned duration field is not 0 or negative.
Validate percent complete	Validates the value in the percent complete field is between 0 and 100.
Validate work end before work start	Validates that the actual end date is not earlier than the actual start date.
Set top task	Sets top tasks to maintain hierarchy of tasks.
Set top task on children	Sets top tasks on child tasks to maintain hierarchy of tasks.

Planned task hierarchy

The **Task Hierarchy** tool available for Planned Task displays the relationship between parent and child planned tasks.

About this task

Out-of-box, the **Task Hierarchy** tool is available in both Project and Release Management.

Different **Planned Task** tables have different UI actions to launch the task hierarchy:

- To view the hierarchy of a Project, navigate to the use the **Task hierarchy** context menu action.
- To view the hierarchy of a Product in Release v2, navigate to the product and click the **Product hierarchy** related link.
- To view the hierarchy of a Release in Release v2, navigate to the release and click the **Release hierarchy** related link.

Rm2 Hierarchy

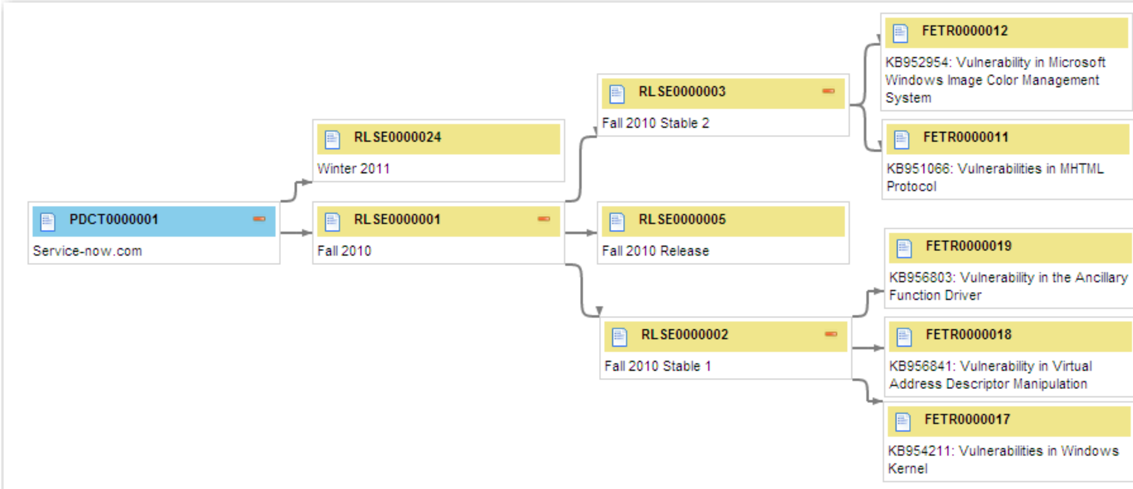

The Task Hierarchy can be added to any planned task table by:

Procedure

1. Navigating to **System UI** > **UI Actions**.
2. Selecting one of the existing Task Hierarchy UI Actions (for example, **Task hierarchy** if Project Management is activated).
3. Change the table to the desired table and rename the UI Action if appropriate, and insert.

   The hierarchy should now be available as a UI Action on the form of the new table.

Configure rollup for planned task fields

You can configure the planned task fields to roll up the field values in the parent entity.

About this task

By default the following fields are configured to roll up for any planned task and planned task derivatives:

Cost, **Budget**, **Effort**, **Actual date**, **Planned date**, **Duration**, **Percent Complete**

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Open the **Planned task rollup** table.
3. In the Related links, click **Show List**.
4. Click **New**.
5. Fill out the Planned task rollup form (see table).
6. Click **Submit**.

Planned task rollup form fields

Field	Description
Child	Select the table you want to include or exclude.
Field	Select the column name you want to include or exclude.

Field	Description
Navigator	The reference column of the parent to which the value must be rolled up.
Parent	Select the parent table.
Application	Auto-generated value.
Rollup	Select the check box to allow roll up. Leave the field unchecked to exclude the field from rollup.

Creating many-to-many task relations

By default, tasks can be related to each other using a parent/child relationship, such as a Problem with a group of child Incidents or a Catalog Request with a group of child Catalog Tasks.

However, it may prove useful to record exactly the nature of the relationship between the task records. When activated, the Many-to-Many Task Relations plugin allows administrators to define relationships between different tasks.

Request many to many task relations

The Many to Many Task Relations plugin (com.snc.task_relations) is included with several plugins. You can request activation of the plugin by itself.

Before you begin

Role required: admin

About this task

- Planned Task
- Field Service Management
- Project Management
- Governance, Risk, and Compliance

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. On the All Applications page, select **Request Plugin** to open the **Activate Plugin** form on Now Support.

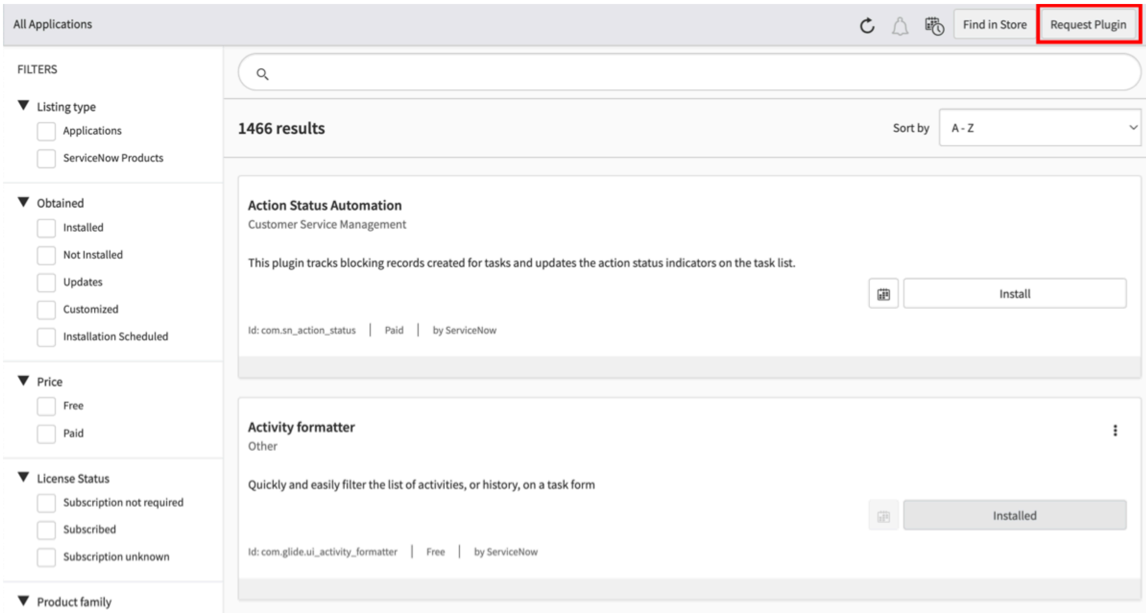

3. On Now Support, select the link to access the Now Support Service Portal Service Catalog.

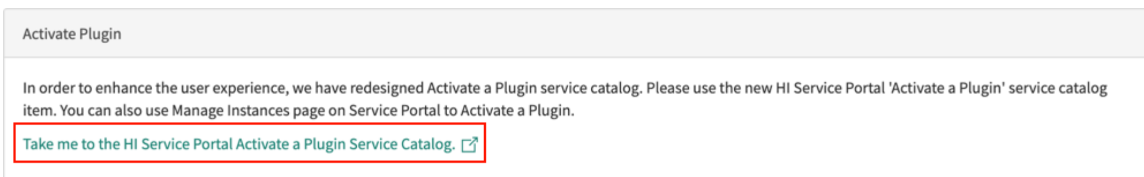

4. Select your instance.
5. Select **Actions > Activate Plugin**.
6. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Instance on which to activate the plugin.
Which plugin would you like to activate	Name of the plugin to activate. **Note:** If the system does not list the plugin you want or if you are activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	The date and time to activate the plugin.

Field	Description
	Note: Plugins are activated in two batches, once in the morning and once in the evening, on every business day in the US Pacific time zone. If the plugin must be activated at a specific time, enter the request in the **Reason/Comments** field.

Example

For example, see the following form to activate the CSM Workspace plugin on an instance named My Instance.

Activate Plugin form

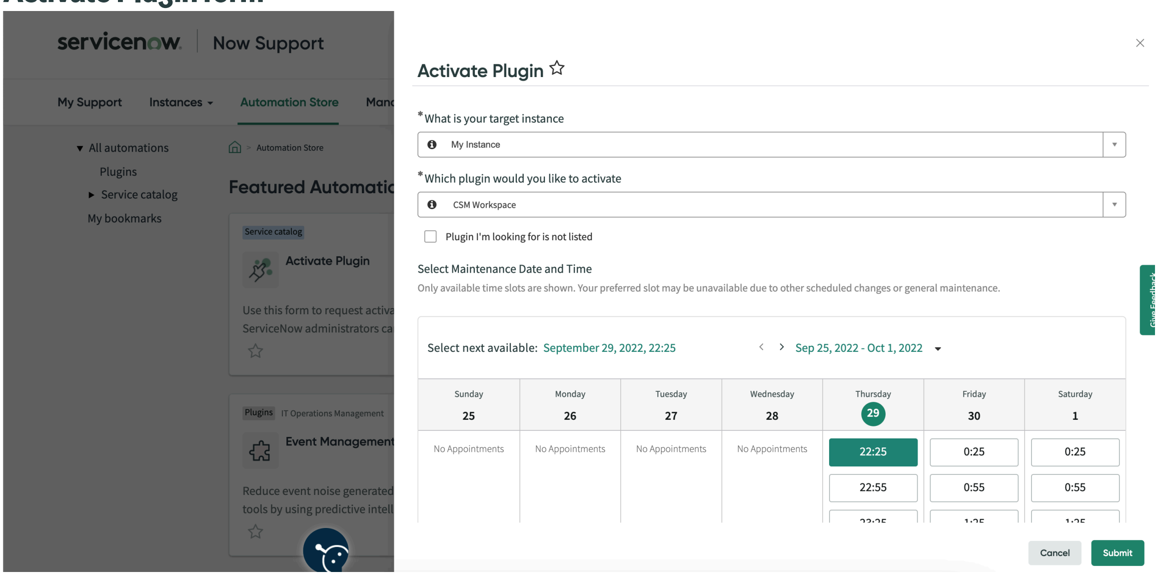

7. Select **Submit**.
   For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

Plugin manifest

When the plugin is activated, the Task Relationships application is available with certain modules.

Task relationships

Module	Description
Knowledge Relation Types	Contains the definitions for the possible types of relationships between knowledge articles and tasks.
Knowledge Relationships	Displays all the defined relationships between existing tasks and knowledge articles.
Relationship Types	Contains the definitions for the possible types of relationships between tasks.

Task relationships (continued)

Module	Description
Relationships	Contains all the defined relationships between existing tasks.

The following Relation Types are available by default:

- Solution is documented in::Documenting solution for
- Caused by::Causes
- Contains::Task of
- Documenting Solution for::Solution is documented in
- Investigated by::Investigates
- Permanent correction for::Permanently corrected by
- Related to::Related to
- Requesting::Requested by
- Solved by::Solves

Define a relationship type

Create type codes that define the relationship between parent and child tasks.

Before you begin

Activate the Many to Many Task Relations plugin (com.snc.task_relations) first. See Request many to many task relations.

Role required: admin

Procedure

1. Navigate to **All** > **Task Relationships** > **Relation Types** and select **New**.
2. Populate the *Parent Descriptor* and *Child Descriptor* fields with a short description of the relationship between the two tasks, such as a parent descriptor of *Caused By* and a child descriptor of *Causes*.
3. Select and hold (or right-click) the form header bar and select **Save**.
   The *Name* field automatically populates with the Parent and Child descriptors.

Define a task relationship allowed from the task relationship type record

Define the Task Relationships Allowed from the Task Relationship Type record.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Task Relationships** > **Relationship Types**.
2. Select a Task Relationship Type record.
3. In the Task Relationship Allowed related list, select **New**.
4. Populate the parent and child tables to define which tables are able to accept the relationship.
5. **Optional:** Define scripts to run in the *Parent* and *Child* script fields.

(Optional) These scripts are run when a parent or child record is run to automatically generate the other task (child or parent). These scripts use the current value of the new record, as opposed to the source record which triggered the script.

6. Select **Submit**.

Modify the displayed field

The list view of the Reference Lookup defines the fields displayed in the Task Relations field and editing interface.

Before you begin

Role required: admin

Procedure

1. Navigate to a form that has a reference to the table whose display values you would like to modify within Task Relations field.
2. Click the magnifying glass to display the Reference Lookup list view.
3. Right-click the list header and select **Configure** > **List Layout**.

   The selected List layout is used as the display value for records within the Task Relations field.

Mark as Solution button

The **Mark as Solution** button is added to the KB popup view and displayed when you search the knowledge base from a task record.

Selecting **Mark as Solution** creates a record in the Task / KB Relationships `[task_rel_kb]` table to associate the KB article with a task. You can disable this button by marking the active field false on the 'solution_button' UI macro.

Task relationships with UI actions

After defining task relationships, you can use UI Actions to define the task relationship when a new task is being created from an old task.

Below are a few examples.

> **Warning:**
> These examples may not work on all instances. They are provided as illustrative examples.

UI Actions examples

Cause an incident

This UI Action allows the change management team to log an incident directly from the change request and records that the change caused the incident.

Create a UI Action on the Change Request [change_request] table and place the following into the script:

```
var inccaus =new GlideRecord("incident");
inccaus.short_description= current.short_description;
inccaus.comments= current.comments.getHTMLValue();
// inccaus.parent = current.sys_id;
inccaus.insert();
CauIncident();
```

```
gs.addInfoMessage("Incident "+ inccaus.number+" created");
action.setRedirectURL(current);
action.setReturnURL(inccaus);


function CauIncident(){
var m2m =new GlideRecord('task_rel_task');
m2m.initialize();
m2m.child= current.sys_id;
m2m.parent= inccaus.sys_id;
m2m.type.setDisplayValue("Caused by::Causes");
m2m.insert();}
```

Cause a problem

This UI Action allows the change management team to record a problem from a change request and record that the change caused the problem.

Create a UI Action on the Change Request [change_request] table and paste the following script:

```
var probcaus =new GlideRecord("problem");
probcaus.short_description= current.short_description;
probcaus.comments= current.comments.getHTMLValue();
// probcaus.parent = current.sys_id;
probcaus.insert();
CauProblem();

gs.addInfoMessage("Problem "+ probcaus.number+" created");
action.setRedirectURL(current);
action.setReturnURL(probcaus);


function CauProblem(){
var m2m =new GlideRecord('task_rel_task');
m2m.initialize();
m2m.child= current.sys_id;
m2m.parent= probcaus.sys_id;
m2m.type.setDisplayValue("Caused by::Causes");
m2m.insert();}
```

Fix a problem

This UI Action allows a change request to be generated from a problem, recording that the change fixes the problem.

Create a UI Action on the Problem [problem] table, and paste the following code:

```
var fixchg =new GlideRecord("change_request");
fixchg.short_description= current.short_description;
fixchg.comments= current.comments.getHTMLValue();
// fixchg.parent = current.sys_id;
fixchg.insert();
FixChange();

gs.addInfoMessage("Change "+ fixchg.number+" created");
action.setRedirectURL(current);
action.setReturnURL(fixchg);
```

```
function FixChange(){
var m2m =new GlideRecord('task_rel_task');
m2m.initialize();
m2m.child= current.sys_id;
m2m.parent= fixchg.sys_id;
m2m.type.setDisplayValue("Fixes::Fixed by");
m2m.insert();}
```

Time Card Management

The Time Card Management feature enables time card users such as task assignees to report and track their time for the assigned tasks.

Video explaining how to set up the time tracking feature for projects, enter time, and approve time cards. Watch this five-minute video to learn about setup of the time tracking feature, entering time and approvals.

Time Card Management works with the Task table to record time worked on various task types, such as projects, incidents, problems, and change requests.

With the Time Card Management feature:

- The time card users can record the time worked on a task using time cards and time sheets. They can submit their time cards and time sheets for approval.
- Time cards and time sheets are routed for approval based on the **Non-project time approver** and **Project time approver** fields in the user time sheet policy. User managers with time card approver role can also log time and submit the time sheet of their resources
- The time card approvers can perform the following tasks:
  - Review and approve or reject the time cards in a submitted time sheet.
  - Recall the approved time sheets or time cards to return them to the users for corrections.
  - Use dashboards to view reports of time card and time sheet exceptions, and categorize time reported by the users.
  - Delegate another time card approver to process time sheets from your resources when you might not have time or will be unavailable. Note the following points for delegation:
    - You cannot delegate approval only for an individual resource.
    - The delegated time card approver cannot further delegate your time sheet approvals.

  The project_manager and the resource_manager roles contain the timecard_approver role in the base system. However, the timecard_approver role can also be used independently without these roles.
- The time card admin can create and manage time sheet policies, and can also approve or reject by exception and process the time sheets. Time card admin can also create or edit time cards of any users if the time cards are in Pending or Rejected state.

Note:
The Time Card Management plugin is required to use time cards. Some of the procedures require the project management feature, which activates time cards automatically. For more information, see Project Management (com.snc.financial_planning_pmo).

Domain separation in Time Card Management

Domain separation provides complete data isolation for domain-specific users. Time card Management is domain separation compliant at the **Data only** level.

Activate Time Card Management

Administrators can activate the Time Card Management plugin (com.snc.time_card). The plugin also activates the Performance Analytics - Content Pack - Project Portfolio Suite Dashboards plugin (com.snc.pa.time_card). However, you need a Performance Analytics license to use the dashboards.

Before you begin

Role required: admin

About this task

> **Note:**
> The Time Card Management plugin also gets activated as part of PPM Standard plugin (com.snc.financial_planning_pmo).

Procedure

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

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

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

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

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

Time sheet policies

Time sheet policies contain the policies to which a time sheet, or a time card must adhere.

By default, the Default time sheet policy is available with the system. As a Time card admin, you can create multiple time sheet policies based on different requirements and workflows in your organization. For example, specify a time sheet policy for each department or team and assign appropriate users to them.

Time sheet policies also provide an ability to specify appropriate approval workflow for project and non-project tasks.

A time sheet policy can be set as a default policy. The default policy is a global time sheet policy which applies to all the users who are not assigned to any other time sheet policy. Only one time sheet policy can be set as a default policy.

Navigate to **Time Sheets** > **Administration** > **Time Sheet Policies** to view the list of timesheet policies.

Create a time sheet policy

As a time card administrator, you can create a time sheet policy to define the requirements for time card users to record their time worked. The time card approvers must review and process time sheets under the applicable time sheet policy.

Before you begin

Role required: timecard_admin

Procedure

1. Navigate to **All** > **Time Sheets** > **Administration** > **Time Sheet Policies**.
2. Click **New**.
3. On the form, fill in the fields.

Time Sheet Policy form

Field	Description
Name	Unique name of the policy.
Allow blank time cards	Option to enable submission of blank time cards for approval. By default, this option is not selected.
Auto create time card on planned task update	Option to create a time card automatically when you update a planned task, if the following conditions are true: ◦ You are a time card user ◦ You are assigned to the task through **Assigned to** or **Additional assignee list** field ◦ The task is not in a pending state. By default, this option is not selected.
Auto fill time card with time worked entries	Option to fill in time cards automatically when you enter time in the **Time worked** field on the Task form. If a time card does not exist for the task, one is created when the time worked is updated if the following conditions are true: ◦ You are a time card user. ◦ You are assigned to the task through the **Assigned to** or **Additional assignee list** field. By default, this option is not selected.
Auto create time cards every week	Option to generate time cards automatically every week through a scheduled job for all users assigned to the time sheet policy. The time cards are generated based on the planned task assignments and hard-allocated resource plans of the users.

Field	Description
	By default, this option is selected.
Update actual hours and cost in resource plan/reports	Option to update actual hours and actual cost of the associated resource plan based on the hours entered in the approved time card. When selected, you can associate a resource plan related to the selected task with the time card in the following ways: ◦ If the task has multiple resource plans, you can select the appropriate plan. ◦ If the task has one resource plan, that plan is selected by default. The hours from the time card entry and the hourly resource rate are used to update the **Actual hours** and **Actual cost** values of the associated resource plan. **Note:** If the task has no resource plan, the associated records update the corresponding project. By default, this option is not selected.
Allow recall	Option to enable the recall action on a time sheet or time card after it is approved and processed. By default, this option is selected.
Recall period allowed (days)	Number of days within which you can recall a time sheet or time card. This field appears when the **Allow recall** option is selected. The default value is 30.
Week starts on	Start day of the week for the time sheet. The default value is Sunday. **Note:** For more information on filtering Time Sheets using this field, see KB0852458 article in Now Support Knowledge Base.
Maximum hours per day	Maximum number of hours that can be entered each day in a time sheet. The default value is 24. Setting this field to -1 enables the resource to enter: ◦ A value from 0 through 24 for the hours worked each day ◦ A negative value, to correct the previously entered timecards.
Maximum hours per week	Maximum number of hours allowed in a week in a time sheet.

Field	Description
	The value for total hours for a week in a time sheet is the sum of hours entered per day of that week. If set to -1, the total hours per week allowed is up to 168 (24 x 7) hours. The default value is 40.
Non-project time approver	Type of approval required when you submit a time card that has a non-project task assigned. ◦ **Auto**: Time card is auto-approved when submitted. ◦ **User Manager**: Time card is routed to the user manager for approval when submitted. The user manager has the time card approver role and is also selected as the manager of the time card user on the User form. ◦ **None**: A user with the timecard_admin role can approve time cards when submitted.
Project time approver	Type of approval required when you submit a time card that has the project task. ◦ **Auto**: Time card is auto-approved when submitted. ◦ **Project Manager**: Time cards for a project task are routed to the respective project manager for approval when submitted. The project manager is picked from the task against which you are submitting time. ◦ **User Manager**: Time cards are routed to the user manager for approval when submitted. ◦ **Both**: Time cards are routed to both the user manager and project manager for approval when submitted. A time card remains in the Submitted state when only one of the approvers approves the time card. ◦ **None**: A user with the time_card admin role can approve time cards. **Note:** ◦ The project manager and the user manager can view only those time cards that are routed to them for approval. ◦ If the user manager or the project manager changes while the time card is still not approved, the system automatically updates the approver of the time card. ◦ If the name of the project manager or the user manager is not populated on the respective form, a user with the time_card admin role can approve time cards.
Allow multiple rate types	Option to provide a rate type in a time card to be used during expense line generation. When this option is selected, the **Rate type** field displays on the Time Card form to enable the selection of a rate type.

Field	Description
	For more information, see Rate type in labor rate card.
Default rate type	Rate type used by default for the time card. If you select none, the Standard rate type option is used. This field appears when the **Allow multiple rate types** check box is selected.
Default Policy	Option to set the time sheet policy as the default. The default policy is a global time sheet policy that applies to all users who are not assigned to any other time sheet policy. Only one time sheet policy can be set as the default. A time sheet policy set as the default policy cannot be deleted. You must first set another policy as the default policy.

4. Click **Submit**.

What to do next

- Set the time sheet policy as the default policy, if required.
- Assign the time sheet policy to users.

Set a time sheet policy as default policy

As a time card administrator, you can set a time sheet policy as the default policy. The default policy is a global time sheet policy which applies to all the users who are not assigned to any other time sheet policy.

Before you begin

Role required: timecard_admin

About this task

Only one time sheet policy can be set as a default policy.

Procedure

1. Navigate to **All** > **Time Sheets** > **Administration** > **Time Sheet Policies**.
2. Open the time sheet policy record.
3. Click the **Set as default** related link.

   Note:
   The related link appears only for a time sheet policy which is not already set as the default policy.

Result

The **Default Policy** option on the record form is selected to indicate that the current time sheet policy is set as the default policy. The **Default Policy** option on the earlier default time sheet policy is cleared.

You cannot delete a default time sheet policy. You must first mark another policy as the default policy to delete the current default policy.

Assign a time sheet policy to a user

As a time card administrator, you can assign a time sheet policy to a user.

Before you begin

Role required: timecard_admin

About this task

You can assign only one time sheet policy to a user.

Procedure

1. Navigate to **All** > **Time Sheets** > **Administration** > **Time Sheet Policies**.
2. Open the time sheet policy record.
3. In the **Users** related list, click **Edit**.
4. On the Edit Members form, move the users to the **User List**.
5. Click **Save**.

Result

The selected users appear on the **Users** related list. The time sheet policy is assigned to the selected users.

Create a rate type

You can create rate types using the Rate Types feature.

Before you begin

Role required: timecard_admin

About this task

The rate type functionality is used to categorize different types of work. For example, Standard versus Overtime. It can also instruct down-stream, third-party product about the nature of the work performed.

For example:

- A user works extra time during the day and must differentiate standard time from overtime.
- A technician may be paid at a higher rate based on different types of work involved in the execution of a task.
- An appliance repair task may be billed at one rate for the first hour and a different rate for the remaining hours.
- A service call may last four hours during which specialized equipment is used for one hour. The company must capture the use of the specialized equipment for purposes ranging from additional billing to legal compliance or warranty tracking.

Procedure

1. Navigate to **All** > **Time Sheets** > **Administration** > **Rate Types**.
2. Click **New** and fill the form.
3. Click **Submit**.

Result

- The rate type is displayed in the **Rate Types** list.
- The rate type if active is also displayed in the **Rate Type** field in the Labor Rate Cards form, Time Worked form, Time Card form, and Time Sheet Portal.

Time Sheet Portal

Time Sheet Portal categorizes and displays all your assigned tasks in a single view for a given week. The portal enables you to record time spent on tasks on a day-to-day basis and submit the time sheet in a single action.

Time Sheet Portal works with the Task table to record the time that you spend on various task types, such as projects, incidents, problems, and change requests. It provides a complete breakdown of the time you enter for the task, task categories, and days in a week. It enables you to:

- Generate time cards for assigned projects or project tasks.
- Copy time cards from a previous time sheet.
- Search and create a time card for a task that is not displayed in the assigned tasks list.
- Edit, delete, or add notes to a time card.
- Submit, approve, reject, and recall a time sheet.
- View notifications for the submitted, rejected, and recalled time sheets or time cards.

Time Sheet Portal comprises of the following sections:

Header

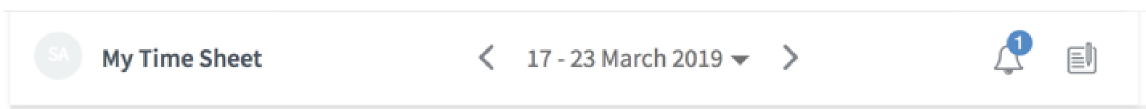

The Time Sheet Portal header:

- Displays a date range of the time sheet for the selected week and provides options to navigate to the time sheet of previous or next week.
- Provides a calendar along with the date range. Clicking a date in the calendar displays the time sheet for the week of the selected date.
- Provides a link to navigate to the time sheet for the current week when a different time sheet is being viewed.
- Provides an option under **My Time Sheet** that enables you to open and edit the time sheets of your resources. The option is available only for a user manager.
- Provides an option to view the time sheet in the Form view using time sheet form view icon ().

- Displays the notifications for the submitted, rejected, and recalled time sheets under the show notifications icon ( ). Each entry displays the number of time cards and their state. Click on a notification item to open the time sheet.

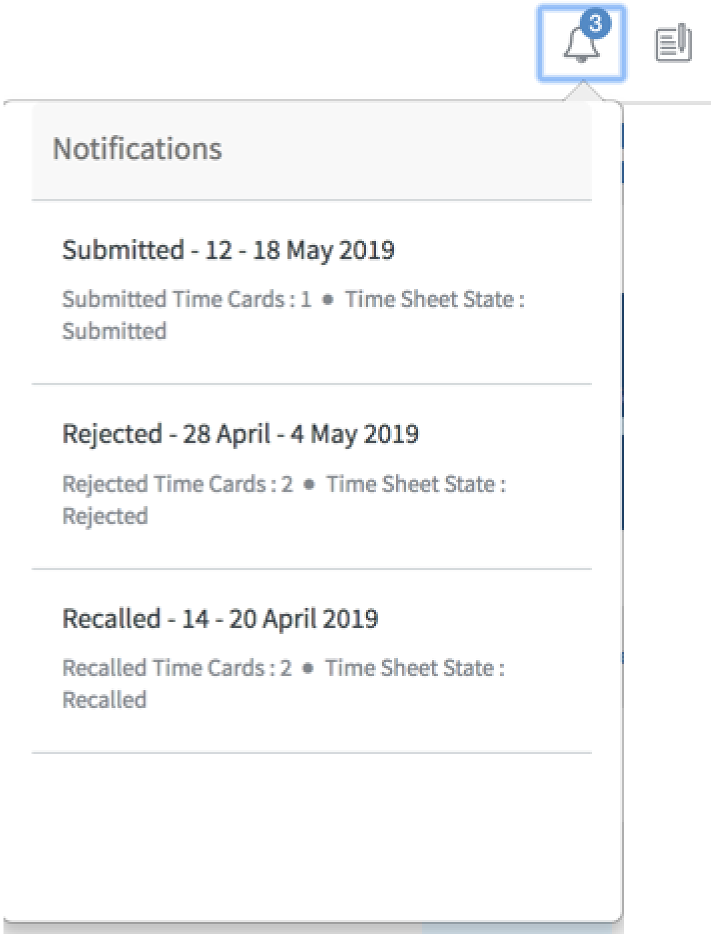

Tasks

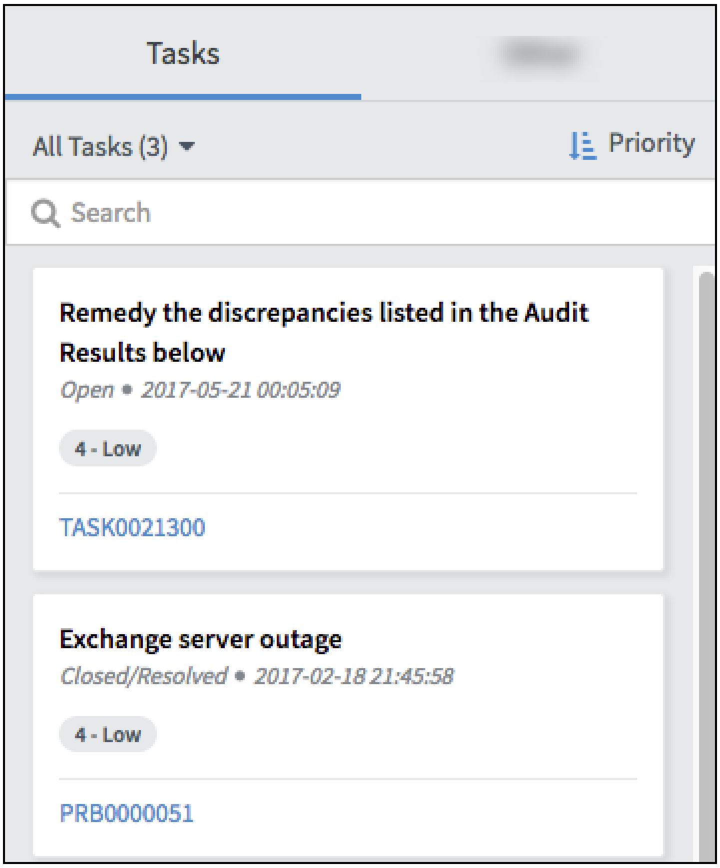

The **Tasks** tab displays your assigned tasks as cards. Each card provides information about a task, such as short description, state, last updated date and time, and priority. Tasks are derived from the following filter conditions:

- All your assigned incidents, problems, or change requests that are active.
- All your assigned incidents, problems, or change requests that are closed in the selected week.
- All tasks for which you are hard-allocated during the selected week.
- All your assigned projects or project tasks that are active during the selected week. Project or project tasks are derived from the following conditions:
  - A project or project task that has started and its actual start date is before the end of the selected week.
  - A project or project task has closed and its actual end date is after the week has begun.
  - When there are no actual dates for a project, planned dates of the project must occur within the selected week.
- All active stories assigned to you.
- All active scrum tasks assigned to you or closed in the selected week.
- All active test plans assigned to you.
- All active defect and enhancements assigned to you or closed in the selected week.

Note:
If the Customer Service Management plugin (com.sn_customerservice) is activated, your assigned tasks that are active and closed in the selected week are also derived from the following tables:

- Case [sn_customerservice_case]
- Work Order [wm_order]
- Work Order Task [wm_task]

On the **Tasks** tab, you can perform the following operations:

- Type a keyword in the search field to view only cards with details that match the keyword.
- Sort cards in ascending or descending order based on priority, irrespective of the task type.
- View cards based on a category, such as projects, project tasks, incident, problem, change, defect, enhancement, test plans, and scrum tasks.
- Create time cards for assigned tasks. See Create time cards and log time through Time Sheet Portal.

  Note:
  Time sheet policies apply when you create or edit a time card.

Group Tasks

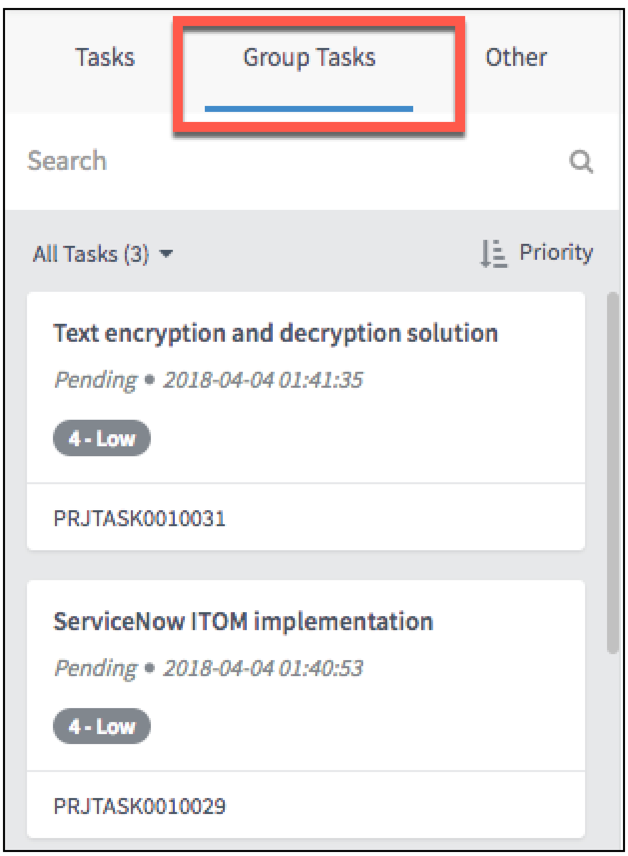

The **Group Tasks** tab displays the tasks assigned to your **Assignment group**. The group tasks are displayed as cards in the tab. A card provides information about the task, such as short description, state, last updated date and time, and priority.

Other

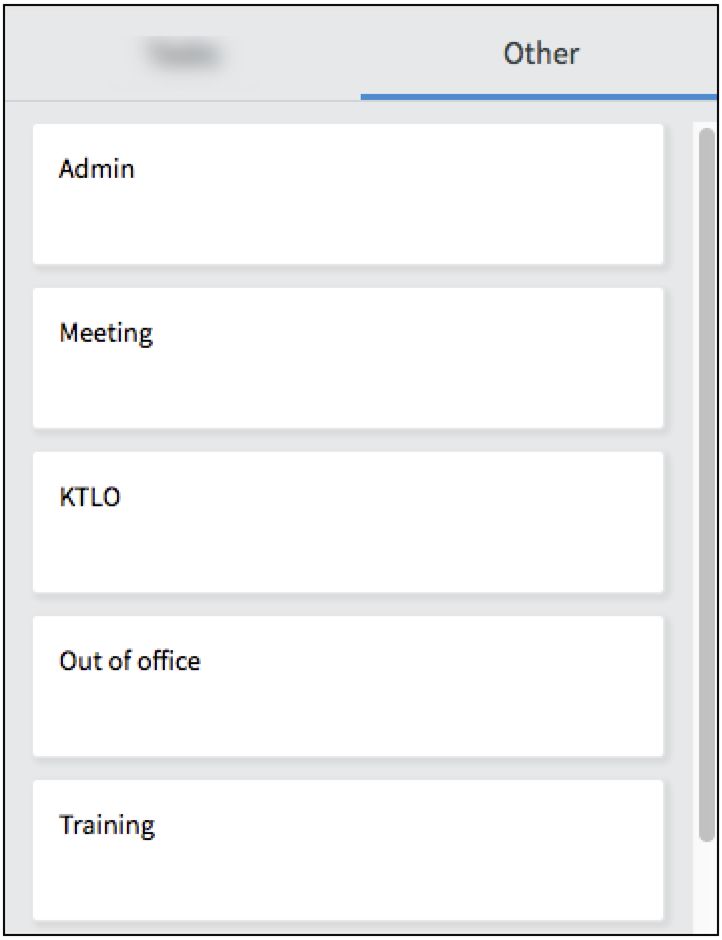

The **Other** tab displays cards used to log time for operational work, such as administration, meeting, and KTLO.

Note:
The list of categories in the time card controls the list of cards displayed on the **Other** tab. If an administrator marks any of the categories as inactive in the time cards, those categories are not displayed on the **Other** tab.

Time Sheet

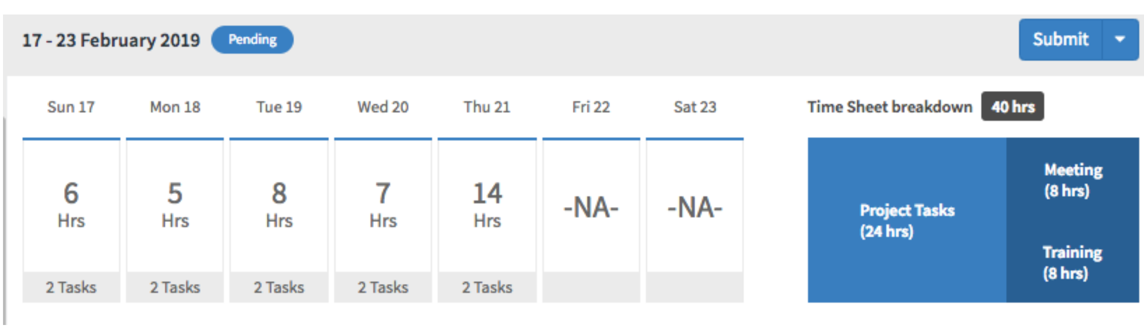

The Time Sheet section:

- Displays the date range of the time sheet for the selected week.
- Indicates the state of the time sheet being viewed.

- Displays the total number of tasks and hours that you have entered for each day in a week. When you click a day in the time sheet, the column for that day is highlighted in both the **Time Sheet** and **Logged Time Cards** sections. The first day displayed in the Time Sheet section is based on the value specified in the **Week starts on** field in the users Time Sheet Policy. For example, if the **Week starts on** field is set to Friday, the Time Sheet section starts with Friday.
- Indicates the total number of hours that you entered in the time sheet. Time Sheet breakdown also displays a breakdown of user-entered hours in different work categories, such as project tasks, admin, meeting, follow on tasks, or change requests.
- Provides options that enable you to submit a time sheet and copy time cards from a previous time sheet. For an approver, options to approve or reject a submitted time sheet and recall a processed time sheet are available.

Logged Time Cards

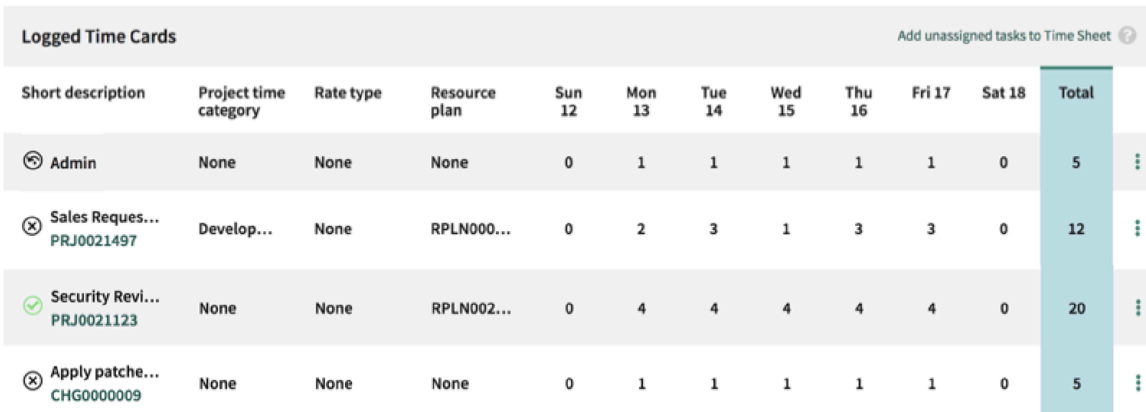

Logged Time Cards — Add unassigned tasks to Time Sheet

Short description	Project time category	Rate type	Resource plan	Sun 12	Mon 13	Tue 14	Wed 15	Thu 16	Fri 17	Sat 18	Total
Admin	None	None	None	0	1	1	1	1	1	0	5
Sales Reques... PRJ0021497	Develop...	None	RPLN000...	0	2	3	1	3	3	0	12
Security Revi... PRJ0021123	None	None	RPLN002...	0	4	4	4	4	4	0	20
Apply patche... CHG0000009	None	None	None	0	1	1	1	1	1	0	5

Each row in the Logged Time Cards section represents a time card entry for a given week. It has details such as icon for state of the time card, short description of the task, project time category, rate type, resource plan, and logged time. Each row of the time card contains options that enable you to edit, delete, or add notes to a time card. For more information, see Create time cards and log time through Time Sheet Portal.

Note:

- The **Rate type** column is visible only when the **Allow multiple rate types** option in the time sheet policy of the user is selected.
- The **Resource plan** column is visible only when the **Update actual hours and cost in resource plan/reports** option in the time sheet policy of the user is selected.

The icon beside the short description of a row indicates a time card state. The merge icon () in a time card row indicates that a duplicate time card for the same task exists in the time sheet and provides option to merge them into a single time card.

When there are no time cards, the Logged Time Cards section displays the **Generate Time Cards** and **Copy from previous time sheet** buttons to create time cards.

Note:
Time Sheet Portal is not designed for use on mobile devices, and may not appear as shown in this document.

Create time cards and log time through Time Sheet Portal

After you create time cards in Time Sheet Portal, log time in the time cards.

Before you begin

Role required: timecard_ user

Procedure

1. Navigate to **All** > **Time Sheets** > **Time Sheet Portal**.
2. Create time cards in the **Logged Time Cards** section using any of the following options.

   Based on the **Allow time card reporting on** field settings in the **Preferences** tab of a project, the **Add to Time Sheet** and **Add selected to Time Sheet** options are displayed in Time Sheet Portal.

   If a duplicate time card exists for a time card, you can merge both time cards to become a single one, except for time cards in the processed, approved, or recalled state. A duplicate time card is identified by comparing the values in the **Short description**, **State**, **Task**, **Category**, **Rate Type**, **Resource Plan**, or **Project Time Category** fields in the time card.
3. In the **Logged Time Cards** section, log time in a time card.

Submit time sheet through Time Sheet Portal

Once you update the time sheet with time worked for a given week, submit it for approval.

Before you begin

Role required: timecard_ user

About this task

When you submit a time sheet, time cards are verified against the assigned time sheet policy. If there is no time sheet policy assigned, time cards are verified against the default time sheet policy. Business rules prevent you from submitting a time sheet in case any violation is found.

Procedure

1. Navigate to **All** > **Time Sheets** > **Time Sheet Portal**.
2. Click **Submit**.
   The time sheet moves to the Submitted state.

Log time and submit time sheets of your resources

As a user manager, you can log time and submit the time sheet of your resources.

Before you begin

A user manager must have time card approver role to create and edit time cards, and submit the time sheet of other time card users. The user manager has either of the required roles and is the one listed as the manager of the time card users.

Role required: timecard_approver

Procedure

1. Navigate to **All** > **Time Sheets** > **Time Sheet Portal**.
2. Click **My Time Sheet** choice list and select the user for whom you want to approve or reject time sheet for under **My Resources**.

The **My Resources** section lists the time card users reporting to you.

> **Note:**
> The choice list with **My Time Sheet** is available only for a user manager.

3. Create time cards and log time through Time Sheet Portal as required.
4. Click **Submit**.

Result

- The time sheet of the selected user moves to the Submitted state. The time sheet policy assigned to the time card user is used for validation when you submit the time sheet.
- All the time cards associated with the time sheet are also submitted automatically and move to the Submitted state.
- The **Approve** and **Reject** buttons appear in the Time Sheet section. The two buttons appear only for a user manager.

What to do next

Once you have submitted the time sheet for your resource, you can click **Approve** or **Reject** to approve or reject it. Only the time cards that are routed to the user manager, are approved or rejected. The routing for approval happens based on the **Non-project time approver** and **Project time approver** fields in the time sheet policy of the time card user.

Time Sheets

A time sheet groups all the time cards for a user for the given week.

With time sheets:

- Time card users can submit all the time for their work week in a single step by using a time sheet.
- Time card approvers can approve all the time cards in a time sheet for a user in a single step by approving the time sheet. They do not need to approve multiple time cards for a given user individually.
- Track the activities of a time sheet, such as who submitted or approved a time sheet, in the Activities section on the Time Sheet form. This time sheet activity audit is useful for tracking when you delegate responsibility for your time card processing to another user. To track the activities, enable the **State** field of the Time Card [time_card] table for auditing the time card activities, if it is not already enabled. For more information, see Include a table field in auditing.

A time sheet can have any of the following states:

Pending

A time sheet has been created, but the user is still making changes before submitting it.

Submitted

A time sheet has been submitted for approval. The approver can approve or reject it.

Approved

A time sheet has been approved. If the time sheet isn't automatically updated to Processed, this means there was no task associated with the time sheet.

Processed

A time sheet has been processed by the system. After a time sheet is approved, the after business rule Create expense from approved time card is triggered. This business rule creates an expense line for the associated task, and then updates the state to Processed.

Rejected

A time sheet has been sent back to the submitter for changes.

Recalled

A time sheet has been recalled by the submitter for modification.

Create a time sheet

As a time card user, you can create a time sheet to group all your time cards for the given week and submit them in a single step.

Before you begin

Role required: timecard_user

About this task

A user can create only one time sheet per week.

Note:
When a user creates a time card for a week, a time sheet is also automatically created for that week.

Procedure

1. Create the time sheet with one of these options.

2. On the form, fill in the fields.

Time Sheet form

Field	Description
Week starts on	Starting date of the week for which the time sheet is created. **Note:** The time sheet policy **Week starts on** controls the start day of the week.
Total Hours	Number of hours the user has worked in that week. This field is automatically populated from the hours recorded for the associated time cards.
User	Name of the user for which time sheet is created.
State	Current state of the time sheet. All new time sheets begin as Pending. Default states: Pending, Submitted, Approved, Processed, and Rejected, Recalled.
Comments	Comments related to the time sheet.
Notes	Any additional information.

3. Click **Submit**.

What to do next

Add time cards to the time sheet. You can use the related links and related list to add time cards. You can also add time cards using the Time Sheet Portal.

Time sheet form related links

Field	Description
Generate Time Cards	Link to generate time cards for the assigned project tasks.
Copy from previous time sheet	Link to open the **Copy Time Sheet** window for you to copy the time cards from previous time sheets.
Time Sheet Portal	Link to open the Time Sheet Portal to view and manage the time sheet.

Time sheet form related lists

Field	Description
Time Cards	List of the time cards that are part of the time sheet. To create a time card, click **New**.

Submit a time sheet

Once the time sheet is updated with time worked, you can submit the time sheet for the week to submit all the time cards for the week together.

Before you begin

Role required: timecard_user

About this task

When a user submits a time sheet, the time sheet policy assigned to the user is used for validation. If there is no time sheet policy assigned to the user, the default time sheet policy is used. Business rules prevent a user from submitting a time sheet in case any violation is found.

Note:
You can also perform this task from the Time Sheet Portal.

Procedure

1. Navigate to the time sheet using one of the following options:
   - To submit a time sheet in the Pending state, navigate to**Time Sheets** > **My Time Sheets** > **Pending**.
   - To submit the time sheet for the current week, navigate to **Time Sheets** > **My Time Sheets** > **Current Time Sheet**.
   - To submit a time sheet in the Rejected state, **Time Sheets** > **My Time Sheets** > **Rejected**. You can make required changes as suggested in the rejection comments to resubmit a rejected time sheet.
2. To submit in the Time Sheets list, open the time sheet.
3. Add comments, if required.
4. Click **Submit Time Sheet**.

Result

- The time sheet moves to the Submitted state.
- All the time cards associated with the time sheet are also submitted automatically and move to the Submitted state.

Approve or reject a time sheet

View, approve, or reject time sheet or time cards for your user, for the given week, in a single step.

Before you begin

Role required: timecard_approver or timecard_admin

About this task

A user manager is the one who has either of the required roles and is selected in the **Manager** field on User form of the time card user.

During out-of-office time, the user manager can Delegate time sheet approvals to another user with the required roles as the time sheet approver. For the specified time period, the delegated user also receives all the time sheets submitted by the time card users to process.

Procedure

1. Navigate to **All** > **Time Sheets** > **Time Sheets** > **Pending Approval**.
2. In the Time Sheets list, open the time sheet to approve or reject.
3. If required, add comments.
4. Click **Approve** or **Reject**.

Result

- The time sheet moves to the Approved or Rejected state.

  Note:
  If a time sheet is rejected, it gets listed in the notifications for the rejected time sheets under the show notifications icon ( ) on Time Sheet Portal.

- All the time cards associated with the time sheet are also approved or rejected automatically.
- If any time cards in the time sheet are pending approval by one of the approvers as defined by the time sheet policy, the time sheet remains in the Submitted state.

Note:

- You can also approve or reject the time cards within a time sheet selectively by selecting one or more time cards and approving or rejecting them.
- If a time sheet is rejected, all the associated time cards in the Submitted state are also rejected automatically. The time cards in the Approved and Processed state for the time sheet remain unaffected.
- If all the time cards in a time sheet are approved, the time sheet automatically moves to the Approved state.
- If a time card in a time sheet is in the Rejected state, then irrespective of the state of other time cards in the time sheet, the time sheet also moves to the Rejected state.

Time cards

Time cards are used to record the time worked on a task by a task assignee.

The time card management feature works with the Task table to record time worked on Projects, Incidents, Problems, and Change Requests.

Task assignees can record time worked in the *Time worked* field on a task record or enter hours directly into their time card. Some tables support automatic time card creation based on start and end date fields.

Track the activities of a time card, such as who submitted or approved a time card, in the Activities section on the Time Card form. This time card activity audit is useful for tracking when you delegate responsibility for your time sheet processing to another user. To track the activities, enable the **State** field of the Time Card [time_card] table for auditing the time card activities, if it isn't already enabled. For more information, see Include a table field in auditing.

You can associate time cards for the project tasks and other task categories, such as meeting and training, with relevant resource plans.

When you approve a time card, the time logged in the time card is saved day-wise in the Time Card Daily [time_card_daily] table. Using daily time logged data, you can generate time card reports by days for any period irrespective of the time sheet period to which the time cards belong. For example, you can create a monthly time card report that includes time cards from the first day to the last day of the month.

Time cards also have an optional approval mechanism for project managers to approve the time cards. Administrators and time card approvers can see all the time cards for the week. All users who are in a role that is responsible for working on tasks also can access their personal time cards. A time card can have any of the following states:

Time Card states

State	Description
Pending	A Time card that isn't submitted for approval.
Submitted	A time card submitted for approvals.
Approved	An approved time card.
Processed	An approved time card with an expense line for the associated task.
Rejected	A rejected time card.
Recalled	An approved time card recalled for any required adjustments.

Create a time card

You can create time cards to log time against the work you have done.

Before you begin

Role required: timecard_user

About this task

The **Allow time card reporting on** field on the project form determines the level at which the time cards for the project tasks can be created.

Note:
Time Cards created from Project form do not populate Resource plans automatically to avoid random allocations. For more information, see KB0814884 article in the Now Support Knowledge Base.

Time cards can be created automatically or manually.

- *Automatic*: Use the following options to automatically generate the time cards:
  - *By updating task*:

    Auto create time card on task update

    Configure time cards to be created when a user updates a task record. The **Auto create time card on task update** setting in the assigned time sheet policy controls this behavior and is set to false by default.

    Auto fill time card with time worked entries

    Configure time cards to be created when a user records **Time worked**. A time card is created if a time card does not exist for the task. The time sheet policy **Auto fill time card with time worked entries** controls this behavior and is set to false by default.

    In **Project Task**, **Incident**, **Problem**, and **Change** records, the **Time worked** field does not appear by default and must be configured on the form.
  - *By scheduled job*: Configure time cards for the project tasks to be generated automatically for users through a scheduled job. Only the admin can configure a scheduled job.

  Note:
  Time cards cannot be created automatically when you use the mobile interface. Use the desktop interface if you want to use the automatic time card feature.
- *Manual*: Create a time card for each task and enter the time manually.

Procedure

1. Create the time card using one of the following options.

2. On the form, fill in the fields.

Time card form

Field	Description
Week starts on	Starting date of the week of the time sheet. **Note:** The time sheet policy **Week starts on** controls the start day of the week. A message is displayed when you try to associate a resource plan with a time card when the week on which the time card starts is outside the resource plan dates.
State	Current state of the time card. All new time cards begin in the Pending state. Different states of a time card are Pending, Submitted, Approved, Processed, Rejected, and Recalled.
Category	Type of task for which the time card is created.
Task	Task that is associated to the time card.
Project time category	Type of activity in the project that time is reported for.

Field	Description
	This field appears when you select **Project/Project Task** or **Task work** in the **Category** field.
User	Name of the user that the time card is created for.
Rate type	Rate type that is considered during the expense line generation. For more information, see Rate type in labor rate card. **Note:** This field appears when the **Allow multiple rate types** option is selected in the assigned time sheet policy of the user.
Resource Plan	Resource plan of the task associated with the time card. When the time card is approved, the hours from the time card entry and the hourly resource rate are used to update the **Actual hours** and **Actual cost** values of the associated resource plan. A resource plan is associated with the time card in the following ways: ◦ If the task has multiple resource plans, you can select the appropriate plan. ◦ If the task has one resource plan, that plan is selected by default. **Note:** ◦ If the task has no resource plan, the associated record updates the corresponding project. ◦ A message is displayed on the time card when the resource plan dates are outside the dates of the time submission week. ◦ This field appears when the **Update actual hours and cost in resource plan/reports** option is selected in the assigned time sheet policy of the user.
Remaining effort	Amount of work time left for the planned task. Remaining effort is the remaining effort in the planned task table and is calculated as: `Remaining effort = Planned effort - Actual effort` **Note:** The field is not available on the Time Card form by default. Configure the form to add this field if required.

3. Click **Submit**.

Result

- The time card is created for the selected time card period.
- If the time sheet for the week does not exist, a time sheet is created for the time card week.

What to do next

After the time card is created, the hours for that task can be incremented automatically from the **Time worked** field in the task record. The time sheet policy **Auto fill time card with time worked entries** controls the way the time is updated in time cards. The policy is set to *false* by default. If

automatic updates for time worked are not configured, the user must manually update the time card.

Related topics

Map a time card category with operational work types

Copy time cards from a previous time sheet

Another option for creating time cards is to copy them from an existing timesheet, which copies all the time cards (for project as well as non-project tasks) from a selected time sheet.

Before you begin

Role required: timecard_user

About this task

You can also perform this task from the Time Sheet Portal.

Procedure

1. Open the time sheet record that you want to create the time cards for.
2. Click the **Copy from previous time sheet** related link.
3. In the **Select Time Sheet** field, select a previous time sheet that you want to copy the time cards from.
4. **Optional:** You can copy the time logged for the tasks in the previous time sheet, select the **Copy time logged** check box.
5. Click **OK**.

Result

Time cards for project tasks

Time cards for all the project tasks that are in progress or planned in the current week are copied from the selected time sheet. If a time card for a project task exists, a duplicate time card is not created during copying.

Time cards for non-task category

Time cards for all non-task category such as meetings and trainings are copied from the selected time sheet. If a time card for a non-task category exists, a duplicate time card is created during copying.

Note:

- If the previous time sheet has rate types and the time sheet policy of the user allows rate types, then rate types are also copied into the current time sheet.
- If the previous time sheet has rate types, but the time sheet policy of the user no longer allows rate types, then rate types are not copied to the current time sheet.
- If the previous time sheet has a rate type that is inactive, that rate type is not copied. If necessary, the user can enter a different rate type before saving the individual time cards.

Auto-generate time cards

As an admin, you can configure a scheduled job to generate time cards automatically for project tasks assigned to time card users.

Before you begin

Role required: admin

About this task

In addition to the manual option **Generate Time Cards**, a scheduled job can be run to automatically generate the time cards for project tasks. A job can be scheduled to run every week, for example, on every Sunday to generate time sheets for all users for the next week. You can configure when to run the scheduled job based on the business process of the organization. By default, the scheduled job is turned off. For more information, see Schedule a script execution .

Note:
The scheduled job auto-generates time cards only for those users who have the **Auto create time cards every week** option set to true in their assigned time sheet policy.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. In the Scheduled Jobs list, select **Auto Generate Time Cards**.
3. Configure the following parameters in the script as per the business requirements:
   - run for (CURRENT_WEEK, NEXT_WEEK, LAST_WEEK)
   - Group Name (includeGroups, excludeGroups)

   Note:
   - If a time card exists for a project task for the time card week, a duplicate time card is not created.
   - Time cards are generated only for those project tasks that are in progress or planned in that week.
   - The **Allow time card reporting on** field on the project form determines the level that the time cards for the project tasks are created at.

Example:
The following example script is executed when a scheduled job is run to generate time cards for the `Current week` for the `Database` and `Hardware` groups:

```
// Clone this schedule job to run separately for different
 groups at different times
// One of the following values
// TimeCardConstants.CURRENT_WEEK
// TimeCardConstants.NEXT_WEEK
// TimeCardConstants.LAST_WEEK
var runFor = TimeCardConstants.CURRENT_WEEK;

// Comma separated group sys ids or group names.
 TimeSheet/TimeCards will be auto generated for time card users
 in the given groups.
var includeGroups = [Database,Hardware];
var excludeGroups = [];

var generator = new TimeCardGenerator();
generator.generateFromConfig(runFor, includeGroups,
 excludeGroups);
```

Submit a time card

As a time card user, once a time card for the week is updated with the time worked, you can submit the time card individually.

Before you begin

A time card in the Pending or Rejected state can be submitted.

Role required: timecard_user

About this task

When a user submits a time card, the time sheet policy assigned to the user is used for validation. If there is no time sheet policy assigned to the user, the default time sheet policy is used. Business rules prevent a user from submitting a time card in case any violation is found.

Procedure

1. Navigate to:
   - **Time Sheets** > **My Time Cards** > **All**
   - The **Time cards** related list in a time sheet.

   Note:
   You can select multiple time cards to submit.

2. In the Time Card list, open the time card to submit.
3. Click **Submit Time Card**.

Result

- The time card moves to the Submitted state.
- The time card is submitted to an approver based on the values set in the **Project time approver** and **Non-project time approver** fields in the time sheet policy assigned to the user.

  Note:
  If the project or user manager is changed after submitting the time card, the time card is automatically redirected to the new project or user manager for approval.

Approve or reject a time card

As a time card approver, you can view and approve or reject a submitted time card.

Before you begin

Role required: timecard_approver or timecard_admin

About this task

You can view only those time cards that are submitted to you for approval or you have been delegated as the time card approver by one of the other user managers. The time sheet policy assigned to the user governs the approval process.

Procedure

1. Navigate to **All** > **Time Sheets** > **Time Cards** > **Pending Approval**.
2. Select the time card to approve or reject.
3. Click **Approve** or **Reject**.

Result

- The time card moves to the Approved or Rejected state.
- Expense line is created for an approved time card.
- For an approved time card, a Time Card Daily record is created.

Note:

- A user manager can also approve the time cards within a time sheet by selecting one or more time cards and approving them. When all the time cards in a time sheet are approved, the time sheet automatically moves to the Approved state.

  A user manager is one who is selected in the **Manager** field on User form of the time card user.

- If a time card is in the Rejected state, the associated time sheet also moves to the Rejected state irrespective of the state of other time cards in the time sheet.

  Note:
  If a time sheet is rejected, it gets listed in the notifications for the rejected time sheets under notifications icon on Time Sheet Portal.

- If **Project time approver** is set to **Both** in the assigned time sheet policy, the time card is in the Submitted state only if both approvers approve it. If one of the approver rejects, the time card state is set to Rejected and the **Approved by** field is cleared, requiring approval of both approvers on resubmission.

Record time worked

The time card retrieves time accrued on a project or spent working on any record in the Task table from the **Time worked** field.

This field does not appear by default on the Project Task, Incident, Problem, and Change forms and must be added by personalizing the form. Time recorded in this field is used to populate an existing time card or to create a new time card if one does not exist. A time sheet policy controls this behavior.

Note:

When time worked is updated, a time card is created only if:

- the user is a time card user, AND
- the user is assigned to the task through **Assigned to** or **Additional assignee list** field.

The *Time worked* field has a counter that acts like a stopwatch for the duration of the time spent in the record. A button in the field can stop and start the counter. By default, the *Time worked* counter is enabled and begins recording the elapsed time when the record is opened. Stop the counter with the stop button and restart it with the play button.

If you are creating time cards from time worked entries, you can ask your admin to add the **Time Worked** related list to display the time worked records on the time card form. You will also notice an informational message on the time card to let you know that changes to time worked records overrides values in the time card. This is displayed using a Create a formatter and add it to a form, which can be added or removed by configuring the form.

Manage costs

When the cost management feature is enabled, time cards can be used to manage the cost of labor in the Cost Management application.

When a time card for a project task is approved, an expense line is generated for the corresponding labor cost associated with the project. If the project has a Rate Models associated, then the hourly rate for calculating labor cost for the expense lines is derived from the rate model. If no rate model is associated to the project, the rate listed in

the *Labor Rate Card* is considered. If no Labor Rate Cards is applicable, the property `com.snc.time_card.default_rate` defines the default hourly rate.

Roles

The timecard_admin role enables users to approve, modify, and delete the time cards of other users.

Related topics

Activate Cost Management

Domain separation and Time Card

Domain separation is supported in Time Card. 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.

For more information about using domain separation with Time Card, see https://community.servicenow.com/community?id=community_article&sys_id=616d706edbb318d066f1d9d9689619c1.

Related topics

Domain separation for service providers

Creating database views for reporting

A database view defines table joins for reporting purposes.

For example, a database view can join the Incident table to the Metric Definition and Metric Instance tables. This view can be used to report on incident metrics and may include fields from any of these three tables.

Several useful database views are installed with the Database View plugin and the Database Views for Service Management plugin. These database views cover most metric reporting needs and greatly reduce the need to define new ones.

Any user who can create a report can use database views as the report source, but ACLs on the underlying tables are honored.

Note:

- The accumulated impact on performance grows as the number of tables that are included in the view and the number of records that those tables contain increases. To maximize the performance of the database view, ensure that the 'where' clauses that are defined in the database view are based on indexed fields.
- A database view is not treated like a custom table, so there is no licensing impact.
- Database view tables are not included in FTP exports.

Limitations

- Database views cannot be created on tables that participate in table rotation.
- It is not possible to edit data in the database view output.
- Database view tables cannot be added as a data preserver in clone requests.
- You can reference a table or database view from a different application scope in a Table Name field. However, if the field belongs to a table that extends sys_metadata, the table or database view must belong to the same application scope as that table.

ACLs and database views

You do not need to create ACLs on fields in the view. The system honors contextual ACLs (ACLs with a condition or script) that exist on the underlying table. Non-contextual ACLs (ACLs with only role checks) are still honored just as with previous releases.

To require explicit read ACLs be added to the database views, set the `glide.security.expander.view.legacy` property to **true**. On upgraded instances, add this system property and set it to **true**.

You can still create additional ACLs on the database views. These ACLs are evaluated last and are always honored.

Database view reserved words

Using the terms may cause unintended or undesirable performance. For more information, see the MySQL reserved words document.

Database view creation

Create a database view to join tables. You can then create a report based on the database view.

Use the **Database Views** module to create a database view that joins tables to create a pseudo-table that combines data from each table in the view. You can create a report on this data by adding the database view as a table in a report. Accessing the database view does not require database access.

Create a database view

Create the database view.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **System Definition** > **Database Views**.
2. Select **New**.

3. Name the view as you would name a new table.
   The application converts capital letters to lower case and spaces to underscores.
4. Provide a label and a plural.

   The **Label** and **Plural** fields define how the database view is labeled in lists and forms. For example, give a database view with the label Duplicate installation the plural `Duplicate installations`.
5. **Optional:** Provide a description of the database view so that other users know its purpose.
6. Select **Submit**.

Result

The database view is created and accessible from the **Database Views** module.

What to do next

Add a table to the database view

Add a table to the database view

Specify the table to join to the database view.

Before you begin

Role required: admin.

About this task

The **Table** field in the View Table form specifies the table to join to the database view. The **Variable prefix** is used by the **Where clause** to define the conditions for the join. These conditions can refer to any field, but typically define the join by matching a field in the table to a field in another table in the database view.

> **Note:**
> The Variable prefix field is mandatory.

When you write the **Where clause**, add the field name to the **Variable prefix** of its table with an underscore. In the following example, **mi_id** refers to the **id** field in the Metric Instance [metric_instance] table (mi) and **inc_sys_id** refers to the **sys_id** field in the Incident [incident] table (inc).

Database View Table form

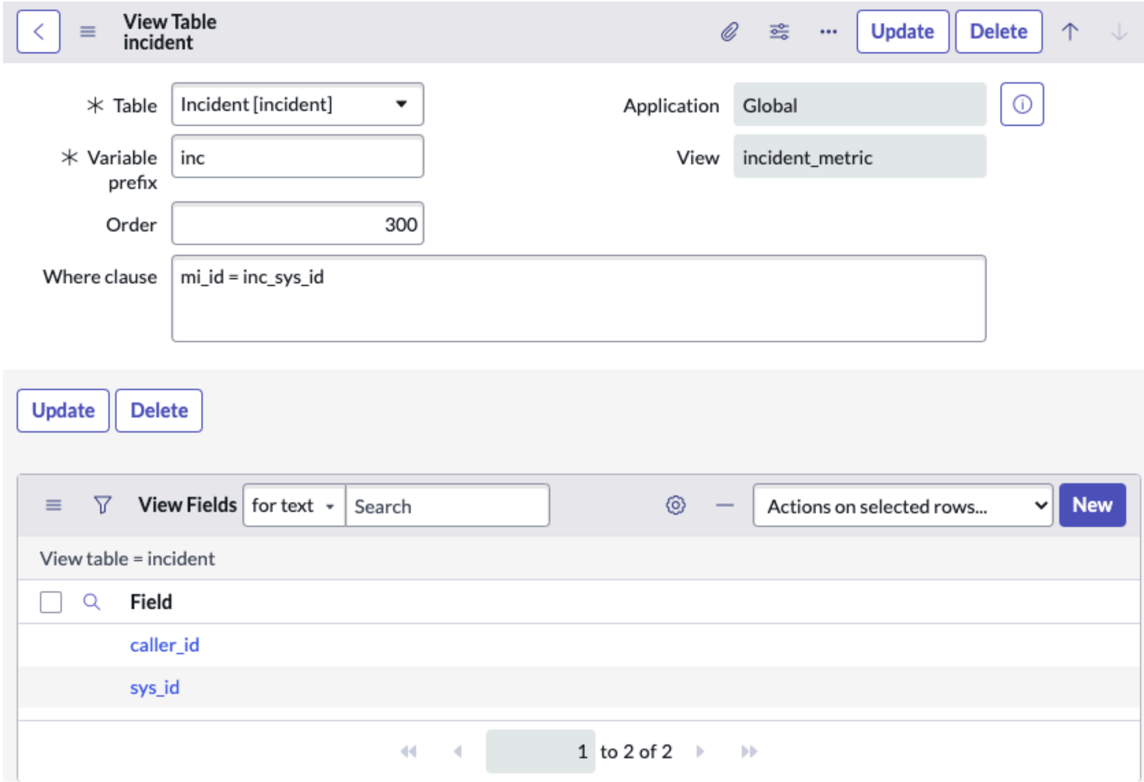

Note:

- Database views cannot be created on tables that participate in table rotation.
- Use only lowercase characters in the **Variable prefix**. Using uppercase characters may prevent you from viewing the database view in a list.

To see an example of using left joins to create a view, see Example left join in creating a database view.

The **Where clause** supports these JavaScript conditional operators:

- =
- !=
- <
- <=
- >
- >=
- &&
- ||

There is no support for LIKE or CONTAINS conditions in the **Where clause** clause of a database view. Link tables based on the sys_id using the = operator to work with the full dataset when filtering lists, running glide queries, or running reports.

Procedure

1. From the Database View form, select **New** on the View Tables related list.
2. Configure the form and add the **Left join** field (a check box) to the form.

3. Select **Save**.
4. Complete the form and select the **Left join** check box.

    - Selecting **Left join** causes the left-hand table in the database view to display all records, even if the join condition does not find a matching record on the right-hand table. Select this check box to view tables that specify a **Where clause**. Selecting **Left join** for view tables without a **Where clause** does not affect the query.
    - Joined tables are ordered left to right from lowest to highest **Order** values.
5. Select **Submit**.
6. Personalize the **View Tables** related list to show the **Left join** column.
   The **Left join** field is **true**.
7. Select a record to view a table.
8. To add an OR to your Where clause, use ||.

   For example, to query all incidents related to RFCs OR all incidents that are the parent of a change request, use the syntax: `inc_rfc = chg_sys_id || chg_parent = inc_sys_id`.

Example left join in creating a database view

This example shows the proper settings when using left-joins to add tables to a database view.

Before you begin

Role required: admin

About this task

The following procedure shows how to create a database view that includes a list of Catalog Tasks and their parents. Most of the steps take place on the View Table form.

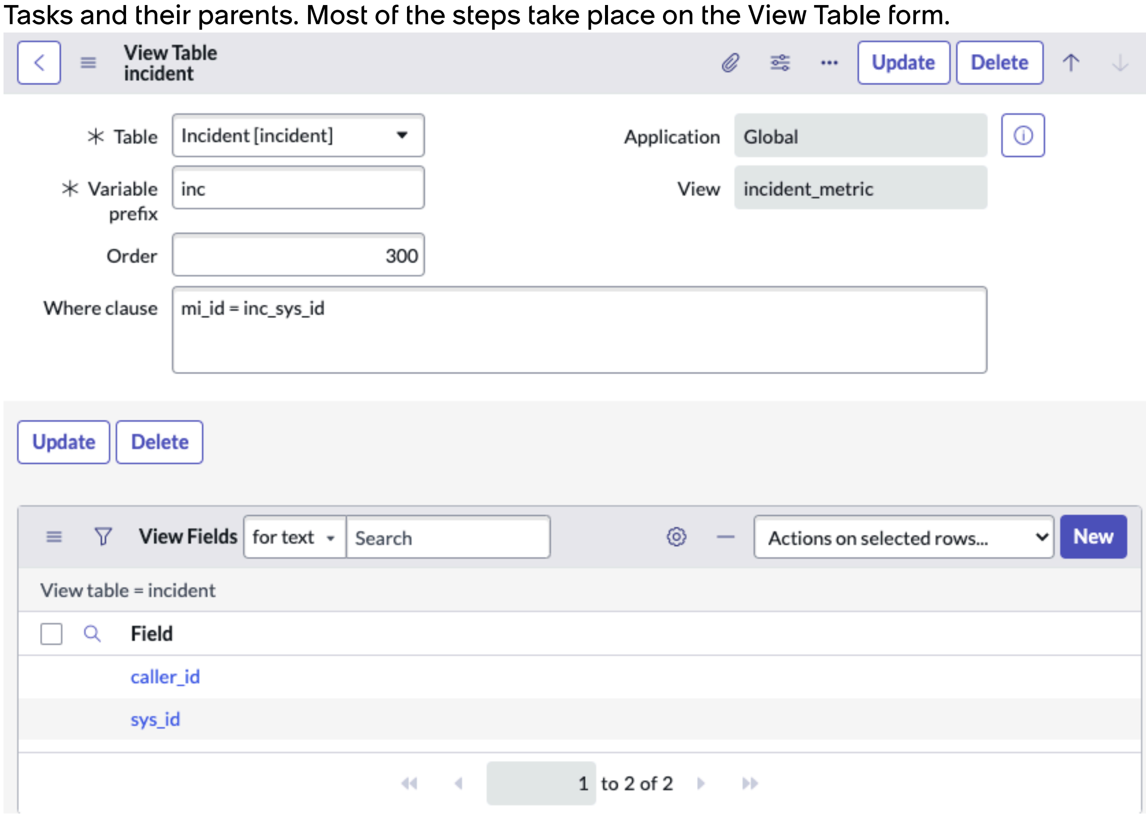

Procedure

1. Navigate to **All** > **System Definition** > **Database Views**.
2. Select the view you want to edit.
   The Database View page appears. In the View Tables related list, specify the tables you want in the database view.
3. In the **View Tables** related list, select **New**.
4. Add a **Left join** check box to the form by Configuring the form layout.
5. Add the Catalog Task [sc_task] table to the database view by filling in the form with the following data and then select **Submit**. Joined tables are ordered left to right from lowest to highest **Order** values.

Field	Value
Table	Catalog Task [sc_task]
Variable prefix	cat
Order	100
Where clause	Leave blank

   The Database View page appears with the Catalog Task table in the View Tables table.
6. In the **View Tables** table, select **New** and add the Requested Item [sc_req_item] table to the database view by filling in the form with the following data and then selecting **Submit**. Joined tables are ordered left to right from lowest to highest **Order** values.

Field	Value
Table	Requested Item [sc_req_item]
Variable prefix	item
Order	200
Where clause	cat_parent=item_sys_id
Left join check box	True

   The Database View page appears with the Requested Item table in the View Tables table.
7. In the **View Tables** table, select **New** and add the Request [sc_request] table to the database view by filling in the form with the following data and then selecting **Submit**.

Field	Value
Table	Request [sc_request]
Variable prefix	req
Order	300
Where clause	cat_parent=req_sys_id
Left join check box	True

   The Database View page appears with the Request table in the View Tables table.

8. In the **View Tables** table, select **New** and add the User [sys_user] table to the database view by filling in the form with the following data and then selecting **Submit**.

Field	Value
Table	User [sys_user]
Variable prefix	user
Order	400
Where clause	cat_opened_by=user_sys_id
Left join check box	False

   - The Database View page appears with the User table in the View Tables table.
   - If the parent record of Catalog Task is a Requested Item, all the fields in the Request table will be blank.
   - If the parent record is a Request, all the fields in Requested Item will be blank.
   - Because of the two left joins, the Catalog Task record returns even if the parent is empty or is not a Requested Item or Request.
   - Because User is not a left join, there must be a matching user in sys_user for the row's Opened By field for the Catalog Task row to return.
9. **Optional:** Perform the following steps to experiment and learn more about how left joins impact database views.
   If you make the left join in:
   - User true, rows that have an empty **Created By** return.
   - Requested Item false, only Catalog Tasks that have a Requested Item as a parent return.
   - Request false, only Catalog Tasks that have a Request as a parent return.
   - Requested Item and Request false, no rows return because nothing can have a Requested Item and a Request as parents.

Specify a field to return

Restrict or specify a field that you want returned by the joined table.

Before you begin

Role required: admin

About this task

You can restrict the fields displayed in the database view output by adding fields to the View Tables related list.

When you restrict fields, you must create a View Field record for the join field from the *Where clause* in the parent record. If you omit a record for this field, it cannot be returned, and the join fails.

Procedure

1. Navigate to **All** > **System Definition** > **Database Views**.
2. Select the database view you want to add a field to return to.
3. In the database view record, select a view table, and then select **New** in the View Fields related list.

4. In the View Field form, select the field you want to display in the database view output.
If no fields are defined in the View Fields related list, all fields are returned. If any fields are defined, then only those fields are returned. For example, you can restrict the display if there are fields that should not be viewable to all users.

View Field form

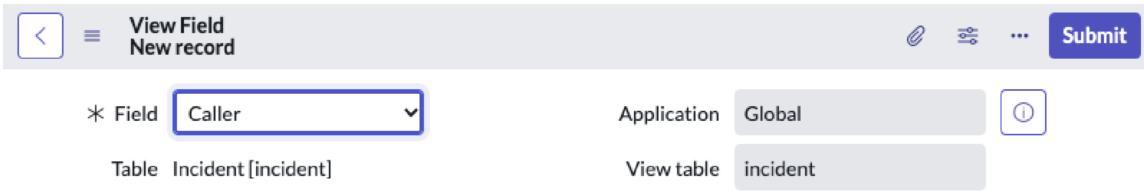

In this example, the *Where clause* uses the *sys_id* field from the Incident table to establish the join. For the join to succeed with a restricted field list, you must include a record for the *sys_id* field in the View Fields related list.

What to do next

Verify your changes using the **Try It** related link in the database view record. See Test the database view.

Relabel a column

Sometimes, two different tables may have fields of the same name that are both important (such as two tables with a sys_updated_on field). For clarity, relabel one of the fields.

Before you begin

Role required: admin.

About this task

To create clear reports, relabel the fields on the Database View [sys_db_view] table without changing the names of the fields.

Procedure

1. Navigate to **All** > **System Definition** > **Language File**.
2. Select **New**.
3. Fill in the form as follows:

New language file fields

Name	Description
Table	Name of the database view
Label	Display label
Plural	Plural form of the display label
Element	Name of the field on the database view in the format variable prefix_fieldname. For example, you may enter inc_name to indicate the Name field on the Incident table. **Variable prefix** is a required field when you define the tables in a database view.

4. Select **Submit**.

Specify the number of records to return

Specify the number of records to return for a database view when the view is used in a script.

A property called *glide.db.max_view_records* controls the maximum number of rows returned when running a GlideRecord query in a script. The default value for this property is 10000. To change this value, add the property to the System Property [sys_properties] table and edit the property's **Value** field, which determines the number of rows to return.

This property only applies when querying a database view table in a script. When the database view table is used in a list or report, this property doesn't apply. Reports or lists based on the database view use all rows in the view.

Test the database view

Verify that the database view works correctly.

Before you begin

Role required: admin.

About this task

Procedure

1. Navigate to **All** > **System Definition** > **Database Views**.
2. Select the database view that you want to test.
3. Test the view by clicking the **Try It** related link.

   Note:
   If you do not see the *Try It* link, the tables necessary for the view do not exist. If this problem occurs, it is possible that you did not activate the necessary plugins to create the supporting tables.

   The database view runs and a list of items is displayed. If you have specified fields in the **View Fields** related list for each table, the output is limited to those fields. If you specified view fields and the view does not display correctly, review the **View Fields** related list for each table and ensure that you include fields from the **Where clause**.

Example: Output of the sample Incident Metrics database view

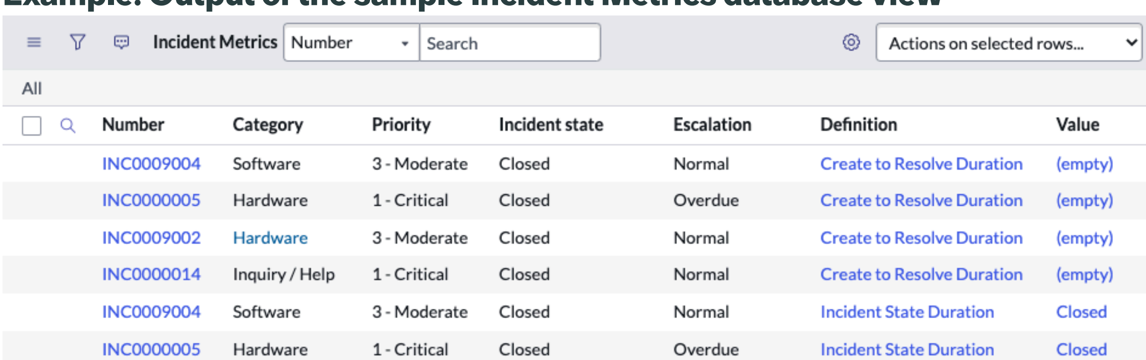

Number	Category	Priority	Incident state	Escalation	Definition	Value
INC0009004	Software	3 - Moderate	Closed	Normal	Create to Resolve Duration	(empty)
INC0000005	Hardware	1 - Critical	Closed	Overdue	Create to Resolve Duration	(empty)
INC0009002	Hardware	3 - Moderate	Closed	Normal	Create to Resolve Duration	(empty)
INC0000014	Inquiry / Help	1 - Critical	Closed	Normal	Create to Resolve Duration	(empty)
INC0009004	Software	3 - Moderate	Closed	Normal	Incident State Duration	Closed
INC0000005	Hardware	1 - Critical	Closed	Overdue	Incident State Duration	Closed

What to do next

To use the database view in a data visualization, select it from the data sources. Database views are listed under tables.

You can create any visualization type using the database view. This example shows the database view visualized as vertical bar.

Displaying function results in a database view

Enhance the display of a database view by adding a function field to the output to display function results.

Function fields perform common database transformations and calculations on fields in the system. Function fields do not store any data, but rather display the results of a database query when the function is executed.

For example, you could identify incidents with ineffective descriptions of fewer than 10 characters by adding a function field that calculates the length of the **Short description** field.

Create a function field to perform a database function

Create a function field that displays the results of a database function, such as a mathematical operation, field length computation, or day of the week calculation.

Before you begin

Role required: admin

You can either select a function field from a physical table or create a function field directly on the database view.

If you create a function field directly on the database view, you must use the variable prefixes from the database view in the function definition to identify the columns. Before you create the function field, take note of the variable prefixes in your database view.

For example, if the variable prefix for the Incidents table is `inc2` and you want to return the length of the description field, the function definition would be `glidefunction:length(inc2_description)`. If the variable prefix for the Metric Instance table is `mi` and you want to concatenate the Definition and Value fields, the function definition would be `glidefunction:concat(mi_definition.name,' ',mi_value)`.

For information on supported functions, see Report on function fields.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Select **New**.
3. Select **Function field**.
4. On the form, fill in the fields.

Function fields form

Field	Description
Table	The table or database view to which the function field is added. If you select a database view, you must include the database view's variable prefixes in the **Function definition** field.
Type	The type of field: string, number, date, and so on.
Column label	The label for the column.
Column name	The name of the column. This field is automatically set to the value of the **Column label** field.
Max length	The maximum length of the return value.

Field	Description
Function field	If selected, creates a field that displays the results of a database function, such as a mathematical operation, field length computation, or day of the week calculation. Once the new function record is saved, you cannot clear the check box to make the field a regular field.
Function definition	The function that the field performs. The format for functions starts with `glidefunction:`, followed by the operation to be performed (for example, `concat`), followed by function parameters. Constants must be enclosed in single quotes. **Important:** If you create the function field on a database view and not a physical table, you must include the database view's variable prefixes in the **Function definition** field.

5. Select **Submit**.

Example: Create a Function field on the Incident Metric database view

In this example, a new Function field, CallerDesc, is added to the Incident Metric database view. The field concatenates the caller name, a hyphen, and the short description.

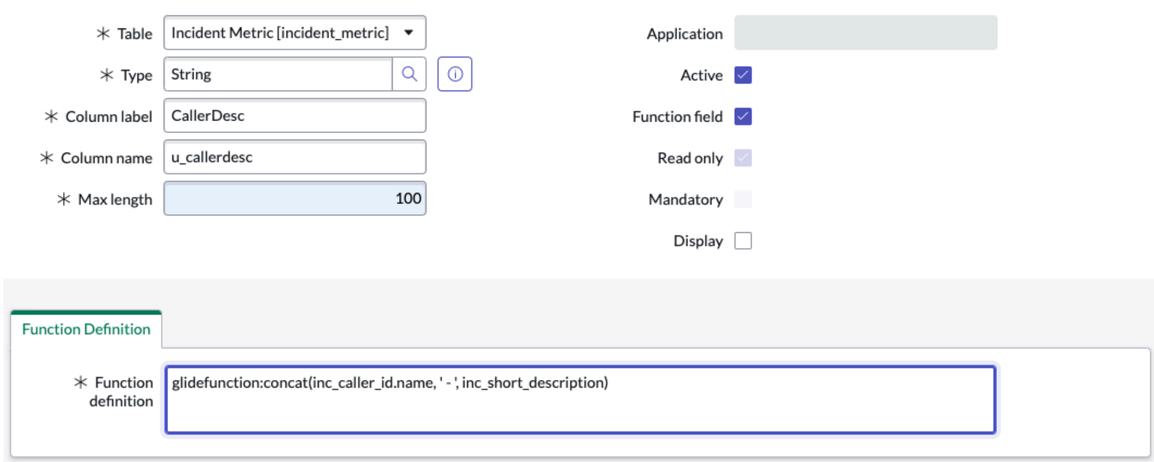

The variable prefix `inc` is included because the function field is defined on the Incident Metric database view and not on a physical table. The prefix is necessary to distinguish which short_description and caller_id columns to use from the tables that are part of the database view.

What to do next

Display function results in a database view

Display function results in a database view

Add a function field to the output of a database view to see function results.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Database Views**.
2. Select a database view.
3. Select the **Try it** related link.
4. Add the function field to the database view output.
   See Configure items on forms or in lists using a list collector.
5. Review the function definition if the expected value is not returned or a returned value indicates the function definition is invalid.
   - If the message **Invalid function** is returned, the function definition is invalid for string field types.
   - If an empty value is returned, the function definition is invalid for other field types.
   - If an unexpected value is returned, review the field definition in the function definition.

Result

The function field operation is executed and the results appear in the database view output.

Use disjunctions in complex queries

ServiceNow performs conjunction statements before disjunction statements in a query.

When you create a complex query, you must use parenthesis around disjunctions where appropriate to ensure proper grouping of query elements. For example, you must use parenthesis in the query `(md_table = 'incident' || md_table = 'task') && mi_definition = md_sys_id && mi_id = inc_sys_id`. Removing the parenthesis from this query returns all records where the md_table value is *incident*.

Database views in the base system

Certain views are included in the base system with the Database Views and Database Views for Service Management plugins.

Database views

Name	Description	Label
change_request_metric	Join change to metric definition to metric instance, creating a view that can be reported on for things like: Changes that were closed by category.	Change Metric
change_request_sla	Join change_request to sla(task_sla), creating a view that can be reported on for things like change request resolved by sla per change category.	Change Request SLA
change_task_metric	Join change task to metric definition to metric instance, creating a view that can be reported on for things like: Change tasks that were closed by change state	Change Task Metric

Database views (continued)

Name	Description	Label
change_task_sla	Join change_task to sla(task_sla), creating a view that can be reported on for things like change tasks resolved by sla.	Change Task SLA
change_task_time_worked	Join change task to task time worked to pull time worked entries associated with incidents.	Change Task Time Worked
incident_metric	Join incident to metric definition to metric instance creating a view that can be reported on for things like: Incidents that were resolved on the first call by category	Incident Metric
incident_sla	Join incident to sla(task_sla) to report on things like incidents resolved by sla per incident category.	Incident SLA
incident_time_worked	Join incident to task time worked to pull time worked entries associated with incidents.	Incident Time Worked
pm_project_metric	Join pm_project to metric definition to metric instance creating a view that can be reported on for things like: Projects that were closed by name or date	Project Metric
pm_project_sla	Join pm_project to sla(task_sla) to report on things like project names by sla.	Project SLA
pm_project_task_metric	Join pm_project_task to metric definition to metric instance creating a view that can be reported on for things like: Project tasks that were closed by name or date	Project Task Metric
pm_project_task_sla	Join pm_project to sla(task_sla) to report on things like release tasks by sla.	Project Task SLA
pm_project_task_time_worked	Join pm_project_task to task time worked to pull time worked entries associated with project tasks.	Project Task Time Worked
problem_metric	Join problem to metric definition to metric instance creating a view that can be reported on for things like: Problems that were resolved on the first call by category	Problem Metric
problem_sla	Join problem to sla(task_sla) to report on things like problems resolved by sla per problem state.	Problem SLA
release_feature_metric	Join release_feature to metric definition to metric instance creating a view that can be reported on for things like: Release Features that were closed by product	Release Feature Metric

Database views (continued)

Name	Description	Label
release_project_metric	Join release_project to metric definition to metric instance creating a view that can be reported on for things like: Releases that were closed by category	Release Metric
release_task_metric	Join release_task to metric definition to metric instance creating a view that can be reported on for things like: Release Features that were closed by feature	Release Task Metric
release_task_sla	Join release_task to sla(task_sla) to report on things like release tasks by sla.	Release Task SLA
sc_request_metric	Join sc_request to metric definition to metric instance creating a view that can be reported on for things like: Requests that were closed by category	Catalog Request Metric
sc_request_sla	Join sc_request to sla(task_sla) to report on things like requests by sla.	Catalog Request SLA
sc_req_item_metric	Join sc_request_item to metric definition to metric instance creating a view that can be reported on for things like: Request Items that were closed by item	Catalog Request Item Metric
sc_req_item_sla	Join sc_req_item to sla(task_sla) to report on things like request items by sla.	Catalog Request Item SLA
sc_task_metric	Join sc_task to metric definition to metric instance creating a view that can be reported on for things like: Catalog tasks that were closed by item	Catalog Task Metric
sc_task_sla	Join sc_task to sla(task_sla) to report on things like tasks by sla.	Catalog Task SLA

Data management

Data is stored and managed in your instance according to a principled structure that administrators can view and configure. Data management functions include importing, exporting, and archiving database data and configuring fields and tables.

Database structure

All of the information in the instances is stored in tables, which consist of a series of records. The record in turn holds a series of fields that hold the individual bits of data and can be viewed either as a list or a form.

Tables can be related to each other in the following ways:

Extensions

A table can extend another table. The table doing the extending (child class) includes all of the fields of the other table (parent class) and adds its own fields. For

instance, the Incident [incident] table has all of the Task [task] table fields (because an incident is a special form of task) and has its own incident-specific tasks.

One-to-Many

Within a table, a field can hold a reference to a record on another table. There are three types of one-to-many relationship fields:

- Reference Field: allow a user to select a record on a table defined by the reference field. For instance, the **Caller** field on the Incident table allows the user to select any record on the User table.
- Glide List: allows a user to select multiple records on a table defined by the glide list. For instance, the **Watch list** field on the Incident table allows the user to select records on the User table.
- Document ID Field: allows a user to select a record on any table in the instance. These fields are much less common, but one example is the **Document** field on the Translated Text [sys_translated_text] table.

Many-to-Many

Two tables can have a bi-directional relationship, so that the related records are visible from both tables in a related list.

Database Views

Two tables can be joined virtually with Database Views to enable reporting on data that might be stored over more than one table.

Data management tools

There are a number of tools that can help manage data within the instance.

Schema map

The Schema Map displays the relationships between tables visually, helping to navigate through the database structure.

The Schema Map provides an interface for viewing the relationships between tables. The inter-table relationships it captures include many-to-many relationships, tables that extend other tables, and tables that reference other tables through reference fields.

Data dictionary tables

Data dictionary tables holds information that defines the database and can be accessed for information on the database schema.

These tables hold important information on the database and its structure:

- `Tables [sys_db_object]`: contains a record for each table.
- `Dictionary Entries [sys_dictionary]`: contains additional details for each table and the definition for every column on each table. Each row represents either a column on a table or a table.
- `Field Labels [sys_documentation]`: contains the human-readable labels and language information.

Data management plugins

Plugin	Description
Data Archiving	Provides the ability to archive records to minimize performance issues.
Database Rotations	Provides tools for managing large tables to minimize performance issues.
Many to Many task relations	Provides the ability to define many-to-many relationships between task tables.

Unique record identifier (sys_id)

A unique 32-character GUID (Globally Unique ID), called a Sys ID (sys_id) identifies each record in an instance. There are various ways a record is assigned a sys_id.

When created within the application, sys_id values are unique. The ServiceNow AI Platform and database should manage all operations on sys_id values. The same sys_id value is never generated twice, which ensures that every record created in every table in every instance in the world has a unique identifier. If two records have the same sys_id value, it occurs as a result of the following situations:

- If a record with the sys_id was copied to the other at the database level outside of the ServiceNow AI Platform.
- If a record with the sys_id was copied using an Update Set or via XML, its sys_id is the same.

Note:
A new record has a sys_id of -1, and once inserted, it is assigned a new sys_id. The sys_id is not meant to show as a field on a form or as a column in a list. To learn more, see the Users are unable to add the sys_id field to a form or list view [KB0690876] article in the Now Support Knowledge Base.

Get the sys_id from the header bar

Users can locate the sys_id of a record using the header bar.

Procedure

1. Navigate to the record.
2. Right click the header bar and select **Copy URL**.

   The sys_id is inside of the URL, after the parameter sys_id=. For example, the following is a URL for an Incident:

   ```
   https://<instance
    name>.service-now.com/nav_to.do?uri=incident.do?sys_id=9d38501
   7c611228701d22104cc95c371
   ```

   Therefore the sys_id is 9d385017c611228701d22104cc95c371.

Get the sys_id from a script

Users can locate the sys_id of a record using a script.

Procedure

- The sys_id value of a record can be found in a business rule (or any other server-side JavaScript) by dot-walking from the GlideRecord.

```
var id = current.sys_id;
```

- The sys_id of a record can be found in client-side JavaScript using g_form.getUniqueValue() as shown in the following example.

```
function onLoad() {
    var incSysid = g_form.getUniqueValue();
    alert(incSysid);
}
```

Get the sys_id from the URL

Users can locate the sys_id of a record by viewing the URL.

About this task

Since the sys_id of a record is always part of the URL for a link to that record, it is possible to retrieve the sys_id by viewing the URL.

Procedure

View the sys_id in the information bar of the browser by hovering over a link to the record. For example, an Incident with the following URL: https://<instance name>.service-now.com/nav_to.do?uri=incident.do?sys_id=23dc968f0a0a3c1900534f399927740e, has this sys_id: *23dc968f0a0a3c1900534f399927740e*.

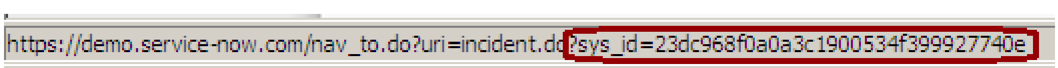

Data archiving

Data archiving involves managing table size growth and archiving old data. It moves data that is no longer needed every day from primary tables to a set of archive tables.

The longer an instance runs, the more likely it is to accumulate data that is no longer relevant. For example, task records from two years ago are typically less relevant than currently active tasks. Old data may eventually cause performance issues by consuming system resources and slowing down queries and reports.

If you cannot delete this data because you need it for auditing or for historical purposes:

1. Create an archive rule that archives the data and removes it from immediate access to free up system resources.
2. Create a destroy rule to delete the data after a specified period.
3. Configure archive and destroy rule properties that control batch processing of the archive and destroy rules.

You can archive records in core tables such as the Task [task] table and records in custom tables that you create on the ServiceNow AI Platform.

To archive Configuration Management Database (CMDB) CI records, use the Working with CMDB Data Manager.

Sample Benefits of Archiving Data

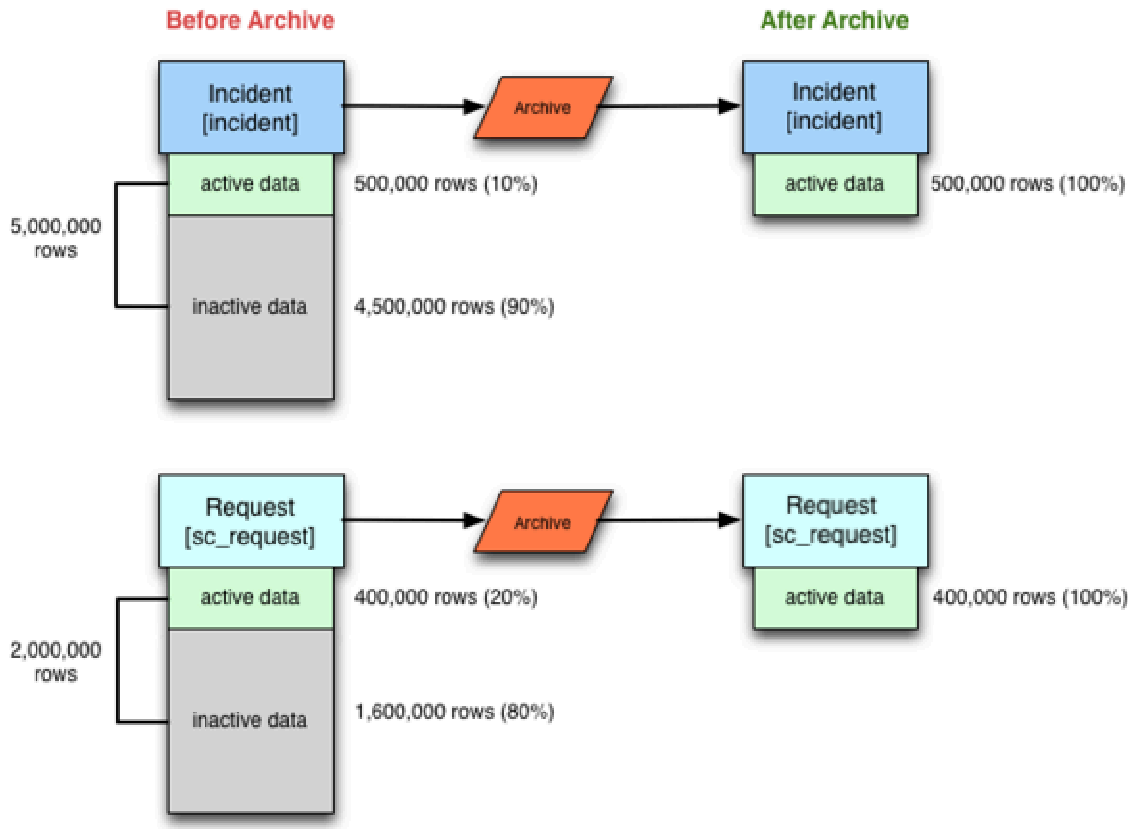

Data archiving supports Domain separation for service providers. For example, incidents that belong to a domain keep their domain designation even after they are archived.

Note:
If you want to archive emails, activate the Email retention plugin and use the archive and destruction rules that come with the plugin. Do not use the archive feature to create your own archiving rules on the email table.

Tables and modules created by an archive rule

The first time you activate an archive rule, it performs the following actions:

- Creates the archive table in the database. The archive table has the same name as the primary table with an "ar_" prefix. For example, if you archive the Incident `[incident]` table, then the archive table is `[ar_incident]`.
- Stores an XML version of each archived record in the sys_archive_log table. This archive log is the same table for all archive rules, and you cannot alter this behavior. It is also the only place where the sys_id is stored together with the display value for reference fields.

For example, for ar_incident `<assigned_to>Fred Luddy</assigned_to>`, the sys_archive_log record is as follows:

```
<assigned_to display_value="Fred
 Luddy">5137153cc611227c000bbd1bd8cd2005</assigned_to>
```

- Converts multiple joined tables into a single flat-file archive table. The archive table no longer consists of a base and extended tables.
- Converts reference field values (values set by references to records in other tables) into string values. The archive record contains the display value of the reference field at the time of the archive.
- Adds a module to the **Archive Tables** list in the **System Archiving** application. The module name is a combination of the word "Archive" plus the display name for the archived table. For example, the archive module for the Attachment `[sys_attachments]` table is **Archive Attachment**. Click the module name to view records in the archive table.
- Creates a list of the archive table using the default list view.
- Creates a form for the archive table using the default form view. The form excludes any dot-walking fields such as **Caller ID.Email**.

Conversion of Multiple Joined Tables into a Flat Archive Table

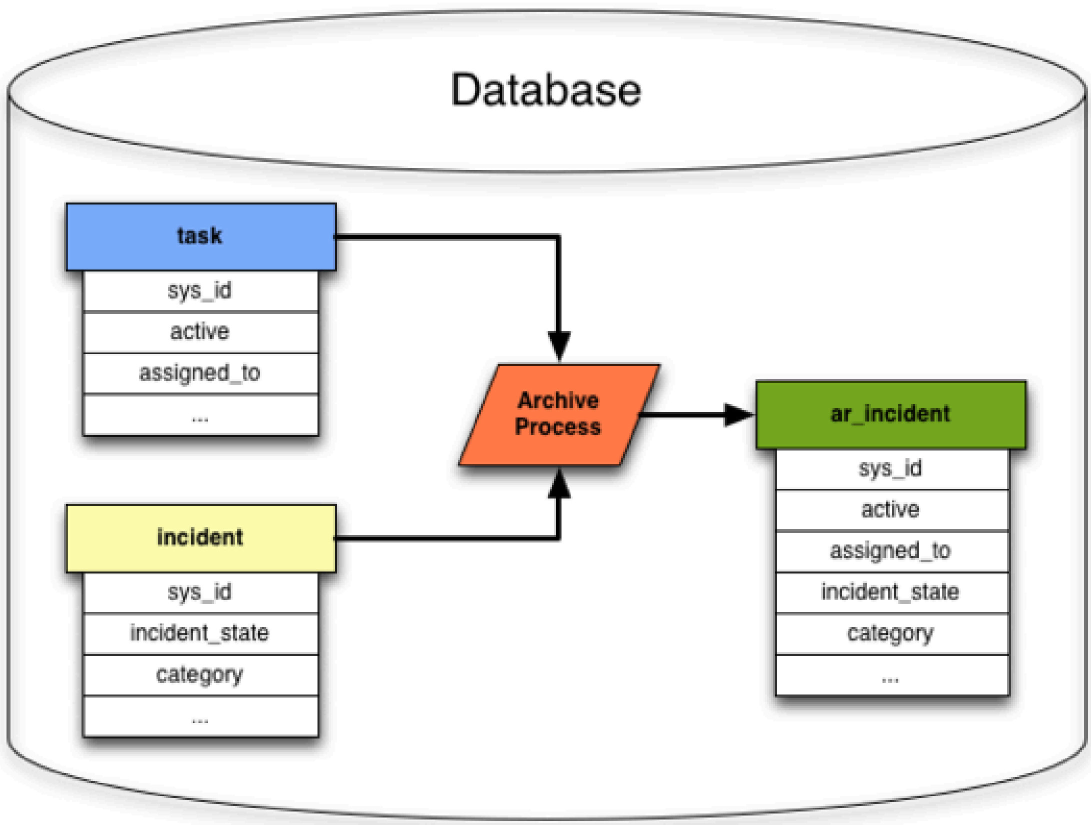

Reference values converted to strings

Archived data is stored as a flat file with no reference fields to other tables. The archive process converts any references to other tables to string values.

In the case of a reference field, the string uses the display value such as the caller's user name. For example, the **Caller** reference field in an incident would display the string *ITIL User*. If the reference was a document ID and the archive rule included the option to archive related document IDs, then the string is the document ID of the related record.

It is important to note that archive records do not receive any future changes to referenced values. For example, if you change the user name for "John Smith" to "John A Smith", all active incident records automatically show the caller as "John A Smith" because of the reference

between the Incident and User tables. However, all archived incident records display the user name that existed at the time of the archive. Any incident for "John Smith" continues referencing this user. Likewise, if you delete a user from the system, current incidents no longer display the deleted user as a caller. However, there can be archived incidents that still display the string "John Smith" as the user because the user existed at the time of the archive.

Query archived data

Archived tables are not optimized for ad hoc queries. They only contain index entries for the display value, creation date, and the primary key of sys_id.

For this reason, do not make on-demand queries against an archived table, such as searching for all priority 1 archived incidents. Instead, only search against the indexed fields. For example, search for incident INC100001 or incidents created on a specific date.

Archive tables and ACLs

By default, archive tables use the ACLs for the unarchived table of the same name. For example, the archived Incident `[ar_incident]` table uses the ACLs defined for the unarchived Incident `[incident]` table.

You can manage access to archive tables explicitly by creating ACLs for specific archive tables and setting the *glide.security.enable_archive_table_acls* property to **true**. The system then follows one of two paths:

1. If one or more active ACLs are defined for an archive table, those ACLs control access to the archive table.
2. If no ACLs are defined for an archive table, the system reverts to default behavior and uses the ACLs for the unarchived version of the table.

> **Note:**
> The two paths are mutually exclusive: If archive table ACLs deny access, the system doesn't attempt to revert to the default behavior.

The read operation read is the only operation evaluated, and other operations are prevented.

The Execution Plan UI is aware of this logic and presents information accordingly. For example, adding the first ACL to an archive table shows that the archive table ACL is "masking" ACLs on the unarchived (original data) table.

If you have existing ACLs on archived tables, they are ignored unless you set the *glide.security.enable_archive_table_acls* property to **true**. Those newly activated ACLs may possibly cause access issues. To prevent this occurrence, the system sets the *glide.security.enable_archive_table_acls* property as follows:

- Instances without the *glide.security.enable_archive_table_acls* property use the default value of **false**.
- Upgraded instances don't install the property. The property must be added manually and set to **true** to enable the archive table ACL behavior.
- Zbooted instances install the property and set it to **true**.

Create an archive rule

The System Archiving application includes several sample archive rules that illustrate the archive features. Related records can be added to an archive rule, and you can view an estimate of the number of records the rule affects.

Before you begin

Role required: admin

About this task

The sample archive rules are inactive by default.

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Rules**.
2. Select **New**.
3. Fill in the fields as appropriate.

Archive Rule form

Field	Description
Name	Required. Enter a unique name that identifies the rule. Since this name is the display field for archive rules, references to archive rules display the name you enter here.
Table	Required. Select the table containing records to archive. ◦ There can only be one archive rule per table (even if the rule is inactive). If there is an existing archive rule for a table, the table no longer appears as an option on the archive rule table list. ◦ You can't change the table in an existing archive rule. To archive records in a different table, create a new archive rule. ◦ You can't select a table in a different application scope. ◦ You can't create an archive rule for a table where `update_synch` is set to true in the dictionary entry for the table. ◦ You can't create archive rules for some internal system tables or peripheral tables. For example: ▪ Sys Audit [sys_audit] ▪ Audit Deleted Record [sys_audit_delete] ▪ Audit Relationship Change [sys_audit_relation] ▪ Attachment [sys_attachment] ▪ Journal Entry [sys_journal_field]
Active	Select this check box to enable the archive rule. Clear the check box to disable the rule. Leave your archive rules inactive until you calculate an estimate of the number of records the rule affects and verify that rule behaves as expected.
Retain references	Select this check box to keep the sys_id for a reference field instead of storing the display name as a string in the archive table. **Important:** The **Retain references** check box becomes read-only after you select it and save the archive rule. If you select this option an existing archive rule, newly-archived records store the sys_ids of the referenced records but existing records will still have string values. See Migrating non-reference fields to reference fields.
Description	Enter a description of the archive rule.

Field	Description
Conditions	Select the fields and values that must be true in order for the archive rule to run. Typically, you would archive inactive records older than a certain date. Ensure the correct records are selected by testing the conditions in a list view before you activate the archive rule. For example, the condition **[Closed] [before] [Last 2 years]** archives records for the current year and the prior year, and not two years ago from today's date. To archive records closed more than two years ago from today, use the relative operator in a condition like **[Closed] [relative] [before] [2] [Years] [ago]**.
Auto Rearchive	If a record is restored, this enables the record to be automatically rearchived after a set period of time.
Auto Rearchive Duration	If you set the Auto Rearchive option, you must set a time interval before the restored record is automatically archived.

4. Select **Submit**.
5. Calculate the estimated number of records to archive by selecting the **Recalculate Estimate** related link.
   The updated count appears in the **Record estimate** field.

What to do next

After verifying that the archive rule is selecting records as expected, activate the archive rule.
If you do not want to wait for the scheduled job to run the archive rule, you can manually start the archive rule by clicking **Run Archive Now**.

Archive related records

Use the **Archive Related Records** related list to add related records to the archive rule.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Rules**.
2. Select the appropriate archive rule.
3. From the **Archive Related Records** related list, click **New**.
4. Fill in the Archive Related Records form.

Archive Related Records form

Field	Description
Archive map	Displays the archive rule to which the related records apply.
Action	Select the action you want the archive rule to take on related records. Choices include:

Field	Description
	◦ **Archive**: archive records that reference the archived record. ◦ **Clear**: remove the reference to the archived record. The record no longer references the archived record and does not appear as a related record in future archives. ◦ **Delete**: delete records that reference the archived record.
Reference	Select the relationship of the records you want to apply an action to. The **Reference** field lists all relations that currently exist for the table being referenced. There are two types of possible relations. Another table has a reference field calling the archived table. For example, if you are archiving problem records, there is a **Problem ID** field in the Incident table that references the related problem records. ◦ The **Archive** action archives the related record in addition to the primary record. For example, if you select the **Problem in Incident** reference, the related record rule also archives any incident record that references an archived problem. ◦ The **Clear** action removes the reference to the primary record. For example, if you select the **Problem in Incident** reference, the related record rule updates any incident record with a reference to the archived problem record by clearing the reference. If the reference is a many-to-many relationship, the related record rule deletes the reference instead of clearing the reference. ◦ The **Delete** action deletes any record that references the primary record. For example, if you select the **Problem in Incident** reference, the related record rule deletes any incident record that references the archived problem record. Another table has a **Document ID** field which might point to the archived table. For example, if you are archiving problem records, there is a **sys ID** field in the Attachments table that may reference the problem record. The list indicates document ID relationships by displaying an asterisk (*) character at the end of the selection name. ◦ The **Archive** action updates the **Document ID** of the related record to point to the archived table. For example, if you select the **Table sys ID Attachment(sys_attachment)*** reference, the related record rule updates the attachment record to change the **Document ID** to refer to the archived table record. ◦ The **Clear** action removes the **Document ID** reference to the archived record. The record no longer references the archived record and does not appear as a related record in future archives. ◦ The **Delete** action deletes any record that references the primary record. For example, if you select the **Table sys ID Attachment(sys_attachment)*** reference, the related record rule deletes any attachment record that references the archived primary record.

Field	Description
	Note: You can't select references from some internal system tables or from peripheral tables. For example: ◦ Sys Audit [sys_audit] ◦ Audit Deleted Record [sys_audit_delete] ◦ Audit Relationship Change [sys_audit_relation] ◦ Attachment [sys_attachment] ◦ Journal Entry [sys_journal_field]
Reference table	Displays the table where the rule looks for related records.
Reference element	Displays the reference field or **Document ID** the rule queries for.
Reference table rule	Select an existing archive rule that applies to the related records you are archiving. For example, if you already have an archive rule for the Incident table, select the existing Incident table rule when archiving records related to incidents records. **Note:** You control the cascade depth by specifying the archive rule for an archive related record. Related records without a specified archive rule aren't cascaded. ◦ Prior to the Washington DC release, if an archive rule was defined for the table matching a related record's table, the system would automatically cascade and process that archive rule and the related records associated with it. ◦ Beginning in the Washington DC release, even if an archive rule exists for a related record's table, you must declare that rule in the **Reference table rule** field to achieve that cascading behavior. If you have multiple related records on multiple tables, you can decide which of those records you want to include in the cascade by defining the reference table rule.

5. Click **Submit**.

Verify the number of records affected

Each archive rule provides an estimate of the number of records the rule affects.

Before you begin

Role required: admin

About this task

This estimate only includes primary records and excludes any related records added to the archive rule. The estimate helps you determine if the archive rule affects the number of records you expect it to. If the estimate is too high or low, change the archive rule conditions and then click the Recalculate Estimate related link.

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Rules**.
2. Select the archive rule you want to estimate records for.
3. Click the **Recalculate Estimate** related link.

Setting the language of archived strings

On internationalized instances, the archive process uses the language of the SYSTEM user to select the display value strings.

If there is no SYSTEM user, the instance uses the default language setting to select the display value strings. You can either create a SYSTEM user with a specific language setting or set the system default language to select the language of archived strings.

Create a destroy rule

Create a destroy rule to delete archived records and related records after a specified amount of time.

Before you begin

Role required: admin

About this task

If you need to amend the destroy rule properties running against your archived data, see Archive rule and destroy rule properties.

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Destroy Rules**.
2. Click **New**.
3. Fill out the form fields.

Field	Description
Name	Provide a descriptive name.
Table	Select the archive table that the system created when you archived records.
Archive	Improve archive performance by linking this destroy rule to an existing archive rule for the archive table that you selected.
Description	Provide a description.
Destroy Related Records	Automatically destroy related records associated with the archived records. **Note:** Peripheral records, which include records in the Journal Entry [sys_journal_field], Attachment [sys_attachment], and Sys Audit [sys_audit] tables are deleted automatically, even if you decide to preserve related records.
Archive Duration	Specify the amount of time that records stay in the archive table before the system deletes them.

4. Click **Submit**.
5. You can check the archive destroy log at any time

Archive rule and destroy rule properties

Rules to prevent the archive process from consuming too many system resources have been automated. You can manually change the system behavior.

The archive process uses a producer and consumer model. Records are queued for archiving in the Archive Job Execution Chunks [sys_archive_run_chunk] table in batches.

- The ArchiveProducerJob creates chunks of sys_ids records that need to be archived. Each chunk contains records pertaining to a single archive rule. The maximum number of records in a chunk is specified by the property *`glide.db.archive.chunk_size`*. The maximum number of records for a single rule across all chunks is specified by *`glide.db.archive.max.rule.records`*.
- The producer starts up 4 (default) ArchiverJob consumers. The exact number of consumers is specified by *`glide.db.archiving.max_consumer_workers`*.
- Consumers read and process chunks of sys_ids. The number of chunks each consumer processes is specified by *`glide.db.archive.max.batches`*. The chunks processed by each consumer may be from different archive rules.
- Consumers repeat reading chunks of sys_ids until none remain or the limit of chunks a consumer can process (specified by *`glide.db.archive.max.batches`*) is reached, after which it shuts down and schedules another consumer.
- If a chunk remains in the running state longer than what's specified in *`glide.db.archive.chunk.max.process.time`*, the system assumes that the consumer abruptly terminated, and it changes the chunk's status to error. This doesn't stop or end any running consumers.

Note:
The archive and destroy properties listed in the table are not used if the Archive Job Execution Chunks [sys_archive_run_chunk] table is deleted.

To manually control how many records the archive rule and destroy processes at one interval, you can change these default settings by adding the following system properties.

Archive System Properties

Name	Description	Type	Default Value
glide.db.archive.chunk_size	Number of unique identifiers added to each record in the Archive Job Execution Chunks [sys_archive_run_chunk] table. The record is formatted as CSV values in the keys field of the Archive Job Execution Chunk (sys_archive_run_chunk.do) form.	Integer	1000
glide.db.archive.chunk.max.process.time	Number of milliseconds a record can remain in the Active state	Integer	600000

Archive System Properties (continued)

Name	Description	Type	Default Value
	in the Archive Job Execution Chunks [sys_archive_run_chunk] table. The time is the difference between the system time and the value of the **Start** field in the Archive Job Execution Chunks [sys_archive_run_chunk] table.		
glide.db.archive.max.rule.records	Maximum number of records to batch per consumer, per rule.	Integer	10000
glide.db.archive.debug	Option that enables verbose logging when copying the column definitions from the original table to the archive table.	true | false	false
glide.db.archive.max.batches	Number of chunks, which are records from the Archive Job Execution Chunks [sys_archive_run_chunk] table, processed by a consumer before it shuts down and schedules another consumer.	Integer	10
glide.db.archiving.max_consumer_workers	The maximum number of concurrent consumers cluster-wide.	Integer	4

Manage archived data

View archived data, change an archive schedule, and restore archived data.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Tables**.
2. Select the module name for the archived table you want to view, for example, **Archive Incidents**.
3. Select a record from the archived table to see the archived record.

Change an archive schedule

All active archive rules are executed by a system-scheduled job set to run every 60 minutes. You can modify the job if you need to change the interval.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Scheduler** > **Scheduled Jobs**.
2. Open the **Archive** record.
3. Modify the **repeat** value.

Restore archived records and related records

When you restore a record, the instance inserts it back into the primary table and flags the record as having been restored in the log. You can also restore the record along with related records.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Log**.
2. Select the archived record to restore.
3. Click the applicable related link:
    - **Restore Record**
    - **Restore Record and Related Records**

    Warning:
    Do not delete archive record log entries. Deleting an archive log entry prevents you from restoring the data for the archived records.

Result

The record is restored to the primary table. The archive record is deleted from the archive table.

What to do next

View the restore date and time in the **Restored** column in the Archive Log related list.

Stop an active archive rule

Stop archiving by deactivating, deleting, or cancelling an archive rule.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Rules**.
2. Select the archive rule that you want to stop.
3. Stop the archive rule from running.
   - Deactivate the archive rule by clearing the **Active** check box. The `Archive - Handle archive chunks` business rule updates the state of the corresponding record in the Archive Job Execution Chunks [sys_archive_run_chunk] table from **waiting** to **not-needed**.
   - Contact Customer Service and Support to delete the archive rule.
   - Cancel the archive rule by calling `GlideArchiver().cancel()` from **System Maintenance** > **Scripts - Background**. This automatically changes the chunks in the Archive Job Execution Chunks [sys_archive_run_chunk] table from **running** to **cancelled**.

Result

No additional chunks are processed.

Data archive table size limits

There are constraints to the size of the table you use to archive your instance data.

Limits on data archive tables

There is a limit on the size of an archive table you can create to store data. The limit is based on the size of the data, not the record count. If you receive an error that an archive table could not be created, it might be related to this size limitation.

Archive table size limits

Table Attribute	Size
Static row size limit	65535 bytes
Row size limit	8126 bytes
Columns	1000 for Oracle 1011 for others

Review the custom columns on the table being archived. Also review the custom fields on the tables the table extends from. Determine if those columns are no longer needed and are not being used by current customizations. If you identify unused columns, you can remove them. Doing this might reduce the row size of the table being archived. If that is unsuccessful, contact ServiceNow® Customer Service and Support to see what other alternatives are available.

Migrating non-reference fields to reference fields

Manually preserve sys_ids in reference fields instead of storing the display name as a string.

Determining tables to migrate

To begin the data migration process, run the following background script. This script creates migration records for both direct table extensions and related archived records in the Archive Rule Reference Migration [sys_archive_ref_migration] table, ensuring that only valid tables are processed.

```
var tables =
 GlideDBObjectManager.get().getAllExtensions(current.table);
    for (i = 0; i < tables.size(); i++) {
        var gr = new GlideRecord('sys_archive_ref_migration');
```

```
        gr.addQuery('table', 'ar_' + tables.get(i));
        gr.query();
        if (!gr.next()) {
            if (GlideTableDescriptor.isValid('ar_' +
tables.get(i))) {
                var gr2 = new GlideRecord("sys_archive_ref_migration");
                gr2.initialize();
                gr2.setValue('rule', current.sys_id);
                gr2.setValue('table', 'ar_' + tables.get(i));
                gr2.setValue('reference_migration_progress', 'waiting');
                gr2.insert();
            }
        }
    }
    //Also get insert related records tables as well
    var map = new GlideRecord('sys_archive');
    map.addQuery('table', current.table);
    map.query();
    if (map.next()) {
        var id = map.getValue('sys_id');
        if (!(id === undefined)) {
            var related = new GlideRecord('sys_archive_related');
            related.addQuery('archive_map', id);
            related.addQuery('action', 'archive');
            related.query();
            while (related.next()) {
                if (!GlideTableDescriptor.isValid('ar_' +
related.getValue('table'))) {
                    gs.log('Related Record table: ' +
related.getValue('table') + ' not created yet');
                    continue;
                }

                var gr3 = new GlideRecord("sys_archive_ref_migration");
                gr3.initialize();
                gr3.setValue('rule', current.sys_id);
                gr3.setValue('table', 'ar_' +
related.getValue('table'));
                gr3.setValue('reference_migration_progress', 'waiting');
                gr3.insert();
            }
        }
    }
```

Copying archive references

For each table that goes through the migration process, the Archive Reference Copy (RefCopyJob) job identifies the sys_id of the reference fields and updates the display value to be the correct sys_id. The job configures 10k records at a time unless there are more than 10k records with the same timestamp. The migration progress relies on the archived timestamp.

Changing field types from string to reference

After the tables associated with an archive rule completely migrate, the ArchiveRefJob job executes. This job changes the sys_dictionary types of the archive table from string to reference.

Fixing node failures for RefCopyJob and ArchiveRefJob

If node failure occurs while these jobs are running, it leaves the status of the data migration in an improper state. If the *RefCopyJob* fails, it can leave a table in a migrating status. You can verify this condition by checking if rows in the sys_archive_ref_migration are stuck in **migrating** status for an extraordinary amount of time. Update the specific row status from **migrating** to **waiting** and the *RefCopyJob* continues the data migration on the table when the job runs again.

The node can also fail when the *ArchiveRefJob* terminates prematurely. Check if tables have fields that are reference fields and some that are still string type fields. The job might have stopped in the middle of changing the field types. You can resolve this condition by setting up a trigger job to execute in a background script, which starts the process over again:

```
GlideRecord trigger = new GlideRecord('sys_trigger');
trigger.initialize();
trigger.setValue('state', 0);
trigger.setValue('trigger_type', 0);
trigger.setValue('next_action', new GlideDateTime());
trigger.setValue('job_context', 'fcRuleId=' + ruleId);
trigger.setValue('name', 'Job Reference Migration' + ' Node - ' + new
 GlideClusterSynchronizer().getSystemID());
trigger.setValue('trigger_class', 'com.glide.db.auxiliary.job.ArchiveRefJob');
trigger.insert();
```

System clone

Use the System Clone application to copy everything in a database from one instance to another.

Cloning is typically used to copy a production instance to a pre-production instance to test changes. Cloning data comes from the most recent nightly backup.

Note:
A new, unified admin experience based on the existing clone engine is now available in the Clone Admin Console. The Clone Admin Console provides enhanced visibility for cloning data between instances, one of our most-used automations, as well as a number of other enhancements.

Clone process

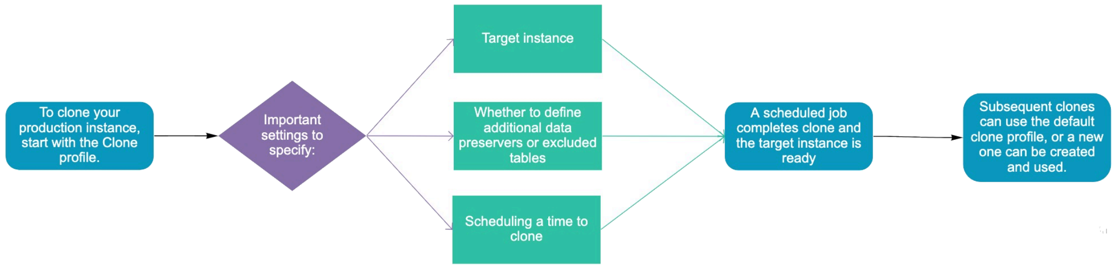

In response to a clone request, the ServiceNow platform performs the following tasks:

1. Generates a file to preserve operational data on the target server.

   Note:
   This file contains the data preserved by data preservers.

2. Copies the database schema from the source instance to the target instance.
3. Creates tables in the target instance database using the source instance table definitions.
4. Copies data from the most recent nightly backup of the source instance to the target instance database.

   Note:
   Certain exclusions are automatic, large tables are normally excluded. These include audit, log, and email tables. MetricBase tables aren't excluded by default.

5. Briefly disables UI traffic and requests to the target instance server.
6. Displays the message **Clone in progress...** to any user accessing the target instance.
7. Restores operational data preserved from the target instance.
8. Runs any post-clone cleanup scripts on the target instance.
9. Briefly suspends all email functions on the target instance.
10. Queues an event to regenerate text indexes.
11. Enables UI traffic and requests to the target instance server.

During a clone, the target instance may be intermittently unavailable. After clone completion, you have up to 24 hours to contact Customer Service and Support and request a rollback of the target instance to its pre-clone state. You're notified when the rollback is complete.

Note:
If the source instance has a clone depth level of >=5 then the clone isn't allowed.

If the source instance purpose is DART (Data Access for Responsible Training) then the clone isn't allowed and an error message will be displayed.

Clone to an instance on a different version

The System Clone application can target an instance running a different instance version from the source.

A central web service controls the clone processing and automatically modifies the target instance version to match the source instance version. This matching process starts up to eight hours before the time specified in the **Date and time** field on the System Clone form. This web service also confirms that there's enough disk space on the target instance for the clone to proceed.

When cloning from a backup, the target instance doesn't need additional time to upgrade or downgrade. The ServiceNow platform performs any version changes during a brief window where the target instance is unavailable, after it copies data from the source instance backup.

Clone from a backup

The clone uses data from the most recent, nightly backup of the source instance when cloning. Backups that are used for cloning are a maximum of 36 hours old. System Clone begins the initial preparation, including selecting the latest backup to use, only at the date and time processing is scheduled to start.

If cloning from a source backup fails, the system uses the legacy clone engine instead. The legacy clone engine can't preserve data from extended tables, relationships, hierarchies between tables, and dot-walked queries. You may want to restore the target instance from a backup and then reschedule the clone in such cases.

After cloning from a backup, the target instance is unavailable for several minutes before the clone is marked as complete in the source instance. If the source and target instances are on different versions of the ServiceNow AI Platform, the target instance is modified to match the source instance version during this time.

When starting a clone from a backup, the date and time the backup was taken, as well as periodic progress messages, appear in the **Clone Log** related list.

System clone backup log

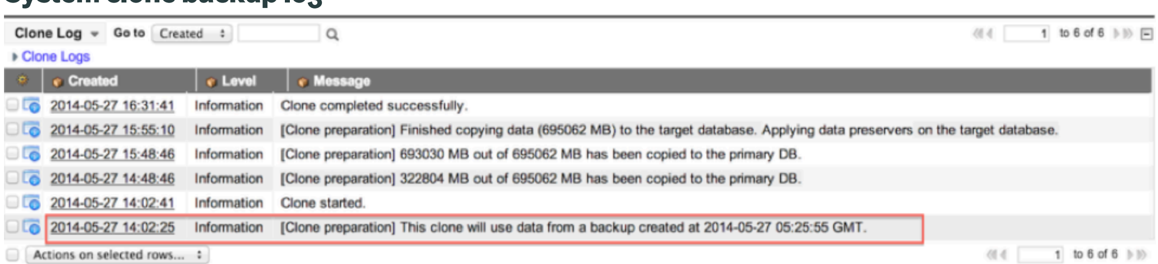

Created	Level	Message
2014-05-27 16:31:41	Information	Clone completed successfully.
2014-05-27 15:55:10	Information	[Clone preparation] Finished copying data (695062 MB) to the target database. Applying data preservers on the target database.
2014-05-27 15:48:46	Information	[Clone preparation] 693030 MB out of 695062 MB has been copied to the primary DB.
2014-05-27 14:48:46	Information	[Clone preparation] 322804 MB out of 695062 MB has been copied to the primary DB.
2014-05-27 14:02:41	Information	Clone started.
2014-05-27 14:02:25	Information	[Clone preparation] This clone will use data from a backup created at 2014-05-27 05:25:55 GMT.

Clone over production instances

As long as the system property `glide.db.clone.allow_clone_target` is TRUE, an instance can serve as a clone. Modifying data on the source instance during a clone can cause a data mismatch between records or duplicate record entries.

Deprecated Clone Options

The clone option check box **Preserve users and related tables** has been removed from Clone Options in the Utah release. In some cases, past customizations to the clone request page or related to clone, may cause this field to remain on your form.

> **Important:**
> Checking this deprecated field won't have any effect on the user, role, or related tables during your clone.

.

Request a clone

Request a clone to copy data from a production instance to a non-production instance or to copy data between non-production instances.

Before you begin

Role required: admin

About this task

https://player.vimeo.com/video/1006365658?h=e210d21863&badge=0&autopause=0&player_id=0&app_id=58479https://player.vimeo.com/video/1006365658%3Fh=e210d21863&badge=0&autopause=0&player_id=0&app_id=58479https://player.vimeo.com/video/1006365658%3Fh=e210d21863&badge=0&autopause=0&player_id=0&app_id=58479

The ServiceNow AI Platform uses data from the most recent, daily backup of the source instance when cloning. Backups that are used for cloning are a maximum of 36 hours old. System Clone begins the initial preparation, including selecting the latest backup to use, only at the date and time processing is scheduled to start.

If cloning from a source backup fails, the system uses the legacy clone engine instead. The legacy clone engine cannot preserve data from extended tables, relationships, hierarchies between tables, and dot-walked queries. You may want to restore the target instance from a backup and then reschedule the clone in such cases.

Note:

- For instances that use an Oracle database, see KB0538884 - System Clone Support for Oracle Customers.
- You can automate cloning by scheduling it in the Options pane.
- You cannot initiate a clone request when the source instance is in debug mode. If the source instance is in debug mode, contact Customer Service and Support to turn off debug mode before you submit a clone request.

Procedure

1. Log in to the instance that you want to clone.
   This instance becomes the source instance of the clone request.
2. Clone Target Registration and Authentication record for each target instance that you want to receive clone data.
3. Verify the list of tables that are excluded from cloning and add or remove tables to exclude from the target instance.
4. Verify the list of tables and system properties that you want saved on the target instance with the following.

   You can use data preservers. You can also create or modify data preservers, as needed.

   The legacy clone engine does not support data preservers for these records:

   - Tables that extend the Task table
   - Relationships

- Hierarchies
- Dot-walked queries

If you are preserving any data that the legacy clone engine does not support, verify that there is a recent backup of the target instance available. If the clone-from-backup-process fails for any reason, you can restore the target instance from the backup.

5. Preserve any unpublished applications on the target instance.
6. Navigate to **System Clone** > **Request Clone**.
7. **Optional:** Specify a predefined clone profile.
   A clone profile stores target and clone options. The clone profile automatically populates your clone request with your selected profile settings. See clone profiles for clone requests.
8. In the **Target instance** field, select the target instance that you want to receive the cloned data. Create a separate clone request for each target instance that you want to receive clone data.
9. In the **Clone Scheduled Start Time** field, select the time that you want the cloning to start.
   You can schedule multiple clone requests for the same source instance. For example, create one clone request to copy data to non-production instance A and another clone request to copy data to non-production instance B. The scheduling engine determines whether multiple clone requests against the same source instance can occur simultaneously or whether they must occur sequentially.
   The system verifies the scheduled start time and either accepts the date-time value that you selected or suggests an available date-time value. The validation process prevents scheduling conflicts with other automations using the same target instance.
10. In the **Email upon completion** field, enter your email address so that you can receive alerts after the cloning finishes, is canceled, or has an error.
11. Click the **Options** arrowhead so that it turns downward and use the following table to make appropriate selections in the Options pane.

Clone options

Field	Description
Exclude tables specified in Exclusion List	Prevents cloning records from tables on the source instance which are listed under#**System Clone** > **Exclude Tables**. If a table is on the Exclusion List, clone will exclude the records on the table as well as records on the child tables. When excluding tables, their table schema and hierarchy will still be cloned to your target instance. As a result, your target instance will have empty but usable tables after the clone. **Note:** ServiceNow out-of-the-box table exclusions will still be excluded and are not affected by this setting. These include tables containing auditing, license usage, logging, and notifications. You can choose to disable this setting if you need the data from your excluded tables. This is called a Full Clone and is recommended for upgrade or deployment testing. The legacy clone engine does not support this option. The default setting is that tables specified in the Exclusion List are excluded from a clone.

Field	Description
Exclude audit and log data	Prevents the cloning of audit and log records from the source instance. This creates empty but usable audit and log tables on the target instance. **Note:** If you exclude audit and log data from your clone, the Activity Stream for records won't match the source instance. This is because the Activity Stream relies on the audit table to generate the history. The default setting is that audit and log data are excluded from a clone.
Exclude attachment data	Prevents the cloning of certain files like ◦ Video files ◦ Image files ◦ Other large binary files from the sys_attachment table. **Note:** ServiceNow out-of-the-box attachments and other system-relevant files (e.g. catalog item images, theme images, and icons) are not affected by this setting. The following ServiceNow out-of-the-box attachments and other system-relevant files are not affected by this setting and are not excluded from a clone: ◦ Table name values that start with ZZ_. ◦ Table names: sys_certificate, sc_cat_item,sys_upgrade_manifest, ecc_agent_jar, ecc_agent_mib, sys_store_app, or invisible.sys_store_app. The default setting is to exclude attachment data.
Preserve theme	Preserves the theme and CSS elements on the target instance. As a result, your target instance will keep its theme, its look and feel after a clone. The default setting is to preserve the theme on the target instance.
Lock settings for this clone request	If you use a clone profile, this option locks the settings and options at the time of the clone request. Any subsequent changes to the clone profile, regardless of when the clone runs, do not affect the clone request. This option is not selected by default.
Amount of data copied from specific tables	Limits the number of days of historical data for the task table and its child tables (e.g. Incident, problem, and change tables) to 90 days. To reduce clone time, consider excluding large tables altogether. When excluding tables, your target instance will have the same table schema and hierarchy (i.e. empty usable tables) as the source instance. The default setting is to clone all data from the task table and its child tables to the target instance.
Preserve In Progress Update Sets	Preserves the last 90 days of in-progress update sets in the global application scope. This option allows you to keep in-progress, global update sets created within the last 90 days on your target instance.

Field	Description
	Note: This option does not preserve your in-progress scoped update sets. Update sets that are older than 90 days will not be saved. It is recommended to review and commit your update sets before cloning. The default is to not preserve update sets.
Clone frequency	This option allows you to schedule recurring clones from your source to your target instance. It allows you to define the clone frequency and the maximum number of occurrences. By default, the clone frequency is set to None. For more information about scheduling cloning, see Schedule cloning.

12. Click **Submit**.
    If there are no issues with the clone request, the system displays the Authenticate Target modal.

    Trouble?

    If your clone request preserves large tables, a warning message displays the large tables from your preserver list. You can cancel your clone request and remove the tables from your preserver list, or you can specify a reason and continue the clone request.
13. In the **Username** and **Password** fields, enter the username and password for an administrator account on the target instance and then click **Authenticate**.
14. Review the clone settings and click **OK**.
    An email is sent to the supplied address after the clone finishes, is canceled, or has an error.

What to do next

You can:

- Schedule automatic clonings.
- Cancel your clone request.
- View the cloning history of completed clonings.

Clone Target Registration and Authentication

A clone target record specifies the instance URL and credentials used for cloning.

Before you begin

- Provide credentials for the target instance for a user with the admin role. Use a local, user account, not an LDAP or SSO user account. The target instance credentials must exist in the User [sys_user] table as a user record or as part of an LDAP integration. Clone requests cannot redirect authentication requests to a single sign-on identity provider.
- Verify the system property *glide.db.clone.allow_clone_target* is set to **True**. By default, this property is enabled on instances whose name ends in Dev, Test, Stage, UAT, or QA.
- If the target instance uses IP range based authentication, it must allow the IP range 10.0.0.0/10.255.255.255 to communicate on a local network.
- Role required: clone_admin or admin.

Procedure

1. Navigate to **All** > **System Clone** > **Clone Targets**.
2. Click **New**.
3. Enter the URL for the receiving instance (target).

   The system validates the instance allows clone targets and that High Availability Cloning is active. Production and demonstration instances fail these validation checks.

   #### Clone target invalid

   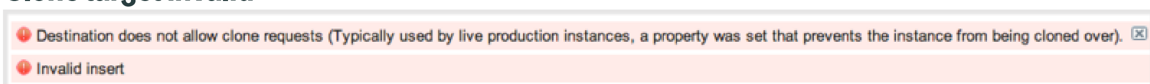

4. Enter the basic authentication credentials for a user account with the admin role on the target instance.

   

   Note:
   You cannot request cloning multiple targets from the same source. Instead, make a separate request for each target.

   The system validates the user credentials have admin access to the target instance.

5. Click **Submit**.
   The system checks connectivity and validates the user credentials against the target instance.

Exclude a table from cloning

Exclude a table to create an empty but usable table on the target instance.

Before you begin

Role required: clone_admin or admin

About this task

Note:
When using Exclusions (Exclude Tables) the data is removed during our exclusion substage prior to the node repoint. The data is removed before the instance becomes available to users.

The **System Clone** > **Exclude Tables** module lists the tables that aren't copied during a system clone. On the clone request form, the option "Exclude tables specified in Exclusion List" is on by default. It verifies that the exclusions are processed during a clone. By default, the system excludes certain required tables such as logging, auditing, notifications, workflow contexts, and license usage.

Note:
Excluded workflow context data includes records stored in the wf_context table, and in related tables with names starting with a prefix of wf_. This also includes the workflow scheduler table. This helps to prevent the occurrence of workflow timer syncing issues that might take place due to the length of the cloning process if workflow contexts were included.

`sys_db_object` and `sys_db_object.*` can't be added to the exclusions.

When excluding child tables of the Task table, which use the table per hierarchy extension model, child tables are excluded. You don't need to exclude the parent Task table together with the child table.

You can use wildcards to exclude several tables from cloning with one entry. For example, `sys_script.*` excludes all tables starting with sys_script.

Note:
The wildcard has to include the dot '.' before the asterisk.

To preserve the existing data on the target instance, see Create a data preserver. Data on tables that reference the table, such as business rules, isn't excluded. To avoid unexpected results, we (strongly) advice against altering out-of-box preservers.

To preserve or exclude related tables, you must preserve and exclude all of the following tables. If any of the tables are missed, the customer may not be able to log in to the instance after a clone.

- sys_user
- sys_user_role
- sys_user_group
- sys_user_grmember
- sys_group_has_role
- sys_user_role_contains
- customer_contact (if the Customer Service Management plugin is active)

Procedure

1. On the source instance, navigate to **System Clone** > **Exclude Tables**.
2. Select **New**.
3. Enter the table **Name**.
   Entering a parent table results in the clone process also excluding its child tables. For example, excluding the Task table would also exclude the Change, Incident, and Problem tables.
4. Select **Submit**.

Data preservation on cloning target instances

You can use data preservers to protect data on the target instance from being overwritten. If you have custom applications, you must also manually preserve unpublished application content.

Data preservers

Sometimes, it is necessary to preserve some data on an instance targeted for cloning. For example, if the target is a MID Server, you must not overwrite the MID Server [ecc_agent] table. Preserved data is stored on the target instance before cloning begins and is restored on the target instance after cloning.

> **Warning:**
> You must define data preservers on the source instance. Defining them on the target instance does not preserve the data.

Data preservers typically preserve system settings and themes, such as:

- Instance-specific authentication settings
- Bookmark [sys_ui_bookmark]
- Recent Selection [sys_ui_recent_selection]
- User Preference [sys_user_preference]

> **Note:**
> A clone does not support preserving data from a database view.

Do not use data preservers to transfer large sets of data, such as user groups. If you must preserve table data, such as users, groups, and roles, consider exporting the records to a file and importing them after cloning.

Data preservers for Multi-SSO

The system automatically creates the necessary data preservers for cloning when you activate Multiple Provider Single Sign-On integration.

Name	Table	Conditions
Certificate	X.509 Certificates [sys_certificate]	None
Core Instance Properties	System Property [sys_properties]	• [OR] [Name] [is one of] [glide.authenticate.external, glide.authenticate.external.logout_redirect] • [OR] [Name] [starts with] [com.snc.integration.saml_esig] • [OR] [Name] [is one of] [glide.smtp.port, glide.smtp.auth, glide.smtp.encryption] • [OR] [Name] [starts with] [glide.authenticate.multisso] • [OR] [Name] [is] [glide.authenticate.sso.redirect.idp] **Note:** The properties `glide.smtp.port`, `glide.smtp.auth`, and `glide.smtp.encryption` are deprecated.
Digest Properties	Digest Properties [digest_properties]	None
Identity Providers	Identity Providers [sso_properties]	None
SAML2 Update1 Properties	SAML2 Update1 Properties [saml2_update1_properties]	None

Note:
Although you can modify these data preservers, it is good practice not to. The Digest Properties [digest_properties], Identity Providers [sso_properties], and SAML2 Update1 Properties [saml2_update1_properties] tables are required for multiple source, single sign-on (SSO) to function properly. If multiple source, single sign-on is disabled on the target instance, you can safely remove all three data preservers. Remove them at the same time, as the system terminates the clone with an error message when you attempt to clone with one or two of these tables being preserved.

Data preservers for SAML

Preserving SAML SSO-related settings can prevent the target instance from using the wrong issuer and audience parameters when making authentication requests to your IdP. To preserve SAML settings, create data preservers for the following tables:

- System Property [sys_properties]: to preserve SAML properties.
- X.509 Certificates [sys_certificate]: to preserve SAML certificates.
- User [sys_user]: to preserve SAML users.

You also need to preserve properties and users that are involved in SAML.

Preservation of unpublished applications

You cannot use data preservers to save unpublished applications. Instead, application developers must choose how they want to preserve unpublished applications.

The cloning process does not preserve version differences for applications in development. Instead, the system clone only copies the application version installed on the source instance onto the target instance. If the target instance had a development version of the same application, the application will be editable after the clone, but it will be at whatever version was installed on the source instance. If the application was missing from the source instance, the cloning process deletes the application from the target instance.

Create a data preserver

Data preservers maintain specified data on a target instance.

Before you begin

Role required: clone_admin or admin

About this task

Sometimes, preserving certain data on a target instance is desirable. For example, when using a MID Server, you can avoid overwriting the MID Server [ecc_agent] table. Preserved data is stored in a dynamically generated list on the target instance before the clone and restored on the target instance after the clone is complete. You define data preservers on the source instance.

Data preservers are primarily intended to preserve system settings and themes, such as instance-specific authentication settings. Do not use data preservers to transfer large sets of data, such as user groups. If you must preserve table data such as users, groups, and roles, consider exporting the records to a file and importing it after the clone is complete.

Consider whether to preserve the data in the following tables.

- Bookmark [sys_ui_bookmark]
- Recent Selection [sys_ui_recent_selection]
- User Preference [sys_user_preference]

If you set a data preserver on a table where the source instance has more records than the target instance, the data preserved on the target instance also includes the additional records from the source instance.

For example, assume that the data preserver is already in place.

- In the source instance, the sys_temp table contains 100 records.
- In the target instance, the sys_temp table contains 20 records.

After the clone, the sys_temp table in the target instance contains 100 records.

- The 20 records in the target sys_temp table are preserved successfully (per the data preserver specification). These records were part of the 100 records in the source sys_temp table.
- The source sys_temp table brings over the remaining 80 records to the target sys_temp table.

To resolve this issue and to preserve only the records in the target table, create an exclude table record for the target table, in addition to setting the data preserver on the source table.

Important:
Configure preservers on the source instance.

Procedure

1. On the source instance, navigate to **System Clone** > **Preserve Data**.
2. Click **New**.
3. Enter the table label as the **Name**, for example, User Preference for the [sys_user_preference] table.
   The data preserver must have a table name or it cannot be submitted.
4. Select the **Table** to be preserved.
   The data preserver must have a table selected or it cannot be submitted.
5. Select the **Theme** check box if the data being preserved is a UI property.
6. Define the data to be preserved using the Condition Builder .
   You can use conditions to define particular records you want to preserve during a clone. For example, to only preserve particular system properties, you can add conditions for each property name you want to preserve.

   Note:
   The condition to match regular expressions [match regex] is no longer supported.

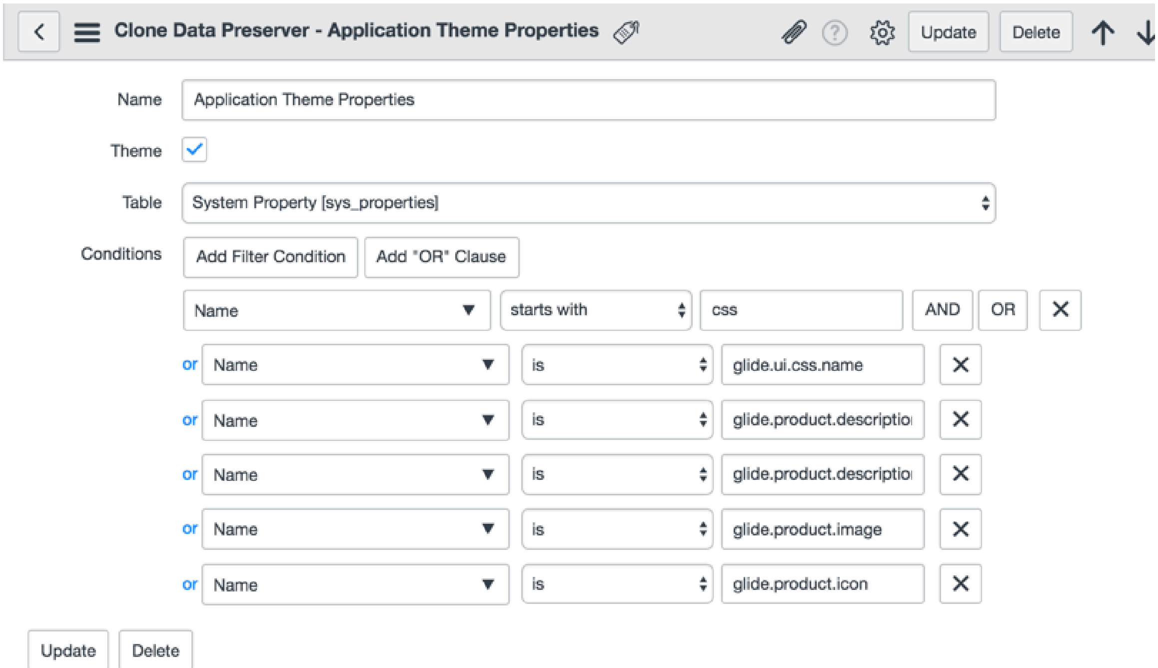

Warning:
If the clone from backup fails for some reason, the clone process fails over to the legacy clone engine. The legacy clone engine cannot preserve data from extended tables, relationships, hierarchies between tables, and dot-walked queries. You may want to reschedule a system clone or manually transfer data in such cases.

7. Click **Submit**.

   If you want to delete the data preserver later, make sure not to modify or delete the following data preserver records:

   - Core Instance Properties
   - Semaphores
   - Email Accounts

   Note:
   DB views cannot be preserved.

   Preservers cannot be empty, and users will not be able to submit the clone if preservers are empty.

Preserve SAML properties

If you want a clone target instance to keep its existing SAML integration, you must edit the **Core Instance Properties** data preserver to include the SAML properties.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Clone** > **Preserve Data**.
2. Select **Core Instance Properties**.
3. Add the following **Conditions**.

- [OR] [Name] [is one of] [glide.authenticate.external, glide.authenticate.external.logout_redirect, glide.authenticate.failed_requirement_redirect]
- [OR] [Name] [starts with] [glide.authenticate.sso.saml2]
- [OR] [Name] [starts with] [com.snc.integration.saml_esig]

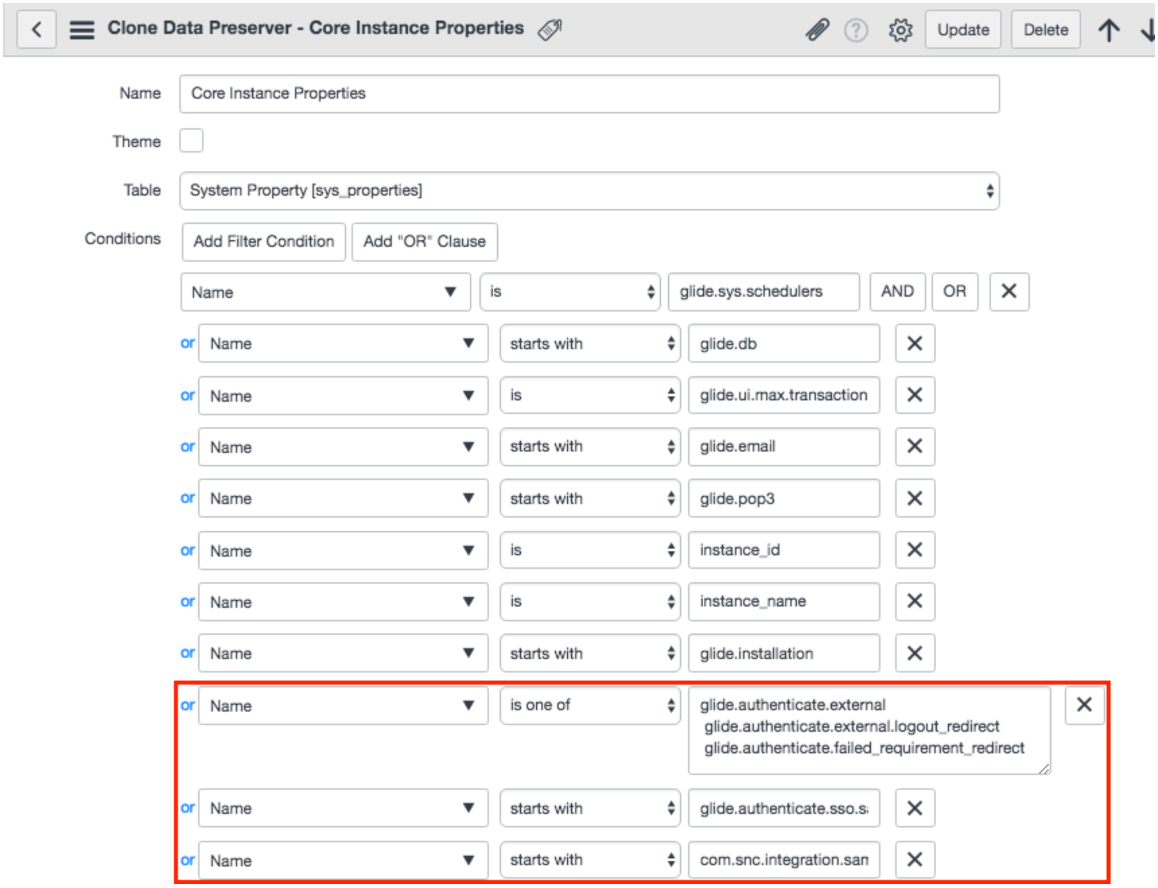

Note:
Ensure the **Theme** check box is cleared so these properties are preserved regardless of whether you preserve the instance theme.

4. Click **Update**.

Preserve applications and customizations in development during a system clone

Manually preserve a copy of each application and customization that you currently have in development before you can clone the application version to the target (development) instance.

Before you begin

Ensure that you have write access to the application record.

Ensure that you have access to a source control repository.

Role required: admin

About this task

The cloning process does not preserve version differences for applications and app customizations in development. Instead, the system clones only the copies of the application and app customization versions that are installed on the source instance onto the target instance. If the target instance had a development version of the same application, the

application is editable after the clone, but is at the version that was installed on the source instance. If the application was missing from the source instance, the cloning process deletes the application from the target instance.

Procedure

1. To preserve the application on the clone target instance, do one of these actions:

   Version differences between instances

<table>
<tr><th>Application version state</th><th>Action to take</th></tr>
<tr><td>The application version on the clone target instance is different than the source instance version.</td><td rowspan="2">Export each application from the clone target instance. The choices include:
◦ Link each application to a source control repository.
Note:
If the application is already linked to a source control repository, commit the latest version to it.
◦ Publish each application to an update set.</td></tr>
<tr><td>The application is available only on the clone target instance.</td></tr>
<tr><td>The application version on the clone target instance is the same as the source instance.</td><td>None. The system clone process copies this application version onto the target instance during the clone.</td></tr>
</table>

2. Request a system clone of the source instance to the target instance.

   Example
   For example, clone your production instance over your development instance.

3. After the clone process finishes, log in to the clone target instance.

4. **Note:**
   If source control is linked, then post clone, the platform will automatically retrieve applications and customized applications. If this is disabled via `glide.source_control.post_clone_import_enabled` you will need to manually retrieve by doing the following.

   If you saved each application to a source control repository, use one of these actions to retrieve them from the source control repository:

   Note:
   For what to expect after application customization post clone, see Results post cloning for application customizations.

   Retrieve applications from a source control repository

Application installation state	Action to take on clone target
The application and customization were previously installed on the source instance.	Apply remote changes from the source control repository.

Application installation state	Action to take on clone target
The application was never installed on the source instance.	Delete the repository configuration (sys_repo_config) and import the customization from the source control repository.

Remote changes after clone

Field	Description
glide.source_control.post_clone_import_enabled	To disable the apply remote changes automation, set to **False**. The default is **True**.
glide.source_control.post_clone_import_delay_time_sec	To provide a delay time, which will delay processing of the queue, provide a value. The default is zero.
glide.source_control.post_clone_import_pause_refresh_time_sec	To provide a time interval in which the refresh repository job will not run, provide a value. The default is three hours (10800).

5. If you saved each application to an update set, do one of these actions to retrieve them from the update set:

Retrieve applications from an update set

Application installation state	Action to take on clone target
The application was previously installed on the source instance.	**a.** Delete the application version that was cloned from the source instance. **b.** Load the update set that contains the current application version.
The application was never installed on the source instance.	Load the update set that contains the current application version.

Result

The applications previously in development are available for further development on the clone target instance.

Example: Preserve the Marketing Events application

Let's say that your company previously created version 1.0 of a custom application called Marketing Events. You have already published version 1.0 of the Marketing Events application to the application repository and installed it on your production instance.

Over time, users have submitted enhancement requests for the application, and you decide to develop version 2.0 of the Marketing Events application on a non-production instance to address these requests. As development nears completion, you want to update your non-production instance to the latest copy of production for some comprehensive testing.

Because you previously used a source control integration to develop version 1.0 of the Marketing Events application, you have already linked the Marketing Events application to a source control repository. You commit version 2.0 of the Marketing Events application to the source control repository.

You schedule a clone of the production instance over the development instance. After completion, you log in to the development instance and see that it has version 1.0 of the Marketing Events application, because that was the version installed on the source instance.

Because the application was already installed on the source instance, you apply the remote changes from the source control repository to receive the latest application version. The development instance now has version 2.0 of the Marketing Events application and is available for further development and testing.

Clone profiles for clone requests

A clone profile enables you to store predefined target and clone options. The clone profile automatically populates your clone request with your selected profile settings.

Clone profiles

Navigate to **System Clone** > **Clone Profiles** to view your clone profiles. With a clone profile, you can:

- Create a profile with specific target instance and option settings, tables to exclude, data to preserve, and cleanup scripts to run
- Create a clone request directly from the clone profile
- Apply the clone profile to a clone request

You can create, edit, and delete clone profiles from the **Clone Profile** view. The system profile is a read-only clone profile that you can't delete. It shows a predefined list of table exclusions, data preservers, and cleanup scripts. This list is based on the plugins that your instance contains.

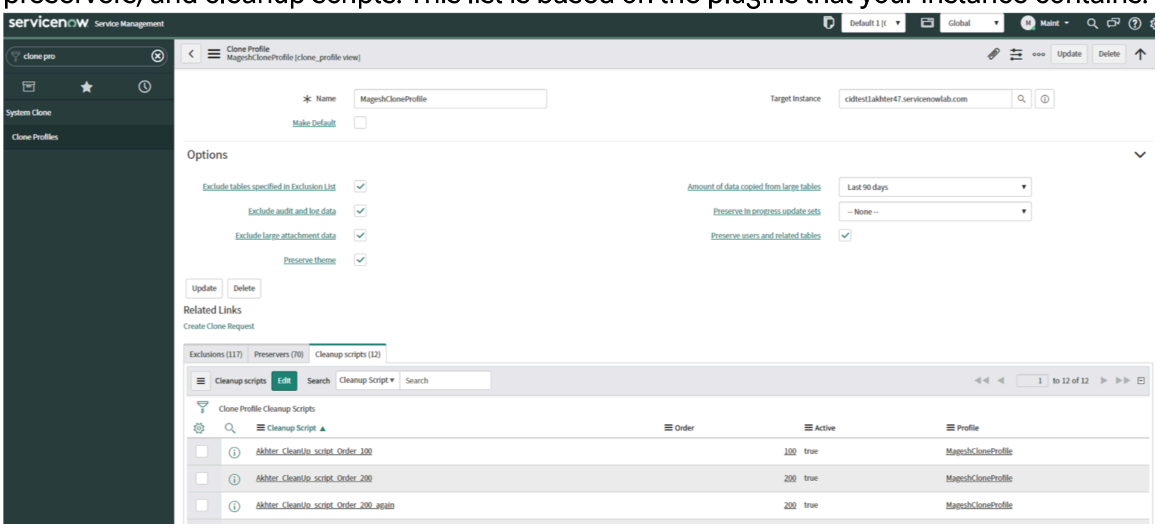

To set a new clone profile as the default profile used when you request a clone, select the **Make Default** option. Verify that this is the correct clone profile you want to use for the clone scenario you are requesting.

If you create a new data preserver, exclusion, or cleanup script, it is not automatically added to your clone profiles. To add a preserver, exclusion, or cleanup script, open the clone profile, select **Additional actions** > **Configure** > **Form Layout**, and move the new preserver to the selected list.

You are not required to use clone profiles. If you leave the clone profile field empty when scheduling a clone, the system uses the exclude tables, data preservers, and cleanup scripts configured under **System Clone** > **Clone Definition**.

Cancel a clone

You can cancel requested, scheduled, and active clones without negatively impacting system stability or usability. Canceling a clone restores the target instance to the pre-clone state, retaining all original data.

Before you begin

Role required: admin

About this task

Active system clone

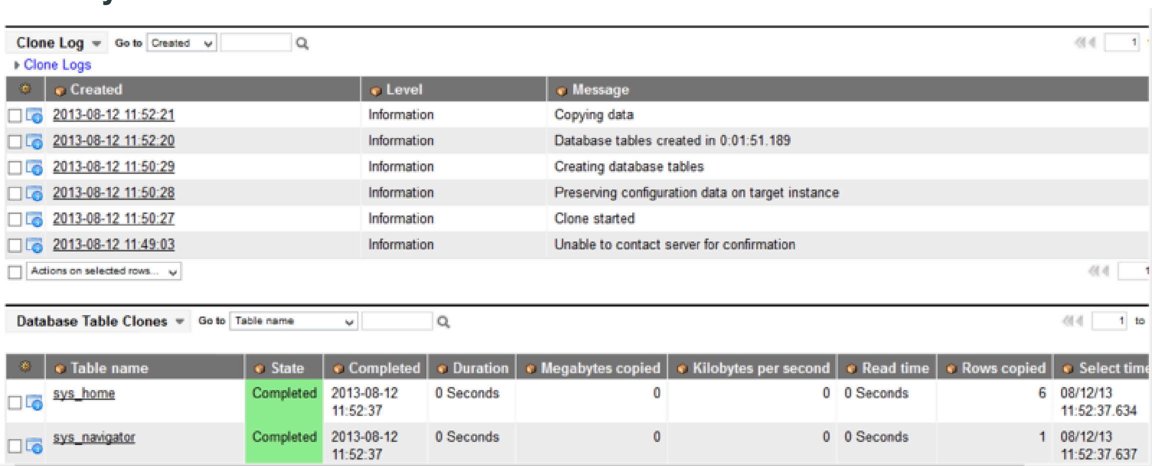

Clone Log — Go to Created

Clone Logs

Created	Level	Message
2013-08-12 11:52:21	Information	Copying data
2013-08-12 11:52:20	Information	Database tables created in 0:01:51.189
2013-08-12 11:50:29	Information	Creating database tables
2013-08-12 11:50:28	Information	Preserving configuration data on target instance
2013-08-12 11:50:27	Information	Clone started
2013-08-12 11:49:03	Information	Unable to contact server for confirmation

Actions on selected rows...

Database Table Clones — Go to Table name

Table name	State	Completed	Duration	Megabytes copied	Kilobytes per second	Read time	Rows copied	Select time
sys_home	Completed	2013-08-12 11:52:37	0 Seconds	0	0	0 Seconds	6	08/12/13 11:52:37.634
sys_navigator	Completed	2013-08-12 11:52:37	0 Seconds	0	0	0 Seconds	1	08/12/13 11:52:37.637

After starting a clone, the **Clone Log** and **Database Table Clones** related lists appear on the form. These related lists show general log messages, and the details of individual tables respectively.

The duration a clone remains active varies upon the amount of data being cloned, and whether the source and target instance are in the same physical location. If a clone takes longer than anticipated, ServiceNow Technical Support can identify additional details about the clone progress.

Procedure

1. Navigate to **All** > **System Clone** > **Live Clones** > **Active Clones**.
   The system displays the list of currently active clones.
2. Select the clone you want to cancel.
   The system displays the System Clone record.
3. From **Related Links**, click **Cancel Clone**.
   The **Cancel Clone** modal displays.

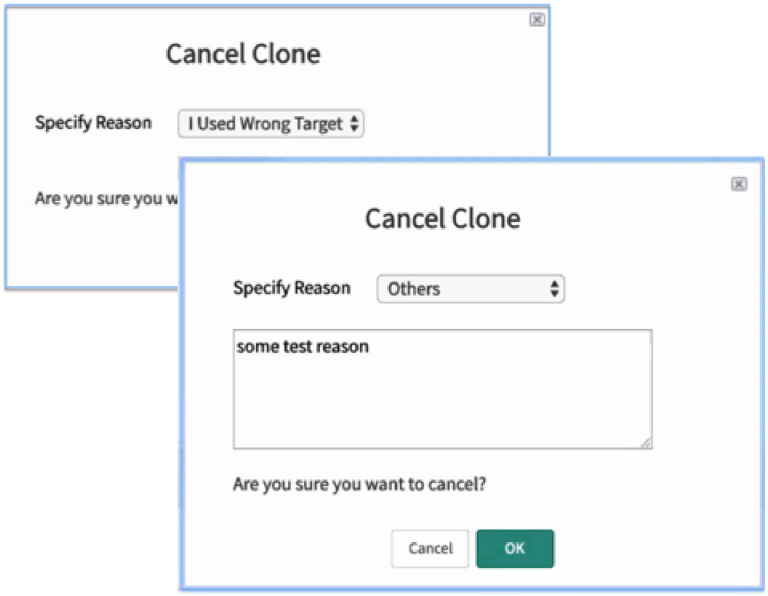

4. In **Specify Reason**, select a reason for the cancellation and click **OK**, or select **Others**, enter a reason in the text box and click **OK**.
   The system stops any current clone activities and sets the **State** to **Canceled**.

What to do next

If you want to restart a canceled clone, you must create a new clone request.

Schedule cloning

You can use System Clone to schedule automatic cloning, which is the easiest way to keep your cloned instances up to date.

Before you begin

Role required: admin

About this task

Instead of manually cloning instances, you can schedule cloning that happens automatically. You create a cloning schedule in the same interface you use to create a clone. This topic assumes that you created a clone but not a cloning schedule for it.

Procedure

1. Click **System Clone** > **Active Clones** > **<one-of-your-clones>**.

2. **Note:**
   (Optional) if the **Options** panel is not already displayed.

   Click the **Options** arrowhead so it turns downward.
   The **Options** panel appears.

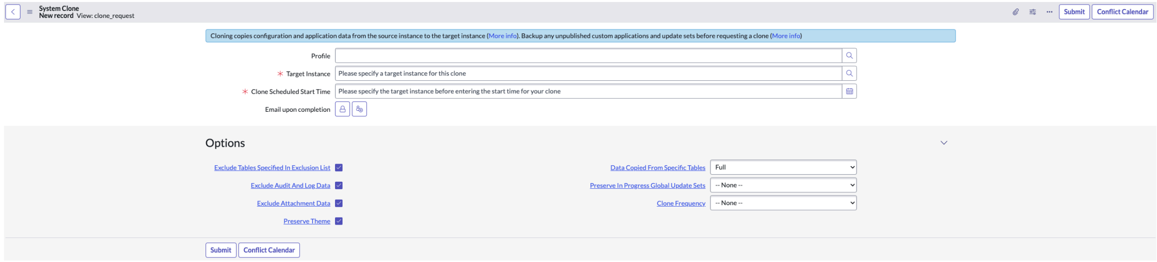

3. **Note:**
   A target instance must be selected or an error message appears.

   Select the **Conflict calendar** to view a calendar with their current clone time and potential conflicts if you want to schedule for a different time.
   The conflict calendar appears in a new 

   tab.
4. Enter values in the following fields to schedule automatic clonings.

5. Click **Submit**.
   The system displays the **Authenticate Target** modal.
6. Enter the **Username** and **Password** for an administrator account on the target instance and click **Authenticate**.
   The System Clone form displays for the target clone.
7. To see the cloning schedule for this target, click the **Recurring Clones** tab.
   Each line in the list shows a separate, scheduled cloning session.
8. To see log messages for past clonings on this target, click the **Clone Log** tab.
9. To see cloning schedules for all the clones in the system, click **System Clone** > **Live Clones** > **Clone History**.
   The **System Clones** page lists all the cloning instances in the system along with their scheduled clonings.

Modify cloning schedules

You can cancel scheduled clonings but not modify them.

Before you begin

Role required: admin

About this task

This topic assumes that you have a cloning schedule that you want to modify. Cloning schedules cannot be modified. Instead, you have to cancel scheduled clonings and create a new cloning schedule.

Procedure

1. Click **System Clone** > **Active Clones** > **<one-of-your-clones>**.
   The **System Clone**record opens for your cloning instance.

2. To delete one or more scheduled clonings, on the **Recurring Clones**tab, select the check box next to the row(s) you want to cancel and on **Actions on selected rows**, click **Delete**.
3. If you deleted all future clonings, you can create a new cloning schedule.
   You cannot create more than one cloning schedule for a target instance.
4. To stop a cloning that is already in progress, under **Related Links**, click **Cancel Clone**.
   The system stops any current cloning activities and sets the **State** to **Canceled**. Future, scheduled clonings are not deleted.

View clone status

You can view the status of a clone to make sure the cloning process isn't stuck.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Clone** > **Live Clones** > **Active Clones**.
   The system displays the list of currently active clones.
2. Select the clone you want the status of.
   The system displays the System Clone record of that clone.
3. Under **Related Links**, click **Check Clone Status**.
   The system updates the **State** field and produces a log entry in the **Clone Log** that shows the status of the clone.

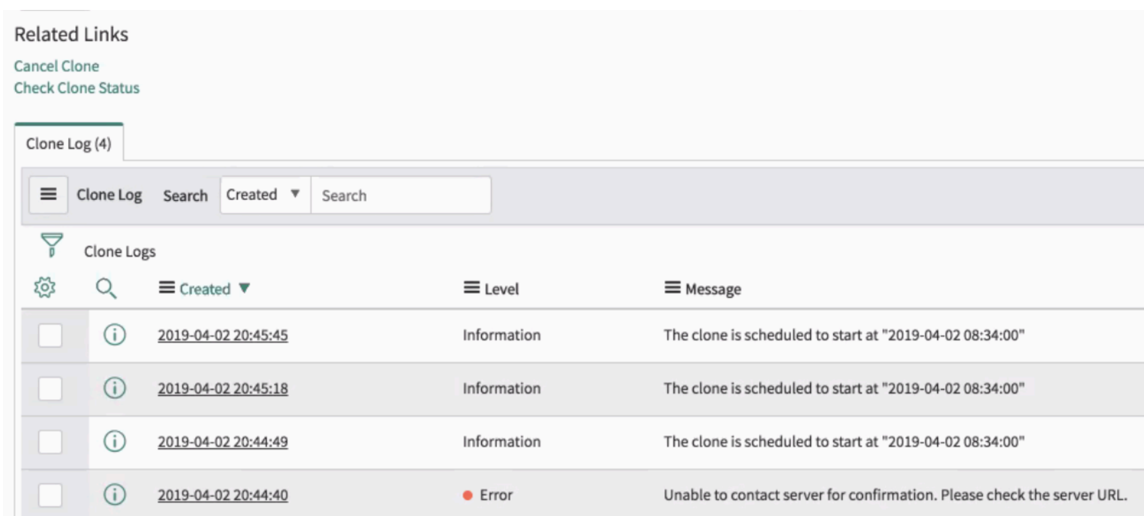

   If a cloning status is failed, you might roll back the cloning and schedule a new cloning.

View clone history

You can view the status and history of any system clone request.

Before you begin

Role required: clone_admin or admin

About this task

The System Clone [clone_instance] table stores records for all previously and currently scheduled clones.

servicenow

Procedure

1. Navigate to **All** > **System Clone** > **Live Clones** > **Clone History**.
   Clone history also displays the **State** for current and past clones. Clones in the **Draft** state do not appear on the Clone History table.

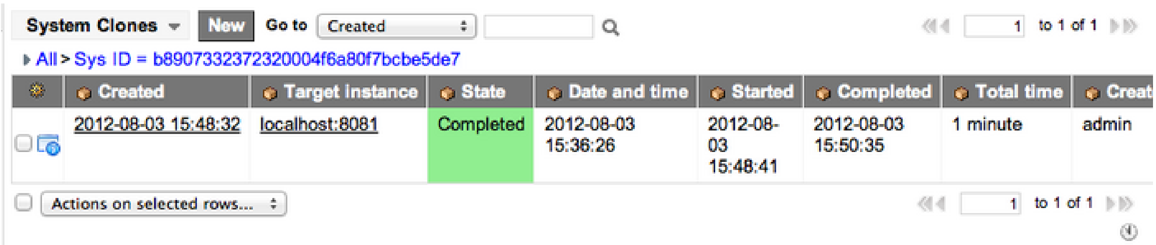

Clone states

Clone state	Description
Requested	The clone was requested and is awaiting approval.
Scheduled	The clone is ready to begin at the scheduled time and date.
Active	The clone is currently running.
Completed	The clone completed successfully.
Canceled	A user canceled the request.
Hold	The server rejected the clone request. This can happen either because the clone was not ready to proceed by the scheduled time or because additional clone requests were submitted before the first one completed.
Error	The clone encountered an error while running. Contact technical support for help resolving this issue.
Draft	The clone is scheduled to be created. This state never appears in the Clone History table.

2. Select a System Clone record to view its history.

Roll back a clone

Return a clone target to its state before the latest cloning. Roll back cloning if something went wrong in the cloning process.

Before you begin

Role required: admin

Rollbacks remove the latest cloning updates on a cloning target. You can only roll back the latest cloning updates. Rollbacks must occur no more than seven days after the cloning.

Note:
Rollbacks must occur no more than two days after the cloning for users with sharded instances.

Procedure

1. Navigate to **All** > **System Clone** > **Live Clones** > **Clone History**.
2. Select the clone to roll back.
3. Under **Related Links**, click **Rollback**.

The **Rollback Clone** modal displays.

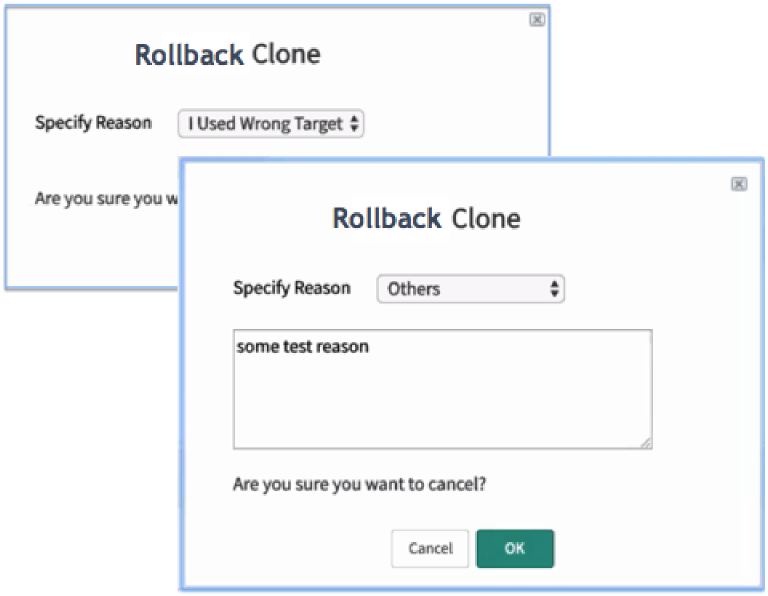

4. In **Specify Reason**, select a reason for the rollback and click **OK**, or select **Others**, enter a reason in the text box and click **OK**.
   The system rolls back the clone and sets the **State** to **Rollback Requested**, and, if successful, to **Rolled Back**, or, if unsuccessful, to **Rollback Failure**.

   You can retry a rollback but if you get repeated failures, contact ServiceNow Customer Support.

Result
The instance contains the same data it had before the clone was applied.

Post-clone cleanup scripts

Cleanup scripts automatically run on the target instance after the cloning process finishes.

Use cleanup scripts to modify or remove bad data. Cleanup scripts run after data preservers and the clone are complete.

You can add new post-cloning scripts on the source instance to perform any action that can normally be accomplished through script includes or business rules. To add a script, navigate to **System Clone** > **Clone Definition** > **Cleanup Scripts** and click **New**.

Note:
Cleanup scripts will always run on clones created using a Clone profile, this is necessary as they're available OOB and ensure a healthy state post clone. After the clone completes, all cleanup scripts are combined together as a scheduled job named **Execute Clone Cleanup Script: Execute Cleanup Scripts Sequentially** and run till completion in the global scope. The cleanup scripts are ordered based on the order field. Cleanup scripts must be defined on the **Source** instance.

You can make post-clone scripts active or inactive to control whether these scripts run or do not run. You can also set an order number on each script, which enables you to set the order that the active scripts run, with lower numbers having a higher priority.

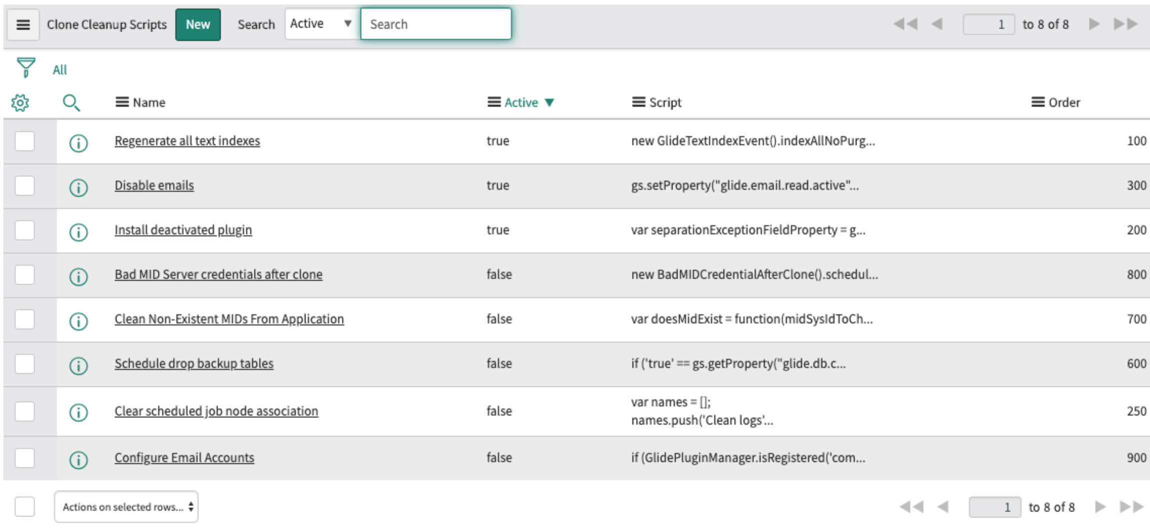

Note:
All cleanup scripts run in the global scope irrespective of the scope in which you have configured the cleanup script. An example of how to run a scoped script is to

- Create a script include in the desired scope with the cleanup logic in it.
- Have a Restricted Caller Access configuration in place to allow access to your scoped script, be sure you are in the scope of the script include used to create the RCA.

Field	Description
Source Scope	Global
Source Type	Scope
Status	Allowed
Target Scope	The scope of the script include.
Target Type	Script Include
Target	The script.
Operation	Execute API

- Call the script include from the clone cleanup script.

The following post-clone cleanup scripts perform various actions on the target instance.

Post-clone cleanup scripts

Script	Description
Bad MID Server credentials after clone	Runs a script include called `BadMIDCredentialAfterClone` on a cloned instance to detect bad MID Server user credentials. This script include creates scheduled jobs that log MID Servers in the **Down** state to the MID Server Issue [ecc_agent_issue] table after an instance clone.
Clear scheduled job node association	Resets any scheduled jobs that were active on the source instance to the **Ready** state. This script also clears the value of the **System ID** and **Claimed by** fields on all scheduled jobs.

Post-clone cleanup scripts (continued)

Script	Description
Configure Email Accounts	Migrates email accounts that existed on the source instance to the target instance if they are not enabled there. This script also migrates the email properties to the target instance.
Disable emails	Disables email on the target instance. A default data preserver maintains other email settings from the target instance.
Install deactivated plugin	Enables the Domain Separation plugin for instances that use this feature.
Regenerate all text indexes	Rebuilds text indexes on the target instance after a clone. Text indexes are not cloned from the source to the target instance.
Schedule drop backup tables	Schedules the deletion of the data that is contained in the target instance database prior to the clone. This original data is preserved for 24 hours following a clone to enable you to roll back an instance to the pre-clone state. If the target instance is downgraded as part of the clone, backup data is not available.

Clone Admin Console

The Clone Admin Console provides a streamlined experience and enhanced visibility for cloning data between instances. You'll find your existing clone profiles, exclusions, preservers, and scripts within the new console. It has added guidance on the clone request page, an on-demand backup option, and provides more visibility into clone progress.

https://player.vimeo.com/video/1006365658?h=e210d21863&badge=0&autopause=0&player_id=0&app_id=58479https://player.vimeo.com/video/1006365658%3Fh=e210d21863&badge=0&autopause=0&player_id=0&app_id=58479https://player.vimeo.com/video/1006365658%3Fh=e210d21863&badge=0&autopause=0&player_id=0&app_id=58479

Cloning

Clone copies data and metadata from one ServiceNow instance (source instance) to another ServiceNow instance (target instance). Cloning is typically used to copy the data and metadata from a production instance to a non-production instance to test changes. Cloning data comes from the most recent, daily backup. To learn more about clone, see System clone.

Notes regarding the Clone Admin Console

1. Legacy clones and new clones are stored separately. Clones requested via the Clone Admin Console are stored on a new table and displayed within the new console.#Legacy clones aren't shown in the console. Clones initiated via the legacy Request Clone page are stored on the legacy Clone History table.
2. Since requests submitted from the clone admin console are stored in a separate table, request numbers in this new table may coincide with request numbers in the legacy table (clone history).
3. Your existing clone settings and configurations remain the same. You can find your existing clone instances, clone profiles, exclusions, preservers, cleanup scripts within both the new Clone Admin Console, and within their original lists and forms.

4. Access and clone-related permissions: The Clone Admin Console requires the clone admin or the admin role on the source and the admin role on the target instance to place a clone request.
5. The new experience uses the same clone engine as the legacy clone.#Known (backend) clone behavior is the same between the legacy and the new user interface.

Overview of the Clone Admin Console

The following sections reflect the different parts of the Clone Admin Console.

Home

The home page displays all current clones in your instance. You can search for your clones by using the **search bar** and entering in the number for your clone.

Note:
The dashboard shows any clones requested through the app. You can't view clones requested through the Now Platform in the dashboard.

Filter options are available to identify the clones you want to view with the filter drop-down button. You can apply the following conditions to be displayed on your dashboard:

- Canceled
- Completed
- Error
- Requested
- Rollback failure
- Rollback requested
- Rolled back
- Rolling back
- Running

You can choose to view your clones in either a grid or list view by selecting the appropriate button option.

Configurations

The **Overview** tab displays the current number of **Clone instances** and **Clone profiles** in your instance.

On the **Clone instances** tab, you can view all currently available clone instances. Once you add an instance to this list, you can use it as a clone source or clone target for your clones. To add your non-production instance to your clone instances list, select **New**.

The **Clone profiles** tab displays all of your available **Clone profiles**. You can think of clone profiles as templates for clone, which you can customize, save and reuse to achieve consistent outcomes with each of your clones. To learn more about Clone Profiles see Clone profiles for clone requests.

Note:
The profile **System Profile** is available out-of-the-box and can't be modified.

Custom profiles you create use the default Exclusions, Preservers, and Scripts from the System Profile as a basis.

Note:
When you're creating a custom profile, all existing custom exclusions and preservers are automatically added to it.

You can create as many custom clone profiles as you'd like and edit them as needed. To change the definitions of a clone profile, such as exclusions, preservers or cleanup scripts, select the number under the definition and select the **Edit** button on the page.

Note:
Certain out-of-the-box system items can't be removed from the exclusions, preservers, or scripts list.

Definitions

The **Definitions** page displays an overview for the following clone definitions.

- **Exclusions**
- **Preservers**
- **Cleanup Scripts**

The **Exclusions** tab lists the tables that aren't copied during a system clone. By default, the system excludes tables for logging, auditing, notifications, workflow contexts, and license usage. You can add **Exclusions** by selecting New and adding the table you want to exclude.

The **Preservers** tab displays a list of all of your available data preservers. **Preservers** protect data on the target instance from being overwritten. Preserving data in a table only preserves the specified table and parent tables.

Note:
Child tables must be individually preserved.

To create a preserver see Create new clone preserver.

Importance of saving your in-flight dev work. During a clone, in-progress development work is over-written (for example work-in-progress update sets, scoped apps which only exist on the target instance but not on the source instance). If you have in-progress update sets, you must export them prior to the clone and reimport them after the clone is finished. If you have custom applications which are not yet deployed to production, you must install the respective plugins or apps after the clone is completed. To learn more about clone and app development best practices, see Whitepaper here.

Tip for dealing with large quantities of in-progress updates sets:

1. On your target instance, which has the work in progress (WIP) update sets, filter to obtain a list view with all your WIP update sets.
2. In a separate window/tab, create a new update set.
3. Batch all incomplete update sets from step 1 - i.e. set the parent field on your list of WIP update sets to the new batch update set
    - How to mass-update: You can mass-update the list by holding down the shift key + holding a left click of your mouse while you're selecting the whole column with your mouse.
    - Double click on any one of the fields in the selected column – type the desired value and update all of them to that value.)

4. Set the parent update set to "Completed" and name it "CLNREQ_DATE" (naming suggestion to make it easier to identify).
5. Export the parent update set.
6. Back on your source instance, retrieve the update set (Important: Don't commit!) The parent update set is now added to the retrieved update sets on your source instance.
7. Verify you wait 24 h until your updates set is captured by the daily backup (alternatively, you can select an on-demand backup on the clone admin console.)
8. Perform the clone as usual from your source instance > target instance.
9. Because your parent update set has been retrieved in the backup, your target instance will have your parent update set along with your dev work which is in flight after the clone.
10. Back on your target instance, you can set the parent update set back to work in progress, and unparent it.
11. On the source instance: Retrieved parent update set can now be cleaned up (i.e. deleted).

The **Cleanup scripts** tab displays a list of all of your available scripts. You can write cleanup scripts to automate post-clone steps. To create a cleanup script see Create cleanup scripts.

Request clone

The app features a clone request page, with guidance and explanations for how the various clone settings affect your clone. The new request page also features a new scheduling calendar to help to prevent timing conflicts with ServiceNow maintenance windows.

To learn more about how to request a clone see Request a clone in Clone Admin Console.

Request a clone in Clone Admin Console

Request a clone to copy data from a production instance to a non-production instance or to copy data between non-production instances.

Before you begin
Role required: admin

Configure your target instance before requesting your clone. Do the following

- Navigate to the **Configurations** > **Clone instances**.
- Select **New**.
- Fill in the new clone instance form.

Note:
You can create multiple clone profiles, utilize a reusable clone template, and clone profiles allow you to select the correct exclusions and preservers for your clone.

Configure a clone profile. Do the following

- Navigate to **Configurations** > **Clone profiles**.
- Fill in the new clone profile form.

Note:

Please note that requests initiated via the classic page `clone_instance.do` will not show up on the new clone console dashboard. However, they can still be found on the classic clone history page `clone_instance_list.do`.

Procedure

1. Navigate to **All** > **Clone Admin Console** > **Clone Dashboard**.
2. Select **Request Clone**.
3. On the form, fill in the fields.

Clone request form

Field	Description
Source Instance	The instance you want to clone.
Target Instance	The instance you want to clone the source instance to. **Note:** You can add an instance or select an existing instance without leaving the page.
Clone Profile	The clone profile to be used for your clone.
Clone Scheduled Start Time	The start time to begin cloning your instance.
Email upon completion	The email to be notified of clone completion.

4. **Note:**
   Additional settings can be configured for your clone request, however, these options may significantly increase your clone completion time.

   Select the following options you want to configure for your clone.

   - The amount of data copied from large tables

     Note:
     By excluding this data, this can extend the time of your clone as the data is copied over from the source, and all data older than 90 days is removed.

   - Apply On-demand Backup
   - Clone frequency
   - Exclude tables specified in the Exclusion list

     Note:
     When a table is excluded, this also excludes all of its child tables. For example, excluding the task table, also excludes its child tables such as Incidents, problems, and change tables.

   - Exclude audit and log data

     Note:
     When **sys_audit** is excluded it will break the activity stream for records. The activity stream is built from the **sys_audit** table, so if it's empty all of the updates will show as the username who created the record and display the created time.

   - Exclude attachment data

- Preserve theme
- Number of days of in-progress Global Update Sets to be preserved

5. Select **Continue**.
6. Review your clone request summary and select **Confirm and Submit Clone Request**.

Create cleanup scripts

Use cleanup scripts to modify or remove data that is not appropriate for your clone.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Clone Admin Console** > **Clone Dashboard**.
2. Select **Definitions** from the secondary navigation.
3. Select the **Cleanup scripts** tab.
4. Select **New**.
5. Enter a name for your script.
6. Select the order number for your script.
7. Check **Active** when your script is ready.
8. Enter your script in the **Script** field, and select **Save**.

Create new clone preserver

Create clone preservers to protect data on the target instance from being overwritten.

Before you begin

Role required: admin

Preserving large amounts of data can significantly increase your clone duration. Only preserve tables if necessary.

The **Preservers** tab displays a list of all of your available data preservers. **Preservers** protect data on the target instance from being overwritten. If you have custom applications, you must also manually preserve your unpublished application content. To create a preserver, do the following.

Procedure

1. Navigate to **All** > **Clone Admin Console** > **Clone Dashboard**.
2. Select **Definitions** from the secondary navigation bar.
3. Select the **Preservers** tab.
4. On the **Preservers** task page, select **New**.
5. Enter the table label as the **Name** for your preserver.

   Note:
   The data preserver must have a table selected or it can't be submitted.

6. Select the **Table** to be preserved.
7. Select the **Theme** check box if the data being preserved is a UI property.

8. Define the data to be preserved using the **Condition builder**, and select **Save**.
   The success message displays.

Database rotation

Database rotation involves managing table size growth and archiving old data.

Data within an instance grows quickly, and as these tables grow in size they require management. Database Rotation preserves instance performance and averts risk associated with querying growing data sets utilizing three techniques. These techniques are based on the concept of managing large quantities of data by separating whole sets into individual tables. After this task is performed, each technique handles data in a different manner:

- Table Rotation works by rotating among a small set of tables, and deleting and reusing the old tables for new data.
- Table Extension works by periodically starting a new table and allowing old tables to be easily archived and removed from the system.
- Table Sharding works by enabling records to be grouped in a single table by document ID.

Table rotation plugins

- **Database Rotations Plugin** activates Table Rotation and Extension without any tables automatically included (com.snc.db.rotation)
- **Database Rotations Default Tables Plugin** applies Table Rotation and Extension to specific tables (com.snc.db.rotation_default_tables)

Database Rotation

Functionality	Tables applied to
Table Rotation	`syslog` `sys_querystat` `ecc_queue` `ecc_event` `cmdb_metric` `sysevent`
Table Extension	`sys_email`

Table rotation and extension

The System Definition Table Rotation module allows you to define a new table rotation, a new table extension, a new table shard, or modify an existing one.

Table rotation groups

Name	Base cleaned	Clean base rotation	Duration	Initialized	Rotations	Shard by	Shards	Type
cmdb_ire_incomplete_payloads	false	2019-10-19 12:18:33	1 Day	true	7			Rotation
ecc_agent_command_audit_log	true	(empty)	1 Day	true	7			Rotation
ecc_agent_host_service	true	(empty)	60 Days	true	3			Rotation
ecc_event	true	(empty)	1 Day	true	7			Rotation
ecc_queue	true	(empty)	1 Day	true	7			Rotation

- **Name**: auto-generated from table name
- **Duration**: overall time parameter for function
- **Initialized**: sets function as active (true) or inactive (false)
- **Rotations**: number of tables to be created within Duration
- **Type**: indicates Extension (archiving), Rotation (deletion), or Shard (doc ID) functionality

When you define a new rotation, a schedule is created and new data is subsequently written to one of the tables in the rotation group. You will notice the that the group includes the original table plus a number of additional tables. Be aware that deleting a rotation will delete the additional tables and all the data, therefore the rotation should not be deleted if the data is needed.

See these topics for more information on managing tables:

- Table Rotation
- Table Extension

Activate database rotation

For new instances, database rotation is active by default. If you are upgrading from a previous version, you can activate the Database Rotations plugin if it is not already active.

About this task

Note:
Deployment of this plugin should be executed in partnership with a ServiceNow representative.

If it is not already active, you can activate the **Database Rotations** plugin if you have the admin role.

Procedure

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

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

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

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

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

What to do next
Do not activate the **Database Rotations Default Tables** plugin. Instead, specify the tables manually, after consulting a ServiceNow representative.

Table cleaner

Table cleaner deletes older records automatically and prevents data from growing exponentially.

Table cleaner is a scheduled job that runs once per hour (by default) to delete older, expired, or unwanted records from tables. Table cleaner prevents tables from growing to an unmanageable size and improves query performance.

You can view a list of table cleaner rules in the Auto Flush [sys_auto_flush] table by typing `sys_auto_flush.list` in the filter navigator. The Auto Flush table displays rules for base system tables and their corresponding record ages. Multiple rules can appear for a single table if each rule has a unique condition defined.

Which records are deleted

Each table cleaner rule specifies the following:

- The target table (to delete records from).
- A **Matchfield** value that corresponds to a Date/Time column in the target table. Ideally, the **Matchfield** is a date field that represents how long the record has been active (for example, sys_created_on) or how long since the last update (sys_updated_on).
- An **Age in seconds** value that determines when the delete is triggered.
- Options for cleaning related records in peripheral or audit tables.
- One or more optional conditions to filter which records are deleted. For example, you might specify that only records where 'active = false AND state =closed' are deleted.

When the **Matchfield** date is farther in the past than the **Age in seconds** value, the table cleaner job deletes the record.

Slow rule handling

When the table cleaner job runs, each table cleaner rule runs several queries as part of the process. If there's no index on a rule's match field or on significant portions of its condition, rule processing can be slow because its queries are running inefficiently on large amounts of data.

If a table cleaner rule has a query that takes longer than 30 seconds to complete, the entire table cleaner job is stopped. By default, table cleaner waits two days before including that rule in the table cleaner job again, which enables the table cleaner job to run without disruption in the meantime. You can configure the duration of the waiting period by adding a system property. See Table cleaner properties.

Disabling table cleaner

You can prevent an administrator from creating a table cleaner rule or running the table cleaner on a specific table by adding the Disable Table Cleaner attribute to the table's dictionary record. Some internal system tables have the Disable Table Cleaner attribute added by default.

Table cleaner limitations

Note the following limitations.

- Table cleaner rules aren't supported for tables configured with table rotation or table extension. Some tables in your instance might have legacy table cleaner rules that were established before table rotation or table extension were enabled. These legacy rules can be safely ignored.
- Performance depends on the size of the table and the conditions you specify. For example, if you use a custom column without an index in a large table, performance is severely degraded. Performance also depends on the number of rows to be deleted.
- Table cleaner spends a maximum of 20 minutes to delete records from a single table. If queries are slow, the volume of records deleted in the 20-minute period may be small.
- Table cleaner does not call `DBDelete.setWorkflow()`. This means `DBDelete` objects run with `workflow=false` (false is the default value for a Java Boolean), hence Delete business rules and workflows are not triggered.

Activate table cleanup

Activate table cleanup for tables with older, expired, or unwanted records.

Before you begin
Role required: admin

Procedure

1. Open the Auto Flush [sys_auto_flush] table by typing `sys_auto_flush.list` in the filter navigator and select **New**.
2. Fill in the fields on the form.

Auto Flush form

Field	Description`
Tablename	The name of the target table.
Matchfield	The Date/Time field in the target table that is used to monitor duration. Defaults to sys_created_on, which deletes records based on the amount of time since their creation date.
Age in seconds	The amount of time the system waits before deleting records. When the **Matchfield** value becomes older than the **Age in seconds**, the record is deleted.
Active	Option to activate or deactivate table cleaning for this table.
Application	The application scope to which this table cleaner job applies.

Field	Description`
Clean peripheral	If selected, related records in the Journal Entry [sys_journal_field] table are also deleted. If cleared, the system deletes records from the target table, but not any related journal records in this tables.
Clean audit	If selected, related records in these audit tables are also deleted: ◦ Sys Audit [sys_audit] table ◦ Audit Relationship Change [sys_audit_relation] table Note that the audit records that are created by table cleaner in the Audit Deleted Record [sys_audit_delete] table are preserved. If cleared, the system deletes records from the target table, but not any related audit records in these tables.
Cascade delete	If selected, this option deletes all matching records plus any records referring to them. If cleared, the system deletes matching records, but not records referring to them.
Conditions	Condition builder for specifying filter conditions that define the records to be removed.

3. Select **Submit**.

Result

Table cleaner runs automatically and deletes records when they meet the specified record age and any conditions that you set for them.

Table cleaner properties

These system properties control table cleaner behavior.

To set table cleaner system properties, navigate to the System Properties [sys_properties] table and add the property.

System properties

Property	Description
Configure the number of days table cleaner must wait before attempting to reprocess a slow rule. `glide.db.tablecleaner.days_before_slow_rule_reattempt`	When a table cleaner rule has a query that takes longer than 30 seconds to complete, the entire table cleaner job is stopped. By default, table cleaner waits two days before including that rule in the table cleaner job again. You can change the default waiting period by adding this property.

System properties (continued)

Property	Description
	• Type: integer • Default value: 2
Configure a time limit for a table cleaner rule. `glide.db.tablecleaner.chunk_delete_max_time_spent`	Limits how long the cleaner spends trying to batch-delete records per rule during a single run. The value is given in seconds and defaults to 1200 (20 minutes), which means the table cleaner runs hourly for a maximum of 20 minutes. You can change the default time limit by adding this property. • Type: integer • Default value: 1200

Deleting records safely

You can safely delete records from a table without using scripts and without deleting the table by creating and executing delete jobs. For example, you might want to delete incident records on a test instance without deleting the incident table itself.

After creating a delete job, you can preview the number of affected records before you schedule or execute the job.

If you want to restore deleted records after the job has run, use the rollback option in the delete job.

If you must delete a large number of records from a table rather than a selection of records, use the table cleaner option. For details, see Table cleaner.

Mark records for deletion

Mark records for deletion according to one or more criteria by creating a delete job.

Before you begin

Role required: admin

Procedure

1. Either choose to mark specific records for deletion or mark all records in a list for deletion.
   - To mark only specific records for deletion, navigate to **All** > **System Data Management** > **Delete Jobs**.
   - To mark all records in a list for deletion, in any list of records, right-click a column heading, and then select **Data Management** > **Delete All with preview**.
2. Select the table that contains the records you want to delete.
3. **Optional:** Limit the number of records in the delete job by adding one or more conditions that records must meet.

a. Select a field, operator, and field value.

Example
For example, **[Category] [is] [Software]**.

b. **Optional:** Use the OR and AND operators to add conditions.

c. **Optional:** Select **New Criteria** to add another set of conditions.

Note:
Limiting the number of records that are added to the delete job can help prevent the table from being locked when the job is executed.

4. **Optional:** View the records that match the conditions.
   a. Select **Preview**.
   b. Select the link with the number of matching records.
   The list of records appears in a new browser tab.
5. Return to the Delete Job tab.
6. Select **Continue** to save the current job information.

Result
The job is saved and the form displays options to schedule the job or execute it now.

What to do next
Preview affected records for deletion.

Preview affected records for deletion

Before you run a delete job, preview the count of affected records to see the impact of executing the delete job.

Before you begin
Role required: admin

About this task
Deleting records in a table also deletes records from tables that extend or reference the source table. Preview the affected records before you execute a delete job and learn about other tables that are impacted.

Up to three levels of cascaded records are deleted when the job is executed. For example, if the preview identifies incidents that match the conditions, it will also delete any problems that reference those incidents, defects that reference those problems, and test records that reference the defects. If sys_attachment and sys_attachment_docs table records are associated with the incidents, those records might not appear in the preview but they are deleted as well.

Procedure
1. Navigate to **All** > **System Data Management** > **Delete jobs**.
2. Select a delete job record.
3. Select the **Preview Cascade** related link to view the count of records from the source table that match the conditions in the delete job.

The record count appears in the **Preview record count** column. Counts for records in other tables that reference those source records are listed as well, although those counts are only close estimates.

4. Review the count of affected records by table to ensure that the records are safe to delete.

What to do next

Schedule or execute a job to delete records

Schedule or execute a job to delete records

Schedule a date and time to execute a delete job or execute the job immediately.

Before you begin

Role required: admin

About this task

Consider scheduling the delete job to run during non-business hours to minimize the potential performance impact on your users. Deleting all records in a table temporarily locks the table, which prevents inserts and updates. If you want to delete all records from a table, use the table cleaner option instead. For more information, see Table cleaner.

Procedure

1. Navigate to **All** > **System Data Management** > **Delete jobs**.
2. Select a delete job record.
3. Determine whether to schedule the delete job for a later time or run it right away.

Result

The records are scheduled for deletion or deleted immediately. If you want to restore the deleted records, see Rollback a delete job.

Rollback a delete job

Rollback a completed delete job to restore the deleted records.

Before you begin

Role required: admin

About this task

If you want to restore records deleted through a delete job, execute a rollback operation.

Procedure

1. Navigate to **All** > **System Data Management** > **Delete jobs**.
2. Select a delete job that has been executed.
3. Select the **Rollback** related link to restore the deleted records.
4. When prompted, confirm that you want to rollback the delete job.

Result

The rollback job is executed immediately and the deleted records are restored. If you want to delete the restored records, create a delete job using the same conditions, and then schedule the delete job or execute it immediately.

Updating records safely

Update several records simultaneously without using scripts by creating and executing a batch update job. For example, you might want to update the Assigned to value for multiple incidents on a test instance without running a script.

To update multiple records safely, create an update job, and then preview the number of affected records before you schedule or execute the job. For more information, see Mark records for updating.

If you must restore the updated records to their previous state, use the rollback option in the update job. For more information, see Rollback a delete job.

Mark records for updating

Mark records for updating according to one or more criteria by creating an update job.

Before you begin

Role required: admin

Procedure

1. Either choose to mark specific records for updating or mark all records in a list for updating.
   - To mark only specific records for updating, navigate to **All** > **System Data Management** > **Update Jobs**.
   - To mark all records in a list for updating, in any list of records, right-click a column heading, and then select **Data Management** > **Update All with preview**.
2. Select the table that contains the records you want to update.
3. **Optional:** Limit the number of records in the update job by adding a condition or conditions that records must meet.

   a. Select a field, operator, and field value.

      Example
      For example, **[Category] [is] [Software]**.

   b. **Optional:** Use the OR and AND operators to add conditions.

   c. **Optional:** Select **New Criteria** to add another set of conditions.

   Note:
   Limiting the number of records that are added to the update job can help prevent the table from being locked when the job is executed.

4. **Optional:** View the number of records that match the conditions.

   a. Select **Preview**.

   b. Select the link with the number of matching records.
5. In the Fields & values lists, select the field that you want to update and a new value.
6. **Optional:** Select additional fields and values to update more fields.
7. Select **Continue**.

Result

The job is saved and the form provides options to schedule the job or execute it now.

What to do next

[Schedule or execute a job to update records](#)

Schedule or execute a job to update records

Schedule a date and time to execute an update job or execute the job immediately.

Before you begin

Role required: admin

About this task

Consider scheduling the update job to run during non-business hours to minimize the potential performance impact on your users. Updating all records in a table temporarily locks the table, which prevents inserts and updates.

Procedure

1. Navigate to **All** > **System Data Management** > **Update jobs**.
2. Select an update job record.
3. Determine whether to schedule the update job for a later time or run it right away.

Result

The records are scheduled for update or updated immediately. If you want to revert the updates to the records, see Rollback an update job.

Rollback an update job

Rollback a completed update job to revert the updates to the records.

Before you begin

Role required: admin

About this task

If you want to revert the updates to records made through an update job, execute a rollback operation.

Procedure

1. Navigate to **All** > **System Data Management** > **Update jobs**.
2. Select an update job that has been executed.
3. Select the **Rollback** related link to revert the updated records.
4. When prompted, confirm that you want to rollback the delete job.

Result

The rollback job is executed immediately and the updated records are reverted to their prior versions. If you want to update the records again, create an update job using the same conditions, and then schedule the update job or execute it immediately.

Domain separation and Data Management

Domain separation is supported in Data Management. 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

- There is business logic to ensure data goes into the proper domain for the application's service provider use cases.
- In the application, the user interface, cache keys, reporting, rollups, aggregations, and so on, all consider domain at run time.
- The owner of the instance needs to be able to set up the application to function normally across multiple tenants.

Use case: As a service provider when I use chat to respond to a tenant-customer's message, the client must be able to see my response.

Related topics

Domain separation for service providers

Data export

Export data from the ServiceNow AI Platform in a variety of formats.

Key benefits

- Export an individual record to a PDF or XML file directly from a form. See Export data from a record.
- Export multiple records to a CSV, Excel, PDF, XML, or Google Sheets file directly from a list. See Export data from a list.
- Export multiple records from a table using the CSV, Excel, PDF, or XML processor in a URL with optional parameters. See Export directly from a URL.

You can also export data from your instance using the following options:

- Export multiple records from a table via a web services request. For example, you can create an external application or process to automate the retrieval of data from an instance via web services such as REST or SOAP. For more information, see Web services .
- Push all the data that you want to export to an external file by creating an export set. For more information, see Export sets .

Associated roles

You can export form data to PDF or XML, depending on your role.

- The PDF options are available to all users.
- The XML option is only available to users with the admin role.

You can export records in a list to Excel, CSV, XML, JSON, PDF, or an Export Set depending on your role.

- The Excel, CSV, JSON, and PDF options are available to all users.
- The XML option is only available to users with the admin role.
- The Export Set option is only available to users with the export_set_admin or admin role.

Export formats

Available export formats

Export format	Description
CSV	Export table records as a comma-separated value text file. Use this option to export the currently displayed fields in the list or form as a text file. Configure the list or form to add or remove fields from the export. When exported to CSV, dot-walked fields appear using their full field name, such as u_assignment_group.parent. **Note:** By default, ServiceNow exports all CSV files in Windows-1252 encoding. If you need to export translated data, set the *glide.export.csv.charset* system property to UTF-8.
Excel	Export table records as a Microsoft Excel spreadsheet. Use this option to export the currently displayed fields in the list or form as an Excel spreadsheet. Configure the list or form to add or remove fields from the export.
XML	Export table records as an XML document. Use this option to export all data from a table or all data for a particular record. The XML file has an XML element for each column in the table. See Exporting and importing data via XML. **Note:** Image field data is not preserved when exporting to XML.
PDF	Export table records as a Portable Document Format file. Use this option to export the currently displayed fields in the list or form as a PDF file. Configure the list or form to add or remove fields from the export. **Note:** The PDF export processor prints data from left-to-right, which can produce rendering errors when displaying right-to-left language data.
JSON	Export the table records as a JSON string. The JSON file has an element for each column in the table.

Column headers

The file type you export to, such as CSV or PDF, determines how column headers are exported. In general, formats meant for human consumption such as PDF display the column label, whereas formats meant for machine use, such as CSV, use the database field_name.

Header format by file type

File type	Column header format
Excel	Label
PDF	Label
CSV	field_name
XML	field_name

Header format by file type (continued)

File type	Column header format
JSON	field_name

Export data from a record

Export a record to PDF or XML.

Before you begin

Role required: none for exporting to PDF. Exporting to XML requires the admin role.

About this task

When exporting to PDF, any field hidden by UI scripts (UI policies, UI actions, client scripts) and any fields visible from the current view are exported. When exporting to XML, all fields are exported, regardless of the view.

A PDF is generated only if the form has been submitted. An unsubmitted form generates an empty page.

Procedure

1. Navigate to the record that you want to export.
2. Select the Additional actions menu ().
3. Select an export option.
   - To export the form data to PDF in portrait mode, select **Export** > **PDF (Portrait)**.
   - To export the form data to PDF in landscape mode, select **Export** > **PDF (Landscape)**.
   - To export all fields in the record to XML, select **Export** > **XML (This Record)**.

   When exporting to XML, the system prompts you to save the file, or the browser automatically saves the file to the downloads folder specified in the browser preferences.

Example:

Export options

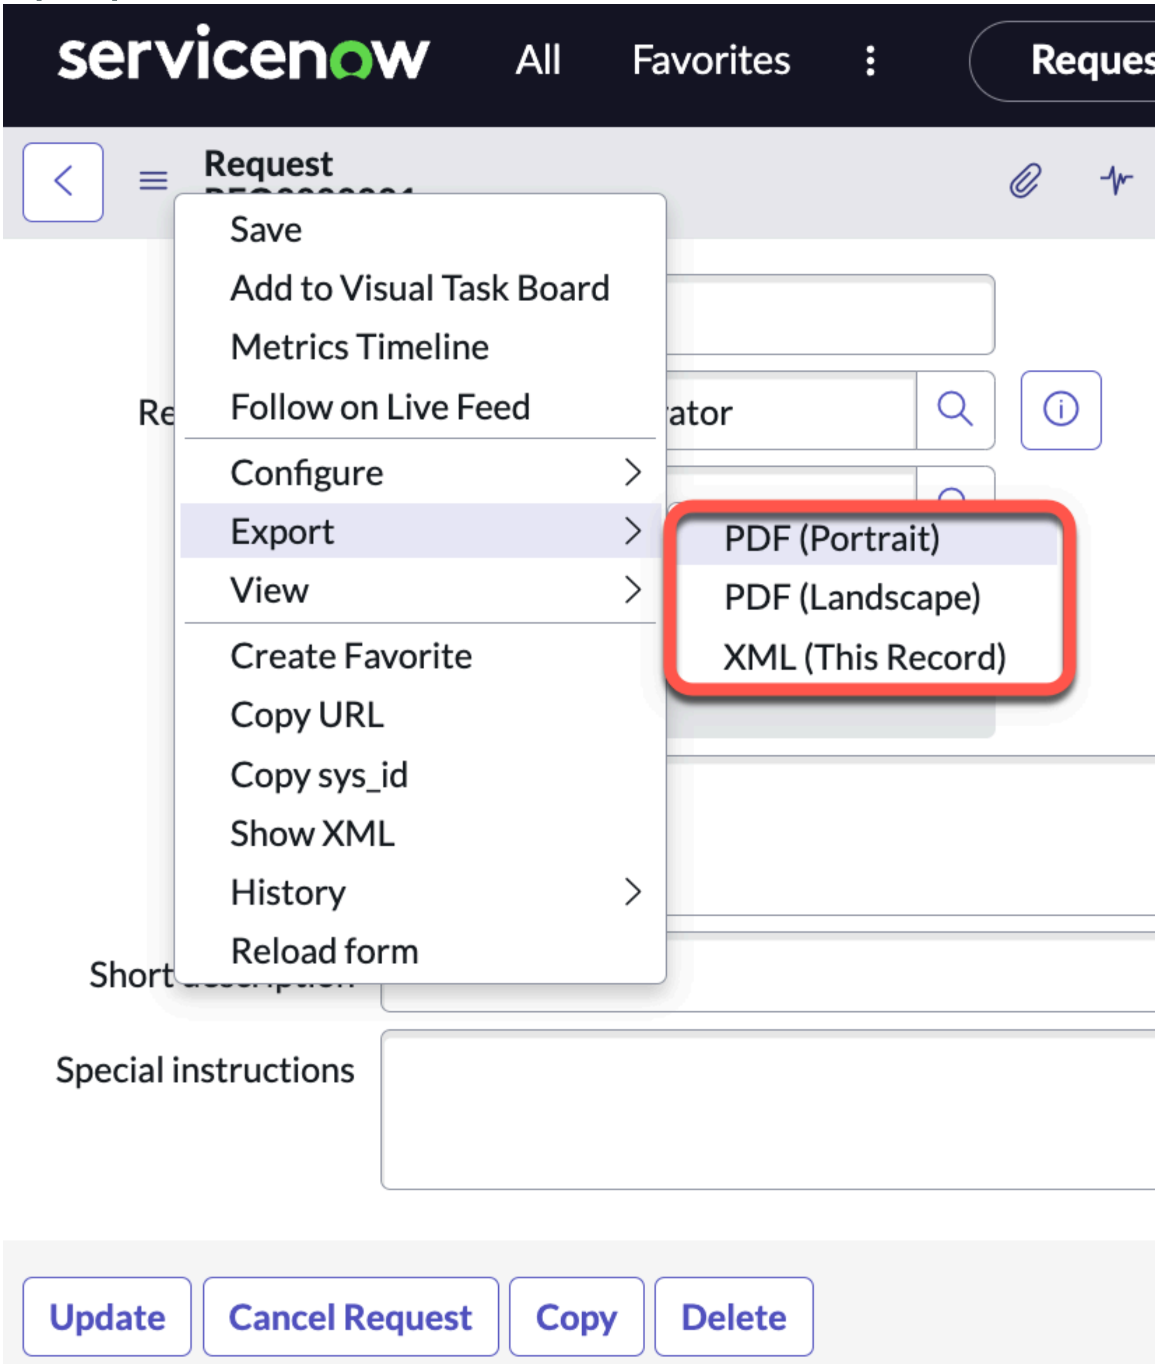

Export data from a list

Export a list of records to a variety of file formats.

Before you begin

Role required: none. Exporting to XML requires the admin role.

About this task

The **Export** option is available in the column heading context menu in List v2. In List v3, it's available in the list title menu.

Note the following important details:

- Data exported from a list may be different from what a user sees in the list view. This generally happens when a table has ACLs that are dependent on the user's current session. To avoid blocking user activity on the browser, list export is executed from a background thread and doesn't have access to user session data.
- To export records in an embedded list, export the record containing the list. See Export data from a record.
- To control how line breaks appear in exported CSV data, use the glide.csv.export.line_break system property. See glide.csv.export.line_break.
- To remove unwanted HTML tags from list data, see the blog Rendering HTML in Exported Lists by a ServiceNow employee in the ServiceNow Community.
- The platform provides a default upper limit for data exports. See Export limits.

Procedure

1. Navigate to the list of records that you want to export.
2. Select and hold (or right-click) any column header in the list.
3. Select an export option.
    - To export the list to Excel, select **Export** > **Excel (.xlsx)**.
    - To export the list to CSV, select **Export** > **CSV**.
    - To export the list to XML, select **Export** > **XML**
    - To export the list to JSON, select **Export** > **JSON**.
    - To export the list to PDF in portrait mode, select **Export** > **PDF** > **PDF Portrait**
    - To export the list to PDF in landscape mode, select **Export** > **PDF** > **PDF Landscape**
    - To export the list and the associated form for each record in portrait mode, select **Export** > **PDF** > **PDF Detailed Portrait**.
    - To export the list and the associated form for each record in landscape mode, select **Export** > **PDF** > **PDF Details Landscape**.
    - To export the list to a file using an export set, select **Export** > **Export Set**.

    For details on using export sets, see Export sets.
4. When the export completes, select **Download**.

servicenow

Example:

Export options

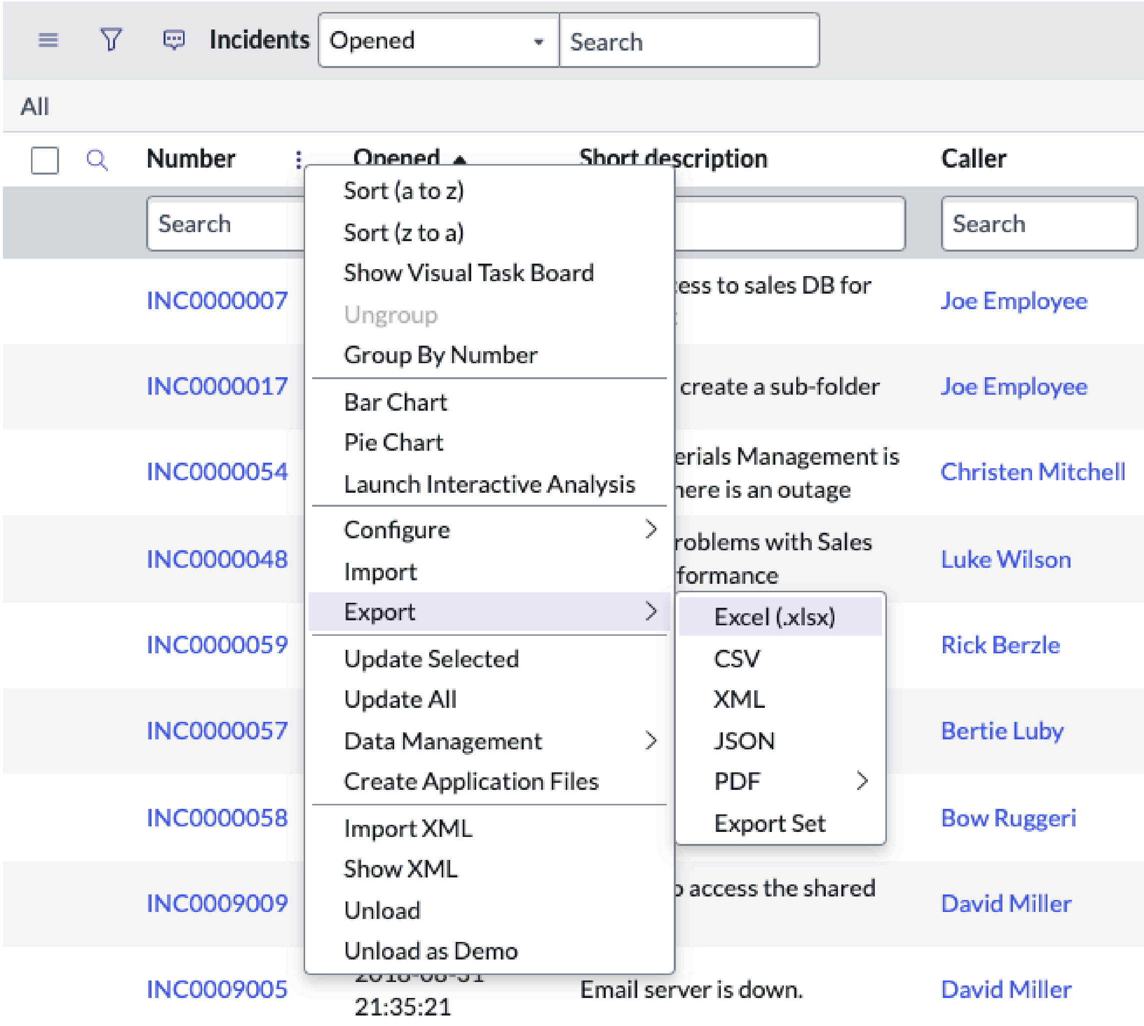

Export directly from a URL

Dynamically export data from a script or web service by building a URL query.

Before you begin

Role required: None

Note:
ServiceNow enforces basic authentication for direct URL access. The data extracted from the URL contains only the fields to which the currently authenticated user has read access.

About this task

You can use URL queries to generate filtered lists of records programmatically before exporting them. URL queries are useful for scripts that generate a list of records or when no user can manually add the filter from the UI.

You must be familiar with the ServiceNow table and column names to export data directly from the URL. See Navigate to a record or a module using a URL.

You can control whether to export the display value or raw value of fields and the field label or field name for headers using query parameters or system properties.

- For details on using query parameters, see Query parameters for display value and header.
- For details on using system properties, see System properties for display value and header.

Query parameters, if present, override system properties.

Procedure

1. Specify the instance URL.
   For example:

   ```
   https://<instance name>.service-now.com/
   ```

2. Specify the table form or list to export.
   For example:

   ```
   incident_list.do
   ```

3. Specify the export format processor to use for the export.
   You can specify CSV, XLS, JSON, XLSX, or XML. For example:

   ```
   ?CSV
   ```

   Note:
   The XML processor in a URL export doesn't export attachments or journal fields. To export attachments and journal fields like the **Work Notes** or **Additional Comments** fields, export the records from a list by using the **Export** > **XML** option in the context menu.

4. **Optional:** Query for specific records by adding the *sysparm_query* parameter.
   The required syntax is `sysparm_query=[column name][operator][value]`. For example:

   ```
   https://<instance
    name>.service-now.com/incident_list.do?XML&sysparm_query=prior
   ity=1
   ```

   Note:
   All queries use a column name, an operator, and a value. See Condition Builder for a list of available operators.

5. **Optional:** Sort the records by column name using the *sysparm_orderby* parameter.
   The required syntax is `sysparm_orderby=[column name]`. For example:

   ```
   https://<instance
    name>.service-now.com/incident_list.do?XML&sysparm_query=prior
   ity=1&sysparm_orderby=assigned_to
   ```

   Note:
   You can sort by only one column using *sysparm_orderby*. To sort by multiple columns, use multiple orderby entries. For example:

   ```
   sysparm_query=ORDERBYassigned_to^ORDERBYpriority
   ```

6. **Optional:** Export fields from a specific view using the *sysparm_view* parameter.
   The required syntax is `sysparm_view=[view name]`. For example:

   ```
   https://<instance
    name>.service-now.com/incident.do?CSV&sysparm_view=ess
   ```

   This URL exports incident fields visible from the Self Service (ess) view.

- By default, exporting data from a list exports only the fields that are visible from the current view. If you want to export fields from another list view, you can switch views from the UI. See View management.
- If you're exporting CSV or Excel data and don't specify a view, the export uses the default list view.
- If you're exporting XML data, all fields are exported by default unless you specify a particular view.

7. **Optional:** Export specific fields from the table using the *sysparm_fields* parameter.
   The required syntax is `sysparm_fields=[field1],[field2]`. For example:

   ```
   https://<instance
    name>.service-now.com/incident_list.do?CSV&sysparm_fields=sys_
   id,number
   ```

8. **Optional:** Export all fields from the table, including the sys_id, using the *sysparm_default_export_fields* parameter.
   The required syntax is `sysparm_default_export_fields=all`. For example:

   ```
   https://<instance
    name>.service-now.com/incident_list.do?CSV&sysparm_default_exp
   ort_fields=all
   ```

   The *sysparm_default_export_fields* parameter has no effect on XML exports unless you include the *sysparm_view* parameter to specify a non-default view.

Example: Sample URLs

URL	Description
https://<instance name>.service-now.com/incident_list.do?CSV	Export all incidents to a comma-separated value text file.
https://<instance name>.service-now.com/incident_list.do?CSV&sysparm_query=sys_id%3E%3Db4aedb520a0a0b1001af10e278657d27	Export a particular incident to a comma-separated value text file.
https://<instance name>.service-now.com/incident_list.do?CSV&sysparm_orderby=sys_id	Export all incidents to a comma-separated value text file and sort the list by sys_id.
https://<instance name>.service-now.com/incident_list.do?XML&sysparm_query=priority=1&sysparm_orderby=assigned_to	Export a list of all priority 1 incidents to an XML file and sort the list by **Assigned to** field.

Call URL export programmatically

Dynamically export data from a script or web service by calling a URL export from any programming language.

Before you begin

Role required: none

About this task

The following procedure includes code samples that demonstrate calling a URL export in C# for a .Net framework call:

Procedure

1. Add the following Imports:

```
using System.Net;
using System.IO;
```

2. Call the Download method:

```
static void Main(string[] args)

     {
        // Call to DownloadFile method supplying the URL and
 location to save CSV file locally
        int read =
 DownloadFile("https://
<instance>.service-now.com/incident_list.do?CSV&sysparm_query=p
riority=1&sysparm_orderby=assigned_to",
                     "c:\\test\\incident.csv");
    }
```

3. Create a Download method as follows:

```
public static int DownloadFile(String url,
                                       String localFilename)
    {
        // Function will return the number of bytes processed
        // to the caller. Initialize to 0 here.
        int bytesProcessed = 0;
        // Assign values to these objects here so that they can
        // be referenced in the finally block
        Stream remoteStream = null;
        Stream localStream = null;
        WebResponse response = null;
        // Use a try/catch/finally block as both the WebRequest
 and Stream
        // classes throw exceptions upon error
        try
        {
            // Create a request for the specified remote file
 name
            WebRequest request = WebRequest.Create(url);
            // Create the credentials required for Basic
 Authentication
            System.Net.ICredentials cred = new
 System.Net.NetworkCredential("user_name", "password");
            // Add the credentials to the request
            request.Credentials = cred;
            if (request != null)
            {
                // Send the request to the server and retrieve
 the
                // WebResponse object
                response = request.GetResponse();
                if (response != null)
                {
```

```
                   // Once the WebResponse object has been
retrieved,
                   // get the stream object associated with
the response's data
                   remoteStream =
response.GetResponseStream();
                   // Create the local file
                   localStream = File.Create(localFilename);
                   // Allocate a 1k buffer
                   byte[] buffer = new byte[1024];
                   int bytesRead;
                   // Simple do/while loop to read from stream
until
                   // no bytes are returned
                   do
                   {
                       // Read data (up to 1k) from the stream
                       bytesRead = remoteStream.Read(buffer,
0, buffer.Length);
                       // Write the data to the local file
                       localStream.Write(buffer, 0,
bytesRead);
                       // Increment total bytes processed
                       bytesProcessed += bytesRead;
                   } while (bytesRead > 0);
               }
           }
       }
       catch (Exception e)
       {
           Console.WriteLine(e.Message);
       }
       finally
       {
           // Close the response and streams objects here
           // to make sure they're closed even if an exception
           // is thrown at some point
           if (response != null) response.Close();
           if (remoteStream != null) remoteStream.Close();
           if (localStream != null) localStream.Close();
       }
       // Return total bytes processed to caller.
       return bytesProcessed;
   }
```

Break up a large export

If the number of records to be exported exceeds the actual export limit, you may want to break the export into smaller increments that do not place a significant performance load on the platform.

Before you begin

Role required: none

servicenow

Procedure

1. Create a filtered list of records that you want to export by following the steps in Export directly from a URL.
2. Write down the number of records returned.
3. If the record number is higher than the defined threshold, issue a `sysparm` query for the first 10,000 records using the following syntax:

   ```
   https://<instance
    name>.service-now.com/syslog_list.do?XML&sysparm_orderby=sys_i
   d&sysparm_record_count=10000
   ```

   This exports the first 10,000 records in order, sorted by the *sys_id* number.

4. Find the next record in order, such as 10,001.
   You can find the next Sys ID value by creating a database view on the table and adding the sys_id column to the view. You don't need to specify a where clause. After you create the database view, view the records and their Sys ID values by selecting **Try it**. You can sort by the sys_id column and enter 10,001 to skip to that row.
5. Right-click the row and copy the sys_id of the next record you want to export.
6. Access the next series of records with a *greater than or equal to* query run against the *sys_id* of record 10,001.

   The following example shows a query that uses a sys_id of b4aedb520a0a0b1001af10e278657d27. Use the syntax shown in this query to export the next set of records.

   ```
   https://<instance
    name>.service-now.com/syslog_list.do?XML&sysparm_query=sys
   _id%3E
   %3Db4aedb520a0a0b1001af10e278657d27&sysparm_orderby=sys_id&sysp
   arm_record_count=10000
   ```

   Note:
   URL queries use typical percent encoding. In this example, the greater than sign (>) is encoded as %3E and the equal sign (=) is encoded as %3D.

7. **Optional:** Continue issuing this query, using the starting sys_id for the next set of records until you have exported all the necessary records.

Exporting and importing data via XML

Migrate data from one instance to another by exporting and importing XML files.

Key benefits

- Directly export records as XML from any list or form.
- Transfer all fields and values exactly from one instance to another.
- Perform unscheduled data imports as needed.
- Save time by not building import sets or transform maps.

Use cases

- Export smaller batches of records on an infrequent basis by exporting and importing data via XML.
- Export and import data using XML files for records created in a development instance that must be migrated with an update set as part of a migration procedure. Examples of these records include lookup tables, unit test records, and other information required to support production. Typically, this information is only migrated once and the overhead of an import set is not justified.

Important considerations

- The XML export preserves all record field values including system generated field values such as:
  - Sys ID
  - Creation date
  - Update date
- Image field data is not preserved when exporting to XML.
- The XML import directly inserts records into the target table. During the import the system:
  - Ignores any business rules that normally apply to the table.
  - Does not provide any opportunity to transform incoming data.
  - Automatically matches a reference field's display value to the local Sys ID for some tables.
- Data export and import via XML is helpful for copying records from one system to another but does not remove the need to know the relationships between tables. If you need validate or transform data or reconcile the sys ID values of reference fields, use an XML import set instead. See Export sets.

Export data to XML

Export data from a source instance to an XML file.

Before you begin

Role required: admin

Procedure

1. Log in to the source instance (the instance that should send the data).
2. Export a record or a list of records to XML.

What to do next

Import data from XML

Import data from XML

After you have successfully exported data from the source instance to XML, import the XML file directly to the target instance.

Before you begin

Role required: admin

About this task

Importing XML doesn't trigger business rules or update the instance cache.

Procedure

1. Log in to the target instance (the instance that should receive the data).
2. Navigate to any list in the system.
   Any list can be used because the XML file contains the destination table name for the records.
3. Select and hold (or right-click) any column title and select **Import XML**.
4. In the Import XML form, select **Choose File** and select the previously exported XML file.
5. Select **Upload**.

Result

Data is imported into the target instance.

Trouble?

If the data doesn't import, navigate to **System Definition** > **Tables & Columns** and verify that the table from which the data was exported also exists in the target instance. If the table doesn't exist, you can move it using an update set. For details on using update sets, see Get started with update sets.

Automatic matching of display values

During the import of XML records, the system attempts to match some reference field display values to a local sys_id value.

If the system finds an existing record with a matching display value on the local instance, the import uses the sys_id of the existing record rather than the sys_id of the imported record.

For example, suppose you export an incident record that is assigned to the user John Smith. In the exported XML file there is an entry such as:

```
<incident>
 ...
 <assigned_to display_value="John
 Smith">7712173d2ba80200c5244f74b4da159a</assigned_to>
 ...
</incident>
```

This user already exists on the target instance but has a different sys_id value such as:

```
<sys_user><name>John Smith</name>
 ...
 <sys_id>18cab8de2be80200c5244f74b4da15f7</sys_id>
 ...
</sys_user>
```

Since the display value matches an existing record, the system uses the local instance's existing sys_id value for the reference field such as:

```
<incident>
 ...
 <assigned_to display_value="John
 Smith">18cab8de2be80200c5244f74b4da15f7</assigned_to>
 ...
</incident>
```

The system can match display values for the following tables.

- User [sys_user]
- Group [sys_user_group]
- Role [sys_user_role]
- Group Roles [sys_group_has_role]

Exporting currency fields to Excel

Exporting currency fields to Excel applies Account formatting and can be configured to convert all values to US dollars or to export values in the user's session currency.

When exporting currency fields to Microsoft Excel, the cells containing currency data use Accounting formatting. This formatting allows you to perform numeric operations on those values as well as view the currency symbol, such as $ or €.

When exporting currency fields the type of currency, such as $ or €, is preserved by default. The conversion rates for non-USD currencies are stored on the Exchange Rates [fx_rate] table.

You can choose to export all currency values either in US dollars or in the user's session currency.

- To export all currency values in US dollars, set the property *glide.excel.fixed_currency_usd* to **true**.
- To export all currency values in the user's session currency, set the property *glide.excel.convert_to_user_currency* to **true**.

If both *glide.excel.fixed_currency_usd* and *glide.excel.convert_to_user_currency* are set to **true**, *glide.excel.fixed_currency_usd* overrides *glide.excel.convert_to_user_currency*.

Currency export behavior and symbols

The behavior of the currency field in Excel depends on the **Symbol** of the currency.

- When the currency symbol is $, the Excel cell type is set as **Currency**.
- When the currency symbol is 1 character (such as ¥ or ¢) or more than 1 character (such as KR), the Excel cell type is set to **Custom**.

Related topics

Currency administration

Single-currency mode

Control default currency field configuration and use in an instance

Enable export debug logging

When the property `glide.export.debug` is true, the instance logs export processing including database query time and the time taken to write data to the file.

When the property `glide.export.debug` is *true*, the instance logs export processing including database query time and the time taken to write data to the file. Debug logs are indicated by the text *Export API*. Prolonged use of this property can affect performance, so it is best to use it while debugging export processing, and then set the property back to *false*.

```
7/17/14 15:53:48 (500) EB39A310EB022100C46AC2EEF106FED9 Maximum
 record count for this instance is: 10000: , request is for: 0
 Cap the Record count to Maximum Record Count
```

```
07/17/14 15:53:48 (522) EB39A310EB022100C46AC2EEF106FED9 Export
 API - ExportProcessor : Processing EXCEL export request,
 ExportParameters:TableName=incident, Query=active=true,
 Limit=10000, SortBy=null
07/17/14 15:53:48 (527) EB39A310EB022100C46AC2EEF106FED9 Export
 API - ExportProcessor : Export using background thread
07/17/14 15:53:48 (528) EB39A310EB022100C46AC2EEF106FED9
 #748 /poll_processor.do -- total transaction time: 0:00:03.357,
 total wait time: 0:00:00.000, session wait: 0:00:00.000,
 semaphore wait: 0:00:00.000, source: 0:0:0:0:0:0:0:1%0
07/17/14 15:53:48 (550) SYSTEM Enabling elevated role:
 security_admin
07/17/14 15:53:48 (588) SYSTEM Export API - ExcelExporter : 29
 rows retrieved from database duration_milliseconds=2
07/17/14 15:53:49 (534) NONE New transaction
 EB39A310EB022100C46AC2EEF106FED9 #751 /poll_processor.do
07/17/14 15:53:49 (544) EB39A310EB022100C46AC2EEF106FED9
 #751 /poll_processor.do Parameters -------------------------
    sys_action=poll
    sysparm_processor=poll_processor
    job_id=a61a2314eb022100c46ac2eef106fe0a
07/17/14 15:53:49 (547) EB39A310EB022100C46AC2EEF106FED9
 #751 /poll_processor.do -- total transaction time: 0:00:00.013,
 total wait time: 0:00:00.000, session wait: 0:00:00.000,
 semaphore wait: 0:00:00.000, source: 0:0:0:0:0:0:0:1%0
07/17/14 15:53:49 (740) SYSTEM Export API - ExcelExporter : Rows
 written to file duration_milliseconds=1150
```

Data export reference

Reference topics provide details on exporting data.

Default values for column headers and column values

Default values are used for column headers and column values, unless overridden by query parameters, Export Set fields, or system properties.

The following table describes the default values used if you do not use query parameters, Export Set fields, or system properties to control output format.

Output format	Column headers	Column values
CSV	**Use field name**	**Use raw value** if *glide.export.csv.raw.value* set, otherwise **Use display value**.
Excel	**Use field label**	**Use display label**
JSON	N/A	**Use raw value**
XLSX	**Use field label**	**Use display value**
XML	N/A	**Use raw value**

Field types affected by export controls

Different field types are affected differently by export controls.

The following table describes how different field types are affected by export controls.

Field type	Display value	Raw value
Date/Time	CSV: User timezone date/time with user format.	CSV: UTC timezone date/time with system format.
	Excel/XLSX: User timezone date/time with user format if *glide.excel.use_user_date_format* set, otherwise system format.	Excel/XLSX: UTC timezone date/time with user format if *glide.excel.use_user_date_format* set, otherwise system format.
Date	Date with user format.	Date with system format.
	Excel/XLSX: Date with user format if *glide.excel.use_user_date_format* set, otherwise system format.	Excel/XLSX: Date with user format if *glide.excel.use_user_date_format* set, otherwise system format.
Time	CSV: Time with user format.	CSV: UTC time, assuming 1970-01-01 as the date with system format.
	Excel/XLSX: User timezone time with user format if *glide.excel.use_user_date_format* set, otherwise system format.	Excel/XLSX: UTC time, assuming 1970-01-01 as the date with user format if *glide.excel.use_user_date_format* set, otherwise system format.
References	Display value of the referenced record.	Sys_id of the referenced record.
Choice	Label of the selected choice.	Value of the selected choice.
Currency	Value with the currency symbol.	USD value without the currency symbol.
Price (calculated)	Value with the currency symbol.	USD value without the currency symbol.
Price (fixed)	Value with the currency symbol.	USD value without the currency symbol.

System properties for display value and header

Use system properties to export the display value or raw value of fields and the field label or field name for headers.

System properties control the output for the CSV, Excel, JSON, and XLSX file formats. System properties define default behavior if overrides such as query parameters are not specified. System properties apply to exported files downloaded with URL parameters, exported from list view, or generated by export sets.

The following table describes each of the system properties that control display values and headers.

Format	Property	Default Value
CSV	*glide.export.csv.raw.value*	**false** **Note:** Setting the property as false also crops the value's string to 32000 characters.

Format	Property	Default Value
	`glide.export.csv.column_header_label`	**false**
Excel	`glide.export.excel.display_value`	**true**
	`glide.export.excel.column_header_label`	**true**
JSON	`glide.json.return_displayValue`	**false**
XLSX	`glide.export.xlsx.display_value`	**true**
	`glide.export.xlsx.column_header_label`	**true**

Example: CSV

Example:

glide.export.csv.raw.value = **false** (default)

glide.export.csv.column_header_label = **false** (default)

number	opened_at	short_description	caller_id	state	assignment_group	assigned_to
INC0000060	12/12/16 7:19	Unable to connect to email	Joe Employee	Closed	Network	David Loo
INC0000059	8/10/16 9:14	Unable to access team file share	Rick Berzle	New		
INC0000058	8/10/16 9:37	Performance problems with email	Bow Ruggeri	New		
INC0000057	8/10/16 9:14	Performance problems with wifi	Bertie Luby	New		
INC0000055	9/20/17 21:47	SAP Sales app is not accessible	Carol Coughlin	In Progress	Service Desk	Beth Anglin

Example:

glide.export.csv.raw.value = **true**

glide.export.csv.column_header_label = **true**

Number	Opened	Short description	Caller	State	Assignment group	Assigned to
INC0000060	12/12/16 15:19	Unable to connect to email	681ccaf9c0a8016400b98a06818d57c7	7	287ebd7da9fe198100f92cc8d1d2154e	5137153cc611227c000bbd1bd8cd2007
INC0000059	8/10/16 16:14	Unable to access team file share	5137153cc611227c000bbd1bd8cd2006	1		
INC0000058	8/10/16 16:37	Performance problems with email	f298d2d2c611227b0106c6be7f154bc8	1		
INC0000057	8/10/16 16:14	Performance problems with wifi	7a82abf03710200044e0bfc8bcbe5d27	1		
INC0000055	9/21/17 4:47	SAP Sales app is not accessible	46c1293aa9fe1981000dc753e75ebeee	2	d625dccec0a8016700a222a0f7900d06	46d44a23a9fe19810012d100cca80666

Example: Excel

Example:

glide.export.excel.display_value = **true** (default)

glide.export.excel.column_header_label = **true** (default)

Number	Opened	Short description	Caller	State	Assignment	Assigned to
INC0000060	2016-12-12 07:19:57	Unable to connect to email	Joe Employee	Closed	Network	David Loo
INC0000059	2016-08-10 09:14:29	Unable to access team file share	Rick Berzle	New		
INC0000058	2016-08-10 09:37:45	Performance problems with email	Bow Ruggeri	New		
INC0000057	2016-08-10 09:14:59	Performance problems with wifi	Bertie Luby	New		
INC0000055	2017-09-20 21:47:23	SAP Sales app is not accessible	Carol Coughlin	In Progress	Service Desk	Beth Anglin

Example:

glide.export.excel.display_value = **false**

glide.export.excel.column_header_label = **false**

number	opened_at	short_description	caller_id	state	assignment_group	assigned_to
INC0000060	2016-12-12 15:19:57	Unable to connect to email	681ccaf9c0a8016400b98a06818d57c7	7	287ebd7da9fe198100f92cc8d1d2154e	5137153cc611227c000bbd1bd8cd2007
INC0000059	2016-08-10 16:14:29	Unable to access team file share	5137153cc611227c000bbd1bd8cd2006	1		
INC0000058	2016-08-10 16:37:45	Performance problems with email	f298d2d2c611227b0106c6be7f154bc8	1		
INC0000057	2016-08-10 16:14:59	Performance problems with wifi	7a82abf03710200044e0bfc8bcbe5d27	1		
INC0000055	2017-09-21 04:47:23	SAP Sales app is not accessible	46c1293aa9fe1981000dc753e75ebeee	2	d625dccec0a8016700a222a0f7900d06	46d44a23a9fe19810012d100cca80666

Example: JSON

Example:

glide.json.return_displayValue = **false** (default)

```
{"records":
[{"parent":"","made_sla":"false","caused_by":"","watch_list":"","upon_reject":"","sys_upda
ted_on":"2017-10-11
20:16:07","child_incidents":"","hold_reason":"","approval_history":"","number":"INC0000001
","resolved_by":"6816f79cc0a8016401c5a33be04be441","sys_updated_by":"admin","opened_by":"6
81ccaf9c0a8016400b98a06818d57c7","user_input":"","sys_created_on":"2016-02-10
```

Example:

glide.json.return_displayValue = **true**

```
{"records":
[{"parent":"","made_sla":"false","caused_by":"","watch_list":"","upon_reject":null,"sys_up
dated_on":"2017-10-11
13:16:07","child_incidents":"","hold_reason":"","approval_history":"","number":"INC0000001
","resolved_by":"System Administrator","sys_updated_by":"admin","opened_by":"Joe
Employee","user_input":"","sys_created_on":"2016-02-10
```

Example: XLSX

Example:

glide.export.xlsx.display_value = **true** (default)

glide.export.xlsx.column_header_label = **true** (default)

Number	Opened	Short description	Caller	State	Assignment group	Assigned to
INC0000060	2016-12-12 07:19:57	Unable to connect to email	Joe Employee	Closed	Network	David Loo
INC0000059	2016-08-10 09:14:29	Unable to access team file share	Rick Berzle	New		
INC0000058	2016-08-10 09:37:45	Performance problems with email	Bow Ruggeri	New		
INC0000057	2016-08-10 09:14:59	Performance problems with wifi	Bertie Luby	New		
INC0000055	2017-09-20 21:47:23	SAP Sales app is not accessible	Carol Coughlin	In Progress	Service Desk	Beth Anglin

Example:

glide.export.xlsx.display_value = **false**

glide.export.xlsx.column_header_label = **false**

number	opened_at	short_description	caller_id	state	assignment_group	assigned_to
INC0000060	2016-12-12 15:19:57	Unable to connect to email	681ccaf9c0a8016400b98a06818d57c7	7	287ebd7da9fe198100f92cc8d1d2154e	5137153cc611227c000bbd1bd8cd2007
INC0000059	2016-08-10 16:14:29	Unable to access team file share	5137153cc611227c000bbd1bd8cd2006	1		
INC0000058	2016-08-10 16:37:45	Performance problems with email	f298d2d2c611227b0106c6be7f154bc8	1		
INC0000057	2016-08-10 16:14:59	Performance problems with wifi	7a82abf03710200044e0bfc8bcbe5d27	1		
INC0000055	2017-09-21 04:47:23	SAP Sales app is not accessible	46c1293aa9fe1981000dc753e75ebeee	2	d625dccec0a8016700a222a0f7900d06	46d44a23a9fe19810012d100cca80666

Query parameters for display value and header

Use query parameters to export the display value or raw value of fields and the field label or field name for headers.

Query parameters control the output for CSV, XLS, JSON, XLSX, and XML file formats. Query parameters, if present, override system properties.

The following table describes each of the query parameters that control display values and headers.

Query parameter	Description
sysparm_display_value	Exports the display value (**true**) or raw value (**false**) for the CSV, XLS, XLSX, and XML file formats.

Query parameter	Description
	See Field types affected by export controls for more information.
displayvalue	Exports the display value (**true**), raw value (**false**), or both (**All**) for the JSON file format. See Field types affected by export controls for more information.
sysparm_export_column_header_label	Exports the CSV, XLS, and XLSX file formats with the field label (**true**) or field name (**false**).

Example: sysparm_display_value

Example: *sysparm_display_value* **= true**

Number	Opened	Short description	Caller	State	Assignment group	Assigned to
INC0000060	2016-12-12 07:19:57	Unable to connect to email	Joe Employee	Closed	Network	David Loo
INC0000059	2016-08-10 09:14:29	Unable to access team file share	Rick Berzle	New		
INC0000058	2016-08-10 09:37:45	Performance problems with email	Bow Ruggeri	New		
INC0000057	2016-08-10 09:14:59	Performance problems with wifi	Bertie Luby	New		
INC0000055	2017-09-20 21:47:23	SAP Sales app is not accessible	Carol Coughlin	In Progress	Service Desk	Beth Anglin

Example: *sysparm_display_value* **= false**

Number	Opened	Short description	Caller	State	Assignment group	Assigned to
INC0000060	2016-12-12 15:19:57	Unable to connect to email	681ccaf9c0a8016400b98a06818d57c7	7	287ebd7da9fe198100f92cc8d1d2154e	5137153cc611227c000bbd1bd8cd2007
INC0000059	2016-08-10 16:14:29	Unable to access team file share	5137153cc611227c000bbd1bd8cd2006	1		
INC0000058	2016-08-10 16:37:45	Performance problems with email	f298d2d2c611227b0106c6be7f154bc8	1		
INC0000057	2016-08-10 16:14:59	Performance problems with wifi	7a82abf03710200044e0bfc8bcbe5d27	1		
INC0000055	2017-09-21 04:47:23	SAP Sales app is not accessible	46c1293aa9fe1981000dc753e75ebeee	2	d625dccec0a8016700a222a0f7900d06	46d44a23a9fe19810012d100cca80666

Example: displayvalue

Example: *displayvalue* **= true**

```
{
  "records": [
    {
      "parent": "",
      "made_sla": "false",
      "caused_by": "",
      "watch_list": "",
      "upon_reject": null,
      "sys_updated_on": "2017-10-11 13:16:07",
      "child_incidents": "",
      "hold_reason": "",
      "approval_history": "",
      "number": "INC0000001",
      "resolved_by": "System Administrator",
      "sys_updated_by": "admin",
      "opened_by": "Joe Employee",
```

Example: *displayvalue* **= false**

```
{
  "records": [
    {
      "parent": "",
      "made_sla": "false",
      "caused_by": "",
      "watch_list": "",
      "upon_reject": "",
      "sys_updated_on": "2017-10-11 20:16:07",
      "child_incidents": "",
      "hold_reason": "",
      "approval_history": "",
      "number": "INC0000001",
      "resolved_by": "6816f79cc0a8016401c5a33be04be441",
      "sys_updated_by": "admin",
      "opened_by": "681ccaf9c0a8016400b98a06818d57c7",
```


Example: *displayvalue* **= All**


```
{
  "records": [
    {
      "parent": "",
      "dv_incident_state": "Closed",
      "caused_by": "",
      "watch_list": "",
      "dv_sys_mod_count": "21",
      "dv_opened_by": "Joe Employee",
      "dv_sys_id": "9c573169c611228700193229fff72400",
      "upon_reject": "",
      "sys_updated_on": "2017-10-11 20:16:07",
      "approval_history": "",
      "number": "INC0000001",
      "dv_work_end": "",
      "dv_made_sla": "false",
```


Example: sysparm_export_column_header_label

Example: *sysparm_export_column_header_label* **= true**

Number	Opened	Short description	Caller	State	Assignment group	Assigned to
INC0000060	2016-12-12 07:19:57	Unable to connect to email	Joe Employee	Closed	Network	David Loo
INC0000059	2016-08-10 09:14:29	Unable to access team file share	Rick Berzle	New		
INC0000058	2016-08-10 09:37:45	Performance problems with email	Bow Ruggeri	New		
INC0000057	2016-08-10 09:14:59	Performance problems with wifi	Bertie Luby	New		
INC0000055	2017-09-20 21:47:23	SAP Sales app is not accessible	Carol Coughlin	In Progress	Service Desk	Beth Anglin

Example: *sysparm_export_column_header_label* **= false**

number	opened_at	short_description	caller_id	state	assignment_group	assigned_to
INC0000060	2016-12-12 07:19:57	Unable to connect to email	Joe Employee	Closed	Network	David Loo
INC0000059	2016-08-10 09:14:29	Unable to access team file share	Rick Berzle	New		
INC0000058	2016-08-10 09:37:45	Performance problems with email	Bow Ruggeri	New		
INC0000057	2016-08-10 09:14:59	Performance problems with wifi	Bertie Luby	New		
INC0000055	2017-09-20 21:47:23	SAP Sales app is not accessible	Carol Coughlin	In Progress	Service Desk	Beth Anglin

Export limits

The platform provides a default upper limit for data exports.

The purpose of the upper limit is to avoid creating performance issues when a table is excessively large. If you must export more records than the threshold permits, Break up a large export into separate manageable chunks.

servicenow

Export limit properties

You can set the number of records to return during an export using the `sysparm_record_count` URL parameter.

However, the system analyzes the following settings to determine whether an export limit should be applied.

1. First, the platform checks the property that defines the format-specific export limit. Each format can have a different limit. Although this property can be set to any value, exceeding the default export limit can impact system performance. You may want to set the property at or below the default limit and have users Break up a large export to export large amounts of data.
2. If the format-specific property is not set, the system checks the property for the general export limit. This property can also be set to any value, but exceeding the default export limit can impact system performance.
3. If both the format-specific export limit and the general export limit property aren't set, the system enforces the default export limit.

Important:
To set these properties, navigate to **System Properties** > **Import Export**.

Default export limit

Format	Format-specific export limit	General export limit	Default export limit
XML	glide.xml.export.limit	glide.ui.export.limit	10,000
CSV	glide.csv.export.limit	glide.ui.export.limit	10,000
EXCEL (XLSX)	glide.xlsx.export.limit	glide.ui.export.limit	10,000
EXCEL (XLS)	glide.excel.export.limit	glide.ui.export.limit	10,000
EXCEL (XLSX)	glide.xlsx.max_cells	N/A	500,000
EXCEL (XLS)	glide.excel.max_cells	N/A	500,000
PDF	glide.pdf.max_detail_pages **Note:** The number of rows can be set from 0 to 250. If no value is specified, the default is 250. If a value greater than 250 is specified, the default value of 250 is used.	250	250
PDF	glide.pdf.max_rows	N/A	1,000

Default export limit (continued)

Format	Format-specific export limit	General export limit	Default export limit
	Note: The number of rows can be set from 0 to 5,000. If no value is specified, the default is 1,000. If a value greater than 5,000 is specified, the default value of 1,000 is used.		
PDF	glide.pdf.max_columns **Note:** Only 25 header labels fit on a page.	N/A	25
JSON	glide.json.export.limit	glide.ui.export.limit	10,000

A warning threshold property called `glide.ui.export.warn.threshold` controls how the records are exported. When you export enough records from a list to exceed the threshold, a choice dialog appears. You can choose to wait for the export to complete or email the exported records as an attachment. The warning threshold can be changed in the system property. The email attachment must not exceed the maximum Email size limits.

Export limit examples

Example	Property	Breaking export limits
Exporting to CSV	• glide.csv.export.limit = 20,000 • glide.ui.export.limit = 10,000 • com.glide.processors.XMLProcessor.max_record_count = 20,000 • Default export limit for CSV = 10,000	If you export 30,000 records, the number of records returned from the database is 20,000. The number of records specified for export exceeds the value set in the `glide.csv.export.limit` property.
Exporting to Excel (XLSX)	• glide.xlsx.export.limit = no entry • glide.ui.export.limit = no entry • Default export limit for Excel (XLSX) = 10,000	If you export 30,000 records, the number of records returned from the database is 10,000. The number of records specified for export exceeds the default export limit for Excel.
Exporting to PDF	• glide.pdf.max_rows = 1,000 • Default export limit for PDF = 1,000 • Maximum export limit for PDF = 5,000	If you export 5,000 records, the number of records returned from the database is 1,000. The number of records returned is limited because the number of records specified for export exceeds the default value in the `glide.pdf.max_rows` property.

Excel export threshold

Excel exports are intended for relatively small exports with fewer than 500,000 cells. CSV can handle larger exports.

If an exported spreadsheet has more than 500,000 cells, the export process stops. You receive the Excel file at that point. The following message displays in the bottom row: `Export stopped due to excessive size. Use CSV for a complete export.`

The Excel export cell threshold is customizable using the *`glide.xlsx.max_cells`* integer property (or *`glide.excel.max_cells`* if using XLS).

Note:
Increasing this threshold may cause a memory issue in your instance. The threshold is set at an appropriate level to prevent resource issues.

The export puts the information into the Excel document with 32,000 rows per spreadsheet.

Roll back and delete recovery

With rollback contexts, you can roll back certain actions such as a patch upgrade, plugin activation, and background script executions, and you can recover record deletions and all related changes.

The roll back and delete recovery features are available on instances that use MySQL and MariaDB databases. Instances that use Oracle databases only support roll back. Instances that use SQL Server do not support roll back or delete recovery.

Roll back and delete recovery database support

Database type	Roll back	Delete recovery
MySQL	yes	yes
MariaDB	yes	yes
Oracle	yes	no

Deleted Records module

This module works on records in audited tables. Recovery of cascaded deleted records must be done within seven days of the record deletion. After seven days, only data records and references on tables that audit deletions can be recovered, which is the same functionality as prior releases.

To find this module, navigate to **System Definition** > **Deleted Records**.

Delete Recovery module

This module works for any deleted record. This recovery must be done within seven days of the record deletion.

To find this module, navigate to **Rollback & Recovery** > **Delete Recovery**.

Script Execution History module

This module works on scripts executed using the **Scripts - Background** module. This history only includes seven days of script executions.

To find this module, navigate to **Rollback & Recovery** > **Script Execution History**.

Rollback contexts

Rollback contexts contain everything necessary to roll back a software upgrade or plugin activation. They include deleted records, patch updates, Scripts-Background script executions, database actions, and plugin activations. A rollback context is created for each patch upgrade within a family, and each plugin activation, provided that the plugin supports rollback contexts.

To use rollback contexts, activate the Restore Deleted Records (com.snc.undelete) and Delete Recovery (com.glide.delete_recovery) plugins.

Rollbacks are typically performed on pre-production instances where functionality must be restored before you can find the root cause of a problem in the upgrade. Rolling back deletes data, which can often make it hard if not impossible to discover the problem that made the rollback necessary.

A rollback context is created when:

- *`GlideRecord.delete()`* or *`GlideRecord.deleteMultiple()`* delete records.
- There is a patch upgrade.
- You activate a plugin that supports rollback contexts.
- A script executes using the **Scripts-Background** module, and rollback was enabled by selecting the **Record for Rollback?** check box.

Rollbacks do not impact other database activities. If a database activity modifies a record that is part of a rollback context, the rollback does not affect that record.

Because rollback contexts contain a significant amount of data, rollback contexts are deleted after 10 days. Therefore, rollbacks must occur within 10 days of the latest upgrade or plugin activation. If you need to retain a rollback context for more than 10 days, you can do so by adding a system property. See Rollback context properties.

> **Note:**
> Do not roll back a rollback context until checking with Customer Service and Support. A roll back deletes data and may remove evidence of the upgrade or activation issue preventing debugging of the problem.

To find this module, navigate to **Rollback & Recovery** > **Rollback Contexts**.

If any of the following operations occur during a rollback, a rollback context is not created:

- Tables or columns are dropped from the schema.

  > **Note:**
  > Index drops are okay.

- A table is truncated.
- A table or column is renamed.
- A column is re-parented or promoted.
- A column type changes.
- A column width is decreased.

The rollback process on Now Support does the following:

- Updates the reported WAR to the rolled-back version and the assigned WAR remains set to the version before the rollback.
- Sets the property `glide.war.no_upgrade` on the instance is set to the version before the rollback.
- Displays the message, "Desired war matches reverted war specified by property [`glide.war.no_upgrade`]. Upgrade script will NOT run".
- Changes the state to **Expired** and the rollback purges all stored data.

Roll back patch upgrades or plugin activations

Use the **Rollback Contexts** module to roll back patch upgrades and plugin activations.

Before you begin

Role required: admin

Review this information and requirements before performing rollback:

- Rolling back is possible only between patches in the same software release. For example, Rolling back is possible from Orlando patch 2 to patch 1, but not from Orlando to New York.
- The rollback must be done within 10 days of the most recent upgrade or plugin activation. If you need to retain a rollback context for more than 10 days, you can do so by adding a system property. See Rollback context properties.
- You can only roll back to the previous version.

  For example, rolling back from patch 3 to patch 1 requires two rollbacks.
- You must have the Restore Deleted Records (com.snc.undelete) and Delete Recovery (com.glide.delete_recovery) plugins activated.
- You can only roll back one active instance at a time.

  Rolling back contexts on multiple nodes is not supported.
- A rollback deletes data and may remove evidence of an upgrade or activation record, which makes it difficult to debug the problem that made the rollback necessary.
- To roll back to a patch version before the Assigned WAR, submit a request for Customer Service and Support to do the rollback.

Procedure

1. Navigate to **All** > **Rollback & Recovery** > **Rollback Contexts**.
2. Select the context you want to rollback.
3. Roll back by selecting the **Rollback** related link.

   Note:
   The rollback link only works if the app was the last to be updated on the instance. If another app has been updated on the instance, the rollback won't work.

4. Confirm the rollback.
   An indicator appears showing the progress of the rollback. `glide.war` updates to the rolled-back version. The log shows the stages of the rollback process.
5. After the rollback is complete, log out of your instance and log back in again.

6. Open a new tab and see if the context state changed to **Rolled back**.
   If not, contact Customer Service and Support.

Use the Deleted Records module to restore a deleted record

You can recover deleted records that are in audited tables.

Before you begin

Role required: admin

About this task

The instance can track deletions on any table, and references on audited tables, with a few exceptions. Record deletions are not tracked in these circumstances:

- Record deletions are not tracked on tables with the `no_audit_delete=true` dictionary attribute.
- Not all tables with a sys prefix track deletions by default. You can Enable auditing for a system table.
- References are restored only if the reference field is on an audited table and the Restore Deleted Record plugin is activated.
- References that use an **Image** field type are not restored.

Procedure

1. Navigate to **All** > **System Definition** > **Deleted Records**.
2. Open the deleted record that you want to restore.

   Note:
   You can only restore one deleted record and its associated references at a time.

3. Do one of the following to restore the record.

Option	Description
Undelete With Related	Recovers the record and all cascaded deletes and other database actions that re sulted from the delete. This option appears when a rollback context is available for the delete.
Recover entire operation	If this record was deleted as part of another delete, all records from the parent delete are recovered including all cascaded deletes and other database actions that resulted from the parent delete. If this is the top level delete, then this is the same as **Undelete with Relat ed**. This option appears when a rollback context is available for the delete.
Undelete Record	Recovers only this record.

Option	Description
	Important: The undelete operation performs an in sert, which can trigger business rules. This might lead to additional inserts and updates on related records in other ta bles, potentially resulting in unexpect ed outcomes. For guidance on avoiding this behavior, contact Customer Service and Support.

An instructions page appears, explaining the process and its limitations.

4. Click **Restore deleted record**.
   A Progress page shows the progress of the restore process, and shows a count of the references restored.

Result

When the process completes, a Restore Summary itemizes the changes and provides links for viewing the restored record or returning to the deleted records list.

Use the Delete Recovery module to restore a deleted record

You can recover a deleted record and all related changes. The recovery must be done within seven days of the record being deleted.

Before you begin

Role required: admin

About this task

Procedure

1. Navigate to **All** > **Rollback & Recovery** > **Delete Recovery**.
2. Click the item from the **Delete Recoveries** list to be recovered.
   The **Delete Recovery** form is shown.
3. Under Related Links, click **Rollback...**.
4. In the **Rollback entire operation** form, enter `yes`, and click **OK**.
   A Progress page shows the progress of the restore process, and shows a count of the references restored.

Result

When the process completes, a **Restore Summary** itemizes the changes and provides links for viewing the restored record or returning to the **Delete Recoveries** list.

Use the Script Execution History module to roll back a Scripts-Background execution

You can roll back the database actions of a script executed using the Scripts-Background module.

Before you begin

Role required: admin

About this task

When you executed the script using the Scripts-Background module, the **Record for Rollback?** item must have been checked.

If you are still in the **Scripts-Background** module after executing a script, you can use the **Script execution history and recovery available here** link to get to the **Script Execution History** form.

Procedure

1. Navigate to **All** > **Rollback & Recovery** > **Script Execution History**.
2. Click the script execution to be rolled back.
3. In the **Related Links**, click **Rollback Script Execution**.

Rollback context properties

Change the default expiration period for different types of rollback context records.

Because rollback contexts contain a significant amount of data, they're deleted after 10 days by default. The Clean Expired Rollback Contexts scheduled job runs daily to delete expired records in the Rollback Context [sys_rollback_context] table. If you need to retain a rollback context for more than 10 days, you can do so by adding a system property.

To add a system property, navigate to the System Properties [sys_properties] table and add a property for the type of rollback context record that you want to preserve. You can determine which property to add for a given rollback context by matching the property name to the Type column on the Rollback Context [sys_rollback_context] table.

Rollback context properties

Property	Description
The number of days to retain the rollback context for an app install `glide.rollback.expiration_days_app_install`	By default, you have 15 days to roll back an application installed from the ServiceNow® Store before the rollback context expires. You can extend the expiration period by updating this property. • Type: integer • Default value: 15
The number of days to retain the rollback context for a batch app install `glide.rollback.expiration_days_batch_app_install`	By default, you have 10 days to roll back apps installed using the Batch Installation feature before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain the rollback context for demo data `glide.rollback.expiration_clean_demo_data`	By default, you have 10 days to roll back before the rollback context expires. You can change the expiration period by adding this property and setting a new value.

Rollback context properties (continued)

Property	Description
	• Type: integer • Default value: 10
The number of days to retain the rollback context for recorded delete operations `glide.rollback.expiration_days_delete`	By default, you have 10 days after a delete operation completes to roll back before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain the rollback context created for validating Flow Designer flow execution before and after an upgrade `glide.rollback.expiration_days_flow`	By default, you have 10 days to roll back before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain the rollback context for IRE changes recorded from CMDB Integration Studio. `glide.rollback.expiration_days_inst_preview`	By default, you have 10 days after IRE changes are recorded from CMDB Integration Studio to roll back before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain rollback contexts in which the Type is set to Other. `glide.rollback.expiration_days_other`	By default, you have 10 days to roll back before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain the rollback context for a plugin activation `glide.rollback.expiration_days_plugin`	By default, you have 15 days to roll back plugins activated in Application Manager before the rollback context expires. You can extend the expiration period by updating this property. • Type: integer • Default value: 15
The number of days to retain the rollback context for an anonymization job.	By default, you have 10 days after data is redacted by a Data anonymization job to roll

Rollback context properties (continued)

Property	Description
`glide.rollback.expiration_days_reboot`	back before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain the rollback context for a background script execution `glide.rollback.expiration_days_scripts_bg`	By default, you have 10 days to roll back script execution before the rollback context expires. The rollback context tracks all INSERT, DELETE, and UPDATE statements executed by the script and recovers the data by undoing the SQL statements. You can extend the expiration period by updating this property. • Type: integer • Default value: 10
The number of days to retain the rollback context for changes recorded during test executions `glide.rollback.expiration_days_test_execution`	By default, you have 10 days after test execution to roll back before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10
The number of days to retain the rollback context for an upgrade `glide.rollback.expiration_days_upgrade`	By default, you have 10 days after the latest upgrade to roll back the upgrade before the rollback context expires. You can change the expiration period by adding this property and setting a new value. • Type: integer • Default value: 10

ServiceNow AI Platform forms, fields, and lists

Data collection is easy with the ServiceNow AI Platform. Customize your forms with personalization, templates, tailored UI, file attachments, and more.

Form administration	Field administration	List administration
Configure forms and customize their functions to best suit your enterprise demands.	Create and maintain individual data items in your records on the ServiceNow AI Platform.	Control who views and uses data in your instance, and how they use it, with ServiceNow administrative options.

Form administration

Create and modify forms with a wide variety of tools on the ServiceNow AI Platform. Give users roles to configure forms with Form Personalization. Include supplemental info with Annotations, Attachments, and Formatters. Populate forms automatically with Form Templates, to add data based on configurable settings. Form UIs are fully customizable, with desired UI actions, UI policies, and more advanced form configuration options.

Field administration

Fine-tune every detail of your data collection. Field administration includes the ability to create individual pieces of data, modify them as needed. Parent-to-child field relationships and vice-versa are automatically maintained and updated as you make changes in your tables. Multiple Field types are available on the ServiceNow AI Platform, to tailor records on your instance as needed.

List administration

Give users controlled access to data with List administrative tools by ServiceNow. Delegate read-only or read-and-write privileges to users, and set up or modify UI controls for data in your lists.

For more information about how users interact with the ServiceNow user interface, lists, forms, and filters, see the ServiceNow User Interface Overview on ServiceNow University.

Form configuration

Form configuration involves changing the form layout and the related list layout. The form layout changes what appears on the form and the related list layout changes which related lists appear at the bottom of the form.

Important:
Try configuring your forms with Form Builder instead. Form Builder has all of the Form Designer and Form Layout options. It also has other robust features that help you configure each aspect of the form. For more information, see Accessing Form Builder.

Understanding form views; configuring layout; designing forms;configuring related lists

You change aspects of a form such as showing or hiding fields, embedding lists, and adding sections and annotations. This video highlights form layout and design.

Using the form designer

Administrators or users with the personalize_form role can use the form design feature to quickly create new or change existing form views.

Important:
Try configuring your forms with Form Builder instead. Form Builder has all of the Form Designer and Form Layout options. It also has other robust features that help you configure each aspect of the form. For more information, see Accessing Form Builder.

With form design, you can accomplish the same tasks that you can with form layout in a graphical tool called the form designer. You can also access field properties and add information from macros or previously scripted UI elements.

Form design is an alternative to configuring forms that combines several configuration options into one tool. It is available by default for new instances and requires Core UI . Administrators who upgrade from a previous version must activate the Form Designer (com.glide.ui.ng.fd) plugin to use the feature.

Access Form Designer by opening an existing form, selecting the Additional actions icon ( ), and then navigating to **Configure** > **Form Design**.

Accessing Form Designer from an existing form

Parts of the UI form design interface

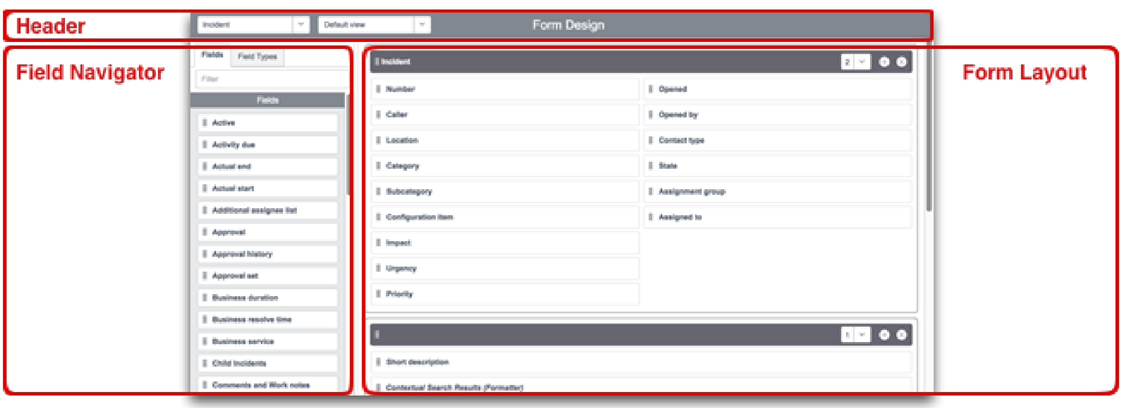

The page header

The page header displays the current table and view being designed. Each form is tied to a unique table and view combination. For example, there can only be one form for the incident table default view. Select a different view to modify different form layouts for a table.

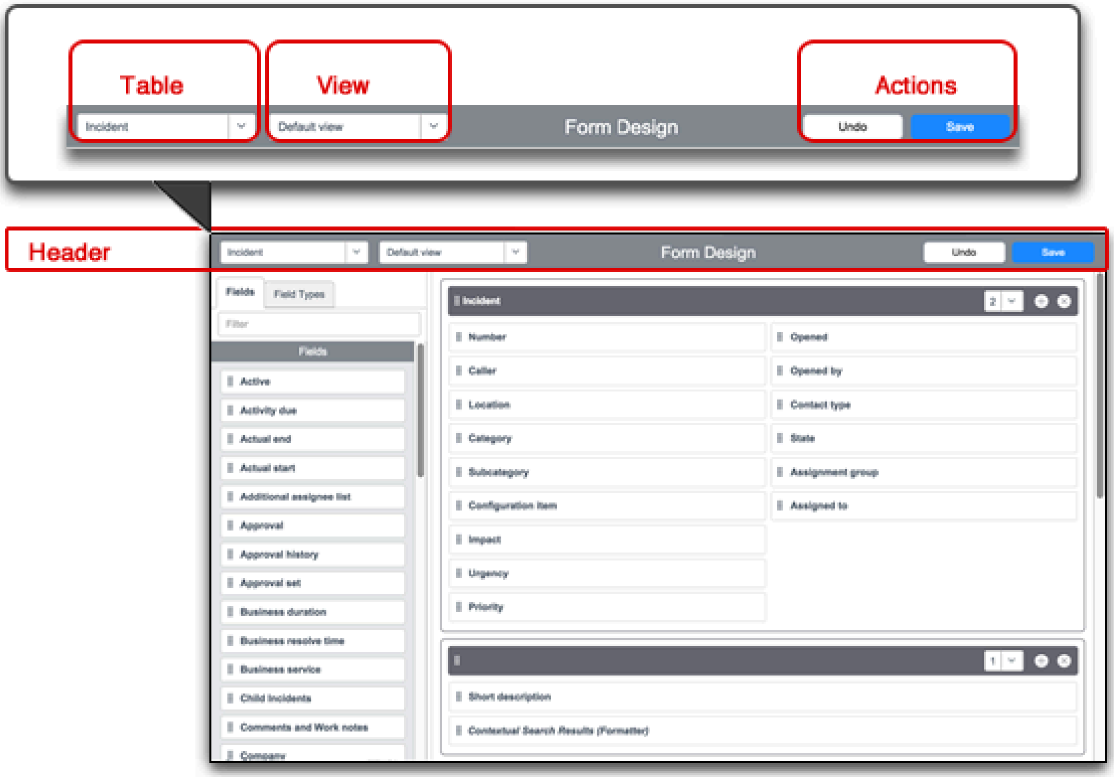

The field navigator

The field navigator allows you to manage fields on the form. It contains the following elements:

- **Filter**: Allows text search on the currently selected tab.
- **Fields**: Displays a list of existing fields you can add to the current form.
- **Field Types**: Displays a list of field types you can add to the current form. Adding a field type to the form layout creates fields in the selected table when the form layout is saved.

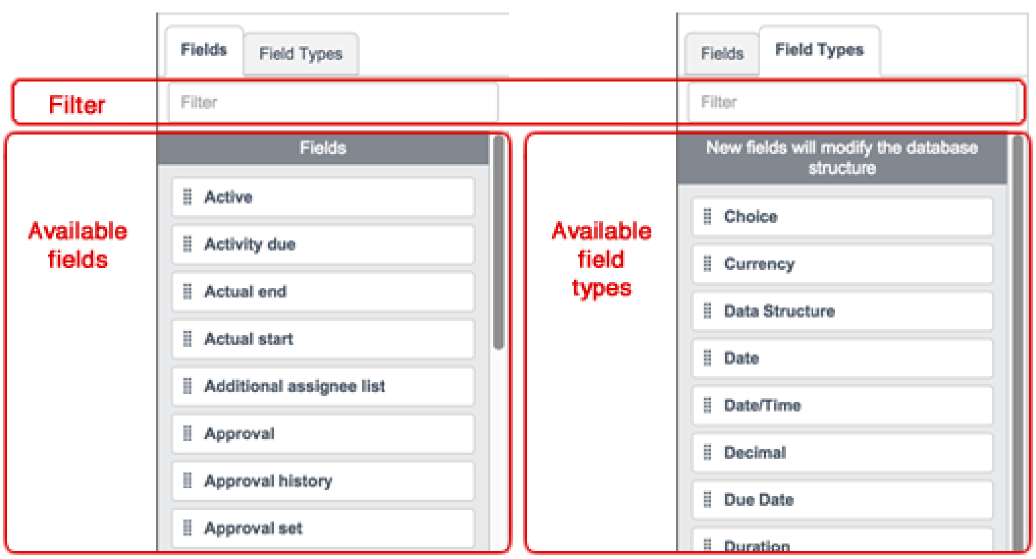

Field Navigator

Field properties are accessed by clicking the gear icon that appears when you point to a field in the form layout.

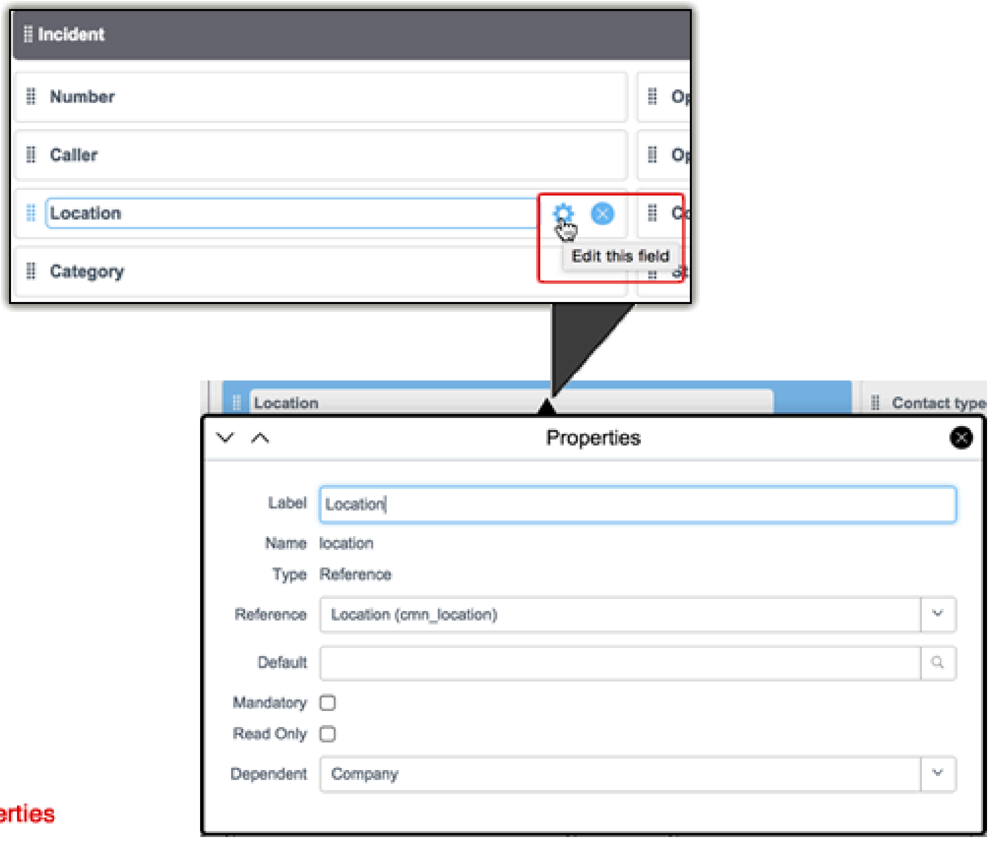

Field properties

Form layout

The form layout displays a separate cell for each field, section, annotation, UI element, and related list that is on the form. Each cell has an identifying label and a handle icon to aid in selecting and moving it. The location of the cell in the form layout represents its relative position on the form. Only fields visible on the form layout are displayed on a form.

Form layout

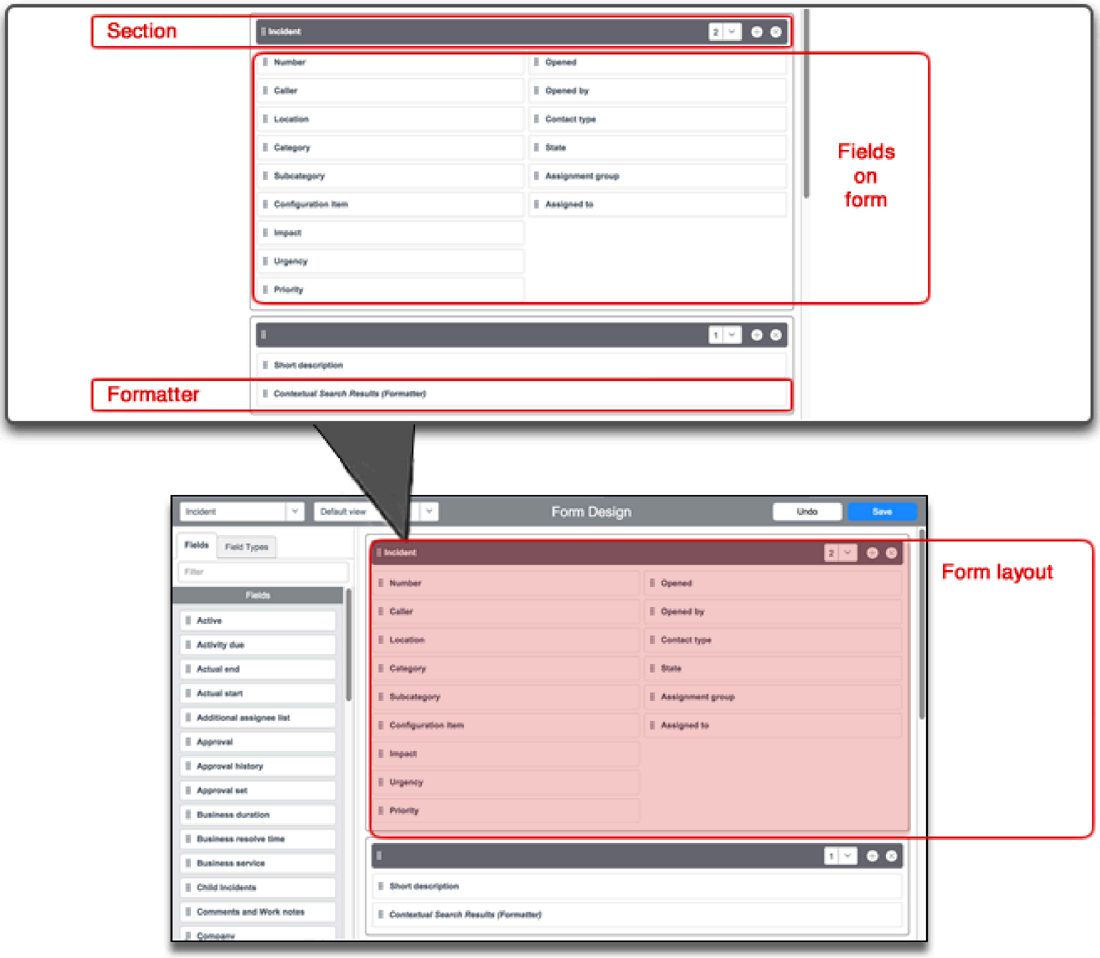

Show or hide fields with the form designer

You can move fields to and from the form workspace to display the fields on the form.

Before you begin

Role required: personalize_form

About this task

Use the filter to quickly find fields in the field navigator.

Procedure

1. Navigate to the form designer.
2. From the list of fields in the **Field Navigator**, select the field you want to display from the list of available fields.
3. Drag the field to the form layout.
4. Drag the field to the location you want it to be on the form.
5. To remove a field from the form layout, select the field and click the **Remove this field** (X) button beside the field label.
6. Click **Save**.

Add a form annotation with the form designer

Use annotations to provide instructions or additional information about fields on a form.

Before you begin

Role required: personalize_form

Procedure

1. Navigate to the form designer.
2. Click the **Field Types** tab.
3. Select the **Annotations** field type.
4. Drag the annotation to the appropriate location on the form.
5. Point to the **Annotations** field and click the gear icon to open Properties dialog box.
6. In the Properties dialog box, enter a value for **Annotation Text**.
   You can use HTML tags to format the annotation text.
7. Click **Save**.

Add a new section with the form designer

Move sections to and from the form layout to display or remove the sections on the form.

Before you begin

Role required: personalize_form

About this task

Sections can have a one- or two-column layout.
By default, the first section on a form is always a read-only section displaying the label of the table. You cannot remove this section.

Procedure

1. Navigate to the form designer.
2. Select an existing section.
3. Click the (+) button beside the section label to add a section.
4. Enter a label for the section.
5. Drag sections to reorder them.
   The label for the first section on the form designer is displayed as the form title.
6. To delete a section, click the (x) button beside the section label.
7. Click **Save**.

The following GIF shows how you would use form designer to add and then delete a form section.

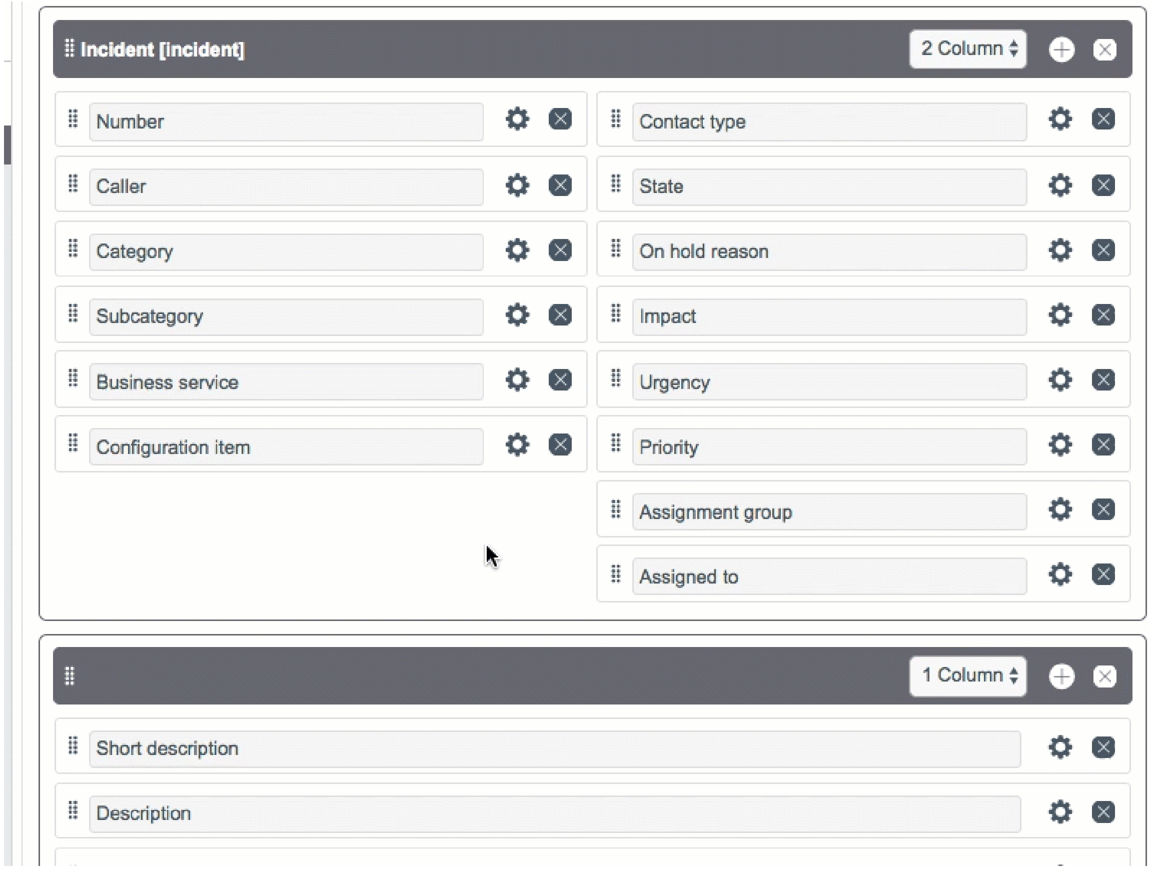

Add a formatter with the form designer

Use formatters to add information from macros or previously scripted UI elements to a form.

About this task

Formatters are used to add information from macros or previously scripted UI elements to a form.

Procedure

1. Navigate to the form designer.
2. Click the **Fields** tab.
3. In the filter, enter the string `Format`.
4. Select the formatter you want to add to the form.
5. Drag the formatter to the location you want it to be on the form.
6. Click **Save**.

Add a new field with the form designer

You can add a field to a form and the associated table.

Before you begin

Role required: personalize_form

About this task

Adding a field creates a dictionary entry for the new column in the current table. You can create a custom table and then use the form designer to add fields to the form.

Use the filter to quickly find field types in the field navigator.

Procedure

1. Navigate to the form designer.
2. Click the **Field Types** tab.
3. Select the cell for the field type to create.
4. Drag the field type to the location you want it to be on the form.
5. Enter a label for the field.
6. **Optional:** Point to a field in the form design area and click the gear icon to access properties.
7. Click **Save**.

Configuring the form layout

Administrators or users with the personalize_form role can configure the form and related list layout.

> **Important:**
> Try configuring your forms with Form Builder instead. Form Builder has all of the Form Designer and Form Layout options. It also has other robust features that help you configure each aspect of the form. For more information, see Accessing Form Builder.

Show or hide fields on a form

Configure form layout to change what appears on the form. You can perform tasks like changing what fields are visible, embedding lists, and adding sections and annotations.

Before you begin

Role required: personalize_form

Procedure

1. Navigate to a form.
2. Click the form context menu icon ( ) and select **Configure** > **Form Layout**.
3. Using the slushbucket, select the fields and the order in which you want them to appear. Available items that appear in green followed by a plus (+) sign represent related tables. To access fields on these tables, use dot-walking.

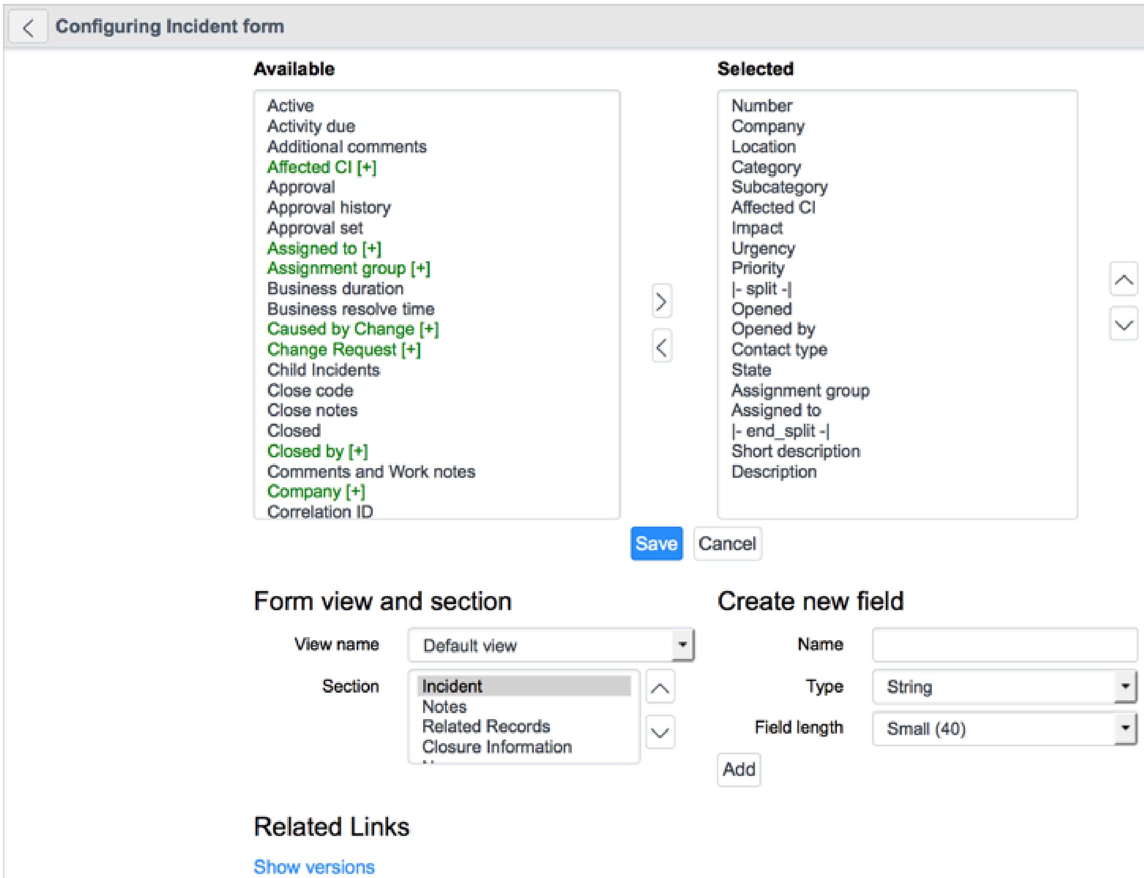

In Core UI, Form Layout is not responsible for the order of the Additional Comments and Work Notes fields. For more information on how to determine the order of these fields, see Customize activities.

4. Click **Save**.

> **Warning:**
> Do not add the same field to more than one section of a form unless the field displays read-only data. Having two or more instances of an editable field can cause data loss and prevent the proper functioning of UI and data policies.

What to do next

Through form configuration, you can also add new fields to the form, which creates columns to the table that underlies the form. See Add and customize a field in a table for more information.

Add a related list to a form

You can configure related lists to appear on forms and in hierarchical lists.

Before you begin

Role required: personalize_form

About this task

Related lists display records in another table that have a relationship with the current record.

Procedure

1. Open the form.
2. Click the form context menu icon and select **Configure** > **Related Lists**.

3. Using the slushbucket, select the related list to display on the form.
4. Click **Save**.
   Related lists appear at the bottom of the form.

Add an annotation to a form

Highlight form elements by displaying blocks of colored text or separators between form elements.

Before you begin

Role required: personalize_form

Procedure

1. Navigate to a form.
2. Click the form context menu icon ( ) and select **Configure** > **Form Layout**.
3. Under **Form view and section**, select a section.
4. In the Available list, double-click *** Annotation** to move it to the **Selected** list.

   The Annotation Details section appears.

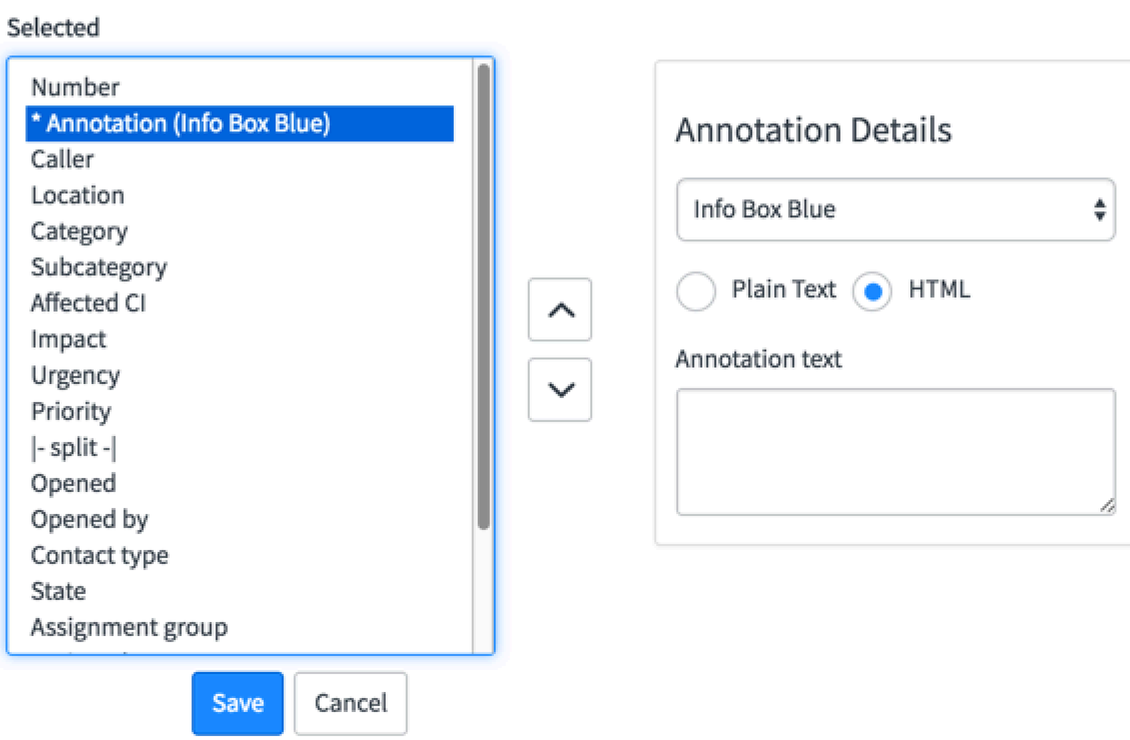

5. Place it above the field to annotate.
6. In Annotation Details, select the type of annotation.
   The available types of annotations include **Info Box Blue**, **Info Box Red**, **Line Separator** (any text you add appears beneath the line), **Section Details**, **Section Separator**, and **Text**.

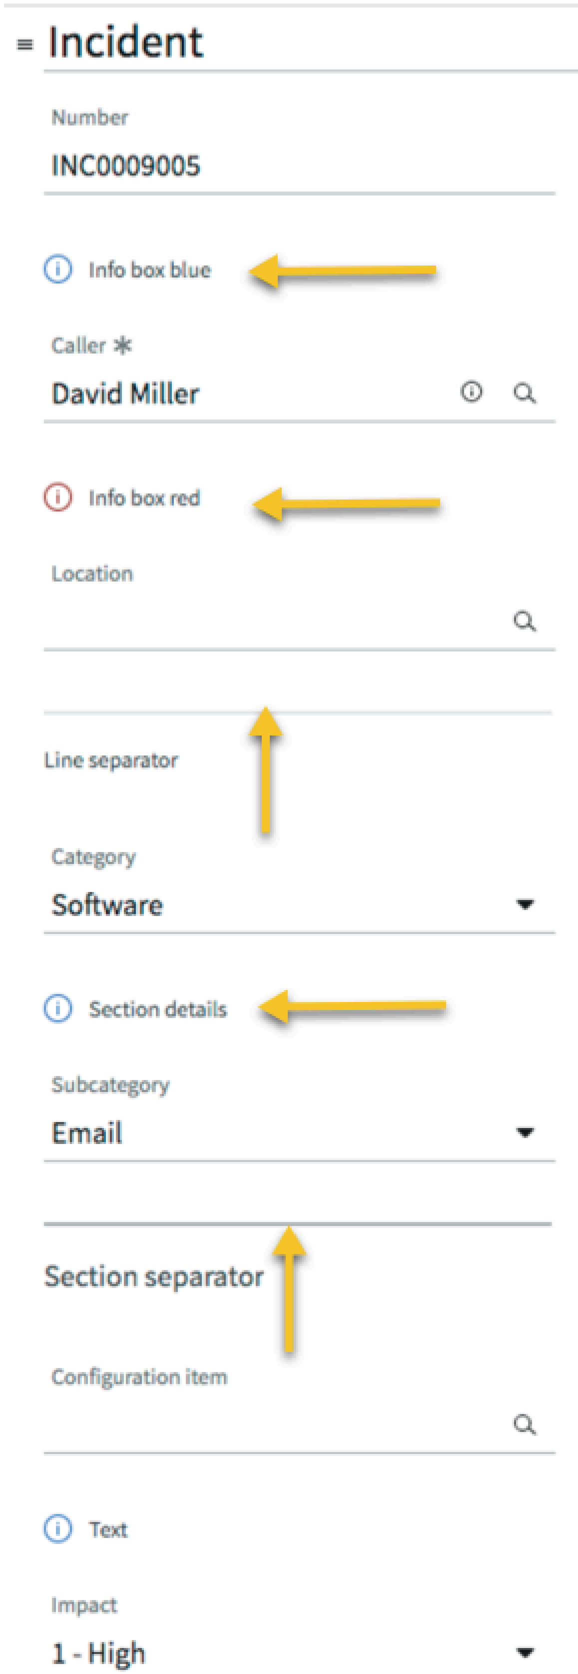

7. Select either **Plain Text** or **HTML** for the data type and enter the annotation in the text box.

   All valid HTML tags are supported.

For example, select the **Text** annotation type, the **HTML** data type, and enter Select the primary location:. The text appears red on the form.

8. Click **Save**.

Toggle annotations

Users can toggle annotations on a form in Core UI . The *glide.ui.show_annotations* user preference controls the visibility of annotations for each user.

Before you begin

Role required: none

Procedure

1. Navigate to a form.
2. Select the more options icon () to see the annotations icon.
   The icon is light gray and cannot be toggled when no annotations are available.

Create a form section

Create sections on forms to help group related fields together.

Before you begin

Role required: form_admin

About this task

Users can expand or collapse form sections to show or hide the fields they need. If you have tabs enabled, each form section appears on a separate tab. The default Change Request form is an example of a form with multiple sections as shown in the following image.

Change Request form sections

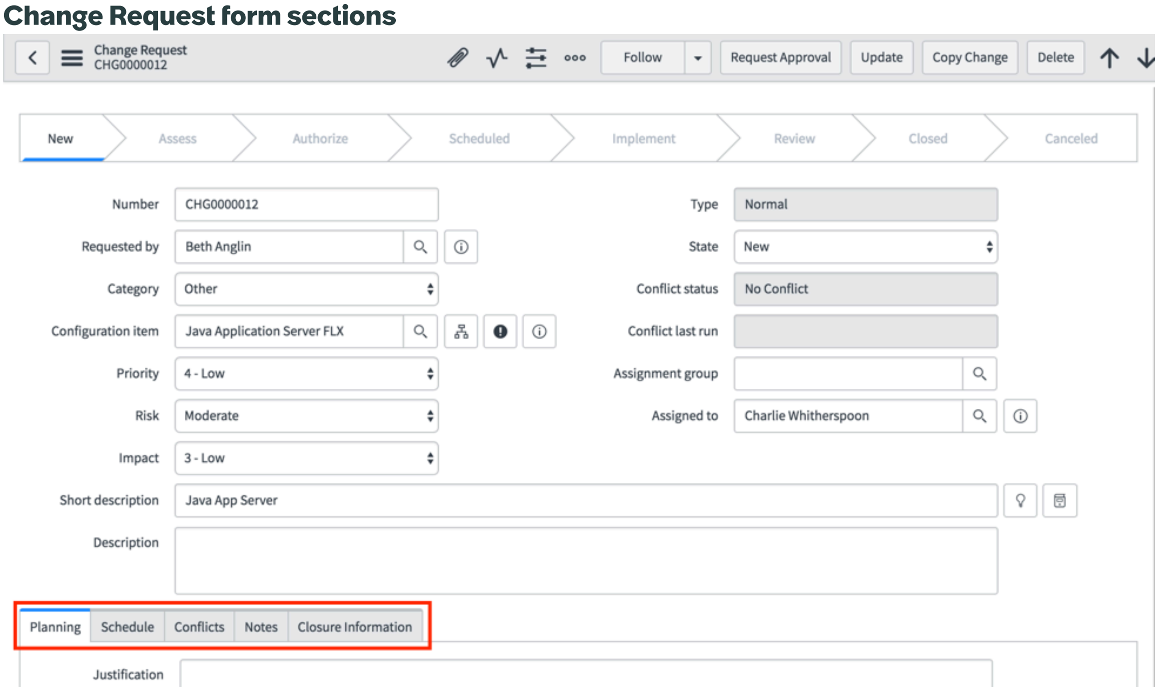

Procedure

1. Right-click the form header and select **Configure** > **Form Layout**.
2. In the **Form view and section** area below the slushbucket, click **New** in the **Section** list.
3. In the **Section caption** field, give the new section a title, then click **OK**.
4. Select a form section name and click the up or down arrow to change the section order.

   > **Note:**
   > The caption for the first section on the form becomes the form title.

5. Add fields to the new section using the slushbucket.
6. Click **Save** when you are finished.
   The new section appears on the form with the fields you selected. If you do not add any fields to a section, the section stays empty.

   If you want to delete the form section later, do the following:

   a. Navigate to **System UI** > **Form Sections**.
   b. Filter for the table that contains the section you want to delete. For example, the Incident table. The name of the form section to be deleted is displayed in the **Caption** field.
   c. Select the check box beside the form section to delete, and then select **Delete** from the **Actions on selected rows** menu.
   d. Click **OK** to confirm the deletion.

Move form splits

Change the place where fields split on a form by moving the split section indicators in a slushbucket.

Before you begin

Role required: admin

Procedure

1. Navigate to a form.
2. Click the form context menu icon ( ) and select **Configure** > **Form Layout**.
3. Move any of the following split section indicators:
   - |- begin_split -|
   - |- split -|
   - |- end_split -|
4. Click **Save**.

Display tabbed forms

Tabbed forms offer a useful way to make forms and related lists take up less space by reducing the scrolling that must be done to navigate the form.

Before you begin

Role required: personalize_form

About this task

Tabbed forms

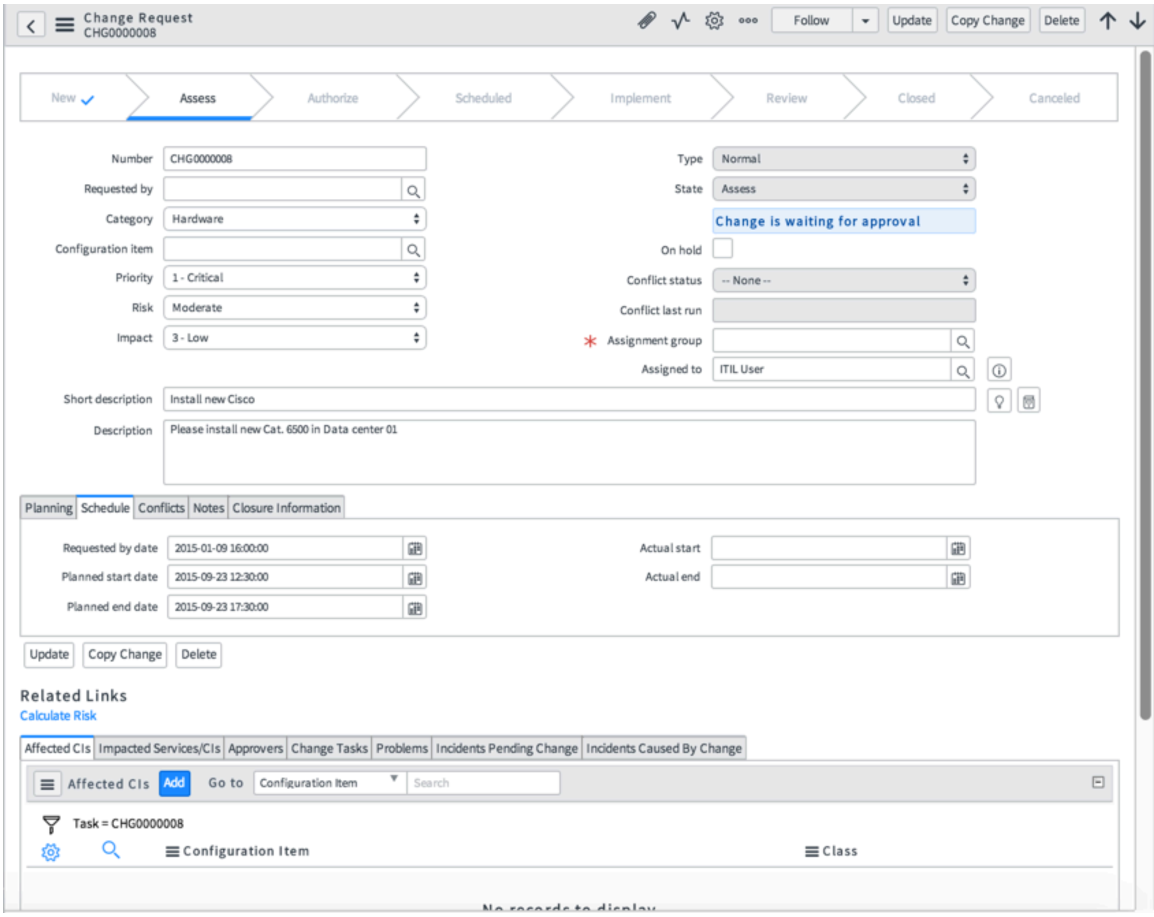

Form sections and related lists are tabbed separately, each with their own tab line. A user always sees the first form section. All sections after that can be tabbed. Tabs are enabled by default for new instances.

Tabbed forms are enabled by default for new instances. A system user preference with the name **tabbed.forms** specifies whether the tabbed UI is used by default for all users. Users can change this preference as described in this procedure.

Procedure

1. Click the gear icon in the banner frame.
2. Enable **Tabbed forms**.

Embed a list within a form

You can embed lists within a form. When a list is embedded in a form, any changes made to the contents of the list are saved when the form is saved.

Before you begin

Role required: personalize_form

About this task

This allows the embedded list to be treated just like another element on the form. In addition, users can create records in the list view. After a row is added to the list, double-click any cell in the list to edit its value. You can add the same lists to a form as embedded lists or as related lists, depending on the path selected in the form context menu.

You can also modify embedded list controls, such as the name of the related list and the roles required to create records in the embedded list.

Procedure

1. Select a record from the list.
2. Right-click the form header and select **Configure** > **Form Layout**.
   The slushbucket shows the available fields and the lists that can be embedded in the form. Lists appear in red at the bottom of the **Available** list.
3. Select a list and move it to the **Selected** column.
4. Use the up and down arrow buttons to position the list in the form.
5. Click **Save**.
   Lists adjust to fit the frames of the adjacent fields.

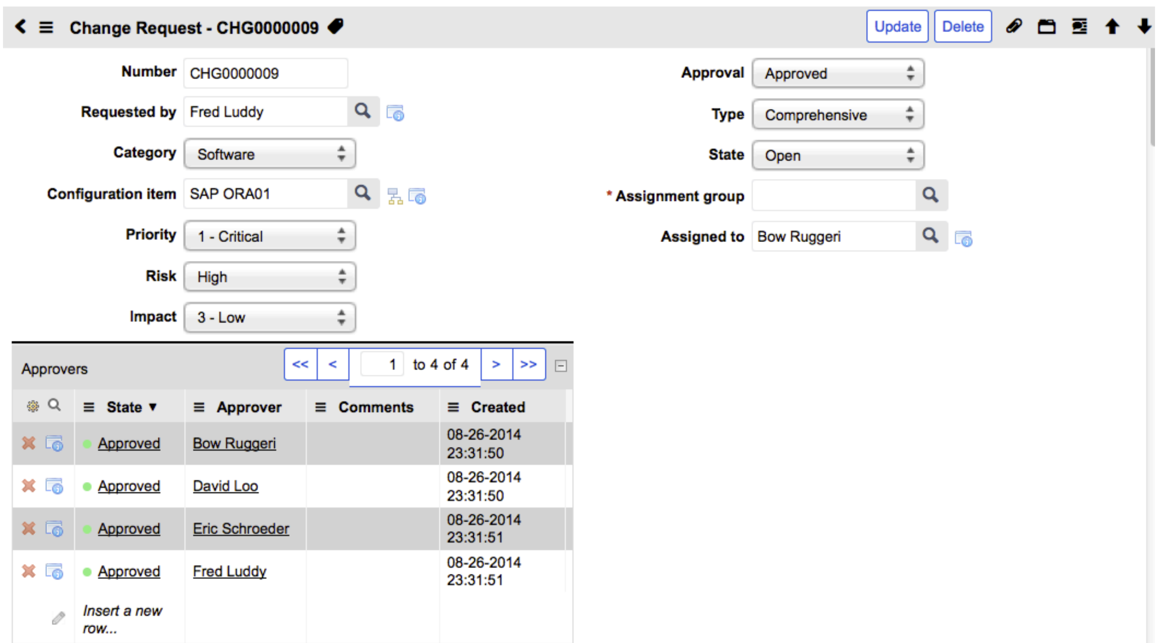

Configure customer updates indicator

You can configure the indicator for customer updates, which are changes that update sets, scripts, service catalog items, and other configuration tables track.

Before you begin

Role required: admin

About this task

The customer updates indicator icon ( ) may appear on the header of forms that have customer updates. Clicking the customer updates indicator opens the update set records for the item.

You can configure this indicator to appear for all or for specific administrators using the `owned_by_indicator.form` user preference.

Procedure

1. Navigate to **All** > **User Administration** > **User Preferences**.
2. Perform one of the following actions.

Option	Description
Enable the indicator for all administrators	Set the *owned_by_indicator.form* preference to **true**.
Enable the indicator for an individual admin istrator	**a.** Make sure the *owned_by_indicator.form* system user preference is set to **false**. **b.** Create a user preference with the following values: ▪ **Name:** `owned_by_indicator.form` ▪ **User:** Select the administrator for which to enable the preference. ▪ **Value:** `true`

Add a chart to a form

Add a chart to a form to show your users graphical data related to the form. For example, you could show an incident SLA-related chart on the Incident form so your users can see how well service level agreements are being met.

Before you begin

Role required: personalize_form

About this task

Note:
The following report types are not supported on forms: List, Pivot, Multilevel Pivot, Calendar, and Single Score.

Chart on a form

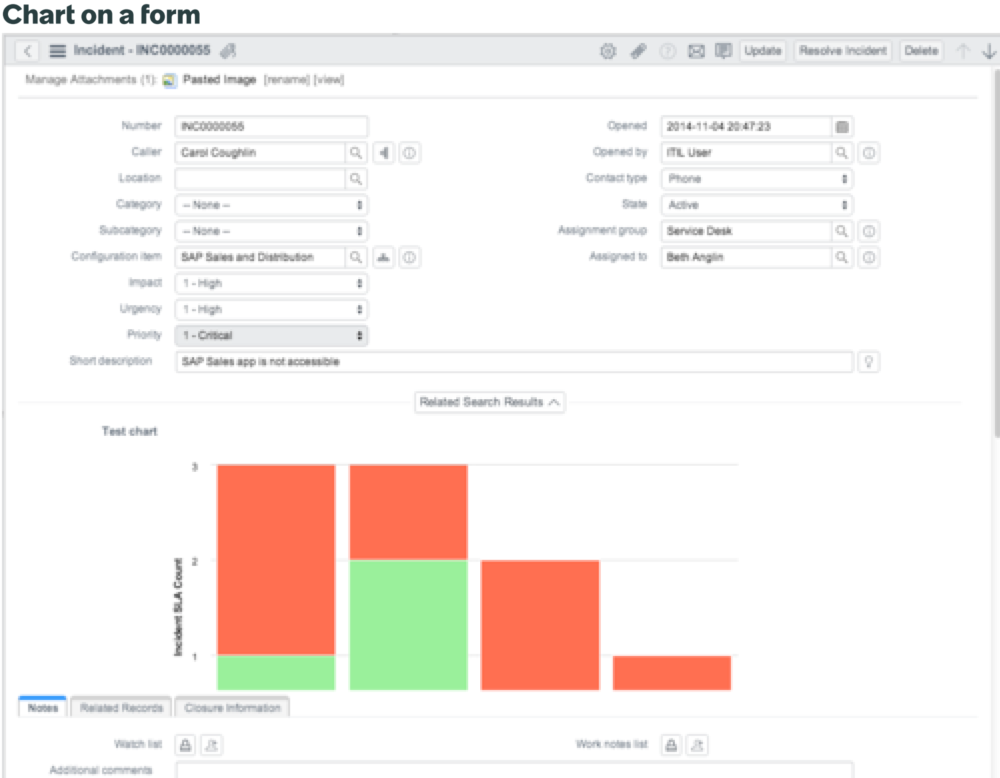

Procedure

1. Navigate to **All** > **System UI** > **Forms**.
2. Select the form you want to configure.
3. Click the form context menu icon ( ) and select **Configure** > **Form Layout**.
4. Using the slushbucket, select *** Chart**.
5. Enter a **Label** in the chart details.
6. Click **Save** to return to the form.
7. Click **Configure chart**.
8. Select a chart in the **Report** field and configure other options as desired.
9. Click **Update**.

Form administration

Administrators can configure several settings that control form functionality.

Form personalization

Control how users configure forms. You can also change the role that is necessary to configure a form and disable the form configurations that specific users already made.

See Administering form personalization for instructions.

Form annotations

Choose the types of annotations that are available on forms and save localized versions of form annotations.

See Administering form annotations for instructions.

Attachments

Control several aspects of form attachments, such as attachment size, allowed file extensions, and the roles that are allowed to attach files.

See Administering attachments for instructions.

Formatters

Create and edit formatters, such as the activity formatter, which displays information on the form that is not a field in the record. You can also configure formatter settings such as the maximum number of formatter entries.

See Formatters for more information.

Form templates

Create templates that automatically populate form fields based on settings you configure.

See Form templates for more information.

UI actions

Add buttons, links, and context menu items on forms to the UI more interactive, customizable, and specific to user activities.

See UI actions for more information.

UI policies

Configure the policies that dynamically change information on a form.

See Create a UI policy for more information.

Other advanced options

Configure advanced form features, such as form focus, form splits, and derived fields.

See Advanced form configuration for instructions.

Administering form personalization

Administrators can configure several aspects of form personalization, which allows users to customize the layout for any form view.

Administrators can manage this function using the following options.

- Activate or deactivate form personalization globally.
- Control user access to form personalization based on roles.
- Manage the personalized forms of users.

Note:
Personalizing a form in this way modifies the form for you only. To make changes to a form that are visible to all users, you must configure the form.

Activate form personalization

Form personalization is activated for new instances. To activate form personalization for upgraded instances, an administrator must activate the Form Personalization (com.glide.ui.personalize_form) plugin.

Before you begin
Role required: admin.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Form Personalization (com.glide.ui.personalize_form) plugin using the filter criteria and search bar.

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

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

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

Change form personalization role requirements

By default, the itil role is required to personalize forms, but you can change this requirement with a system property.

Before you begin

Role required: admin

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. Locate the *glide.ui.personalize_ form.role* property in the System Properties list.
3. In the **Value** field, specify the roles that can access form personalization.

Manage personalized forms

When a user personalizes a form, the system stores the customizations as a user preference record. You can view and manage the user preferences.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **User Preferences**.
2. Filter the list by **[Name] [contains] [personalize]**.
   There is a user preference for each form view each user personalizes. The name format combines the word personalize with the name of the table and the name of the view. For example, if a user personalizes the default view of the Asset [alm_asset] form, the user preference is called *personalize_alm_asset_default*.
3. Delete a user preference to remove the customizations for the user.

Disable form personalization

If you do not want your users to customize forms, you can disable form personalization.

Before you begin

Role required: admin

About this task

Activating the Personalize Forms plugin sets the *glide.ui.personalize_ form* property to true. You can disable form personalization.

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. Locate the *glide.ui.personalize_ form* property in the System Properties list.
3. Set the **Value** field to `false`.

Administering form annotations

Form annotations are additional pieces of information on a form, such as a line or paragraph of text. Use form annotations to provide on-screen instructions to your users.

Form annotations are enabled by default in the base system. To disable them, set the `glide.ui.form_annotations` system property to false.

Support multiple languages for a form annotation

You can store multiple translations of form annotation text.

Before you begin
Role required: admin

About this task
To support multiple languages, use message records to translate annotation text.

Procedure

1. Navigate to **All** > **System UI** > **Messages**.
2. Create a message record for each language you support.
3. On the Message form, set the **Key** field to a unique identifier for the annotation text.
   The annotation text is a good key. The key must be the same for each translation message for the annotation.
4. Select the appropriate **Language**.
5. In the **Message** field, enter the translated annotation text.
6. Edit the form annotation and reference the message key with a *`gs.getMessage`* call.
   For example, if the message key is **Message key text**, enter `${gs.getMessage("Message key text")}` in the form annotation.

Administer form annotation types

You can define the form annotation types to control their appearance.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System UI** > **Form Annotation Types**.
2. Set the **Active** field to **false** for any types you do not want to use.
3. Click **New** to add a type.

Administering attachments

Administrators can configure system properties for how files are uploaded and attached to records. Uploading, viewing, and deleting file attachments triggers a single event that can be used for notifications or in scripts.

Attachment tables

When you store an attachment to any table, a record is created in the Attachment [sys_attachment] table that contains attachment metadata. The file data is stored in the Attachment Document [sys_attachment_doc] table, in 4k chunks. For example, if you attach a 12k file called **My_attachment.pdf**, then there is an Attachment entry with three related Attachment Document entries.

To learn how to upload attachments to records, see Add and manage attachments.

Activity stream thumbnail scaling

The *`glide.ui.activity_stream.scale_images`* property scales large images down and creates thumbnails for the activity stream view of the images. This property is active by default. It applies to images retroactively, so any attachments included in an activity stream in the past also see the effects.

- Maximum dimensions for thumbnails using this property are 525 pixels width by 350 pixels height. The system favors the pixel height first, so you can end up with images that are wider than 525 pixels.
- Attaching image files over 5 MB can create an out of memory error and cause the instance to restart. An out of memory error and restart occur when the system generates the thumbnail when you attach a new image. The error and restart can also occur the first time you load existing images on a record.
- The administrator can add the *com.glide.attachment.max_get_size* system property to control the maximum image scaling dimensions.

Enter the following information into the [sys_properties.list] table:

Field	Description
Name	com.glide.attachment.max_get_size
Type	integer
Value	5242880

The value field represents the size in bytes.

- The base system value is 5242880 bytes (5 MB). You can change the value as necessary.
- With this property set, images that are larger than 5 MB are not scaled and the system creates a link instead. Images smaller than 5 MB are scaled down.

Configure the max image size property for Activity Streams

When attaching an image in an Activity Stream, the system controls the maximum size of the image. Use the *com.glide.attachment.max_get_size* property to control the maximum image scaling dimensions.

Attachment events and logging

A script action or notification can process attachment events. Only one event is created when action is taken on a record with attachments, even when the record has multiple attachments. The following events are provided.

Only one event is created when action is taken on a record with attachments, even when the record has multiple attachments. The following events are provided.

Attachment events

Event	Description
attachment.read	An attachment has been read or downloaded.
attachment.uploaded	An attachment has been uploaded. If multiple attachments are uploaded to a record at one time, only one event is created.
attachment.deleted	An attachment has been deleted. This event is also triggered when the record containing the attachment is deleted. If a record is deleted that contains multiple attachments, a separate event is triggered for each attachment in the deleted record.
attachment.renamed	An attachment has been renamed.

Attachment logging

When an attachment is downloaded, the `attachment.read` event record is written, and you can do something with this event. For example, you can record when and by whom certain attachments are downloaded. For this functionality, **current** is a sys_attachment record, and the event record uses the following parameters:

- **parm1:** File name
- **parm2:** Table name

Debug attachment indexing

To change debugging options for attachment indexing, add these system properties.

- `glide.ts.index.attachment.debug`: When the value is set to **true**, enables log messages for exceptions that occur when indexing attachments (default is **false**).

  You can leave this property enabled during normal operations to capture stack trace information about any exceptions.

- `glide.ts.index.attachment.list_terms.debug`: When the value is set to **true**, logs all indexed terms when an attachment is indexed (default is **false**).

  [Recommended] For optimal performance, set this property to **false** during normal operations. Only enable this property when you are actively debugging an issue.

> **Note:**
> To learn more about the properties that affect attachments processing, see Attachments (instance security hardening) in Instance Security Hardening Settings.

Configure attachment system properties

You can disable the drag-and-drop feature. (Users can still upload attachments by browsing to the file.) You can also limit the attachment file size, restrict who can upload attachments, and restrict what file extensions can be uploaded.

Before you begin

Role required: admin

> **Note:**
> To learn more about the properties that affect attachments processing, see Attachments (instance security hardening) in Instance Security Hardening Settings.

Procedure

1. To disable the drag-and-drop features, navigate to **System Properties** > **UI Properties**.
2. Clear the check box for the **Allow attachment drag and drop in supported HTML5 browsers** property.
3. Click **Save**.

Limit attachment file size

You can specify the maximum size allowable for attachments to avoid issues with the user's active session on the instance. The maximum attachment size for email attachments is configured separately.

Before you begin

Role required: Admin.

About this task

Note:
The maximum attachment size for email attachments is configured separately.

Procedure

1. Navigate to **All** > **System Properties** > **Security**.
2. Enter a value in the **Maximum file attachment size in megabytes** property.
   The system only allows attachment sizes up to 1 GB. By default, this field is blank. If you leave this field blank, the system uses the default maximum limit of 1 GB.
3. Click **Save**.

Require a role to attach files

You can restrict who can upload attachments.

Before you begin

Role required: Admin.

Procedure

1. Navigate to **All** > **System Properties** > **Security**.
2. In the **Attachment limits and behavior** section, locate the **List of roles (comma-separated) that can create attachments** property (*`glide.attachment.role`*).
3. Enter one or more roles separated by commas.
   Only roles listed in this property are able to upload attachments to a record. If no roles are entered, then all roles can upload attachments.
4. Click **Save**.

Restrict attachment file extensions

The *`glide.attachment.extensions`* property restricts the file extensions that users can upload as attachments.

Before you begin

Role required: Admin.

About this task

Restrict attachment file extensions to prevent users from uploading invalid file types and file types that may be more likely than others to contain viruses or malware.

Note:
This property does not restrict attachments based on the actual file type, only based on the extension. To enable MIME type validation, which validates that the MIME type of a file matches the file extension and can block attachments that do not pass this validation, see Upload MIME type restriction (instance security hardening).

Procedure

1. Navigate to **All** > **System Properties** > **Security**.
2. In the **Attachment limits and behavior** section, locate the **List of file extensions (comma-separated) that can be attached to documents via the attachment dialog** property.

3. Enter the file extensions and click **Save**.
If no extensions are specified, then all extensions are allowed. However, if any extensions are specified, all unlisted extensions are restricted. Listed extensions should not include the dot (.) prefix or spaces after commas. For example, enter `xls,xlsx,doc,docx`.

Disable attachments on a table

You can prevent users from adding attachments to records on a specific table.

Before you begin
Role required: admin

Procedure

1. Open a record in the table.
2. Right-click the form header and select **Configure** > **Dictionary**.
3. In the list of dictionary entries, select the first record in the list (the record with no **Column name** entry).
4. Add `no_attachment` to the **Attributes** field, separated by commas from any existing attributes.
See Altering tables and fields using dictionary attributes for more information.

Index attachments on a table

You can enable attachment indexing for a table so text searches can return matches from the record and its file attachments.

Before you begin
Role required: admin

About this task
By default, attachment indexing is enabled for the Knowledge Base. You can enable attachment indexing for other tables. Enabling attachment indexing causes the platform to reindex the selected table, its parent table, and any children of the parent table. For large tables, such as the Task table, reindexing can take several hours and slows down the system until complete. Reindexing is best performed during non-peak times.

Procedure

1. Navigate to **All** > **System Definition** > **Text Index Configurations**.
2. Edit the record for the table that you want to enable attachment indexing for.
3. In the Text Index Table Attribute Maps related list, select **New**.
4. On the Text Index Table Attribute Map form, fill in the fields.

Field	Value
Table	<table name>
Attribute name	Attachment index
Value	true

> **Note:**
> The attachment index attribute only applies to the tables on which you explicitly add it. It does not cascade to child tables. For example, enabling indexing of attachments on the Task table does not enable indexing of attachments on the Incident table. Not all file types are supported for attachment indexing. For a list of supported file types, see Zing can include attachments in search results.

5. Select **Submit**.
   A message appears indicating that the attachment index attribute change will not take effect until you reindex the selected table.
6. Dismiss the reindexing message by selecting **OK**.
7. On the Text Index Configuration form, select the **Generate Text Index** related link.
8. In the confirmation window, perform one of these two steps:
   - To receive an email notification from the system when text indexing for the table is complete, enter your email address and select **OK**.
   - If you do not want to receive an email notification, select **Do not notify me**, then select **OK**.

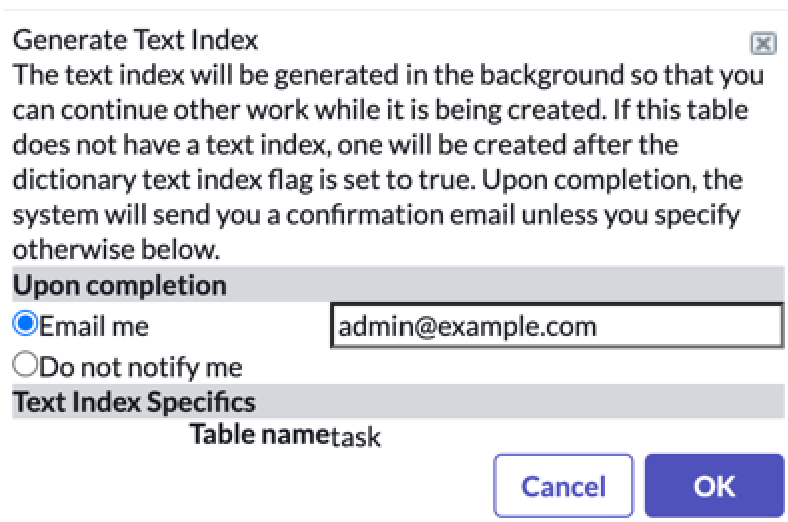

9. Dismiss the confirmation message by selecting **OK**.

Result

The system begins text indexing for the table. When it is complete, attachments can be searched on that table.

Hide the attachment [view] link

Users can open an attachment by clicking either the file name or the **[view]** link. The **[view]** link opens the file from within the browser, which executes JavaScript code as part of the attachment. You can hide the **[view]** link. Users can still view attachments by clicking the file name.

Before you begin

Role required: Admin.

Procedure

1. Add the `glide.ui.disable_attachment_view` and `glide.ui.attachment_popup` properties.
   For instructions on adding properties to the platform, see Add a system property.

2. For the *glide.ui.disable_attachment_view* property, set the **Type** to **true/false** and set the **Value** to **true**.
3. For the *glide.ui.attachment_popup* property, set the **Type** to **true/false** and set the **Value** to **false**.
4. To return to the default behavior (enable the link), set the *glide.ui.attachment_popup* property **Value** to **true**.

Configure attachment icons

Configure the icon that appears beside an attachment of a particular file type.

Before you begin

Role required: admin

Procedure

1. Determine the path of the image file or upload a new image.
2. Navigate to **System UI** > **Attachment Icon Rules**.
3. Open an existing rule or click **New** to create a new rule.
4. Enter the rule details.
5. Click **Submit** or **Update**.

Formatters

A formatter is a form element used to display information that is not a field in the record. Add formatters to a form by configuring the form.

Examples of formatters in the base platform

Formatter	Description
Activity formatter	Displays the list of activities, or history, on a task form. For an example of an activity formatter, see Activity formatter.
Process flow formatter	Displays the different stages in a linear process flow across the top of a record. For an example of a process flow formatter, see Process flow formatter.
Parent breadcrumbs formatter	Provides breadcrumbs to show the parent or parents of the current task. For an example of a parent breadcrumbs formatter, see Parent breadcrumbs formatter.
Approval summarizer formatter	Displays dynamic summary information about the request being approved. For an example of an approval summarizer formatter, see Approval summarizer formatter.

Examples of formatters in the base platform (continued)

Formatter	Description
CI relations formatter	Displays on the CI form a toolbar for viewing the relationships between the current CI and related CIs. For an example of CI relations formatter icons, see CI relations formatter.

Note:
Formatter elements cannot be exported to PDF. When exporting PDF data from a form, any formatter elements added to the form are not displayed in the PDF output.

To create a custom formatter, create a UI macro to define content for the formatter, then create a formatter that refers to the UI macro. You can then add the formatter to a form.

Create a UI macro for a formatter

Create a UI macro to define the content that the formatter displays.

Before you begin

This functionality requires knowledge of Jelly script.
Role required: ui_macro_admin

Procedure

1. Navigate to **All** > **System UI** > **UI Macros**.
2. Click **New**.
3. Complete the form.
4. In the **XML** script field, enter Jelly script to define the content to be displayed by the formatter.

   Note:
   The UI Macro for the formatter represents a row in the UI. It must begin and end with `<TR></TR>` tags.

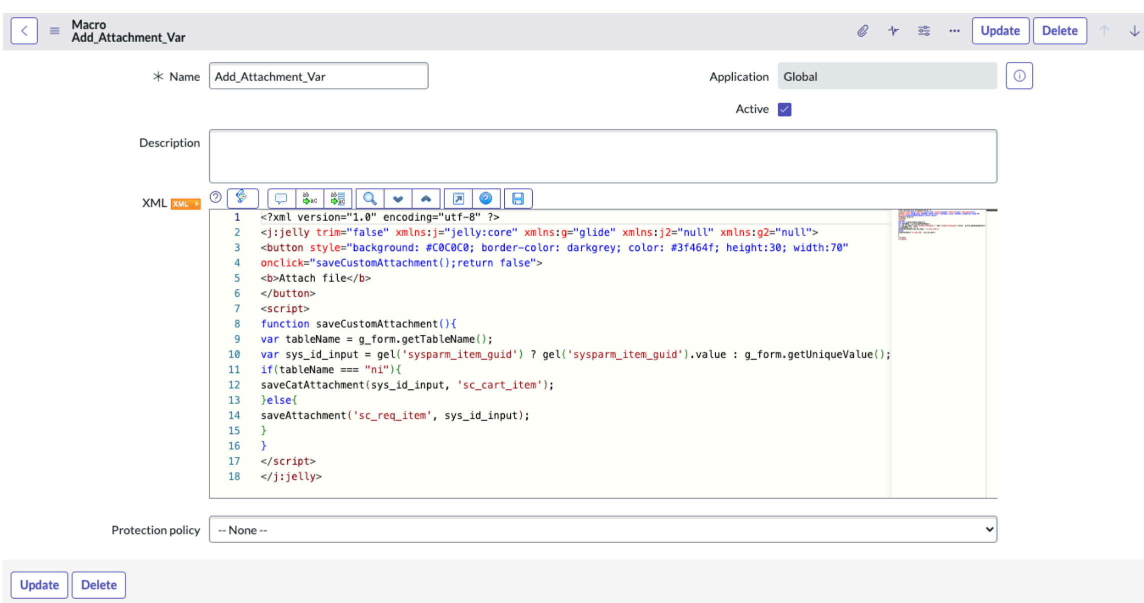

This Jelly script is reproduced below in plain text that you can copy into the Macro form as a basis for your macro script.

```
<?xml version="1.0" encoding="utf-8" ?>
<j:jelly trim="false" xmlns:j="jelly:core" xmlns:g="glide"
 xmlns:j2="null" xmlns:g2="null">
<button style="background: #C0C0C0; border-color: darkgrey;
 color: #3f464f; height:30; width:70"
onclick="saveCustomAttachment();return false">
<b>Attach file</b>
</button>
<script>
function saveCustomAttachment(){
var tableName = g_form.getTableName();
var sys_id_input = gel('sysparm_item_guid') ?
 gel('sysparm_item_guid').value : g_form.getUniqueValue();
if(tableName === "ni"){
saveCatAttachment(sys_id_input, 'sc_cart_item');
}else{
saveAttachment('sc_req_item', sys_id_input);
}
}
</script>
</j:jelly>
```

5. Click **Submit**.

Create a formatter and add it to a form

Create the formatter after creating the UI macro that defines the content.

Before you begin

Role required: admin

About this task

With the UI macro in place, create the formatter that references it, and add it to the form.

Procedure

1. Navigate to **All** > **System UI** > **Formatters**.
2. Click **New**.
3. In the **Name** field, enter a descriptive name for the formatter.
4. In the **Formatter** field, enter the name of the UI macro with `.xml` appended to it.
5. In the **Table** field, select the table in which the formatter will be used.
6. In the **Type** field, select **Formatter**.
7. Click **Submit**.
8. To add the formatter to a form, configure the form design and add the formatter.
   The formatter appears on the form according to your selections.

Activity formatter

The activity formatter provides an easy way to track items not saved with a field in the record, for example, journal fields like comments and work notes.

The activity formatter is enabled by default on the Task [task] table and other tables that extend the Task table, such as the Incident [incident] table. It is also enabled on the Approvals [sysapproval_approver] table.

You can filter the content that appears on the activity formatter and participate in the record feeds on the record. You can create an activity formatter for any form and configure properties that control what fields appear in the formatter. The number next to activities indicates the number of entries in the record feed. The number updates when users use the filter.

Activity type and a time stamp display in the top of each entry.

In Core UI, the activity formatter shows updates in real time so you can see the latest information without refreshing the form. User presence enables you to see when other users are entering comments.

Core UI activity formatter

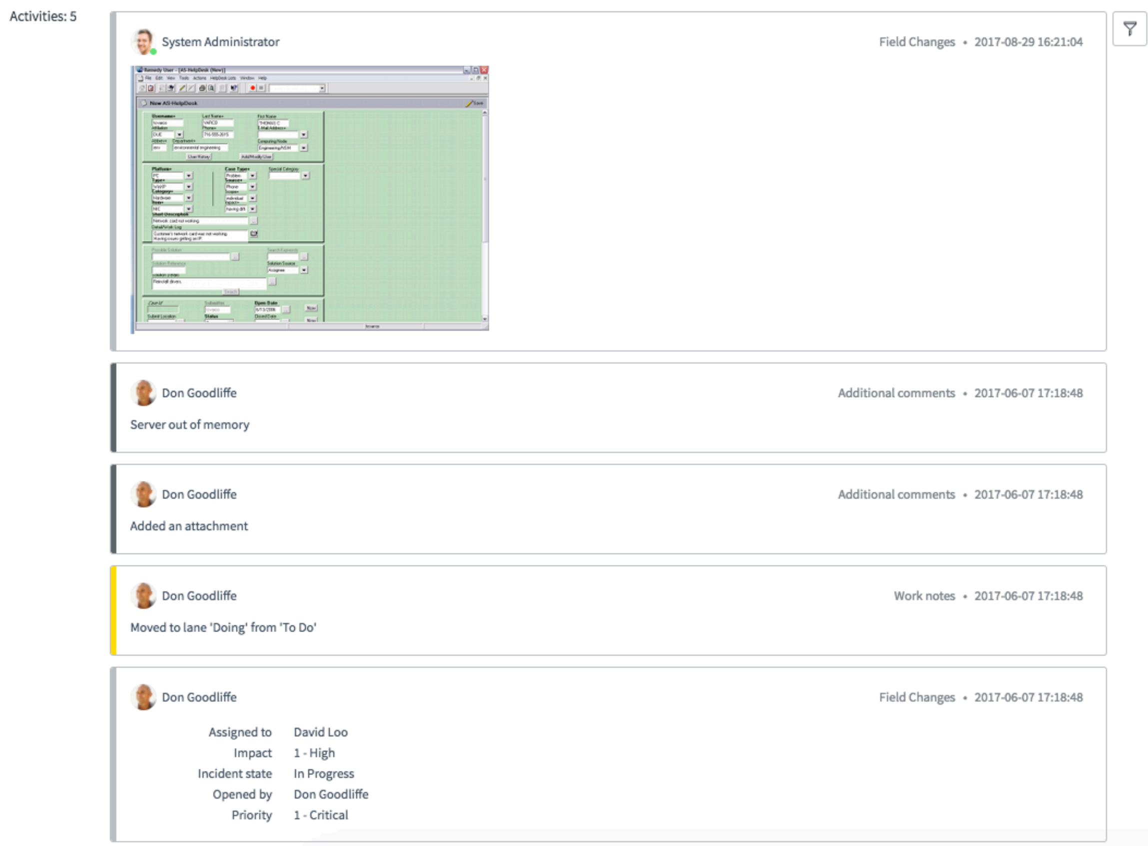

Create an activity formatter

You can create an activity formatter for any audited table.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System UI** > **Formatters**.
2. Click **New**.
3. Enter a name for the formatter, such as `Activities (task)`.
4. Select a **Table**.
5. Enter `activity.xml` in the **Formatter** field.

6. Leave the **Type** as **Formatter**.
7. Click **Submit**.

Result

Note:
You can create more than one activity formatter for a table, however, the system does not allow you to add more than one activity formatter to a form.

What to do next

Add the new activity formatter to forms as needed.

Add the activity formatter to a form

Add an activity formatter to any form to track journal fields on the form.

Before you begin

Role required: personalize_form

Procedure

1. Verify that the table associated with the form is audited.
2. Configure the form layout to add **Activities (filtered)**.

   Note:
   In Core UI, you cannot place another field in between a journal field and the activities formatter. The Activity Stream is built to keep journal fields stacked on top of the activity formatter.

   Note:
   Field styles are not applied to comments and work notes fields used with the activity formatter. Styles for these fields can be set using the `glide.ui.activity_stream.style.comments` and `glide.ui.activity_stream.style.work_notes` system properties located on the sys_properties table.

Enable the Live Feed-Activity toggle

Enable the Live Feed-Activity toggle, which allows users to switch between the activity feed and the document feed for a record.

Before you begin

Role required: admin.

About this task

Use live feed to interact with other users on a record by posting messages and adding attachments to the feed. Use the activity formatter to see an overall summary of activity for the record.

Live Feed/Activity toggle

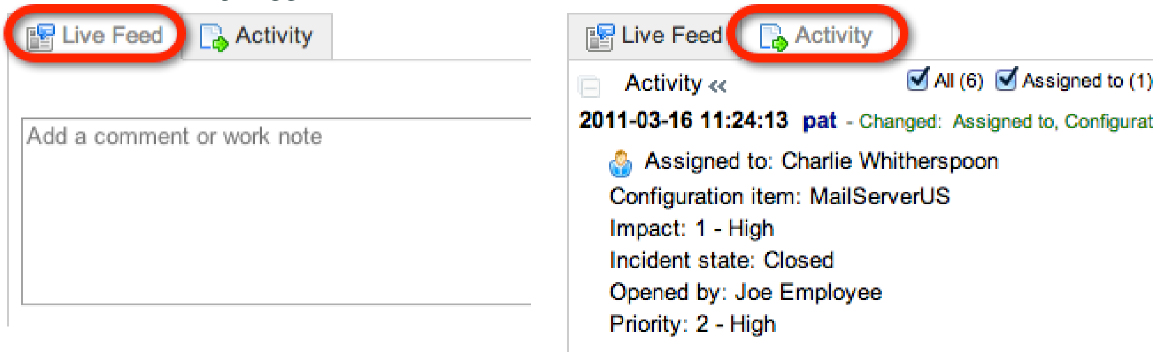

The *glide.ui.show_live_feed_activity property* is enabled. Navigate to **Collaborate** > **Feed Administration** > **Properties** and enable the **Toggle the display of the live feed tab in the activity formatter** option.

Procedure

1. Verify that the live feed and record feed plugins are active.
2. If the activity formatter is not visible, configure the form to add it.
3. Set the *live_ feed* dictionary attribute to true on the form.
   This action adds live feed to the activity formatter.
4. Enable the system property.
   a. Navigate to **Collaborate** > **Feed Administration** > **Properties**.
   b. Enable the **Toggle the display of the live feed tab in the activity formatter** option (*glide.ui.show_live_feed_activity property*).

Customize activities

You can customize which fields appear in the activity formatter. You can add or remove fields from the list of activities that users can select when they open the activity filter.

Before you begin

Role required: personalize_form and personalize_list

About this task

Customize the fields that appear in the activity filter

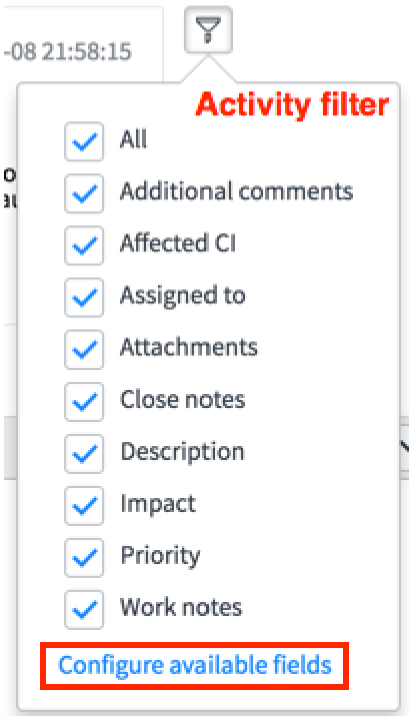

Procedure

1. Select the activity filter icon ( ).
2. At the bottom of the list, select **Configure available fields**.
3. In the slushbucket, select the desired fields in the **Available** column and move them to the **Selected** column.
   The activities appear in alphabetical order, regardless of the order in the **Selected** column.

   > **Note:**
   > In Core UI, **Configure available fields** is responsible for the order of the Additional Comments and Work Notes fields.

4. Select **Save** to add them to the activity formatter and the filter.

What to do next

Administrators can also modify the system property **Incident activity formatter fields**(*`glide.ui.incident_activity.fields`*). Access this property through **System Properties** > **UI Properties**. The system automatically synchronizes the values in the system property and the selections you made.

Configure roles for viewing emails in the activity formatter

The system property *`glide.ui.activity.email_roles`* enables you to control which roles can see emails in the activity formatter.

Before you begin

Role required: admin

About this task

In the activity formatter, users see activity only for fields they have permission to read. For example, self-service users might see the activity formatter on the self-service view of the Incident form, but they do not see work notes, unless security rules have been customized to allow this.

If the **Emails-autogenerated** or **Emails-correspondence**field is included in the activities list, all users see all emails. No determination is made whether an end user, for example, should see an email containing work notes. Configure this property to restrict this capability to specified roles.

> **Note:**
> Email does not appear as an activity until it is sent. If email properties are not configured for outbound delivery, the message can be found by navigating to **System Mailboxes** > **Outbox**.

Procedure

1. Navigate to **System Properties** > **UI Properties**.
2. Locate the property labeled **List of roles (comma separated) that can view emails in the Activity Formatter when "Emails-autogenerated" or "Emails-correspondence" are included.**
3. Add roles to the property, separated by commas.
   These are the only roles that can see email in the activity formatter. All other roles are prevented from seeing email. If no roles are listed, all users can see email. The itil role is on the list by default.
4. Click **Save**.

Configure the max activity size property

Use the glide.max_activity_size property to increase the amount of content that can be visible in the Activity Formatter.

Before you begin

Role required: admin

About this task

The default size for an activity is 100 * 1024, which causes large amounts of content to be hidden in the Activity Formatter. You can change the default value by adding the glide.max_activity_size property.

Procedure

1. Type `sys_properties.list` in the Application Navigator.
2. Select **New** and enter the following values.

Field	Description
Name	glide.max_activity_size
Type	integer
Value	102400 Use this field to increase the number of content allowed in the Activity Formatter. This value should be in bytes, for example 100 * 1024, or 102400.

Configure email reply for the activity formatter in Core UI

The *glide.ui16.emailStreamResponseActions* system property determines whether a user can reply to email using the email reply button in the activity stream in Core UI.

Before you begin

Role required: admin

Procedure

1. In the navigation filter, type `sys_properties.list`.
2. Add the following system property.
   For more information on adding a system property, see Add a system property.
3. Select **Submit**.

Configure activity formatter colors

Use the glide.ui.activity_stream.style.comments and glide.ui.activity_stream.style.work_notes properties to change the color-coding in the Activity Formatter.

Before you begin

Role required: admin

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Add a system property for the **Additional comments** field.
   a. Select **New**.
   b. Enter the following values.

Field	Description
Name	`glide.ui.activity_stream.style.comments`
Type	string

   c. Select **Submit**.
3. Change the color of the **Additional comments** field.
   a. In the System Properties list, select the glide.ui.activity_stream.style.comments record.
   b. Enter the color that you want to use.
      For example, enter `LightGreen`. The default value is transparent.
   c. Select **Update**.
4. Add a system property for the **Work notes** field.

a. Select **New**.

b. Enter the following values.

Field	Description
Name	`glide.ui.activity_stream.style.work_notes`
Type	string

5. Change the color of the **Work notes** field.

   a. In the System Properties list, select the glide.ui.activity_stream.style.work_notes record.

   b. Enter the color that you want to use.
   For example, enter `Blue`. The default value is gold.

   c. Select **Update**.

Process flow formatter

The process flow formatter provides a graphical summary of the stages in a process. The formatter is typically shown at the top of forms that are part of a process.

Each record on the Flow Formatter [sys_process_flow] table represents a process stage and can have a different condition applied to it. When specified conditions are fulfilled, the formatter highlights the current stage and places a check mark next to all previous stages.

These examples show a workflow in Core UI interface.

Core UI process flow for a change request

When any formatter stages are defined for a table, they appear on the form associated with that table in the order specified, assuming the formatter has been added to the form.

Activate the process flow formatter

You can activate the Process Flow Formatter plugin.

Before you begin

Role required: admin

Procedure

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

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

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

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

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

Create a process flow formatter

You can create a process flow formatter stage.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System UI** > **Process Flow**.
2. Click **New**.
3. Complete the form, as appropriate.

Field	Description
Table	Select a table for this process flow formatter.
Name	Enter a name to identify the formatter.
Label	Enter the name to be displayed in the form configuration slushbucket.
Active	Select the check box to ensure that the formatter stage is active. When the check box is cleared, the formatter stage does not appear in the flow display.
Order	Enter a number to indicate where in the process flow the formatter will be displayed. Formatters are arranged with the lowest number on the left and the highest number on the right.
Condition	Use the condition builder to set the conditions under which the formatter is highlighted as current. Any field available in the condition builder, such as SLA or Impact, can be used to trigger a process flow stage.
Description	Describe the process flow formatter. This description does not appear on the actual formatter.

4. Repeat as necessary for each stage.

Parent breadcrumbs formatter

The parent breadcrumbs formatter on the Task table provides breadcrumbs that show the parent or parents of the current task. This formatter can be used also on any table that extends Task.

Before you begin

Role required: personalize_form

About this task

To add the parent breadcrumbs formatter to a form, configure the form and add **Parent Breadcrumbs** to the desired location. The breadcrumbs show only six levels of parents. If more levels exist, the breadcrumbs display an ellipsis ("...").

The **Parent** reference field also has to contain a value for the breadcrumbs to appear. You may need to configure the form to contain the **Parent** field as well.

Breadcrumbs

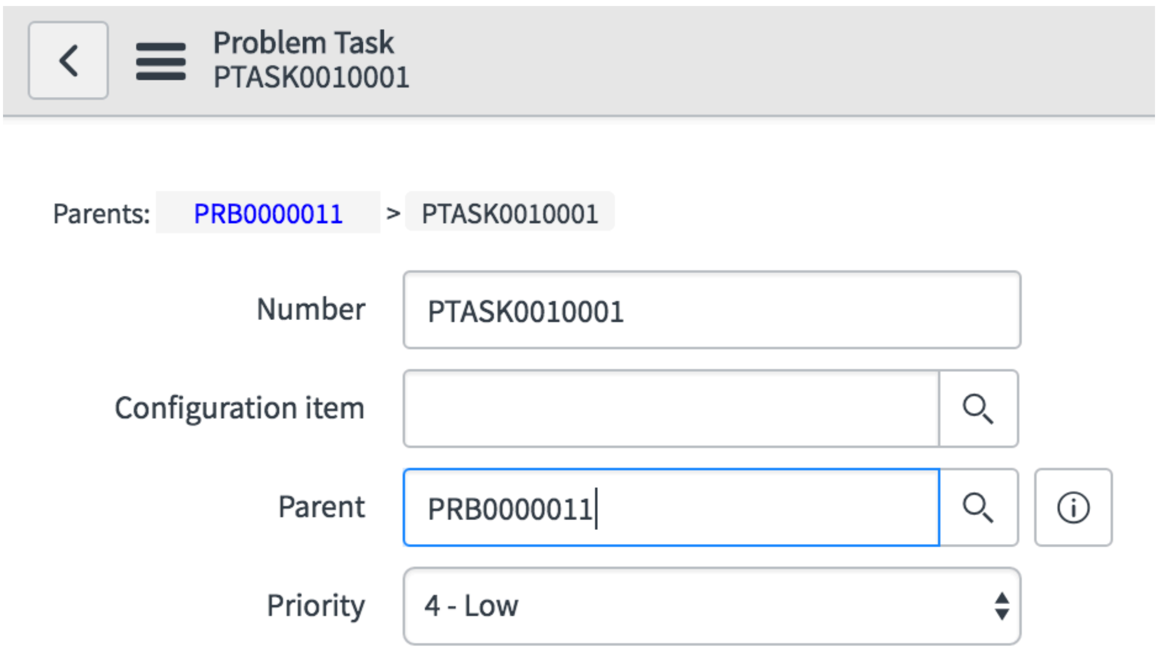

Customize the parent breadcrumbs formatter

You can customize the parent breadcrumbs formatter to control what breadcrumbs appear.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System UI** > **Formatters**.
2. Select **Parent Breadcrumbs**.
3. Click **View UI Macro for this Formatter** to view or modify the underlying formatter.
   By default, the breadcrumb uses the default display field, **gr.getDisplayValue(),** as the link in the breadcrumb. To customize this, add the following line, replacing the **fieldName** parameter with the desired field name (not the field label):

   ```
   pc.setLabelField("fieldName")
   ```

   If a user points to a breadcrumb, the short description for that record appears as a hint by default. To display alternate hints for the breadcrumb, add the following line, replacing the **fieldName** parameter with the desired field name (not the field label):

   ```
   pc.setTitleField("fieldName")
   ```

   Be sure to add these lines in the proper location, as shown in the following example:

```
//parent crumb functions - script include
var pc = new ParentCrumbs(now_GR);
pc.setLabelField("short_description");

//override the default display field to be used for label
pc.setTitleField("number");

//override default short_description hover text
var crumbs = pc.getCrumbs();
```

Use the parent breadcrumbs formatter on non-Task tables

The parent breadcrumbs formatter can be used on non-Task tables as long as the table has a reference to itself through a field called **parent**.

Before you begin

Role required: admin

About this task

To make the formatter available for a different table, duplicate the formatter used by the Task table:

Procedure

1. Navigate to **All** > **System UI** > **Formatters**.
2. Select **Parent Breadcrumbs**.
3. Set the **Table** field to the appropriate table.
4. Right-click the form header and choose **Insert**.
5. Add the new formatter to the appropriate form.

Approval summarizer formatter

The approval summarizer formatter creates the summary at the bottom of an approval form.

The approval summarizer displays different information depending on what is being approved, such as a change request or a service catalog request. Following are two examples.

servicenow.

Summary of a change request

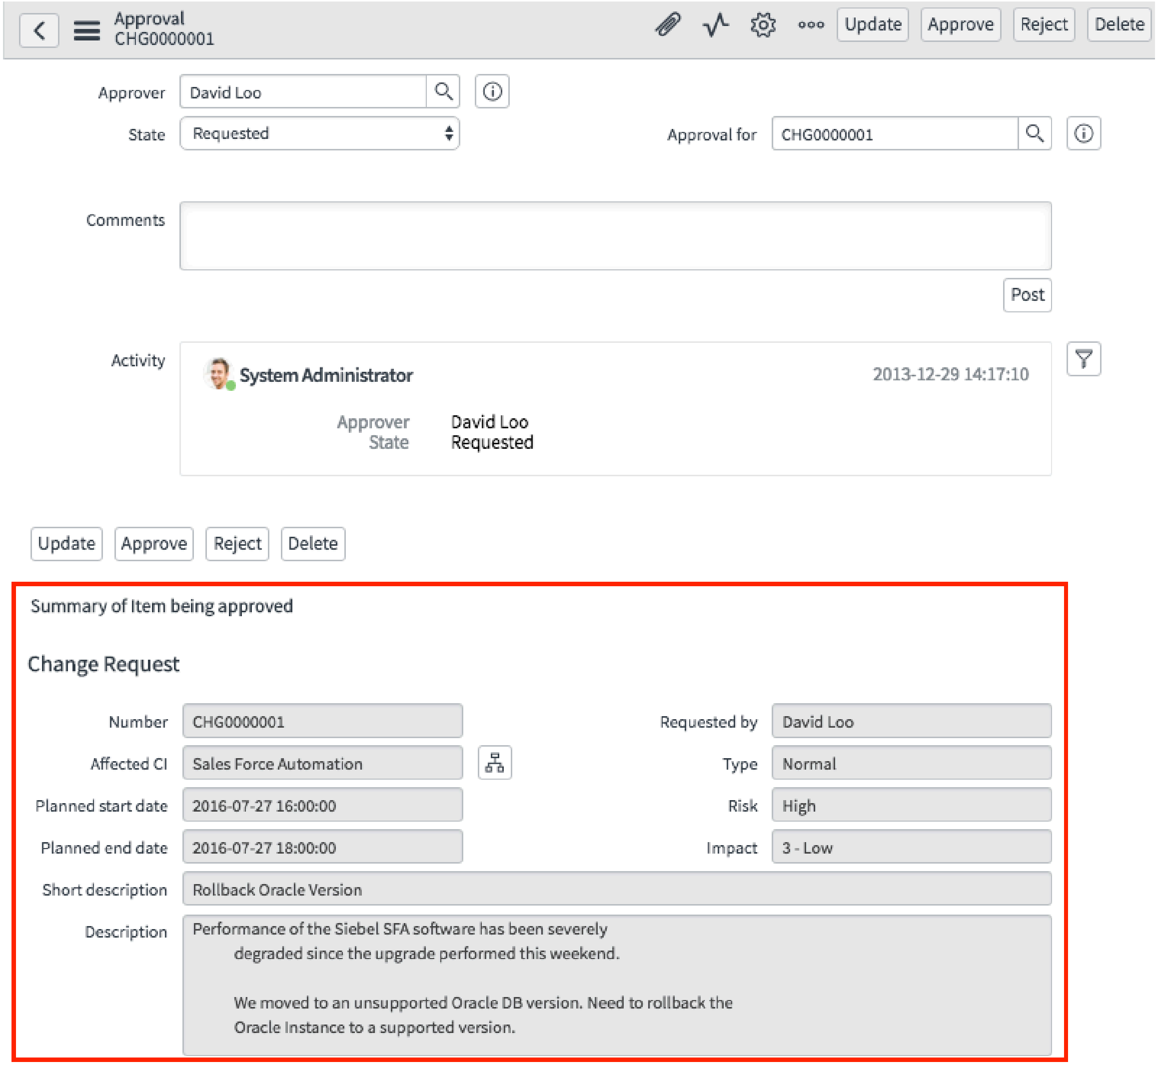

Summary of a catalog request

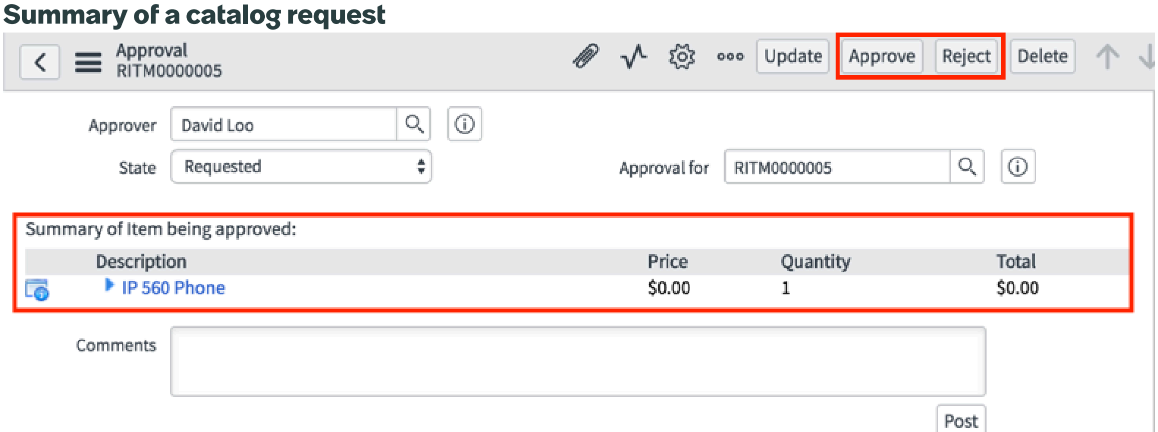

The **Reject** button allows the approver to deny one or more requested items in a multi-item request, before approving the overall request. If a requested item is denied, the workflow for that item never starts. The approver can then choose to **Accept** the item.

Note:
When the overall request is approved, you must ensure this **Reject** button is hidden. If this button is used after request approval, the requested item workflow is canceled, leaving the stage in an inconsistent state. Similarly, the **Accept** button on requested items should only appear before the overall request is approved or rejected.

Override a formatter with macros

The system uses formatters to handle complex rendering of specific form elements.

Before you begin

Role required: ui_macro_admin

About this task

Examples of form elements rendered by formatters in the base platform are:

- Activity formatter: Displays the list of activities, or history, on a task form.
- Process flow formatter: Displays the different stages in a linear process flow across the top of a record.
- Task parent breadcrumbs formatter: Provides breadcrumbs to show the parent or parents of the current task.
- Approval summarizer formatter: Displays dynamic summary information about the element being approved.

A UI macro can override formatters provided in the base system.

Procedure

1. Navigate to **All** > **System UI** > **UI Macros**.
2. Click **New**.
3. In the **Name** field, enter the same name as the formatter you want to override, but omit the .xml extension.
4. Complete the remaining fields on the form.
5. Click **Submit**.

Example:

This example shows the form that defines the existing approval summarizer formatter:

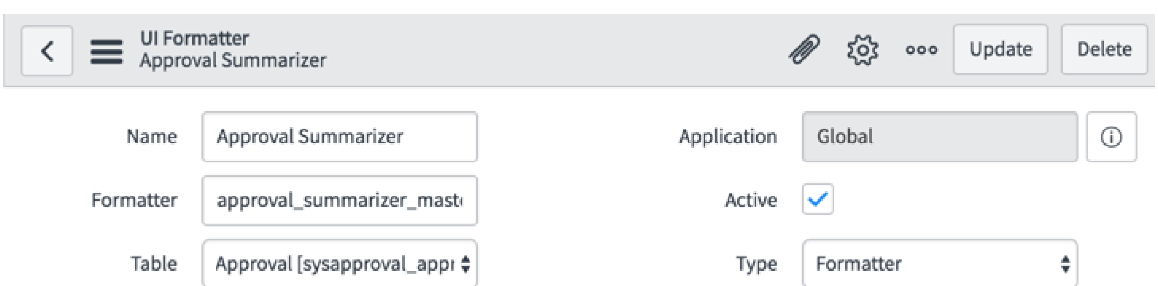

Here is the form for the UI macro that overrides the approval summarizer formatter:

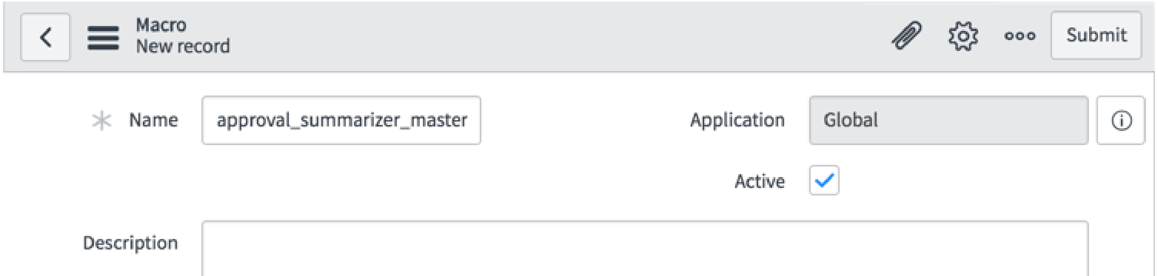

Limit the number of activity stream entries

You can set a system property to limit the number of entries allowed in an activity stream.

Before you begin

Role required: admin

Procedure

1. Enter `sys_properties.list` in the Navigation filter.
2. Search for the property *glide.history.max_entries*.
3. Edit the **Value** property to set the maximum number of entries that users can view in an activity stream.

   Note:
   New entries do not count towards the max number until the form is refreshed.

Option	Description
Name	glide.history.max_entries
Description	Maximum number of entries allowed to display in an activity stream. The default number is 250. New entries do not count towards the max number until the form is refreshed.
Type	integer
Default value	250

Form templates

Templates simplify the process of submitting new records by populating fields automatically.

To use a template, populate the most-used fields for a specific table, save it as a template, and then make the template accessible to your users. Users can manually apply a template when creating records, or an administrator can define scripts to apply templates automatically.

Create templates for the forms that you use frequently, such as incident, problem, and change. There is no limit to the number of templates that you can create or access, but having many templates for each form makes the templates more complex to manage.

You can create table-level access controls to restrict template creation. When applied, a user creating a new record from a template must satisfy the save_as_template access control for every field modified by the template.

Automatically applied templates

You can make a template that automatically applies to new, user-created records on a specific table. To create this kind of template, set the template name to match the name of the table to which the template applies.

For example, imagine you want to apply a template automatically when a user creates a record on the Windows Server [cmdb_ci_win_server] table. Set the **Name** field to **cmdb_ci_win_server** and the **Table** field to **Windows Server [cmdb_ci_win_server]** on the Template form.

Automatic templates are always global. They do not honor the **User** and **Groups** fields on the Template form. Automatic templates do not apply to records created by the system, such as those records generated by business rules, UI actions, or workflows.

Related topics

Hide the template bar for a table

Create a template using the Template form

Create a template record for any table to populate certain fields automatically.

Before you begin

Role required: template_editor_global or admin

About this task

Create a template using the Template form.

Templates allow users to more quickly and uniformly complete records.

The video above shows you how to complete the procedure below to create a template using the Template form.

Procedure

1. Navigate to **All** > **System Definition** > **Templates**.
2. Select **New**.
3. On the form, fill in the fields.

Template form

Field	Description
Name	Name of this template. **Note:** The template name can't match an existing table name.
Table	Table that this template applies to. Select **Global** to make the template available for use with all tables. **Note:** The table list shows only the tables and database views that are in the same scope as the template.
Active	Option for making the template available for use. A template must be active to be used.

Field	Description
User	User who can configure and apply the template. If you define a user, no other users can see the template unless you select the **Global** option.
Groups	Group whose members can configure and apply the template. If you define a group, no other groups can see the template unless you select the **Global** option.
Global	Option to enable any user who can access the templates to view and apply this template.
Short description	Description of the template. **Note:** Adding content to this field doesn't add that content to the **Short description** field of the forms that use this template.
Template	Content that automatically populates records that are based on this template. Select a field from the specified table in the left column and then enter the data to populate automatically in the right column. **Note:** Even though you can select dot-walked fields in the template, they don't apply to fields that are on the form.
Link element	Template that links a child table with the template for the parent table. In the template for the child table, set the value to the field that references the parent table. After you set the value, the child template is explicitly linked to the parent table. **Note:** This field doesn't appear by default. Configure the template form to add the field.

4. Select **Submit**.

What to do next

See Application scope.

Create templates for related task records

Administrators can create a template for a Task table record that also creates one or more related records in the child Task table.

Before you begin

Role required: admin

About this task

Administrators must understand the parent-child relationships between Task tables. For example, the Change Task table is a child of the Change table and the Incident Task table is a child of the Incident table.

Note:
Child templates can only be applied automatically if you first apply the parent template from a module.

Procedure

1. From the parent Task table template, configure the form layout to add these fields.
   - **Next Related Child Template**
   - **Next Related Template**
   - **Link element**
2. Create a template for the parent Task table.
   For example, create this template for the Change table.

Sample Change template

Field	Value
Name	Add server to network
Table	Change [change_request]
Short description	Set up a server on the network
Template	◦ [Short description][Set up server on the network] ◦ [Category][Hardware] ◦ [Assignment group][Hardware]

3. Create a template for the first related task.
   For example, create this template for the Change Task table.

First sample Change Task template

Field	Value
Name	Order server
Table	Change Task [change_task]
Short description	Order server hardware
Template	◦ [Short description][Order server hardware] ◦ [Assignment group][Hardware]
Link element	Change request

4. Create a template for each additional related task.
   For example, create one additional template for the Change Task table.

Second sample Change Task template

Field	Value
Name	Install server on network
Table	Change Task [change_task]
Short description	Install server on network
Template	◦ [Short description][Install server on network] ◦ [Assignment group][Hardware]

Field	Value
Link element	Change request

5. From the parent Task table template, set **Next Related Child Template** to the first related child task.
   For example, in the Add server to network template, select `Order server` in the **Next Related Child Template** field.

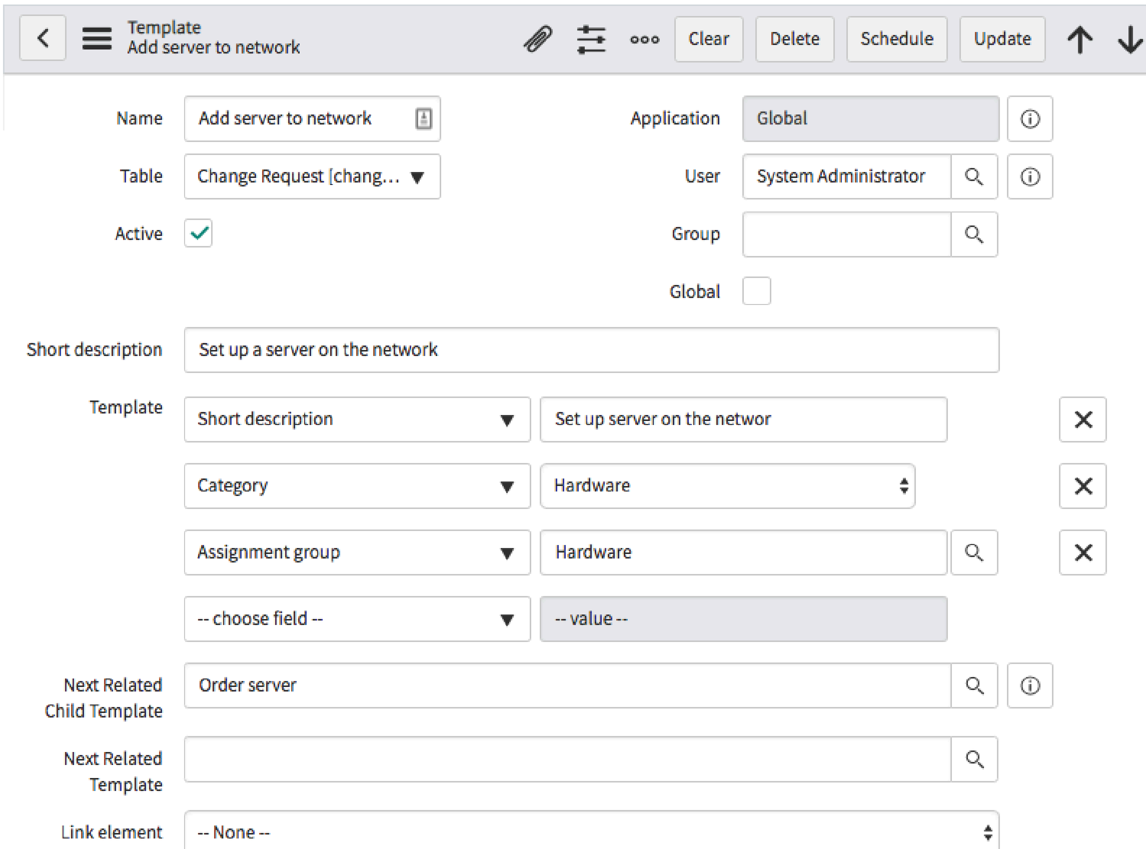

6. For each child related task, set **Next Related Template** to the next related task template.
   For example, in the Order server template, select `Install server on network` in the **Next Related Template** field.

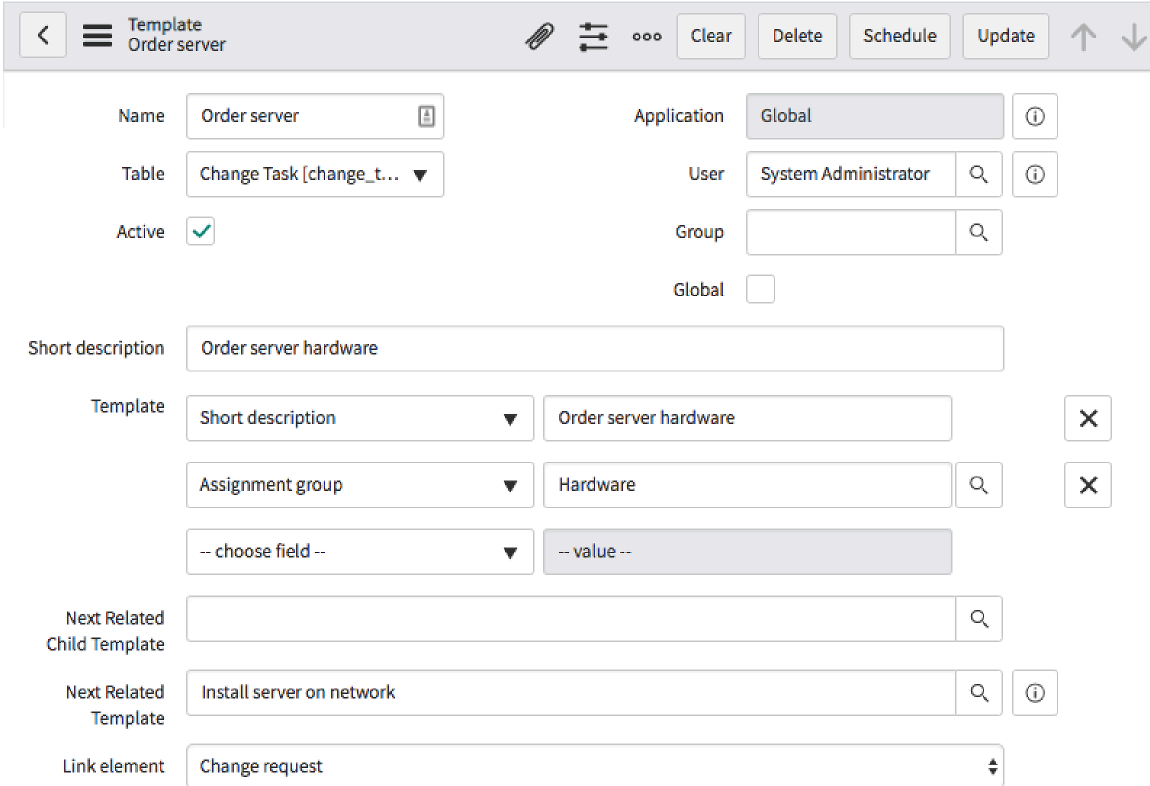

Note:
The last related task template does not have a value for **Next Related Template**. For example, the **Install server on network** template does not have a value in **Next Related Template**.

What to do next

Create a module for the parent Task table template so that child templates can be applied to related task records. For example, create a module for **Add server to network**.

Create a template by saving a form

Save a populated form as a template.

Before you begin

Toggle the template bar so it is visible on forms.
Role required: admin

Procedure

1. Navigate to a form.
2. Complete the form as it should appear when a user applies the template.
3. In the template bar, click the plus icon (+).

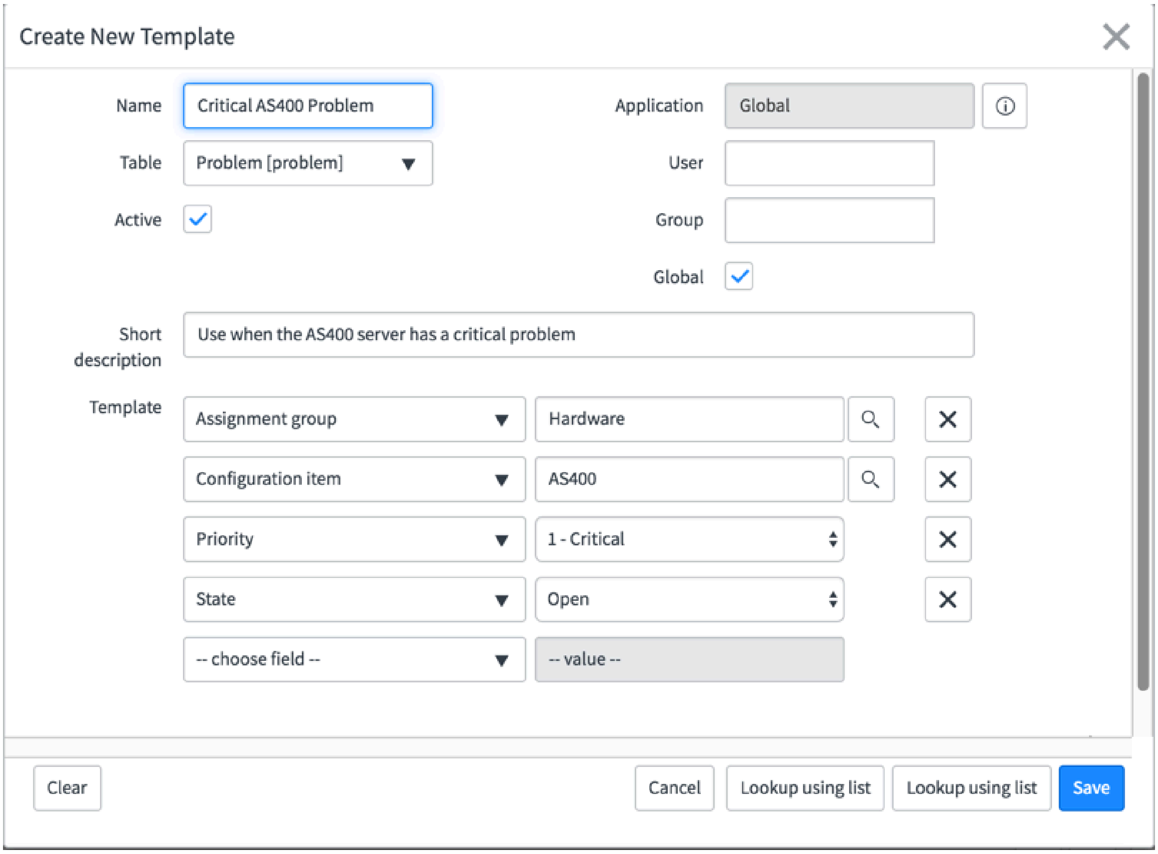

4. Enter a descriptive name to make it easy for a user to select the correct template.
5. Make any additional changes as needed.
6. Click **Save**.

Related topics

ACL rule types

Create records based on a template

You can create and schedule a scheduled job to create records based on a template. For example, you can regularly create a populated task record to perform a weekly backup.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Templates**.
2. Select a template record.
3. Click **Schedule**.
   The Scheduled Entity Generation form appears.
4. In the **Run** choice list, select how frequently to create a record.
5. Complete the schedule information.
6. Click **Submit**.

Create a module for a template

You can create a module to open a form with pre-populated template data.

Before you begin

Role required: admin

About this task

Child templates are only applied if the parent template is applied from a module. Child templates are not applied by applying a template to a new form.

Procedure

1. Point to the application menu, such as **Self-Service**, and click the edit application (pencil) icon.
2. In the **Modules** related list, click **New**.
3. Fill in the **Title** and **Order** fields as appropriate.
4. In the Link Type form section, in the table field, select the table you want to create the form in.
5. In the **Link Type** field, select **URL (from Arguments)**.
6. In the **Arguments** field, enter `<table>.do?sys_id=-1&sysparm_template=<templatename>`
   The `<table>` is the name of the table you selected for the table field, for example `incident`. The `<templatename>` is the name of the template you want to use to pre-populate the form, for example `Incident call type`.
7. Click **Submit**.
8. Refresh the application navigator to view the new module.

Template bar

Use the template bar to apply, edit, and create templates.

When toggled, the template bar appears at the bottom of the form. Select a template title to apply the template to the form. Use the buttons at the bottom-right to add, edit, and access the complete list of templates.

Template Bar in Core UI

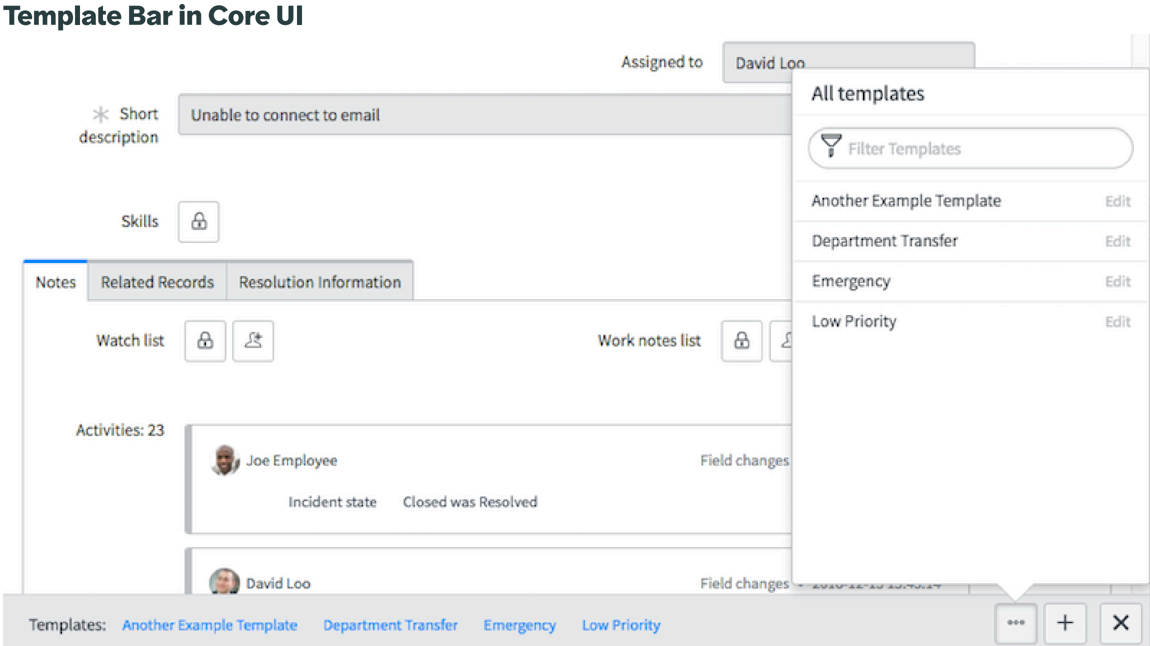

Template Bar Elements

Element	Description
Templates	Available templates appear in blue text on the template bar. Click a template to apply it to the form.
All Templates button ( )	Toggles the All Templates window.
Create New Template button ( )	Displays the Create New Template window.
All Templates popup	This window displays when the **All Templates** button is selected. The window displays a filterable list of all available templates. Click the **Edit** button to the left of any template to edit that template.
Disable Template Bar button ( )	Disables the template bar. To display the bar, see Toggle the template bar

For non-admins, see KB0711956 for more information on enabling the Edit button.

Toggle the template bar

The template bar appears at the bottom of forms. It provides shortcuts to apply, edit, and create templates.

Before you begin

Role required: none

About this task

You can toggle the template bar, which hides or shows it for all forms. The template bar is shown by default.

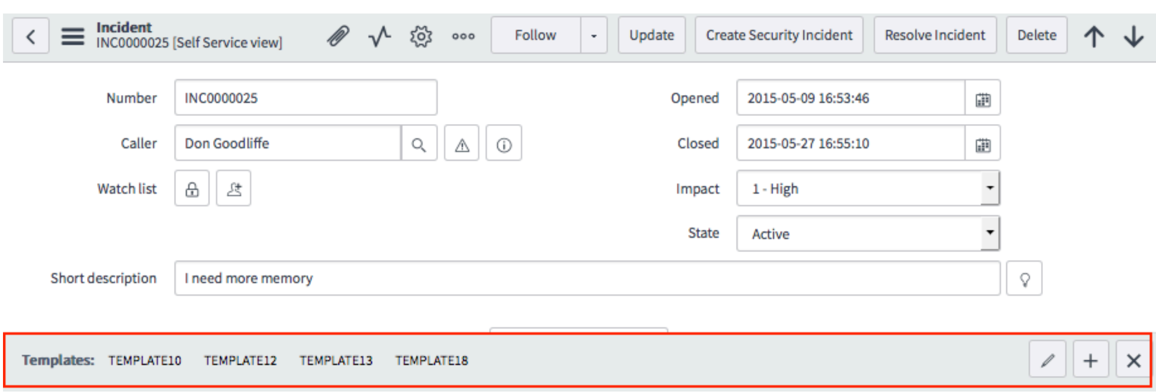

Procedure

1. Navigate to a form.
2. Click the more options icon ( ) in the form header.
3. Select **Toggle Template Bar**.

Scripted templates

You can apply an active template to a record using JavaScript.

Apply a template to current

To apply a template, use the *applyTemplate* method.

```
current.applyTemplate("<templatename>");
```

Apply a template to a GlideRecord

To apply the template to a record other than current, change current to a GlideRecord variable. When using a GlideRecord variable, you may need to initialize it after declaring the variable.

```
var rec1 = new GlideRecord("incident");
rec1.initialize();
rec1.applyTemplate("my_incident_template");
```

Apply a template from a UI action

The following script demonstrates a possible customization to the *Create Change* UI action on the Problem form. After you add this script to the UI action, a user can select the UI action to create a change record with information from both the problem record and the change template.

```
var change = new GlideRecord("change_request");
change.initialize();
change.short_description = current.short_description;
change.description = current.u_details;
change.cmdb_ci = current.u_service;
change.priority = current.priority;
change.requested_by = current.caller_id;
change.assignment_group.setDisplayValue('Change & Release');
change.u_status = 'New';
change.parent = current.number;
if(0 == change.applyTemplate("standard_rfc")) {
    current.rfc = change.insert();
    current.comments = 'Change ' + change.number + ' created.';
}

var mySysID = current.update();

gs.addInfoMessage("Change " + change.number + " created");
action.setRedirectURL(change);
action.setReturnURL(current);
```

Script a template with child templates

When using *applyTemplate* with a template that has one or more child templates, the system creates the parent record before applying the child templates. This behavior ensures that any references or dot-walked fields from the child record to the parent have a valid target.

For example, if a template for the Change Request table has a child template for the Change Task table, applying the Change Request template inserts a Change Request record into the database. It assigns this record as the Change request for the Change Task record, then applies the child template to the Change Task record.

Apply a template using a sys_id instead of a name

To apply a template using a sys_id, use the following method.

```
GlideTemplate.get(template.sys_id).apply(GlideRecord)
```

UI actions

UI actions include the buttons, links, and context menu items on forms and lists. Configure UI actions to make the UI more interactive, customized, and specific to user activities.

Administrators and users with the ui_action_admin role can define UI actions.

UI action controls

You can create a UI action to provide any of these controls:

- A button on a form.
- A context menu item on a form that appears when you open the form context menu or right-click the form header.
- A related link in a form.
- A button in the banner on top of a list.
- A button at the bottom of a list.
- A context menu item on a list that appears when you open the list context menu or right-click the list header.
- A menu item for the action choice list at the bottom of a list.
- A related link at the bottom of a list.

Note:
To hide or restrict the **New** or **Edit** UI action on related list, use list control.

servicenow

Form UI actions

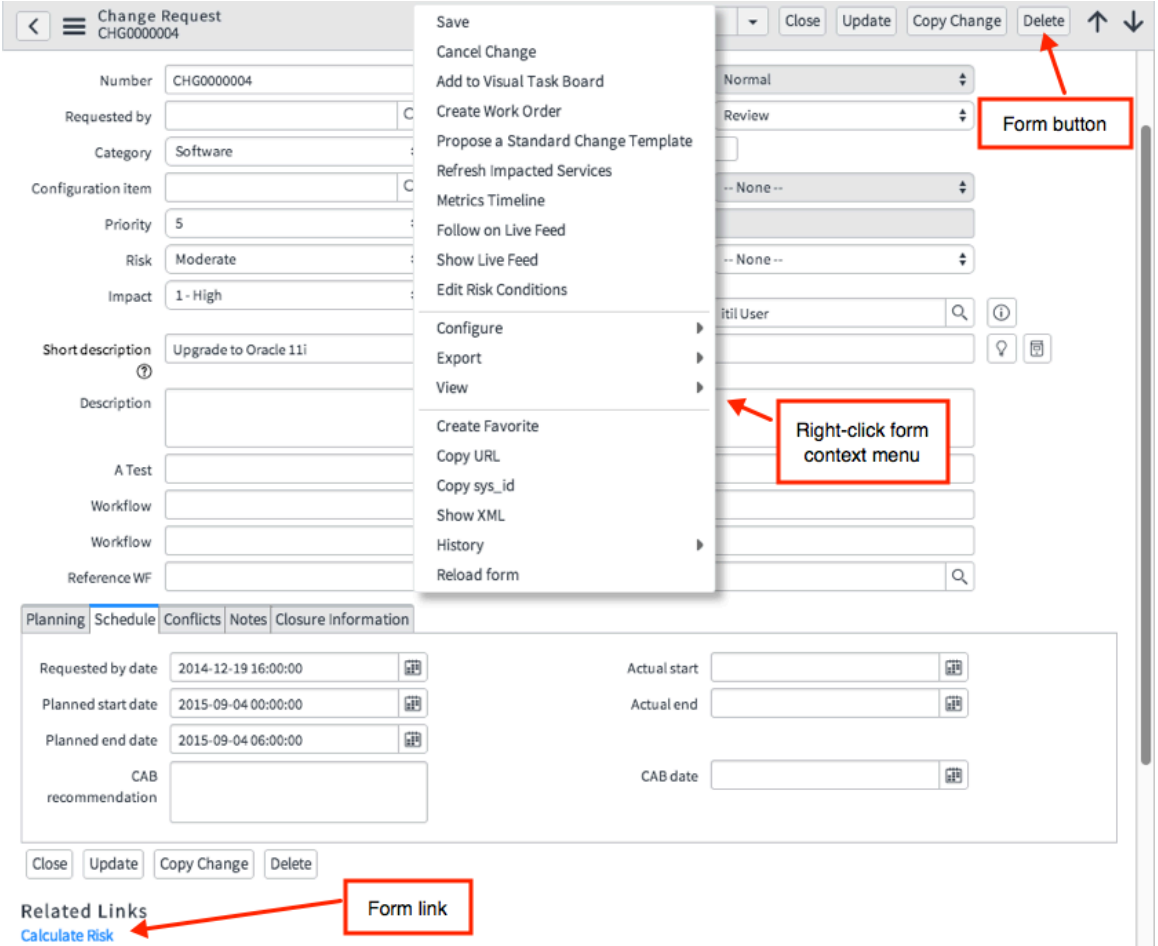

List UI actions

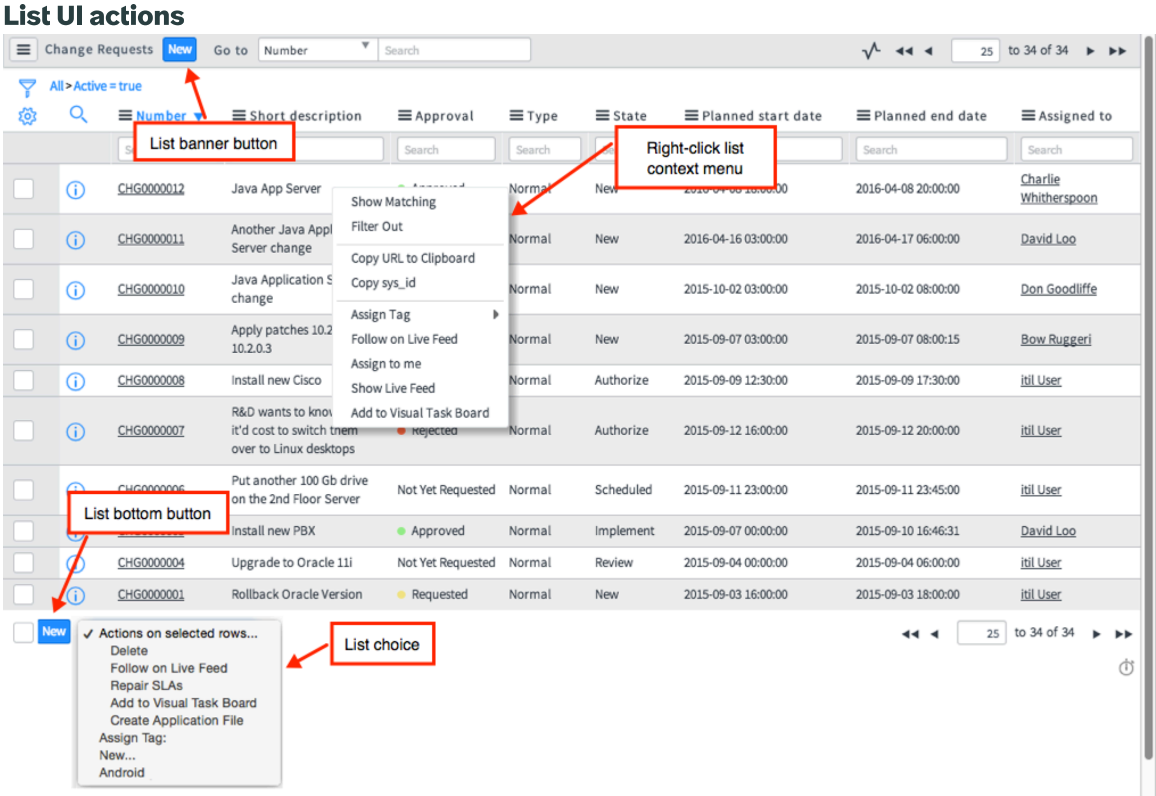

Create a UI action

You can create a new UI action or edit an existing one.

Before you begin

Role required: ui_action_admin or admin

About this task

You can add UI actions to tables and database views that are in the same scope as the UI action and tables that allow UI actions from another application scope to run on them.

Procedure

1. Navigate to **All** > **System Definition** > **UI Actions**.
2. Click **New** or open an existing record.
3. Define the UI action by completing the fields.
   You may need to configure the form to see all the fields.

UI action fields

Field	Description
Name	The text that appears on the button, link, or context menu item.
Table	The table on which the UI action is available. By default, the UI action also appears on tables that extend the selected table (for example, Task actions appear on the Incident table). Select **Global** to make the action available on all tables.
Order	The order in which the UI action appears. The order applies to buttons from left to right and to menu actions from top to bottom.
Action name	A unique name to use when referencing the UI action in scripts.
Active	An option to enable the UI action when selected. To disable a UI action, clear the check box.
Show insert	An option to show a button on new records that have not been inserted.
Show update	An option to show a button on existing records.
Client	An option for the UI action to execute its script in the user's browser, not on the server. When enabled, the **Onclick** field appears above the Condition field.
Form button	An option to put a button on a form.
Form context menu	An option to put an item in a form context menu (right-click the form header).
Form link	An option to put a link in the Related Links section of a form.
Form style	○ Primary—Colors the UI Action blue. ○ Destructive—Colors the UI Action red. ○ Unstyled—Does not color the UI Action.
List banner button	An option to put a button in the banner of a list.

Field	Description
	Note: List banner buttons are not intended to support record-specific conditions, so only the first row is considered when the condition is evaluated to determine whether the button will show for the list. Do not use record-specific conditions (e.g., current.getValue('state') === 'closed') in list banner button UI Actions.
List bottom button	An option to put a button at the bottom of a list. **Note:** List bottom buttons show regardless of condition and are evaluated per record on execution.
List context menu	An option to put an item in a list field context menu (right-click a cell in a list).
List choice	An option to put an item in the action choice list at the bottom of a list. **Note:** List bottom buttons, like list choice actions, show regardless of condition and are evaluated per record on execution.
List link	An option to put a link in the Related Links section at the bottom of a list.
List style	◦ Primary—Colors the UI Action blue. ◦ Destructive—Colors the UI Action red. ◦ Unstyled—Does not color the UI Action.
Overrides	A UI action that this UI action overrides.
Messages	Text strings that the UI action can use as a key to look up a localized message alternative from the Message [sys_ui_message] table. Each message key is on a separate line in the Messages field. The instance looks for a localized message string anytime the UI action makes a getMessage(msg) call where the msg string matches a key in the Messages field.
Comments	Descriptive content regarding this UI action.
Hint	The text that appears when a user points to the UI action control.
Onclick	The name of the JavaScript function to run when the UI action is executed. The function is defined in the **Script** field.
Condition	A JavaScript conditional statement that restricts when a UI action appears. Conditions always run on the server.

Field	Description
	Note: ◦ The current object is not available for conditions on a list context menu. If the **List context menu** option is selected, any use of `current` on these actions is ignored. ◦ You can reference the parent record for the UI action conditions on a related list button. For example, to disable the **New** and **Edit** buttons on the Affected CIs related list for closed changes, copy the global m2m UI actions to the task_ci table and add a condition of `parent.active`. ◦ If you leave one of the fields empty that you specify in your condition statement, that condition defaults to **true**.
Script	The script to run when the UI action is executed. Function names must be unique.
Workspace	
Workspace Form Button	An option to make the UI action appear on the list of UI actions in a workspace.
Workspace Form Menu	An option to make the UI action appear as a list item in the menu associated with UI actions in a workspace.
Format for Configurable Workspace	An option to use the UI action in a Configurable Workspace. If cleared, the UI action applies in Legacy Workspaces.
Workspace Client Script	A script to run when the UI action is executed in workspaces.
Requires role	
Role	The roles required for the UI action to apply. Users must have at least one of the roles for the UI action to apply.
Related lists on the form view	
UI Action Visibility	The views of the form that the UI action applies to. Use this option to restrict the UI action to form views that you specify. A UI action is available for the specified view according to the following rules: **a.** If there are no visibility rules, the action appears on all views. **b.** Any exclude rule on a given view means that the action does not appear on that view. **c.** If there is at least one include rule, then the action appears only on views that are specifically included.
Versions	All versions of the UI action. Use this list to compare versions or to revert to a previous version.

4. Click **Submit** or **Update**.

Note:
If the UI action is enabled to run on the client side, wrap it in a function. Otherwise, the contents of the **Script** field runs when the page loads.

Example:
As part of a UI action script, you can redirect a user to a URL. For example, you might add links to a form or open a new record after it is created from a UI action. To redirect a user to a URL from a UI action, use this syntax in the **Script** field to define the redirect link:

```
action.setRedirectURL ( '');
```

To direct a user to a record, use this syntax, where *new_record* is the variable name for the GlideRecord:

```
action.setRedirectURL (new_record );
```

Override a UI action for an extended table

You can override or remove a UI action for a table that is extended from another table.

Before you begin

This is not applicable to domain separated instances.

About this task

When a UI action is defined for the Task table, it applies to all tasks, including incidents, changes, problems, and any other tables that extend the Task table. Similarly, a global UI action applies to every table. However, you can override a UI action for a specific table. This example demonstrates how to override or remove a UI action on the Task [task] table for only the Incident [incident] table.

Procedure

1. Complete the following steps to override a UI action on the Task table for just the Incident table.

   a. Create a UI action on the Incident table with the same **Action name**.
   If the **Action name** is not defined, update both the new UI action and the UI action to be overridden with the same **Action name**.

   b. Enter a script that is specific to the Incident table.

2. Complete the following steps to remove a UI action on the Task table for the Incident table.

   a. Navigate to the UI action definition for the Task table.

   b. Add the condition *current.getRecordClassName() !='incident'*.

UI policies

UI policies dynamically change the behavior of information on a form and control custom process flows for tasks.

For example, you can use UI policies to make the number field on a form read-only, make the short description field mandatory, and hide other fields. Basic UI policies do not require any scripting, however for more advanced actions, use the **Run scripts** option.

You can also use client scripts to perform all of these actions, but for faster load times use UI policies when possible.

Create a UI policy

Create a UI policy to define custom process flows for tasks.

Adding a UI Policy

Before you begin
Role required: ui_policy_admin

About this task
A UI policy condition evaluates all fields even if they are not visible on the form. This function removes the requirement that a field must be on a form for it to be evaluated.

Note:

- Policies carried over from versions prior to Fuji are evaluated differently. Fields that previously were not evaluated are evaluated.
- UI policies are not supported on search screens.
- UI Policies also apply to forms and lists displayed within Content Management System application.

Procedure

1. Navigate to **All** > **System UI** > **UI Policies**.
2. Click **New**.
   The UI Policy [Advanced view] form opens
3. To change the view, in **Related Links** click **Default view**.
4. Complete the form, as appropriate.
   You may need to configure the form to see all the fields.

UI policy fields

Field	Description
Table	The table for the form to be modified.
Active	The active status of the UI policy. Only active UI policies are applied.
Short description	Short summary of the UI policy.
Order [Advanced view]	The processing sequence, from the lowest to highest number. If two policies conflict, the UI policy with the higher number executes. For inherited UI policies, the extended (child) table's UI policies are executed first. Then the base table UI policies are executed; both from lowest to highest specified value.
When to Apply	
Conditions	The conditions which, if fulfilled, cause the UI policy to be applied. Conditions are built with the condition builder. To set conditions using a script, use a client script instead. Conditions are only rechecked if a user manually changes a field on a form. If the change is made by a UI action, context menu action, or through the list editor, it is not evaluated.
Global [Advanced view]	Option for specifying whether the UI policy applies to all form views. If this check box is cleared, the UI policy is view-specific. By default, the Global UI policy applies to all form views. However, a UI policy can be specific to a view. For example, you can define a UI policy for only the itil view of a form. Use the **View** field to accomplish this.

Field	Description
View [Advanced view]	Option for indicating which form view the UI policy applies to. This field is visible only if **Global** is not selected. If Global is not selected and the View field is left blank, the script applies the default view. For more information on form views, see View management.
Reverse if false [Advanced view]	Option for specifying that the UI policy action is undone when the conditions of its UI policy evaluate to false. In other words, when the conditions are true, the specified actions are taken and when they are false, the actions are undone.
On load [Advanced view]	Option for specifying that the UI policy behavior should be performed OnLoad as well as when the form changes. You can check or clear the **On load** check box in a UI policy to control whether it runs every time a form is loaded when the conditions are satisfied. In this example, an administrator does not want an incident to enter the **Awaiting user info** state unless the user provides an explanation to the customer. The administrator creates a UI policy with the following settings. ◦ In the When to Apply section, adds the condition **[State] [is] [Awaiting user info]** and clears the **On load** check box. This condition means that the UI policy applies only when the state is changed to **Awaiting user info**. ◦ In the **UI Policy Actions** related list, creates a record that makes the **Additional comments** field mandatory when the condition is met.
Inherit [Advanced view]	Option for specifying whether extended tables inherit this UI policy. When a child table has an inherited UI policy from its parent table, the UI policy on the child table always runs first. This event is true regardless of the **Order** of the UI policies. Consider the following example: ◦ A child table has a UI policy with **Order** value **500** that shows the **Urgency** field when its conditions are met. ◦ Its parent table has a UI policy with the same conditions that hides the **Urgency** field. The parent table UI policy has **Order** value **100**. ◦ Although the parent table **Order** field has a lower value, the child UI policy runs first and then the parent UI policy runs. When the conditions are met, the **Urgency** field is hidden.
Script	
Run scripts [Advanced view]	Option for specifying whether advanced behavior can be scripted for both true and false conditions.
Execute if true [Advanced view]	A script that executes when the conditions of the UI policy are fulfilled. ◦ This field is available only if **Run scripts** is selected. ◦ This feature is not available outside the scope or in global scope.
Execute if false [Advanced view]	A script that executes if the conditions of the UI policy are not fulfilled and the **Reverse if false** option is selected.

Field	Description
	○ This field is available only if **Run scripts** is selected. ○ This feature is not available outside the scope or in global scope.
Other fields	
Run scripts in UI type	The UI type for this UI policy: Desktop, Mobile / Service Portal, or Both.
Related List: UI Policy Actions	
Table	[read-only] Field the UI policy action applies to.
Field name	Field on the selected table to which the UI policy performs an action if true. **Note:** If the specified field is not found on the form, the UI policy performs the action on the variable with the same name.
Mandatory	Choice list for specifying how the UI policy affects the mandatory state of the field. Choices are: ○ Leave alone ○ True ○ False
Visible	Choice list for specifying how the UI policy affects the visible state of the field. Choices are: ○ Leave alone ○ True ○ False
Read only	Choice list for specifying how the UI policy affects the read-only state of the field. Choices are: ○ Leave alone ○ True ○ False
Related Links	
Default view or Advanced view	Changes the form view to the default or advanced view. The fields change based on the view.

5. Click **Submit**.

Example: creating a UI policy

Create a UI policy to implement controls in the Incident form when the state changes to **Resolved**.

Before you begin

Role required: ui_policy_admin

About this task

This example demonstrates how to implement the following controls.

- Make a **Close Notes** field mandatory.
- Hide the **Opened by** field.
- Make the **Priority**, **Severity**, and **Urgency** fields read-only.
- Run a client script that displays an alert message.

Procedure

1. Navigate to **All** > **System UI** > **UI Policies**.
2. Click **New**.
3. Supply the following information.

New UI policy

Name	Input
Table	Incident
Conditions	[Incident state] [is] [Resolved]
Reverse if false	Select this check box. If the incident state is not **Resolved**, the UI policy is reversed.
On load	Select this check box to perform the actions when the form is loaded or when the condition changes.

4. Right-click the form header and select **Save** from the context menu.
   The **UI Policy Actions** related list appears.
5. In the related list, click **New**.
6. Provide the following information.

New UI policy actions

Name	Input
Field name	Close notes. This UI action makes the Close notes field mandatory.
Mandatory	True
Visible	Leave alone
Read Only	Leave alone

7. Click **Submit**.
8. Repeat the process to create UI policy actions to hide the **Opened by** field, and to make the **Priority**, **Severity**, and **Urgency** fields read-only.

Client scripts for UI policies

Any scripts you create for UI policies run on the client side.

You can use different options in the UI Policy form to control when and how the UI policy is applied. These options include client scripts, OnLoad execution, and view-specific UI policies. Administrators can use the UI Policy form to create client scripts that run `onChange` when the UI

policy conditions are met (**Execute if true**) or not met (**Execute if false**).To display these scripting fields in the UI Policy form, in the Script section, select the **Run scripts** check box.

For example, to display an alert to the user when the incident **State** field changes to **Resolved**, create the following script in the **Execute if true** field.

```
function onCondition(){
alert('You changed the "Incident state" to Resolved. Please
 enter your comments in the "Close notes" field.');
}
```

Service catalog UI policies control the behavior of catalog item forms. Service catalog UI policies can be applied to a catalog item or a variable set. For more information about using UI policies for service catalog items, see Service catalog UI policy

Advanced form configuration

Administrators can configure advanced form features, such as form focus, form splits, and derived fields.

Allow insert options on task records

The **Insert** and **Insert and Stay** options are disabled by default for task records such as incidents and change requests. You can set a system property to show these options for task records.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Locate the property **Allow the use of the "Insert" and "Insert and Stay" options on task derived tables** (`glide.ui.task.insert`).
3. Select the check box to enable or clear the check box to disable (default) the options for tasks.
4. Click **Save**.

Turn off focusing on the first writable field of a form

Switch the focus on a form to the first element on the page instead of the first writable field to make the form more accessible to users who use screen readers.

Before you begin

Role required: admin

About this task

By default, forms set focus on the first writable field on the form. With first-field focus turned off, the form focuses on the first element on the page instead. Focusing on the first element, instead of the first writable element, helps orient users who use screen readers so they can find all the elements on a form.

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Set the property `glide.ui.focus_first_element` to **false**.

Hide the template bar for a table

Administrators can hide the template bar for a table by creating a system property.

Before you begin

Role required: admin

About this task

Creating a table-specific property and setting it to false prevents users from displaying the template bar on a record for the specified table.

Procedure

1. Navigate to `sys_properties.list`.
2. Add a system property with the following settings.

Property	Value
Name	glide.ui.show_template_bar.<TABLENAME>
Type	true | false
Value	false

3. Click **Submit**.

Deactivate form submission with the Enter key

By default, when you press the `Enter` key in a simple one-line text field, a choice list, or a Boolean field, the form is submitted.

Before you begin

Role required: admin

About this task

You can use a system preference to deactivate this feature if you do not want the `Enter` key to submit the form.

Procedure

1. From the left navigation pane, select **User Administration** > **User Preferences**.
2. Select the **enter_submits_form** preference.
3. Set the value to **false**.
4. Click **Update**.
   The change does not take effect until user preferences are reloaded either at login or when a session is created.

Enable multiple form splits

Form splits enable you to organize fields on a form into columns. Administrators must add a property to enable form splits.

Before you begin

Role required: admin

About this task

Configuration of two form splits

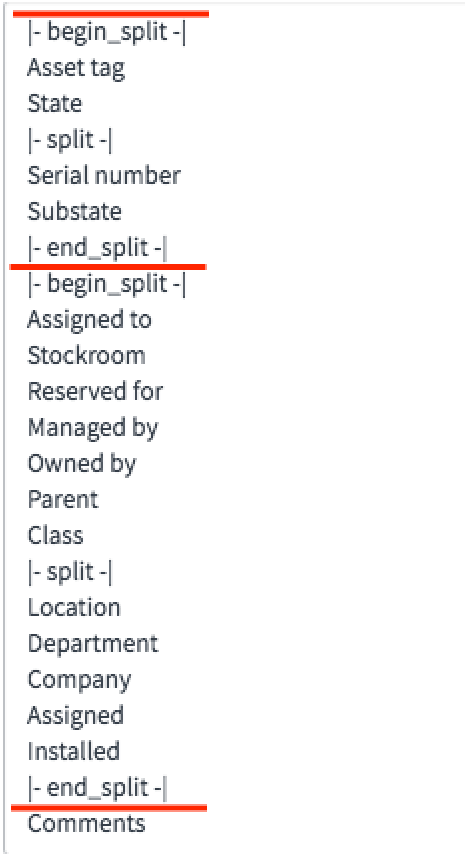

Example of the form with two form splits

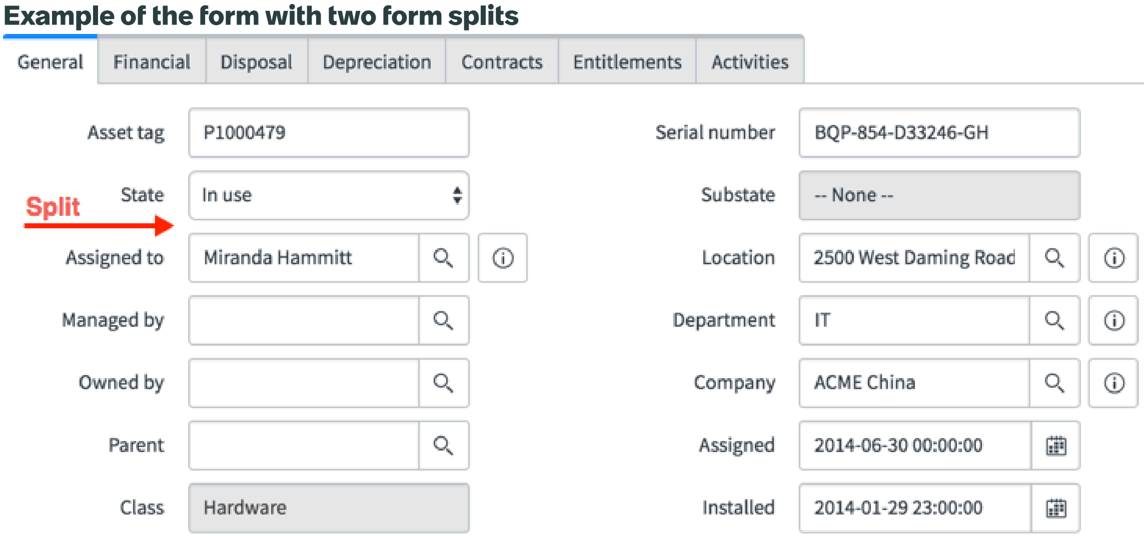

When you organize fields in this manner and the user is viewing the form on a small mobile device, the fields within the first split are listed before the fields in the second split. In this example, the **Asset tag**, **State**, **Serial number**, and **Substate** field are listed before any of the fields below them. You can also create elements that span the form at the top of the form.

Procedure

1. Enter `sys_properties.list` in the Navigation filter.
   The entire list of properties in the System Properties [sys_properties] table appears.

2. If the property does not exist, click **New**.
3. Enter the following information.

Option	Description
Name	glide.ui.form_multiple_splits
Description	Enable multiple form splits
Type	true | false
Value	true

4. Click the form context menu icon and select **Save**.
   The **Categories** related list appears.
5. Click **Edit** and move **UI** to the **Categories List**.
6. Click **Save**.
   The System Property form reopens and the new property appears in the UI Properties page.

Define required fields

You can specify which form fields are required in forms. This action prevents users with the personalize_form role from removing the field by configuring the form layout.

Before you begin

Role required: admin
The Required Form Fields plugin must be active.

About this task

If present on a form, only an administrator can remove required fields. When you configure a form, required fields appear in a gray color and have a tooltip indicating they are required. Required fields are defined in the Required Form Fields [sys_ui_element_required] table.

The expected (although not the only) use case for this feature is as part of a delegated administration scheme. For example, you can grant branch offices the rights to modify forms by granting the personalize_form role, but not allow them to remove certain fields which are critical to overall business processing.

Procedure

1. Navigate to `sys_properties.list`.
2. Locate the property named *`glide.ui.form.enforce_required_fields`* and make sure it is set to true.
3. Navigate to **System UI** > **Required Form Fields**.
4. Click **New**.
5. Select the table and field, and then select the **Required** check box.

   Example
   The following example sets the **Short description** field on the Incident form to be required.

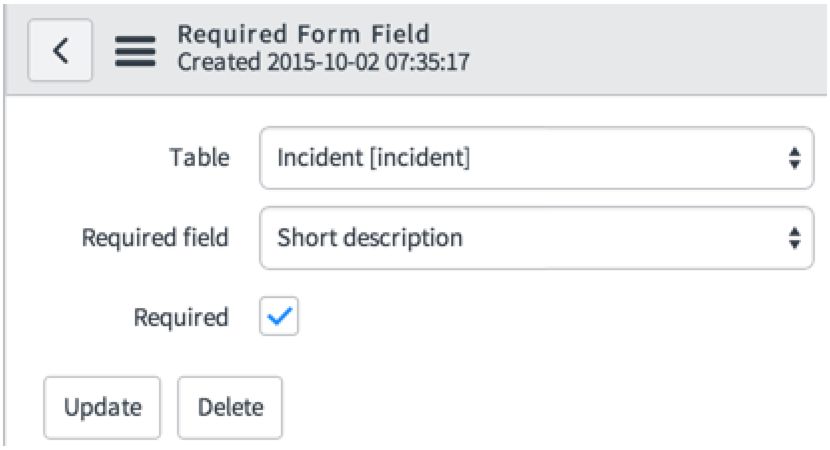

If you decide later that you do not want to make the field required, clear the **Required** checkbox. This action is preferable to deleting the record.

6. Click **Submit**.

Result

If the table specified has extension tables, then the Required Form Field record applies to forms of all extended tables. For example, if an administrator specifies that the **Short description** field is required for the Task table, then this configuration applies to the Incident form, Change Request form, Problem form, and so on.

An extended table can override the Required Form Field rule of its base table. For example, if the **Short description** field is required for the Task table, but not required for the Incident table, it is required for all Task tables except Incident.

Control the label type for derived fields

You can configure the type of label that appears for derived fields.

Before you begin

Role required: admin

About this task

An example would be displaying the email address for the caller when looking at an incident record. The email address is not stored in the Incident table, but is obtained by following the caller reference from the incident to the User table. You control the label type by setting a system property. For example, the field label for an incident email address can be either of the following strings.

- **Caller Email**, which represents the complete label to uniquely identify this field as the email from the caller user record. If you display the email address for the person the ticket is assigned to, its label would be **Assigned to Email**.
- **Email** which is the label for the target field. This label is not unique on the form if, for example, you are also displaying the email address of the person assigned to the incident. However, usually the placement of the field on the form makes it clear what the field represents.

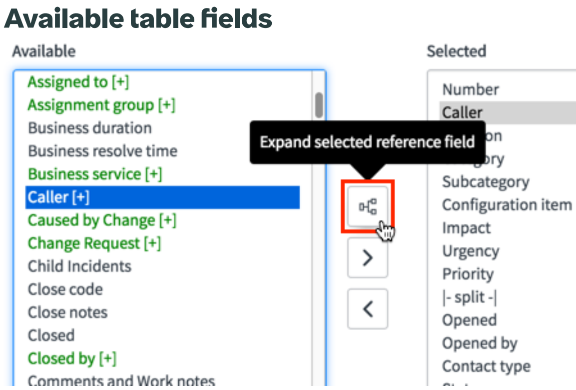

Choose the **Caller** field and click the expander icon.

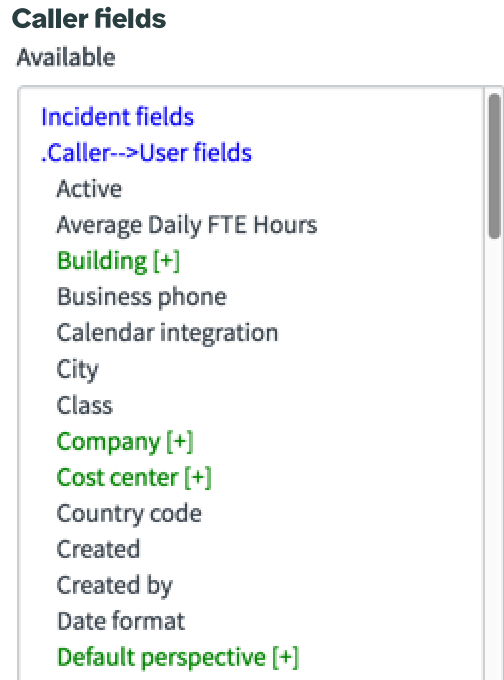

Select the available fields associated with the **Calle**
field.

Procedure

1. Navigate to **All** > **System Properties** > **System**.
2. Locate the **Use short labels for all fields. For example, if a form contains the caller's email address, use the "Email" label rather than the full label of "Caller Email"** (*`glide.short.labels`*) property.
   The default value is **true**, meaning that the short label is displayed in all forms.
3. If you want to display the complete label in forms, clear the check box.
4. Click **Save**.

Field administration

The individual pieces of data in a record are called fields. You enter data in fields on the form or by using the list editor. Administrators can create new or modify existing fields.

You can enter data in fields by using the list editor or by using a form. In form view, fields appear as fields in the form, and in list view they appear as columns of data in the table. Administrators can create new fields or change the type of existing fields.

Field administration considerations

- If you edit a field on a child table that is present on the parent table, it's also changed in the parent table and all other child tables.
- Kanji characters aren't currently supported. If a table column identifier is created using Kanji characters, it's ignored in any update or insert operation from a form.
- When you create a custom field, use one of the supported field types described in Field types. Other field types, such as User Input, are for internal use only and aren't supported for custom fields.

- The Field Name and Table Name field types have a built-in dependency relationship. For example, Field Name fields don't exist independently; they require a Table Name field. Also, Field Name fields always contain a value, which is sys_id by default. You can't set the value for a Field Name field to None, even if you apply the attribute allow_null=true.
- If you need to change the data type for a field, create a new column using the new data type, copy the data from the old column to the new column via a background script, and then rename the original column label.
- Various conditions can cause a field to be read-only. For more information, see the Determining why a field is Read Only (grayed out and not editable) [KB0783470] article in the Now Support Knowledge Base.
- Fields may not display on forms as expected. For more information, see the Troubleshooting fields not appearing on forms [KB0547219] article in the Now Support Knowledge Base.

Field types

These field types are available to administrators when creating fields or changing the type of existing fields.

Field types

Field type	Platform name	Supported in configurable workspace	Description
approval_rules	Approval Rules	No	
audio	Audio	No	Audio field for uploading and embedding `.mp3` or `.ogg` audio files.
auto_increment	Auto Increment	Yes	
auto_number	Auto Number	Yes	
boolean	True/False	Yes	True/False field that holds a true or false value, typically displayed as a check box.
bootstrap_color	Bootstrap Color	No	
breakdown_element	Breakdown Element	No	
calendar_date_time	Calendar Date and Time	No	
catalogue_preview	Catalog Preview	No	
char	Char	Yes	
choice	Choice	Yes	Choice field that restricts values to a list of options, which is ideal for fields with a limited set of valid entries. For example, a priority field or a category field.

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
collection	Collection	No	Collection field type that applies any changes made to the assigned field also to the entire table For example, changes such as attributes or the **Read only** check box that are applied to a table entry with this field type are applied to the table rather than a specific field. This entry is automatically created when a table is created. There shouldn't be more than one entry for a table of this type.
color	Color	No	String field that accepts CSS color declarations,including hex or RGB notation, and displays a preview. For more information on CSS color declarations, go to the Mozilla developer web docs. See the "<color>" article. https://developer.mozilla.org/en-US/docs/Web/CSS/color_value.
color_display	Color Display	No	
composite_field	Composite Field	Yes	
composite_name	Composite Name	No	
compressed	Compressed	No	Compressed field that stores large binary data in a compressed format, reducing database storage requirements. It's used for files like images or documents, automatically compressing and decompressing data to optimize performance. This field type helps maintain efficient storage and database performance, especially when handling large volumes of data.
condition_string	Condition String	Yes	Text field that accepts a plain JavaScript condition statement that is validated automatically for correctness before an update.
conditions	Conditions	Yes	Field that adds the condition builder to a form. You must specify a

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
			dependent field that references the table name.
counter	Counter	No	
css	CSS	Yes	
currency	Currency	Yes	Decimal field with four digits after the decimal point. When displayed on a form, a **Currency** field also includes an additional list of choices for selecting the currency type. If there isn't a default value for the field, then empty **Currency** fields use the reference currency. Adding a value causes the field to use the session currency of the active user. See Change default currency decimal places for how to use two fraction digits. **Note:** After you assign this field type to a field, you can't change it to the **FX Currency** field type.
currency2	FX Currency	Yes	
data_array	Data Array	No	
data_object	Data Object	No	
data_structure	Data Structure	No	Data structure field where you can select one of the following data structures and enter values to organize particular information in the record: • String • Boolean • Integer • Decimal • Object • Array
date	Other date	Yes	Date field that displays a day, which can be selected with a calendar widget.

servicenow

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
datetime	Date/Time	No	Date/Time field that can be selected with a calendar widget. It's used for date-specific entries and is ideal for timestamps and scheduling.
day_of_week	Day of Week	No	
days_of_weeks	Days of Week	Yes	
decimal	Decimal	Yes	Number field with up to two digits after the decimal points. For example, 12.34.
document_id	Document ID	Yes	Document ID field that references a record on a table.
documentation_field	Documentation Field	Yes	Documentation field used by admins to add the label during table or field creation.
domain_id	Domain ID	Yes	Domain ID field that contains a reference to the domain.
domain_path	Domain Path	Yes	String field that contains a domain path, which is a series of three-character codes separated by a slash delimiter that uniquely identifies a domain. For more information on domain paths, see Domain Separation- Advanced Concepts and Configurations [KB0716268] article in the Now Support Knowledge Base.
due_date	Due Date	Yes	Date field that stores a date and time.
email	Email	Yes	
email_script	Email Script	No	Email script field that provides the ability to use field name notation **(${fieldname})**to represent the value of that field. This is a combination of text area and variable picker component and can be used to create text templates that have variables.
expression	Expression	No	
external_names	External Names	No	
field_list	Field List	Yes	
field_name	Field Name	Yes	Reference field for a field name in the table that is selected in a **Table Name**

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
			field type. Make this field dependent on the **Table Name** field.
file_attachment	File Attachment	Yes	File attachment field that can hold one file. All file types are supported.
float	Floating Point Number	Yes	
formula	Formula	No	
geo_point	Geo Point	No	Geo point field for holding longitude and latitude coordinates.
glide_action_list	Action List	No	
glide_date	Date	Yes	
glide_date_time	Date/Time	Yes	
glide_duration	Duration	Yes	
glide_list	List	Yes	
glide_precise_time	Precise Time	No	
glide_time	Time	Yes	
glide_utc_time	UTC Time	Yes	
glide_var	Var	No	
glyphicon	Glyph Icon	No	
graphql_schema	GraphQL Schema	No	
GUID	Sys ID (GUID)	Yes	
html	HTML	Yes	String field with a built-in HTML editor. It supports formatted text with rich text editor and is useful for descriptions that need enhanced readability.
html_script	HTML script	Yes	HTML script field that provides the ability to add variables in the form of field name notation **(${fieldname})** via a variable picker within the HTML editor. This field is used to create rich text content that also has dynamic values.
html_template	HTML template	No	
icon	Icon	No	String field that provides the user access to an icon picker in a form. By default, the favorites icon set is used.

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
image	Image	No	Image field used for uploading and embedding images.
index_name	Index Name	No	
insert_timestamp	Insert Timestamp	No	
int	Int	No	
integer	Integer	Yes	Number field with zero decimal points. **Note:** To search for an integer, input `=` followed by the integer in the search bar. For example, if you want to search for an Order of 100, input `=100`.
integer_date	Integer date	Yes	
integer_time	Integer time	No	
internal_type	Internal Type	No	
ip_addr	IP Address (Validated IPV4, IPV6)	Yes	
ip_address	IP address	Yes	Variable character field that stores IPv4 and IPv6 addresses. See IP address field type for more information.
journal	Journal	No	Journal field that accepts text entries and displays previous entries with a user name and timestamp.
journal_input	Journal Input	Yes	Journal input field that accepts text entries but doesn't display previous entries.
journal_list	Journal List	No	Journal list field that displays the contents of journal fields. You must specify the journal fields as the dependent fields. If a journal list field depends on more than one journal field, the entries are displayed chronologically.
json	JSON	Yes	
json_translations	JSON Translations	No	

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
language	Language	No	
long	Long	Yes	Long field that stores large whole numbers, which can be useful for high-value data. This field handles values that exceed standard integer limits, which is essential in high-volume contexts like serial numbers or counting operations. This field ensures that numeric data is stored accurately without overflow.
longint	Longint	Yes	Longint field that stores large whole numbers, which can be useful for high-value data. This field handles values that exceed standard integer limits, which is essential in high-volume contexts like serial numbers or counting operations. This field ensures that numeric data is stored accurately without overflow.
mask_code	Mask code	No	
metric_absolute	Metric Absolute	No	
metric_counter	Metric Counter	No	
metric_derive	Metric Derive	No	
metric_gauge	Metric Gauge	No	
mid_config	Mid Config	No	
month_of_year	Month of Year	No	
multi_small	Multi Small	Yes	Text area field that stores shorter texts like short descriptions and notes.
multi_two_lines	Two Line Text Area	Yes	
name_values	Name-Value Pairs	No	Name-value pairs field that maps text values. Each mapping is one-to-one, but a single **Name-Value Pairs** field can contain multiple mappings. Each mapping must use a unique name, and the name can't be empty. For example, you can use a **Name-Value Pairs** field to hold header information for a web service request. In this example, the name

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
			of each mapping is the header such as `Content-Type` and the value is the header value, such as `Application/json`. For information on scripting **Name-Value Pairs** fields, see Name-value pairs field type.
nds_icon	NDS Icon	No	
nl_task_int1	NL Task Int1	No	
order_index	Order Index	Yes	Order index field that defines the display order of records in lists or hierarchies. By assigning a numeric value to each record, it enables admins to control the visual sequence, which is helpful for logically organizing tasks, categories, or items. This field type helps improve user experience by making the interface intuitive and orderly.
password	Password (1 Way Encrypted)	Yes	Text field that stores passwords with one-way encryption. One-way encryption stores the password as a secure hash value that can't be decrypted.
password2	Password (2 Way Encrypted)	Yes	Text field that stores passwords with two-way encryption. Two-way encryption stores the password as a secure encrypted value that can be decrypted programmatically within the instance. You can use Password 2 encryption with form variables. To encrypt text fields on forms, see Password2 encryption with KMF. The length for password2 field values must be at least 255 characters.
percent_complete	Percent Complete	Yes	Decimal field that renders a percent complete bar in lists. You can convert any existing **Decimal** field to a **Percent Complete** field.
ph_number	Phone Number	Yes	
phone_number	Phone Number (Unused)	Yes	

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
phone_number_e164	Phone number (E164)	Yes	String field that provides E164-compliant formatting and validation for telephone numbers.
price	Price	Yes	Currency field that enables control over conversions and display. See Price fields for more information.
properties	Properties	No	
radio	Radio	Yes	Radio field that displays mutually exclusive options as radio buttons, letting users select only one choice. It's suitable for simple, definitive decisions, like priority level or yes/no responses. This field type helps improve usability by presenting clear options for single-selection inputs.
records	Records	No	
reference	Reference	Yes	Query field that displays records from another table, which is crucial for relational data structures.
reference_name	Reference Name	No	
related_tags	Related Tags	No	
reminder_field_name	Reminder Field Name	No	
repeat_count	Repeat Count	No	
repeat_type	Repeat Type	No	
replication_payload	Replication Payload	No	
schedule_date_time	Schedule Date/Time	No	
schedule_interval_count	Schedule Interval Count	No	
script	Script	Yes	Text field that accepts JavaScript code input and provides controls, such as syntax checking and formatting. It also provides a list of fields and server APIs. You must specify a dependent field that references the table name for the list of fields.
script_client	Script Client	No	

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
script_plain	Script (Plain)	Yes	Text field that accepts JavaScript code input and provides controls, such as syntax checking and formatting.
script_server	Script Server	No	
short_field_name	Short Field Name	No	
short_table_name	Short Table Name	Yes	Text field that stores a compact version of a table name, often used for efficiency in system configurations, URLs, or logs. It helps map the full table names to shorter identifiers, mainly for internal system use rather than direct user interaction.
simple_name_values	Name-Value Pairs	Yes	
slushbucket	Slushbucket	Yes	
snapshot_template_value	Snapshot Template Value	No	
source_id	Source ID	No	
source_name	Source Name	No	
source_table	Source Table	No	
string	String	Yes	String field that holds a string of characters. Strings of up to 255 characters appear in a single-line text box. Anything 256 characters or over appears as a multi-line text box. **Note:** The database may translate the value that you provide in the **Max length** field to the closest matching database field type. For example, a max string length of 80 maps to the nearest database field type of VARCHAR(100). For Oracle instances, users aren't able to increase the max length of a string field to a value greater than 4000 through the application UI. Changes greater than 4000 require

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
			the CLOB data type in Oracle. If you require a field to be larger than 4000 characters, log an incident to request the change.
string_boolean	String Boolean	No	
string_full_utf8	String (Full UTF-8)	Yes	String field that contains UTF-8 character encoding. This field type is restricted to the same maximum length that can be defined within the application UI as the String data type. Emojis are supported in the String (Full UTF-8) field type.
structure	Structure	No	
sys_class_name	System Class Name	Yes	System field that holds the table name.
sys_class_path	System Class Path	No	
sysevent_name	System Event Name	No	
sysrule_field_name	System Rule Field Name	No	
table_name	Table Name	Yes	String field that lets you select a table. If you use the **Field Name** field type, add this field type and make the **Field Name** field dependent on the **Table Name** field. If you're in a scoped application, you can view only the tables of that scope.
template_value	Template Value	Yes	
time	Time	No	Time field that stores a time in the database as a **Date/Time** field, but only the time part of this field is used.

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
			Note: This field stores the date as 1970-01-01, which doesn't fall within the dates for Daylight Savings Time (DST). As a result, the time isn't adjusted automatically for DST. If a custom feature uses time as a start time, the feature must adjust for DST, if necessary.
timer	Timer	Yes	
translated	Translated	No	
translated_field	Translated Field	Yes	
translated_html	Translated HTML	Yes	HTML text field that displays translations based on the user's language.
translated_text	Translated Text	Yes	Text field that displays translations based on the user's language.
tree_code	Tree Code	No	
tree_path	Tree Path	No	
url	URL	Yes	String field that is a selectable URL field when locked. When configuring, there's an option to allow/disallow links with or without HTTPS. **Note:** In the mobile UI, this field appears as a multi-line text field. The field saves as a single line with the line breaks removed.
user_image	Image	Yes	
user_input	User Input	No	User input field that captures input provided by users, typically in the context of workflows, surveys, or other interactive processes. This field enables users to enter text, select options, or provide other types of input that can then be processed or stored for further use. The data collected through **User Input** fields is often used to guide

Field types (continued)

Field type	Platform name	Supported in configurable workspace	Description
			decision-making in workflows or to gather information for reporting and analysis. It's a flexible field type that can be customized to fit various use cases where user interaction is required.
user_roles	User Roles	Yes	
variable_conditions	Variable Conditions	No	
variable_template_value	Variable Template Value	No	
variables	Variables	No	
version	Version	Yes	
video	Video	No	Video field for uploading and embedding video.
week_of_month	Week of Month	No	
wide_text	Wide Text	Yes	
wiki_text	Wiki	Yes	String field with a built-in wiki text editor that accepts a simplified version of standard wiki text formatting.
wms_job	WMS Job	No	
workflow	Workflow	Yes	List field that displays a stage in a workflow.
workflow_conditions	Workflow Conditions	No	
xml	XML	Yes	

Choice list field type

A choice list is a type of field that lets the user select from a pre-defined set of choices.

Administrators can define the available choices and customize the behavior and appearance of choice lists.

Incident state choice list

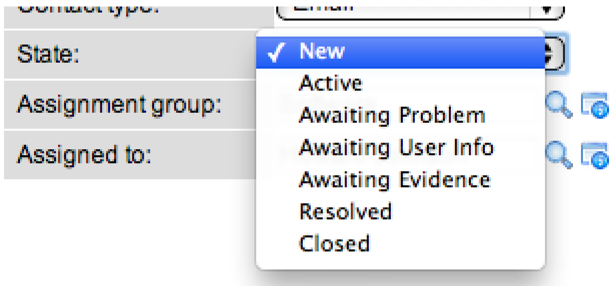

Note:
Choice lists do not support a one to many relationship. Only one choice from a choice list can be selected at a time.

Choice list security

You can use the personalize_choices security role to enable non-administrators modify Choice elements options on all tables.

If more granular control is desired, you can also create a custom ACL (security rule) governing the personalize_choices operation either for a particular field or for all fields (.*) on a particular table. However, access to the personalize_choices operation on a particular field does not confer the ability to add new choices for that field.

To be able to create new choices for a particular field, an ACL that grants personalize_choices access for that field is required. For example, to give the hris_admin role the ability to personalize only the Category field for Human Resources KB articles, you need an ACL granting personalize_choices access to the hris_admin role on the Category field of the Knowledge (kb_knowledge) table.

There are predefined ACLs granting both types of access to the personalize_choices security role, for all fields on all tables. The personalize_choices security role also has read, write, and delete access to the sys_choices table. However, this additional access is not required when making just the Personalize Choices functionality available on a granular basis.

Values to associate with choice labels for scripting

When you write a script that references a choice list, you need to know the value that is associated with each choice.

For example, to check whether the incident_state field is active, you could not use the condition `current.incident_state == "active"` because the value associated with the choice labeled **Active** is the integer **2**. Instead, you would use the condition `current.incident_state == 2`.

The **Type** field on the choice list dictionary entry determines the data type of the values.

To determine the value associated with a choice, right-click the field label and select **Show Choice List**, and then locate the choice for which you need to know the value.

Show choices list values

		Table ▲	Element	Language	Value	Label	Inactive	Sequence
		incident	state	en	1	New	false	1
		incident	state	en	2	Active	false	2
		incident	state	en	3	Awaiting Problem	false	3
		incident	state	en	4	Awaiting User Info	false	4
		incident	state	en	5	Awaiting Evidence	false	5
		incident	state	en	6	Resolved	false	6
		incident	state	en	7	Closed	false	7

The **-- None --** option may not have a sys_choice record associated with it. A choice list field set to **-- None --** evaluates to these values, depending on the script context as listed below.

- For client-side scripts,such as client scripts: "" (empty string)
- For server-side scripts, such as business rules: "0" (string of the number zero)

Related topics

System dictionary

Integer values for default choice lists

Choice provide four default values.

Some common choice lists use integer values that do not match the string labels. For example, the Problem table uses these default values for the **State** field.

State field default values

Value	Label
1	Open
2	Known Error
3	Pending Change
4	Closed/Resolved

These integer values are also used in several default business rules. For example, a business rule on the Incident table sets the active flag to false when the **State** field changes to **7**, which is the default value for the **Closed**. If you change the values of your Incident state options, this business rule may no longer behave as desired or expected.

On the Incident table, the **Active**, **State**, and **Incident state** fields are affected by the following default business rules.

Default business rules

Business rule	Description
mark_closed (incident)	If the incident_state changes to **7** (**Closed**), the **Active** field is set to false

Default business rules (continued)

Business rule	Description
incident reopen (incident)	If the incident_state is less than **7** (**Closed**) and the **Active** field is false, the **Active** field is set to true
mark closed (task)	If the state changes to either **3** (**Closed Complete**) or **4** (**Closed Incomplete**), the **Active** field is set to false
task closer (task)	If the **Active** flag changes from true to false and the state is neither **3** (**Closed Complete**) nor **4** (**Closed Incomplete**), the state is set to **3** (**Closed Complete**)
task reopener (task)	If the **Active** field changes from false to true and the state is either **3** (**Closed Complete**) or **4** (**Closed Incomplete**), the state is set to **1** (**Open**)

Note:
Notice that these business rules do not change incident_state based on a change to either the **Active** field or the **State** field. Changes to incident_state drive the other two fields, not the other way around.

Configure state field choice values

State fields are a subset of choice list fields. Keep the following information in mind when you configure choice values for the state field.

- Use a negative value to add a new active state field.
- Search for and study the business rules that use a state number filter on the **Script** and **Conditions** fields. You can use the Debug tool to trace the order of the business rule execution.
- New values representing inactive states should have a value above 8.

You can define any of the following attributes for a state field by configuring the dictionary. If the attributes are not defined, the system uses the default values. The TaskStateUtil API uses the following attributes. For more information on the TaskStateUtil API, see TaskStateUtil.

Related Attributes

Attribute	Definition
close_states	Semicolon delimited list of state values that are inactive, used to identify whether the task should be set to active or inactive. This is a required attribute to use the TaskStateUtil functionality.
default_close_state	Optional attribute to define the state value of the default close state if you want to define business rules that automatically close a task. Defaults to 3, typically Closed Complete if attribute is not defined.
default_work_state	Optional attribute to define the state value of the default working state if you want to define business rules that automatically set a task for working. Defaults to 2, typically Work in Progress if the attribute is not defined.

State modification examples

Follow these examples for modifying the states of incidents and change requests.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Choice Lists**.
2. At the top of the list, construct a list filter like the following:
   - **Table**: incident
   - **Element**: incident_state
3. Run the filter.

   Notice that the **Closed** state has a value of **7** and the **Resolved** state has a value of **6**. Any state greater than or equal to **7** is assumed to be inactive. Therefore, you should use a positive integer greater than **7** if you want to add a new inactive-type of state. Use a negative value like **-1** or **-2** if you wish to add a new active-type of state field, such as **Awaiting Vendor**.
4. Navigate again to **System Definition** > **Choice Lists**.
5. At the top of the list, construct a list filter like the following:
   - **Table**: change_request
   - **Element**: phase_state
6. Run the filter.

   Notice that the **Complete** state has a value of **8**. Any state greater than or equal to **8** is assumed to be inactive. Therefore, you should use a positive integer greater than **8** if you want to add a new inactive-type of state, such as **Cancelled**. Use a negative value like **-1** or **-2** if you wish to add a new active-type of state field, such as **Pending**.

Troubleshoot change states and business rules

Business rules in the system make assumptions about state values. You can troubleshoot business rules to see the order in which they run and see how it affects changes you make to **State** field values.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Business Rules**.
2. Construct a filter like this one to view the scripts and conditions that pertain to the Resolved incident_state of 6 or the Closed incident_state value of 7:

   The **Script** field contains 7 OR the **Condition** field contains 7 OR the **Script** field contains 6 OR the **Condition** field contains 6 AND the **Table** field is incident AND the **Active** field is true.

What to do next
See Debug Business Rule for information on how to trace the order of business rule execution. You can click **Debug All**, resolve an incident, and then check the trace at the bottom of form to watch the business rules execute. These two line examples show that the mark_closed business rule code is entered `==>` and then exited `<==`.

```
==> 'mark_closed' on incident
<== 'mark_closed' on incident
```

View choice list definitions

The Choice Set [sys_choice_set] table contains a record for every field that uses a choice list.

Before you begin

Role required: personalize_choices

Note:
The personalize_choices role must be explicitly granted to the user; it cannot be an ACL.

About this task

The choice set record is associated with an application file, which allows update sets and team development to track and transfer all choices for a field in a single update record.

Choice list values allow a maximum length of 40 characters. The range of allowable numerical values is [-999, 999].

Procedure

1. Right-click the choice list field label and select **Show Choice List**.
   To view other choice list values, modify the filter at the top of the list.

   Note:
   When you use an ACL to grant personalize_choices on a particular field, **Show Choice List** is not available. It is only available if you explicitly grant the role to the user. **Configure Choices** continues to appear regardless of whether it is an ACL or an explicitly granted user role.

2. Review the items in the list.

   Warning:
   Do not add new choices to the list. To add new choices to a choice list field, use the **Configure Choices** option.

Define an option for a choice list

You can personalize the options that are available in a choice list.

Before you begin

Role required: personalize_choices

Procedure

1. Navigate to a form where the field appears.
2. If the choice list is dependent on another field, enter the choice value that the options depend on.

   For example, on the incident table, the **Subcategory** is dependent on the **Category**. To customize which subcategory choices are available for the hardware category, select **Hardware** in the **Category** field.

3. Right-click the field label and select **Configure Choices**.
4. Use the slushbucket to rearrange the order, add, or remove items or to create new items.
5. Click **Save**.

To dynamically add items to a choice list, use the `addOption` GlideForm method .

Note:
Some business rules may be affected by changes to choice list options (for example, default Incident states).

Reuse a choice list

After defining a set of choice list values, you can reuse the values for another field in a different table.

Before you begin

Role required: personalize_choices

Procedure

1. Right-click an existing choice field (Field A) and select **Configure Choices**.
2. Add the desired choice list values in the **Choices** related list.
3. To reuse the choice list values for another field (Field B) in a different table, right-click the label for Field B and select **Configure Dictionary**.
4. In the **Choice table** field, select the table where Field A resides.
5. In the **Choice field** field, select Field A.

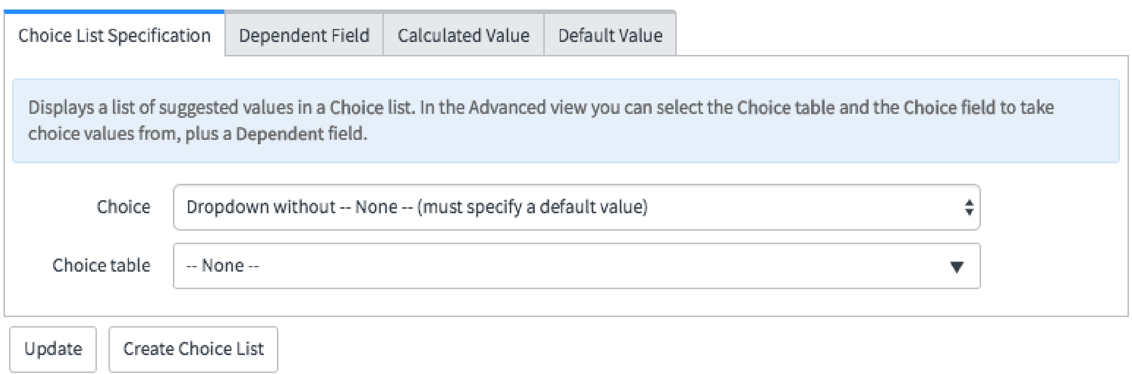

6. Click **Update**.

   The choice list values defined on Field A are displayed on Field B. When you add or remove choice list values on Field A, those changes are also reflected on Field B. After you specify a choice table and a choice field, the field no longer uses the defined choice list.

Remove the None option from a choice list

You can remove the **None** option from a choice list if it is not necessary.

Before you begin

Role required: personalize_dictionary

Procedure

1. Navigate to a form where the field appears.
2. Right-click the field label and select **Configure Dictionary**.
3. Change the **Choice** field value to **Dropdown without -- None -- (must specify a default value)**.

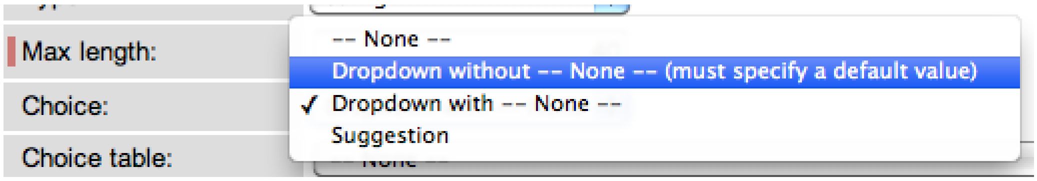

4. Ensure that the **Default** field is populated to determine which choice is displayed by default.

   Note:
   If the field is dependent on another field, the **-- None --** option remains available.

Change the None display value for a choice list

You can change the default display label of the **None** option for a choice field.

Before you begin

Role required: personalize_choices

Note:
The personalize_choices role must be explicitly granted to the user; it cannot be an ACL.

Procedure

1. Navigate to a form on which the field appears.
2. Right-click the field label and select **Show Choice List**.
3. Click **New**.
4. Complete the form.

Choice form fields

Field	Value
Table	Select the table.
Element	Leave the name of the field that is automatically populated.
Language	Enter ISO language code for the label.
Sequence	Leave empty. This field determines the order.
Inactive	Leave cleared.
Label	Enter the label to appear in the choice list. You can use JavaScript, including calls to script includes, to define the label. For example, the JavaScript label in the following example changes the **-- None --** value of the **Time zone** choice list in a user record to use the time zone value of the instance.

Field	Value
	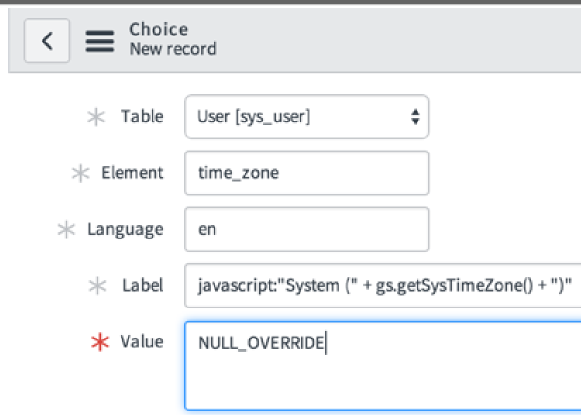
Value	Enter `NULL_OVERRIDE`. **Note:** You must enter `NULL_OVERRIDE` as the value, or the new label appears in addition to the **-- None --** option.
Dependent value	Leave blank.
Hint	Leave blank. **Note:** When field type is set to **List (Glide List)** the hint won't display.

5. Click **Submit**.

Delete all choice list options

You can delete all choices for a choice field from the Choice Set record.

Before you begin

Role required: personalize

About this task

You may want to use this method when you are developing a new application and the business requirements change. If you are updating a choice list that is already in use, consider deactivating the options you no longer use to avoid conflicts with existing data or scripts that may rely on the previous options.

Procedure

1. In the navigation filter, enter `sys_choice_set.list` and press Enter.
2. Open the choice set record for the field.

   For example, to locate the choice set for the incident subcategory, filter by **[Table] [is] [incident] AND [Element] [is] [subcategory]**.

3. Check the box beside the choice set record to delete and select **Delete** from the Actions choice list below the list.

4. Click **Delete** in the confirmation window.
   All choices for the field are deleted.

Create a choice list for another field type

You can create a choice list for a field with another type, such as an integer, string, or reference field.

Before you begin
Role required: personalize_dictionary

About this task
You can use this configuration to standardize data entry and limit available options for a field while still maintaining the original field type.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the dictionary entry for the field.

   Note:
   Reference fields with a large number of records in the reference table cannot be converted to look like choice fields. A reference field with too many records reverts to looking like a reference field.

3. Change the **Choice** value to **Dropdown with --- None ---** or **Dropdown without --- None --- (must specify a default value)**.
4. Right-click the form header and select **Save**.
5. Click **Create Choice List**.

   - The **Choices** related list appears on the dictionary entry form.
   - If records on the table contains data for the field, a choice list value for each unique field value is created. For example, if three records exist on the table and each record has a unique value in the field, then three choices are created.
   - If no data exists in the field, a choice list value of **-- New choice --** is created.

Display invalid choice list values

By default, inactive or invalid choice list values appear in blue text instead of black. You can disable the color indicator for invalid choices.

Before you begin
Role required: admin

About this task
In the following example, the **Network** category has been deactivated, so it appears in blue for records that still contain the inactive value.

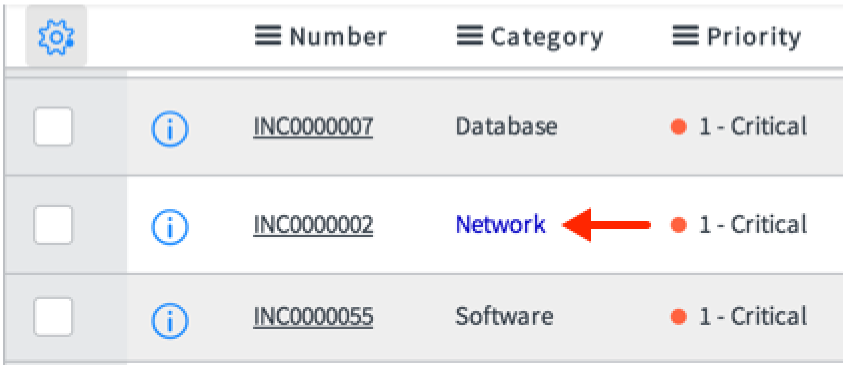

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Clear the check box for the **Display missing choice list entries** property.

Add search option to a choice field

Add a search field to choice fields that have a long list of options.

Before you begin

Role required: admin

Procedure

1. Navigate to a form that contains choice fields.
   For example, incident.
2. From a choice field on the form, for example State, right-click the field and select **Configure Dictionary**.
3. Switch to the advanced view for the dictionary entry form using the context menu by navigating to **View** > **Advanced**.
4. In the Attributes field, type `is_searchable_choice=true`.
   If there are other entries in the attributes field, use a comma to separate the entries.
5. Update the Dictionary Entry form and reload the page containing the choice list.

Result

The choice list contains a search field that users can use to filter the list of choices.

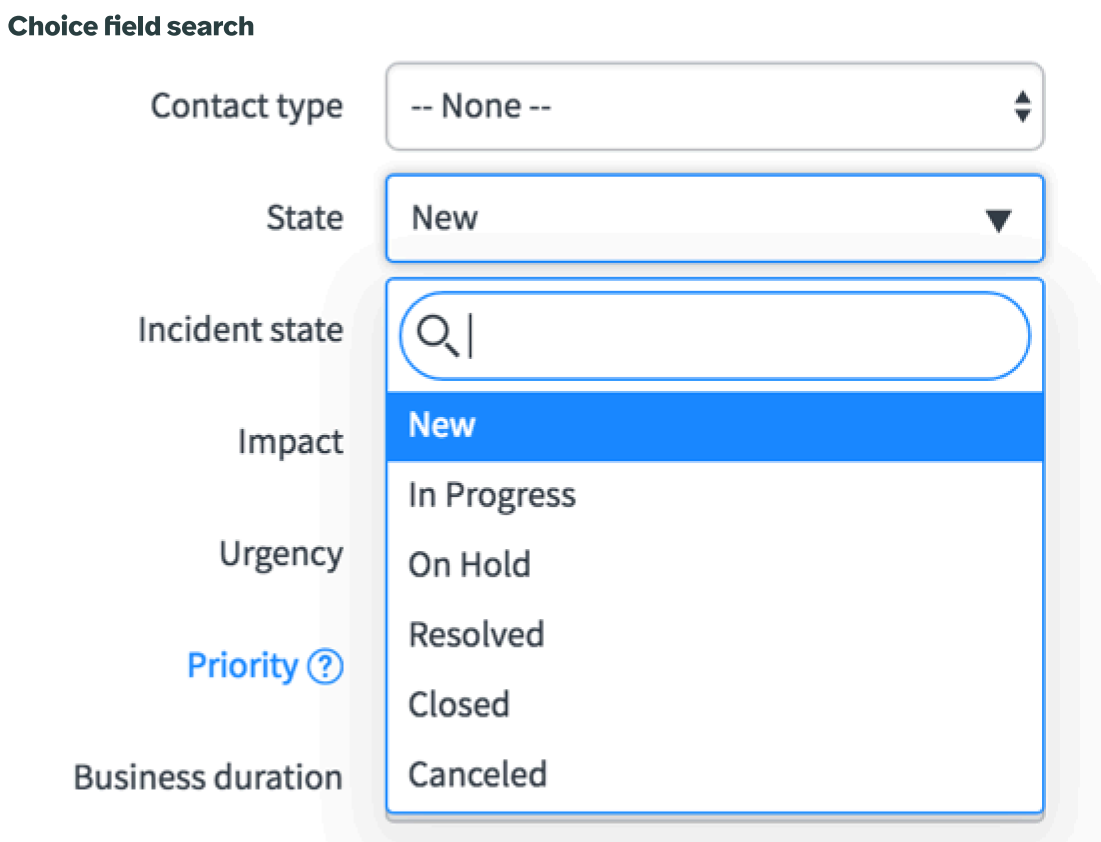

Condition field types

A condition field specifies when to run business logic such as a business rule or workflow.

There are two types of condition field.

Condition field types

Condition field type	Description
Condition string	A text field that accepts a plain JavaScript condition statement. The system validates the condition syntax for correctness before an update.
Conditions	A field that adds a condition builder to a form. Condition builders require specifying a dependent field whose values the system uses to display choice list options. Typically, the dependent field is the **Table** field.

The system evaluates both types of condition field to determine if the conditions are true or false. When true, the system runs the business logic. When false, the system ignores the business logic.

To find dictionary attributes that affect condition fields, see Altering tables and fields using dictionary attributes.

Add the condition count to a condition field

The condition count widget can be activated on condition fields to display a preview of the records that would meet the current set of conditions. For fields where the condition count is activated, the number of records that match the conditions will automatically display. The count

refreshes if the field the condition field depends on, such as Table, is changed. If the Table field is left blank, the widget is hidden.

Before you begin

Role required: personalize_dictionary

Procedure

1. Right-click the field label and select **Configure Dictionary**.
2. Add `show_condition_count=true` to the **Attributes** field.
3. Submit.

Result

In this example, a condition is run on the Incident table to look for incidents where **Category** is **Network**:

Condition count

The list shows that there are 112 records with **Network** as its category.

To refresh the preview, click the update count icon ( ).

To view details of the results, click the number of records to open the list view of the results:

Condition count results

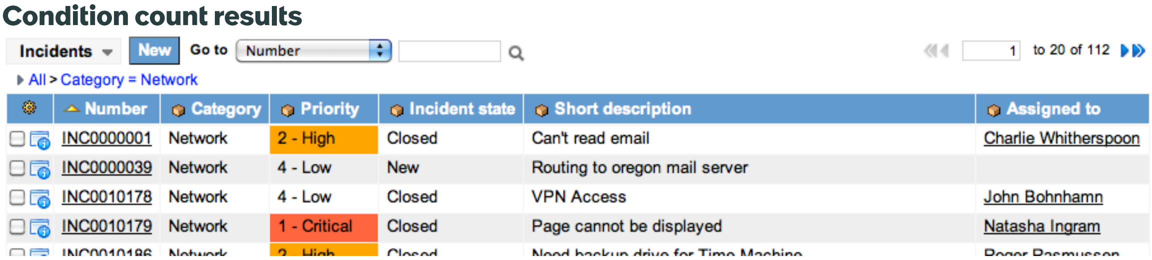

Update a conditions field to use condition builder v2

In Core UI, you can update a conditions field to display the version 2 condition builder.

Before you begin

Role required: admin

About this task

You add a dictionary attribute to the Conditions field to enable condition builder version 2 (v2).

Condition builder v2 configured in Approval Rules form

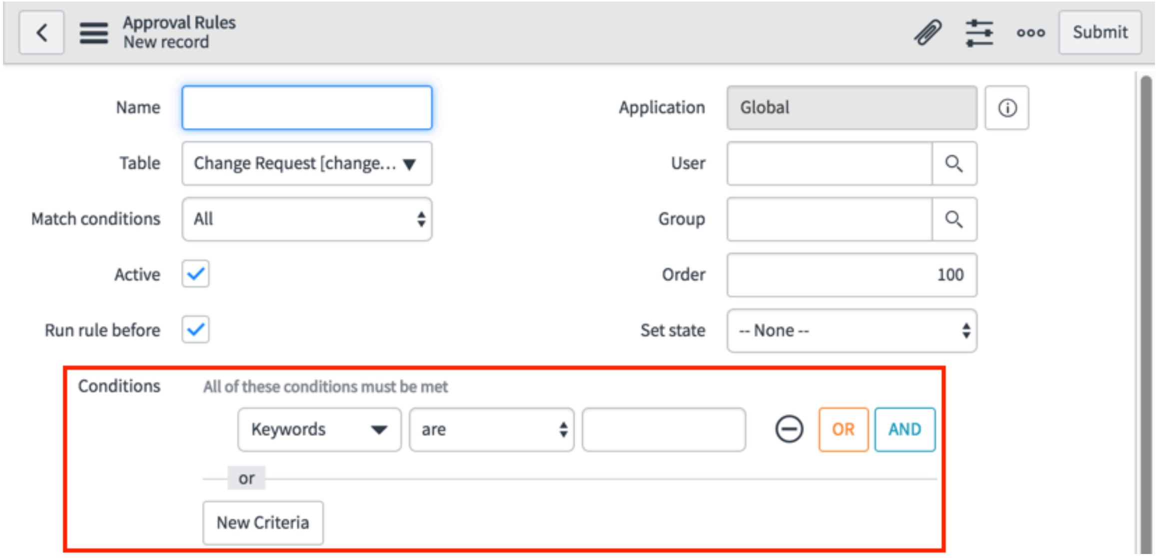

Procedure

1. Open the form with the condition builder to configure.
2. Right-click the **Conditions** label and click **Configure Dictionary**.
3. In the **Attributes** field, enter `condition_builder=v2`.
   If attributes exist, add this string at the end separated by a comma.
4. Click **Update**.
   The form reloads with the v2 version of the condition builder.

Database field type

Several field types are available in the system.

This table shows field types and corresponding MySQL database types. Typically, it isn't necessary to perform any actions at the database level. To learn about changing a field type, see System dictionary.

Database field types

Field types	Options	Dictionary XML type	MySQL DB type
String	small	string	VARCHAR(40)
String	medium	string	VARCHAR(100)
String	large	string	MEDIUMTEXT
String	extralarge	string	MEDIUMTEXT
Decimal		decimal	Decimal (15,2) older instances may have (12,2)
Integer		integer	Integer
True-False		boolean	TINYINT(1)
Date		glide_date	DATE
Date-Time		glide_date_time	DATETIME

Database field types (continued)

Field types	Options	Dictionary XML type	MySQL DB type
Time		glide_time	DATETIME
Duration		glide_duration	DATETIME
Choice		string	VARCHAR(40)
Suggestion		string	VARCHAR(40)
Journal		journal	MEDIUMTEXT
Reference	<reference table>	reference	VARCHAR(32)
List		glide_list	MEDIUMTEXT
Url		url	MEDIUMTEXT
Image		user_image	VARCHAR(40)
Due-Date		due_date	DATETIME

Note:
An example of option types is as follows:

- A string field defined in the platform with a max length of 40 characters → VARCHAR(40).
- A string field defined in the platform with a max length of x, where x = 41-255 characters, inclusive → VARCHAR(x) [examples: String 42 is VARCHAR(42) and String 211 is VARCHAR(211)].
- A string field defined in the platform with a max length of 256 and upwards → MEDIUMTEXT.

Dictionary entry data types

You can only change a dictionary entry's data type when the change does not result in data loss. Use the following guidelines to change a dictionary entry's data type.

Valid data type changes

Condition	Restrictions
The field is empty in all table records.	None. You can convert an empty field from any data type to another without restriction.
The table contains existing data for the field.	You can only convert between logical data types that map to the same physical data type in the database. For example, you can convert a glide duration to a glide datetime since both logical data types map to the DATETIME physical data type in the database.
The field is a string field you are converting to another type of string field.	You can change between string-based data types as long as length changes do not cause any data loss from truncation. For example, you can change from a MEDIUM database type to a VARCHAR(100) database type if none of the existing data is greater than 100.
The field is a string field you are converting to a	You can only convert a string field to a GUID if all of the exiting data in the field are already **Sys ID** values.

Valid data type changes (continued)

Condition	Restrictions
globally unique ID (GUID).	
The field is a GUID field you are converting to a string field.	None. You can convert a GUID field to a string field without restriction.

Document ID field

You can create document ID fields to reference any record on any table.

Before you begin
Role required: personalize_dictionary

About this task
In comparison, a reference field references a record on a specific table. To reference records from any table, two fields need to work together—one to store the table reference and another to store the record reference.

You can select the document ID type when creating a new field.

Procedure
1. Navigate to the form view for the table.
2. Right-click the header and select **Configure** > **Form Layout**.
3. Create a field to store the table name and click **Add**.
   - **Name**: Any label (for example, `Model table`)
   - **Type**: String
4. Create a field to store the record reference and click **Add**.
   - **Name**: Label for your document ID field (for example, `Model ID`)
   - **Type**: Document ID
5. Click **Save**.
6. Right-click the form header and select **Configure** > **Dictionary**.
7. Open the dictionary entry for the document ID field.
8. Under **Related Links**, click **Advanced view**.
9. In the **Dependent** field, enter the column name of the table reference field (for example, `u_model_table`).
10. **Optional:** In **Attributes**, add the *`show_all_tables`* dictionary attribute to display system tables.
11. Click **Update**. When users click the reference lookup for the document ID field, a dialog appears that allows them to select the table and then the record. The sys_id of the selected record is stored in the document ID field and the table name is stored in the table reference field.

12. **Optional:** Configure the form to remove the table reference field.

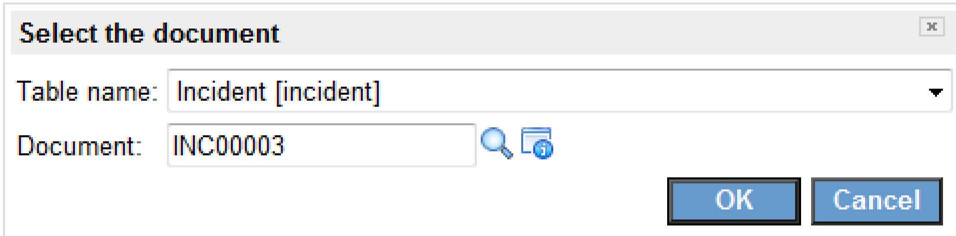

Related topics

Configuring the form layout

Altering tables and fields using dictionary attributes

Function field

Create function fields and scripts in the ServiceNow AI Platform to perform common database transformations and calculations.

Regular fields store a value in the database. Instead of storing data, a function field displays the results of a database query. Function fields do not have an associated database column. Instead, function fields generate a value based on simple computations of other fields and constants. They can be used like any other fields in the system: in forms, lists, query conditions, reports, data visualizations, and so on.

> **Warning:**
> A function field that references another function field in the function will result in an error. Function fields don't hold data or evaluate in run-time.

Example: Use case

You want to identify all incidents with a probably useless short description of fewer than 10 characters.

Instead of querying all records to determine which record meets a given criteria, create a function field or function script that retrieves only the records that meet the criteria.

For a simple computation, it is better to use a function field or function script rather than store a computed value. The value is always calculated on retrieval. Another benefit of using function fields is that the database server performs the transformation rather than the application node.

Function fields versus calculated fields

The values of the calculated fields are stored in the database. The values of function fields or function scripts are not stored in the database, but are calculated at the time of retrieval. Function field values are always up to date.

Function fields versus business rules

Use function fields or scripts when business logic is based on a simple transformation of one or more existing fields. Function fields and scripts let you implement business logic without storing and maintaining the result of a transformation.

Function fields versus filters, query strings, and Rhino

Rather than calling Rhino to perform string transformations, you can use a function field or function script to perform the transformations for you. Function fields and function scripts are more efficient and result in up-to-date values that do not have to be stored or maintained.

Function field limitations

Because function field calculations are not stored in the database, some operations are not relevant.

- Function fields cannot be directly audited or indexed.

  Note:
  To index a function field as if it were a regular field, make sure the individual fields used by the function are indexed. Or for best performance, make sure that there is a composite index including all fields that are used in the function.

- Function fields cannot be encrypted, since a function field value is never stored in the database.
- Function fields cannot be converted to regular fields or vice versa.
- Function fields that you create in the Reporting UI do not support dot-walking. For more information, see Configure function fields in Reporting.
- Security is evaluated on the components of the function and on the calculated value of the field. When used in visualizations, sections that contain information that the user is not permitted to see are hidden from the user.
- Field function names must be unique.

Defining platform functions

Platform functions can be defined in two ways:

1. Use the following application programming interfaces (APIs) to build and use functions in a script.

   Function APIs

APIs	Description
Scoped GlideDBFunctionBuilder	Construct the function to perform a SQL operation.
GlideRecord - addFunction(Object function)	Apply the function to a GlideRecord.
GlideDBFunctionCaseBuilder - Global	Build case statements.

2. Create a field that holds the function definition, as shown in the following example.

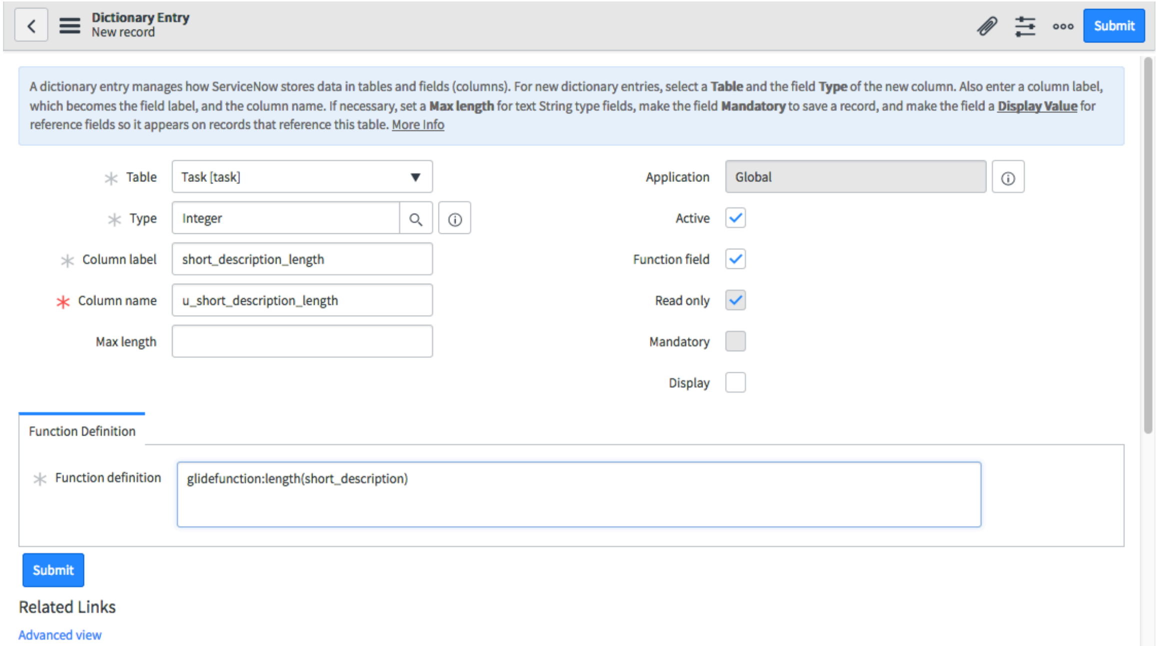

Select the **Function field** check box on the Dictionary Entry [sys_dictionary] form. This action specifies that the field runs a function rather than stores a value.

glidefunction operations

Function fields can perform the following operations using the `glidefunction:<operation>` syntax. When providing a field as an argument, you can dot-walk to related fields. For example, `cmdb_ci.name`.

Note:
The function field type must be compatible with the return type of the glide function. For example, integer, longint, string, and so on. If the returned type cannot be converted to the dictionary type, you get exceptions.

Function field operations

Operation	Description	Example
`add()`	Takes two number fields as input, adds them, and returns the results as a field value. This function also takes numerical values for either input. Place numerical values in single or double quotation marks.	*`glidefunction:add(child_incident`* *`parent_incident)`* Returns 6 if the incident has five child incidents and one parent incident. Possible return types: Decimal, Floating Point Number, Large Whole Number, Whole Number
`coalesce()`	Takes any number of comma-separated fields as input and returns the first non-empty value.	*`glidefunction:coalesce(closed_at`* *`resolved_at,`* *`sys_updated_on)`* If the value of closed_at is empty, the function returns

Function field operations (continued)

Operation	Description	Example
		the value of resolved_at. If the value of resolved_at is also empty, the function returns the value of sys_updated_on. Return type: Text
`concat()`	Takes any number of comma-separated fields and constants as input, concatenates the input, and returns a single string as a field value.	`glidefunction:concat(incident_nu '/', short_description)` Returns "INC0001 / My client needs a new laptop." if the value of the number field is 'INC0001' and the short_description is 'My client needs a new laptop'. Return type: Text
`datediff()`	Takes two date/time fields as input, calculates the difference between the dates in days, minutes, and seconds, and returns the results as a duration field value.	`glidefunction:datediff(closed_at sys_created_on)` Returns the duration of an incident from the creation date to the close date. Example result: 10 days, 8 hours 23 minutes 11 seconds Return type: Duration
`dayofweek()`	Takes two arguments: A date field and a constant of either `'1'` (week starts on Sunday) or `'2'` (week starts on Monday). Returns the results as an integer value that represents the day of the week. The `dayofweek()` function uses UTC dates, but adjusts comparison values based on the instance's time zone.	`glidefunction:dayofweek(resolved '1')`. If resolved_at occurs on a Wednesday, returns 4 if the integer is 1 and returns 3 if the integer is 2. Return type: Whole number
`distance_sphere()`	Takes two locations and returns the distance between them in meters.	• `glidefunction:distance_sphere( '51.473584', '80.206458', '25.767242')` Returns the distance between the two sets of longitude and latitude pairs.

Function field operations (continued)

Operation	Description	Example
		• `glidefunction:distance_sphere(u` `'80.206458',` `'25.767242')` Returns the distance between the u_location geo point field and the longitude and latitude coordinates. • `glidefunction:distance_sphere(l` `location2)` Returns the distance between the location1 and location2 geo point fields. • `glidefunction:divide(distance_s` `'-0.189937',` `'51.473584'),` `'1000')` Returns the distance between the location1 geo point field and London in kilometers. Return type: Decimal
`divide()`	Takes two number fields as input, divides the first by the second, and returns the result as a field value. This function also takes numerical values for either input. Place numerical values in single or double quotation marks.	`glidefunction:divide(u_num2,u_nu` Returns 5 if num2 = 10 and num1 = 2. Possible return types: Decimal, Floating Point Number, Large Whole Number, Whole Number
`greatest()`	Takes two or more values as input and returns the greatest value of the list of arguments.	`glidefunction:greatest(10,100,10` Returns 1000 Possible numeric return types: Decimal, Floating Point Number, Large Whole Number, Whole Number This function also works with date and string fields.
`least()`	Takes two or more values as input and returns the lowest value of the list of arguments.	`glidefunction:least(10,100,1000)` Returns 10

Function field operations (continued)

Operation	Description	Example
		Possible numeric return types: Decimal, Floating Point Number, Large Whole Number, Whole Number This function also works with date and string fields.
`length()`	Takes a string field as input, calculates the field length in characters, and returns the results as a field value.	`glidefunction:length(short_descr` Returns 37 if short_description = "This application is performing a test". Return type: Whole number
`multiply()`	Takes two number fields as input, performs the multiplication, and returns the results as a field value. This function also takes numerical values for either input. Place numerical values in single or double quotation marks.	`glidefunction:multiply(u_num1, u_num2)` Returns 48 if num1 = 8 and num_2 = 6. Possible return types: Decimal, Floating Point Number, Large Whole Number, Whole Number
`position()`	Takes two text fields or two text fields and a whole number as input. One or both of the text fields can also be strings. 1. Returns the position of the first occurrence of the string of the first text field in the second. 2. If a whole number is specified, this function returns the position of the second text field after the position of the whole number. Returns 0 if the first text field is not present in the second (after the position of the whole number if specified).	• `glidefunction:position('e',shor` If the value of the short_description field is "This computer is performing an upgrade", returns 12. • `glidefunction:position('e',shor` If the value of the short_description field is "This computer is performing an upgrade", returns 19. Return type: Whole number
`substring()`	Takes a text field and two whole numbers as input. Returns the first instance of a string that starts at the position of the first whole	`glidefunction:substring(short_de '7', '2')` If the value of the short_description field is

Function field operations (continued)

Operation	Description	Example
	number and is the length of the second.	'We're going to the store', returns 'go'. Return type: Text
subtract()	Takes two number fields as input, subtracts the second from the first, and returns the result as a field value. This function also takes numerical values for either input. Place numerical values in single or double quotation marks.	*glidefunction:subtract(u_num1, u_num2)* Returns 2 if num1 = 8 and num_2 = 6. Possible return types: Decimal, Floating Point Number, Large Whole Number, Whole Number

Create a function field to perform database functions

Create a field that displays the results of a database function, such as a mathematical operation, field length computation, or day of the week calculation. Test in a subproduction instance before deploying to a production instance.

Before you begin

Role required: personalize_dictionary or admin.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Click **New**.
3. Select **Function field**.
4. Fill in these fields on the form.

Fields for function fields

Field	Description
Table	Table to which the function field is added.
Type	Species the type of field: string, number, date, and so on.
Column label	Defines the label for the column.
Column name	Defines the name of the column. Created automatically based on **Column label**.
Max length	Defines the maximum length of the return value.
Function field	If selected, creates a field that displays the results of a database function, such as a mathematical operation, field length computation, or day of the week calculation.

Field	Description
	Once the new function record is saved, you cannot clear the check box to make the field a regular field.
Function definition	Defines the function that the field performs. Starts with `glidefunction:`, followed by the operation to be performed (for example, `concat`), followed by function parameters. Constants must be enclosed in single quotes. For example, the following function definition creates a field that shows the short description, followed by a space, followed by the caller name: `glidefunction:concat(short_description, ' ', caller_id.name)` **Note:** Function parameters support dot-walking. See Dot-walking to data in related tables.

5. Click **Submit**.
   You return to the Dictionary Entry [sys_dictionary] table.
6. Navigate to the table to which you added the function field.
7. Add the function field to the list.
   See Configure items on forms or in lists using a list collector.

Result

Use the function field like any other field. For example, add it to a form, use it in a filter, or use it in the condition builder.
If the function definition is invalid, instead of the expected value you receive the message **Invalid function** if the function field is a string type field, or an empty value for other field types.

Geo point field type

The geo point data type stores longitude and latitude coordinates.

You can store longitude and latitude coordinates by adding a geo point data type field to any table. For example, you can define a geo point field to store a longitude and latitude pair to help technicians find the location of a part in a warehouse.

You can calculate the distance between geo points using distance_sphere() in a function field and use it to order records in a list based on how far away the warehouse is. For example:

```
glidefunction:distance_sphere(geopointField1, geopointField2)
```

See Function field.

Values for geo points are in longitude, latitude format where:

- Both longitude and latitude are restricted to 6 decimal places of precision. For example, “80.206458”, “25.767242”.
- The longitude range is [-180, 180].
- The latitude range is [-90, 90].
- Values entered outside the ranges for longitude or latitude are normalized to values inside their respective ranges.

HTML field type

The HTML editor provides WYSIWYG (what you see is what you get) functionality and HTML source mode editing. Administrators can customize some of the functionality associated with HTML fields.

The HTML editors available depend on your version of the UI.

UI support for available HTML editors

HTML editor	UI support
TinyMCE version 5	Core UI, UI15
htmlArea (legacy)	Basic support in Core UI

Knowledge articles, service catalog item descriptions, release documentation, and HTML content blocks are common examples of HTML fields.

Note:
The TinyMCE editor does not support converting Microsoft Word Shapes into HTML.

While HTML fields can be added to split forms, HTML fields are not intended for this use and can behave unexpectedly due to the limited available width in split pane forms.

If accessibility is enabled, a help icon appears with the HTML field. The help icon indicates how to use keys to access and escape the HTML field toolbar.

Configure a field editor for the HTML field

You can configure HTML fields to use TinyMCE or the legacy HtmlArea editor. This editor provides better stability and more editing functions than the legacy htmlArea editor.

Before you begin

Role required: admin

About this task

There are two options for HTML editors.

- TinyMCE: A What You See Is What You Get (WYSIWYG) field that displays text as readers would see it on the screen. TinyMCE is the default editor.
- htmlArea: The legacy editor, which offers a more basic WYSIWYG interface as well as a mode that shows only HTML markup.

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Locate the property called **HTML field editor to use** (`glide.ui.html.editor`).

3. Select **TinyMCE** or **htmlArea**.
4. Click **Save**.

Configure the HTML toolbar

When HTML fields are configured to use the HTML editor, follow this procedure to configure which buttons are available on the toolbar.

Before you begin

Role required: admin

Procedure

1. Update the toolbar options for your preferred HTML field editor.

Option	Description
htmlArea toolbar	a. Navigate to the System properties [sys_properties] table by entering `sys_properties.list` in the application navigator. b. In the **Name** column, search for `glide.ui.html.toolbar`. c. Select the **glide.ui.html.toolbar** property.
TinyMCE	a. Navigate to the System properties [sys_properties] table by entering `sys_properties.list` in the application navigator. b. In the **Name** column, search for `glide.ui.html.editor.v5.toolbar`. c. Select the **glide.ui.html.editor.v5.toolbar** property.

2. In the **Value** field, enter or remove buttons for each toolbar as a comma-separated list without spaces.

   For the TinyMCE toolbar, use the following buttons. Use a vertical bar ("|") to add a section separator.

   TinyMCE button options

Button purpose	Accepted button names
Formatting	newdocument, bold, italic, underline, strikethrough, justifyleft, justifycenter, justifyright, justifyfull, formatselect, fontselect, fontsizeselect, bullist, numlist, outdent, indent, blockquote, forecolor, removeformat, backcolor, sub, sup
Table functions	tablecontrols
Editing	cut, copy, paste, pastetext, search, replace, undo, redo
Extended functions	link, unlink, cleanup, code, hr, visualaid, charmap, image, media, preview, fullscreen (not supported by Internet Explorer)

   For the htmlArea toolbar, use the following buttons. Use **separator** to add a section separator.

htmlArea button options

Button purpose	Accepted button names
Formatting	formatblock, fontname, fontsize, bold, italic, underline, justifyleft, justifycenter, justifyright, justifyfull, insertorderedlist, insertunorderedlist, outdent, indent, forecolor, hilitecolor
Editing	copy, paste, undo
Extended functions	createlink, inserthorizontalrule, insertimage, insertvideo, inserttable, htmlmode

3. Select **Update**.

TinyMCE toolbar

htmlArea toolbar

Configure TinyMCE to allow deprecated tags

You can set a dictionary attribute on a TinyMCE field to allow the use of deprecated HTML tags, such as <b> and <i>. By default, TinyMCE uses the <strong> and <em> tags for bold and italic formatting.

Before you begin

Role required: personalize_dictionary or admin

About this task

After you set the dictionary attribute, use code view to manually enter deprecated tags. The editor does not validate any tags you enter manually, for example, if you type an incorrect character.

Procedure

1. Navigate to the form with an HTML field that uses TinyMCE.
2. Right-click the HTML field label and select **Configure dictionary**.

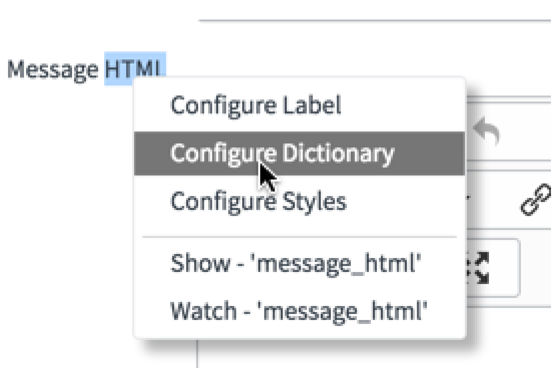

3. In the **Attributes** field, enter `tinymce_allow_all=true`, separated by a comma if needed.

   Dictionary entry attributes can only be added to when the dictionary entry form is in advanced view, as they are not shown in default view.

   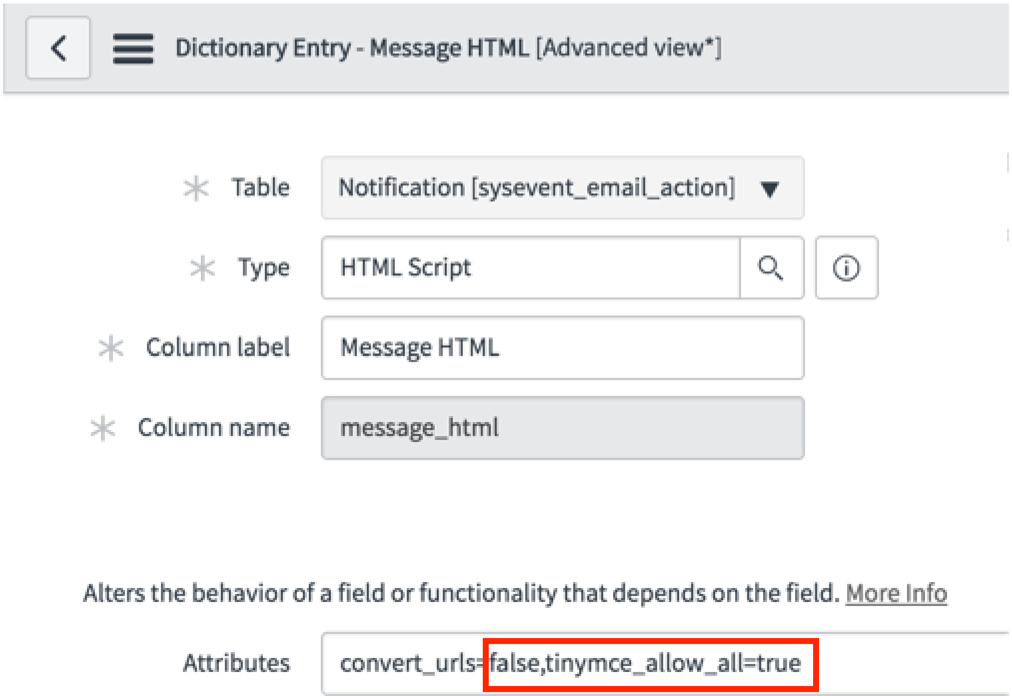

   If other attributes are already listed, use a comma as a separator.

4. Click **Update**.

Configure TinyMCE to allow JavaScript in URLs

You can set a dictionary attribute on a TinyMCE field to allow the use of JavaScript in a URL.

Before you begin

Role required: personalize_dictionary or admin

Procedure

1. Navigate to the form with an HTML field that uses TinyMCE.
2. Right-click the HTML field label and select **Configure dictionary**.

   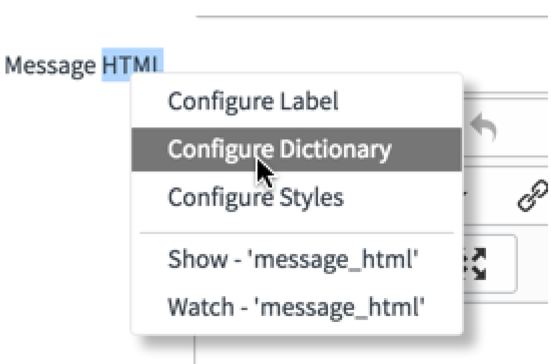

3. In the **Attributes** field, enter `tinymce_allow_script_urls=true`, separated by a comma if needed.

   Dictionary entry attributes can only be added to when the dictionary entry form is in advanced view, as they are not shown in default view.

   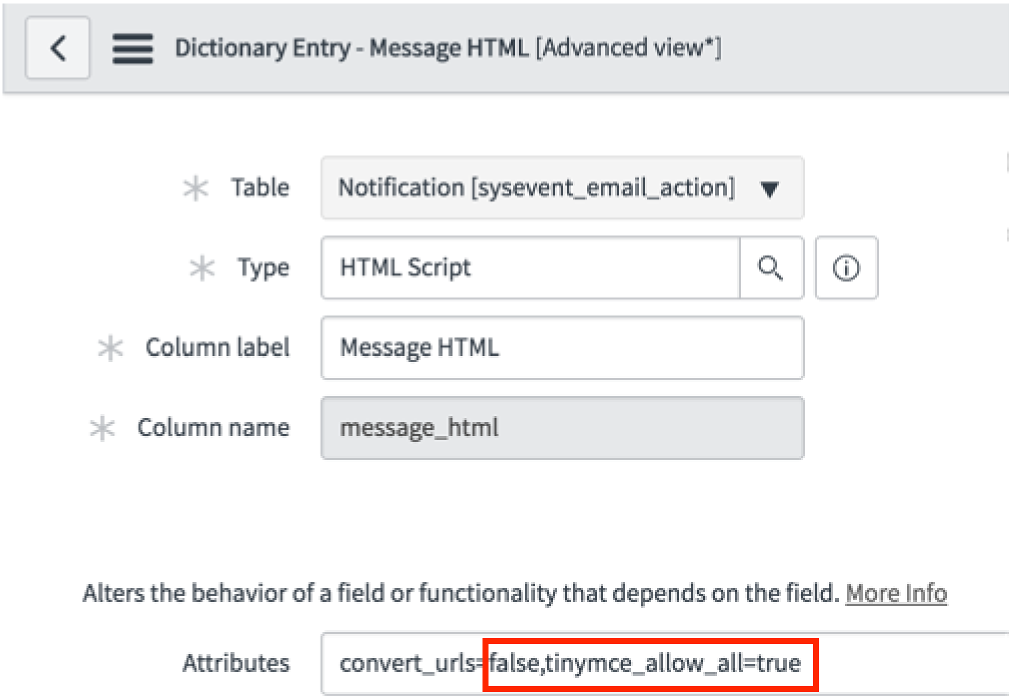

   If other attributes are already listed, use a comma as a separator.

4. Select **Update**.

Formatting icons for the HTML field editors

The formatting table displays how to control the way text appears.

The following table displays icons that are available in the TinyMCE and the htmlArea editors. Names marked with an asterisk (*) are not available with the htmlArea editor.

Formatting table

htmlArea icon	TinyMCE icon	Name	Description
B	**B**	Bold	Applies bold formatting to the selected text or current word. Keyboard shortcut: CTRL + B
I	*I*	Italic	Applies italics formatting to the selected text or current word. Keyboard shortcut: CTRL + I
U	U	Underline	Applies underline formatting to the selected text or current word.

Formatting table (continued)

htmlArea icon	TinyMCE icon	Name	Description
			Keyboard shortcut: CTRL + U
		Align Right	Applies right alignment to the current paragraph. Code reference: `text-align:right`
		Align Center	Applies center alignment to the current paragraph. Code reference: `text-align:center`
		Align Left	Applies left alignment to the current paragraph. Code reference: `text-align:left`
		Justify	Applies justified alignment, which stretches the lines to equal width, to the current paragraph. Code reference: `text-align:justify`
Paragraph		Format	Applies a paragraph style to the current paragraph, such as Paragraph, Heading 1, and Preformatted.
Arial	Font Family	Font Family	Applies a font family to the selected text or current word.
3 (12pt)	Font Sizes	Font Size	Applies a font size to the selected text or current word.
		Insert/ Remove Bulleted List	Applies or removes unordered list tags for the selected paragraphs. Click the arrow beside the button to select a different bullet type.
		Insert/ Remove Numbered List	Applies or removes ordered list tags for the selected paragraphs. Click the arrow beside the button to select a different number type.
		Decrease Indent	Removes indentation from the current or selected paragraphs (removes 30px of left padding; padding cannot be less than 0). Code reference: `padding-left`
		Increase Indent	Applies indentation to the current or selected paragraphs (adds 30px of left padding). Code reference: `padding-left`

Formatting table (continued)

htmlArea icon	TinyMCE icon	Name	Description
		Select Text Color	Applies font color to the current word or selected text. Click the button to use the current color, or click the arrow next to the button to view more colors. Click **More Colors...** to view various color options and the hexadecimal codes.
		Select Background Color	Applies background color to the current word or selected text. Click the button to use the current color, or click the arrow next to the button to view more colors. Click **More Colors...** to view various color options and the hexadecimal codes.

Extended functions

The extended functions available for working with HTML content.

Extended functions

Icon	TinyMCE v4 Icon	Name	Description
		Insert/Edit Link	Configures a link for the selected text. Define the link URL, title (additional information that appears in the tooltip), and the target (same window or new window or tab). Code reference: `<a>`
		Remove link*	Removes the current hyperlink.
		Cleanup Messy Code*	Fixes standard HTML errors for the selected text, such as invalid tags. Clicking this button may change the layout of existing content. If you do not like the results, you can click **Undo** to revert this action.
HTML		Edit HTML Source	Opens HTML source code in a separate window. See Editing in HTML Source Mode.
		Insert Horizontal Line	Inserts a horizontal line at the current location.
		Toggle Invisible Elements*	Shows or hides invisible elements in the article, such as collapsed table borders.
		Insert Special Character*	Inserts a special character (symbol) at the current cursor location. Click the button to view a list of available characters. Point to a character to view the name and HTML code. Click a character to insert it.
		Insert/Edit Image	Inserts an image from the image library or an attachment. You can also add images to the image library with this feature. To learn more, see Embedding Images in HTML Fields.

Extended functions (continued)

Icon	TinyMCE v4 Icon	Name	Description
		Insert/Edit Embedded Media	Embeds a video from the video library or an attachment. You can also add videos to the video library with this feature. To learn more, see Embedding Video in HTML Fields.
		Spell Checker	Checks the spelling of text in the HTML field. Spellcheck is only available in the htmlArea editor.
		Preview*	Opens a preview of the HTML field in a separate window without saving changes.
		Toggle Full Screen Mode*	Expands the HTML field to use the full form view for easier editing. Click the button again to return to standard form view. This feature is only available for the htmlArea editor, and is not available for Internet Explorer.

*These options are not available with htmlArea.

Editing functions in TinyMCE

Several editing functions are built into TinyMCE.

You can add more edit functions to TinyMCE v5 using the `glide.ui.html.editor.v5.toolbar` and properties located on the **UI Properties** page. For more information on configuring the edit functions, see Configure the HTML toolbar.

Editing functions

Icon	TinyMCE v4 Icon	Name	Description
		Cut*	Cuts the selected text. Not supported in all browsers; use keyboard shortcut. Keyboard shortcut: CTRL + X
		Copy*	Copies the selected text. Not supported in all browsers; use keyboard shortcut. Keyboard shortcut: CTRL + C
		Paste*	Pastes the selected text. Not supported in all browsers; use keyboard shortcut. Keyboard shortcut: CTRL + V
		Paste as Plain Text*	Enables paste as plain text without source formatting.

Editing functions (continued)

Icon	TinyMCE v4 Icon	Name	Description
		Paste from Word*	Opens a new window that allows you to copy and paste content from Microsoft Word into the HTML field.
		Find*	Allows you to locate text strings in the HTML field. Search above (up) or below (down) the cursor location.
		Find/ Replace*	Allows you to replace the next (Replace) or all (Replace All) occurrences of a text string in the HTML field. Keyboard shortcuts: CTRL + F, Cmd + F
		Undo*	Reverts the previous edit.
		Redo*	Reapplies the last reverted edit.

*These options are not available with htmlArea.

Table functions in TinyMCE

TinyMCE version 4 uses menus and menu selections to create and edit tables.

TinyMCE version 4

UI element	Action
Table menu 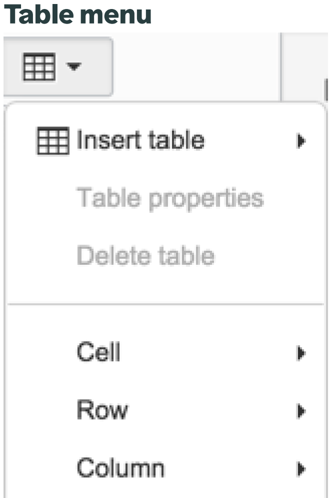 	Click the table icon ( ) to access the TinyMCE version 4 table menu. Use the table menu to: • Insert or delete a table • Modify table properties • Add, move, or delete rows and columns • Modify row and column properties • Split and merge cells

TinyMCE version 4 (continued)

UI element	Action
Table insert grid 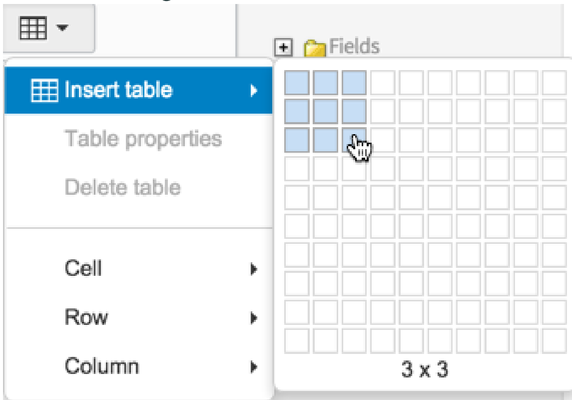	To insert a table in the HTML field, click **Insert table** and highlight squares in the grid to represent the desired number of rows and columns. Click the last highlighted square to insert the table. After you insert the table, you can modify the size by clicking and dragging the handles at the table edges.
Table properties 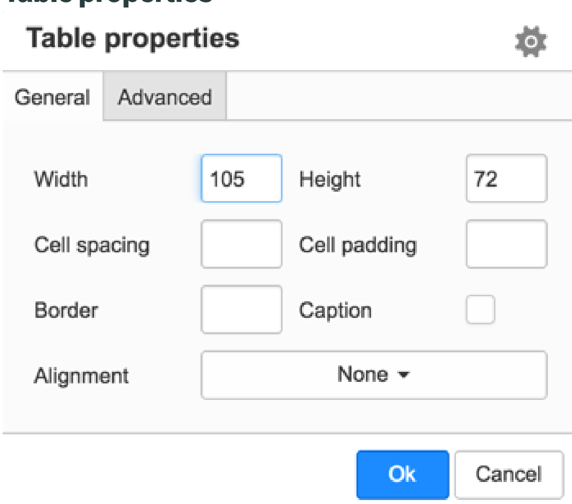	From the table menu, click **Table properties** to open the Table properties dialog box. From this box you can take any of the following actions. • General tab: ◦ Set table width and height ◦ Set cell spacing and padding ◦ Enable borders and captions ◦ Set the table alignment • Advanced tab: ◦ Configure the table style ◦ Select the border color ◦ Select the background color
Cell properties 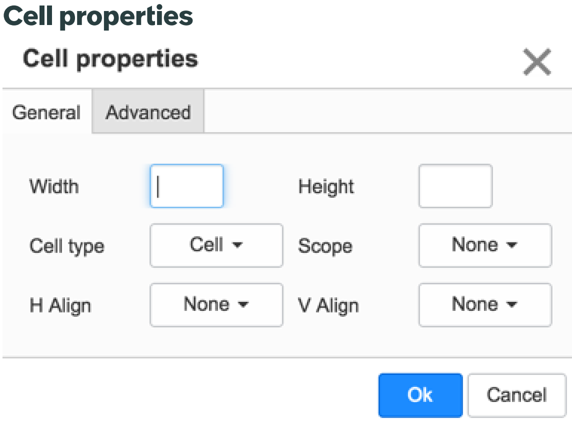	With the cursor in the desired table cell, open the table menu and click **Cell properties**. From this box you can set the following properties for table cells. • General tab: ◦ Width and height ◦ Type and scope ◦ Horizontal and vertical alignment • Advanced tab

TinyMCE version 4 (continued)

UI element	Action
	◦ Configure the cell style ◦ Select the border color ◦ Select the background color
Row properties 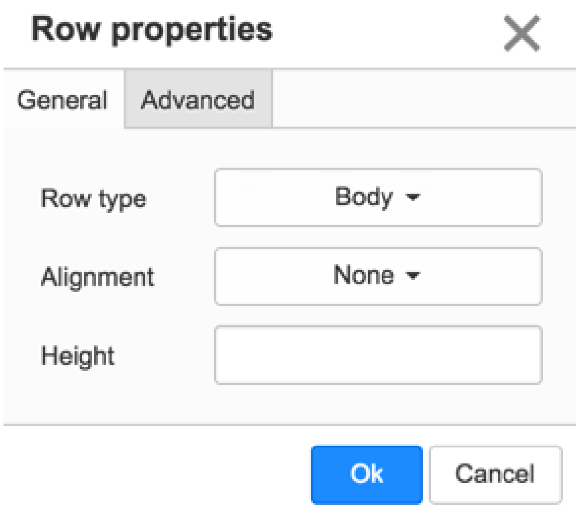 	With the cursor in a table cell in the desired row, open the table menu and click **Row properties**. From this box you can set the following properties for rows. • ◦ Row type ◦ Alignment ◦ Height • Advanced tab ◦ Configure the row style ◦ Select the border color ◦ Select the background color

You can also paste tables into HTML fields from table-based editors.

Customize TinyMCE attributes

Change TinyMCE toolbars and the default height of HTML fields to customize your journal field experience.

Add or remove TinyMCE attributes from the default TinyMCE toolbar. The TinyMCE attributes are applied to TinyMCE toolbars in the platform and workspaces, or to specific plugins in a record or specific tables in a record. You can also apply an attribute to increase the height of HTML fields, enlarging the HTML field that agents work with.

Change the TinyMCE default toolbar

Set the toolbar items in your system properties to enable or disable throughout Workspace.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Update the **Configures the editing toolbar for HTML fields(TinyMCE v5)** (glide.ui.html.editor.v5.toolbar) property to add or remove buttons for the toolbar.

   Note:
   Use a vertical bar ("|") to add a section separator.

Type	Buttons
Default buttons	*bold italic underline undo redo | fontselect fontsizeselect table | forecolor backcolor link unlink | image media code | alignleft aligncenter alignright | bullist numlist fullscreen*
Valid buttons	*newdocument, bold, italic, underline, strikethrough, alignleft, aligncenter, alignright, alignjustify, formatselect, fontselect, fontsizeselect, table, cut, copy, paste, pastetext, searchreplace, bullist, numlist, outdent, indent, blockquote, undo, redo, link, unlink, code, forecolor, backcolor, removeformat, hr, visualblocks, subscript, superscript, charmap, image, media, fullscreen, anchor, toc, codesample, ltr, rtl, emoticons, insertdatetime, nonbreaking, pagebreak, print, visualblocks, visualchars, cut, copy*

3. Select **Save**.

Change the TinyMCE toolbar for a specific table

Set the attributes in the TinyMCE dictionary to determine which TinyMCE attributes show in a specific table.

Before you begin

Role required: admin

Procedure

1. Navigate to a record with an HTML-type field that you want to change.
   For example, select an incident, problem, or knowledge record.
2. Right-click the field label (for example, Article body) and select **Configure Dictionary**.
3. In the Related Links section, select **Advanced View**.
4. In the **Attributes** field, input `editor.toolbar=` followed by the desired toolbar buttons.
   For example, `editor.toolbar=formatselect|bold italic underline strikethrough blockquote subscript superscript removeformat|bullist numlist outdent indent|undo redo|table hr|link unlink|image media code|visualblocks preview fullscreen`.

Note:

- Include all the toolbar items that you want displayed, not just the toolbar items you want to add.
- The configurations made to a field's Attribute field on the associated Dictionary record overrides the value of the System Property glide.ui.html.editor.v5.toolbar.
- Multiple attributes, such as height, toolbar buttons, and toolbar plugins, can be combined within the Attributes field. For example, `editor.height=300,editor.plugins=table colorpicker textcolor link image media codemirror lists advlist fullscreen charmap directionality emoticons hr insertdatetime nonbreaking pagebreak print searchreplace wordcount anchor toc codesample visualblocks visualchars compat3x autolink align_listitems,editor.toolbar= fontselect fontsizeselect | bold italic underline strikethrough forecolor backcolor pastetext removeformat | formatselect searchreplace undo redo | bullist numlist outdent indent alignleft aligncenter alignright | toc table link unlink image media codesample | code fullscreen`.

5. Select **Update**.

Change TinyMCE plugins for a specific table

Set the attributes in the TinyMCE dictionary to enable or disable plugins in a specific HTML field.

Before you begin

Role required: admin

Procedure

1. Navigate to a record with an HTML-type field that you want to change.
   For example, select an incident, problem, or knowledge record.
2. Right-click the field label (for example, Article body) and select **Configure Dictionary**.
3. In the Related Links section, select **Advanced View**.
4. In the **Attributes** field, input `editor.plugins=` followed by the desired plugins, separated with a space.

   For example, `editor.plugins=table colorpicker textcolor link image media codemirror lists advlist fullscreen charmap directionality emoticons hr insertdatetime nonbreaking pagebreak print searchreplace wordcount anchor toc codesample visualblocks visualchars compat3x autolink align_listitems`.

Note:

- Include all the toolbar items that you want displayed, not just the toolbar items you want to add.
- The configurations made to a field's Attribute field on the associated Dictionary record overrides the value of the System Property glide.ui.html.editor.v5.toolbar.
- Allowed plugins include: advlist align_listitems anchor autolink autoresize bbcode charmap codemirror codesample colorpicker directionality emoticons fullscreen hr image insertdatetime link lists media nonbreaking pagebreak preview print readonlynoborder searchreplace table textcolor toc visualblocks visualchars wordcount
- Multiple attributes, such as height, toolbar buttons, and toolbar plugins, can be combined within the Attributes field. For example, `editor.height=300,editor.plugins=table colorpicker textcolor link image media codemirror lists advlist fullscreen charmap directionality emoticons hr insertdatetime nonbreaking pagebreak print searchreplace wordcount anchor toc codesample visualblocks visualchars compat3x autolink align_listitems,editor.toolbar= fontselect fontsizeselect | bold italic underline strikethrough forecolor backcolor pastetext removeformat | formatselect searchreplace undo redo | bullist numlist outdent indent alignleft aligncenter alignright | toc table link unlink image media codesample | code fullscreen`.

5. Select **Update**.

Change the default height of an HTML field

Change the default height of a specific HTML field to expand the size of a journal field.

Before you begin

Role required: admin

About this task

HTML field height is configured per HTML field.

Procedure

1. Navigate to a record with an HTML-type field that you want to change.
   For example, select an incident, problem, or knowledge record.
2. Right-click the field label (for example, Article body) and select **Configure Dictionary**.
3. In the Related Links section, select **Advanced View**.
4. In the **Attributes** field, enter `editor.height=X`, where X is the desired height.
   For example, `editor.height=250`

   Note:

   HTML fields can range from 72 to 2000. HTML fields are by default 64.

5. Select **Update**.
6. To configure the height of form fields dynamically as the text line increases rather than providing a specific height, complete the following steps:

a. Navigate to **All** > **System Properties** > **All Properties**

b. In the search bar, enter `glide.ui.html.editor.v5.enabled_plugins` and select the property.

c. In the value field, add `autoresize`.

d. Select **Update**.
The autoresize plugin is active.

Change the default font size of an HTML field

Change the default font size of a specific HTML field to use a standard font size across forms.

Before you begin

Role required: admin

About this task

HTML font size is configured per HTML field.

Procedure

1. Navigate to a record with an HTML-type field that you want to change.
   For example, select an incident, problem, or knowledge record.
2. Right-click the field label (for example, Article body) and select **Configure Dictionary**.
3. Select the Default Value tab.
4. In the **Default value** field, enter `<p style="font-size:X;"></p>`, where X is the default value.
   For example:
   - To set the font to large, enter `<p style="font-size:large;"></p>`
   - To set the size to 10, enter `<p style="font-size:10pt;"></p>`
5. Select **Update**.

Highlight text in TinyMCE

On the bottom bar of TinyMCE, the path of HTML tags for the text at the cursor position is displayed.

Click a tag in the path to highlight the text affected by the tag.

Insert a line break in the HTML editor

When you use the `Enter` key, the editor creates a paragraph element (`<p>`) tag, which appears as a double space.

To enter a single-line space, use the `Shift` + `Enter` key combination, which inserts a line break (`<br>`) tag.

Add a table to the HTML field

This example uses HTML field controls to format a table in a knowledge article.

About this task

Table example html

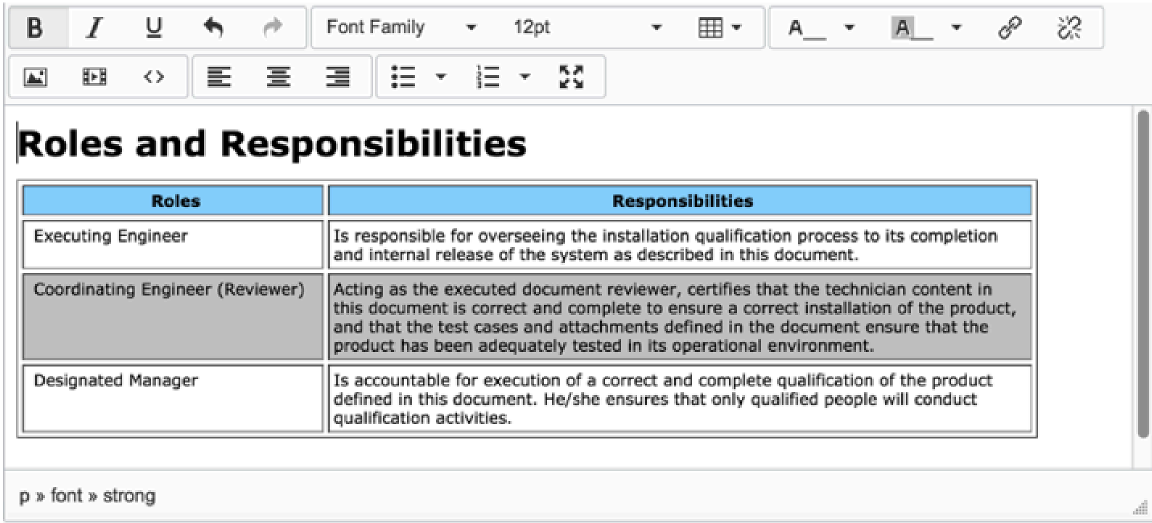

To add the formatted table to a knowledge article:

Procedure

1. Navigate to **All** > **Knowledge** > **Edit** and select the article to edit.
2. In the HTML field, position the cursor in the location for the table.
3. Click the table icon, click **Insert table**, and then select the number of rows and columns.

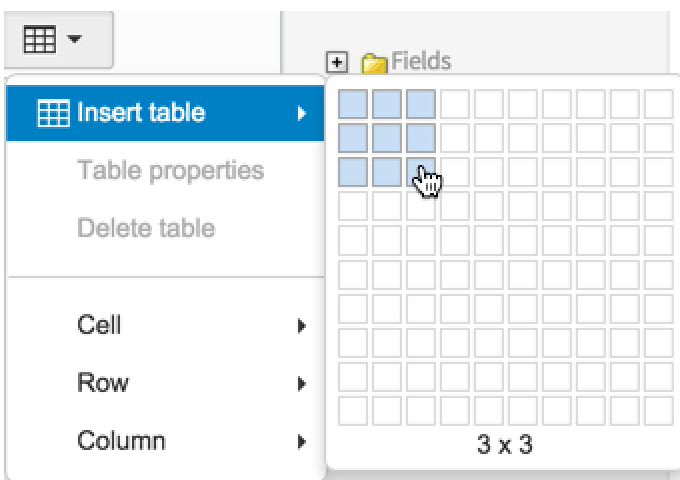

4. Complete the following steps to edit the table properties.

   a. Position your cursor in the table, click the table icon, and select **Table properties**.

   b. Enter the following values on the **General** tab.

      - Width: 75%
      - Cell spacing: 3
      - Cell padding: 3

- Border: 1
- Alignment: Left

c. On the **Advanced** tab, click the text field next to **Border color** and enter `Gray`.
The color picker box to the right turns gray to indicate the color that you entered. You can also click the box and select the color in the palette.

d. Click **Ok**.

5. Complete the following steps to update the header table row.

a. Select the cells in the first table row, click the table icon, and select **Row** > **Row properties**.

b. Enter the following values on the **General** tab.

- Row type: Header
- Alignment: Center

c. On the **Advanced tab**, enter `#87cefa` in the text box beside **Background color** to set it to a light blue.

d. Click **Ok**.

6. To set cell properties, complete the following steps.

a. Select all table cells in the first column except those cells in the header row, click the table icon, and select **Cell** > **Cell properties**.

b. Enter the following values on the **General** tab.

- H Align: Left
- V Align: Top

c. Click **Ok**.

d. Repeat these steps for the table cells in the second column.

7. To set the background color of the middle row, complete the following steps.

a. Position your cursor in the middle table row, click the table icon, and select **Row** > **Row properties**.

b. On the **Advanced tab**, enter `Silver` in the text box beside **Background color** to set it to color #c0c0c0.

c. Click **Ok**.

Repeat this procedure for every other table row.

8. To set column width, complete the following steps.

a. Click the first column of the table, click the table icon, and select **Cell** > **Cell properties**.

b. On the **General** tab, enter `30%` in the **Width** text field.

c. Click **Ok**.

9. Right-click the form header and click **Save**.
10. Enter data in the table cells and then save the article.

Embed images in HTML fields

You can use the HTML field image picker to embed images into HTML fields, and to add images to the image library.

Before you begin

Role required: the role necessary to update the record that contains the HTML field. For example, any user with a role can create a knowledge article and embed an image in it.

About this task

Note:
Administrators and users with the image_admin role manage the image library at **System UI** > **Images**. See Storing images in the database .

Procedure

1. Open the form that contains the HTML field.
2. Click the position where the image is to appear, or to modify an existing image, click the image.
3. Click the insert/edit image icon () on the HTML editor toolbar.
4. In the Insert/Modify image form, enter information in each field.

Insert/Modify Image form

Field	Description
Type	Select the image type. ◦ **Image Library**: Images stored in the db_image table. You can reuse images in the image library in multiple locations. ◦ **Attachment**: Available in the current record only
Image	Begin typing a file name and select an image from the list, or click the reference lookup icon and select an image. To upload a new image, click **New**, click **Choose File**, locate the image, and click **Upload**. If you chose the **Attachment** type, click **Choose File**, locate the image, and click **Attach**.
Tooltip	Enter alternate text that appears when a user points to the image.
Alt	Enter alternate text that can be used to improve accessibility. For example, it could be used with a screen reader. If this field is left blank, it defaults to the text entered in the **Tooltip** field.

Note:
To resize an embedded image, click the image. The sizing frame appears. Drag a sizing point until the image is the desired size. Corner points adjust the size proportionally. Depending on your browser, you may need to highlight the image before the sizing frame appears.

5. **Optional:** To provide additional control over the appearance of an image, click **Advanced options**.

Advanced options

Field	Description
Layout	Select the image **Alignment** (default is **Baseline**) and enter the **Border thickness**.
Spacing	Enter the number of **Horizontal** and **Vertical** pixels around the image.
Size	Enter the **Width** and **Height** of the image (in pixels).

6. Click **OK**.
7. **Optional:** Paste an image into the HTML editor.
   To edit the image, save the entry then select the image. Click the insert/edit image icon and complete the form using the same tables for adding an image.

   (Optional) Pasted images are saved in the system as attachments.

Link to a website in HTML fields

You can insert a link to a website in an HTML field.

Before you begin

Role required: the role necessary to update the record that contains the HTML field. For example, any user with a role can create a knowledge article and link to a website in the article text.

Procedure

1. Move the cursor to the position where the link is to appear.
2. Click the insert/edit link icon () in the HTML toolbar.
3. Enter information in each field.

Insert/Edit Link form

Field	Description
URL	Enter the URL for the link. Copying and pasting is usually the easiest method.
Text	Enter the text you want to display for the link.
Target	Select the target window for the URL. For files, the **None (use implicit)** selection is generally the best choice. However, if you are linking to a complete web page, choose **New window (_blank)** so the browser opens the link in a new tab or window.

4. Click **OK** to insert the link into the field.

Embed videos in HTML fields

You can insert videos into HTML fields. You can also add videos to the video library using the HTML Insert/Modify Video form.

Before you begin

Note:
The Videos [db_video] table is a public table that does not have any security restrictions. Unauthenticated users have full access to videos uploaded to the db_video table.

Role required: The role necessary to update the record that contains the HTML field. For example, any user with a role can create a knowledge article and embed a video in it.

About this task

The following file formats are supported in the base system.

- MPEG-4 video `.mp4`
- WebM Video `.webm`

The larger the file size, the longer it takes to download before the video starts. To reduce file size, you can reduce the bit rate, but doing so reduces the quality of the video. The following bit rates optimize download speed and video quality:

Video resolution	Recommended bitrate
Standard definition (480p)	1.5–2.5 megabits per second (mbps)
High definition (720p)	3–5 mbps
High definition (1080p)	6–8 mbps

Note:
For the best viewing experience the following specifications are suggested:

Optimal viewing settings

Type	Definition
Size	20MB
Bit rate	320KBPS
Resolution	480*320

Internet Explorer and Safari have difficulty streaming videos uploaded to the database. Attach a file rather than embedding if you intend to use one of these browsers or an unsupported file type. For more information on attaching files, see Add and manage attachments.

Note:
Administrators and users with the image_admin role can manage the video library at **System UI** > **Videos**.

Procedure

1. Open the form that contains the HTML field.
2. Click the position where the video is to appear, or to modify an existing video, click the video.
3. Click the insert/edit video icon ( ) on the HTML editor toolbar.
4. Enter information in each field.

Insert/Modify Video

Field	Description
Type	Select the video type. ○ **Video Library**: List of videos stored in the db_video table. You can reuse it. ○ **URL**: from an external source ○ **Attachment**: available in the current record only For an attachment or video library file, select a video from the list or click **New**. For an external URL, enter the URL.
Size	Enter the **Width** and **Height** of the video in pixels.

5. Click **OK**.

Define video file types for HTML fields

You can define the types of video files that can be added to HTML fields.

Before you begin

Role required: image_admin

About this task

Users can add videos to HTML fields. By default, users can add one of the following types of videos to HTML fields: `.mp4`, `.webm`, and `.swf` video file types. You can inactivate video types that you do not want to allow users to add, or add new video types. `.swf` files are only minimally supported. `.mp4` files might be limited by browser type.

Note:
By default, the HTML Sanitizer removes videos from HTML fields. To allow video file types, see Embed videos in HTML fields

Procedure

1. Navigate to **All** > **System UI** > **Embed Object Types**.
2. To deactivate a video type, set the **Active** field to **false**.
3. To add an additional video type, click **New** and complete the form.

   Note:
   If you specify values for the **Codebase** or **Pluginspage** fields, which instruct the browser where to get the plugin, point to `https` pages to avoid warnings from Internet Explorer about unsecure content on the page.

Disable user access to the image library

By default, users can use and upload images to the image library from an HTML field. You can disable access to the image library from HTML fields.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System UI** > **UI Pages**.
2. Select **html_insert_image_dialog**.
3. Locate the following lines of code in the **HTML** field and comment them out.

```
<g:evaluate var="jvar_use_dbimage">
 var imageGR = new GlideRecord('db_image');
  if(imageGR.isValid())
  imageGR.canRead();
</g:evaluate>
```

   and

```
<j:if test= "${jvar_write_dbimage}">
    <j:set var = "jvar_default_insert_image_type" value =
 "dbimage" />
    <g:ui_select_option text = "${gs.getMessage('Image
 Library')}" value = "dbimage" selected =
 "${jvar_default_insert_image_type}" />
</j:if>
```

4. Click **Update**.

What to do next

To remove the **Upload from URL** option in the HTML editor, add a new property named *glide.ui.html.image.allow_url* and set the **Value** to **false**.

Related topics

Add a system property

Paste content into the HTML editor

Paste content from a desktop application to the HTML editor. Content might not paste the same from every application.

Before you begin

Role required: admin

Procedure

1. Navigate to an HTML field.
   For example, the **Text** field of a Knowledge article.
2. Paste in content from another application.
   Google Chrome and Mozilla Firefox do not support paste using the Paste icon. Use the command/control+V keyboard shortcut instead.
3. From the paste formatting options dialogue, choose to **Keep** or **Remove** the text formatting for the pasted content.

Result

The HTML editor has some pasting limitations:

- Pasting images: Pasting images and text from OneNote works in Microsoft browsers. Pasting from OneNote into Chrome or Firefox works but styles and images are not included. You can still select images individually and paste them into the HTML field. Pasting images and text from Microsoft Word into any browser works as expected.
- Pasting tables: Pasting tables into Safari does not always work correctly. Pasting tables into Internet Explorer only works from Excel.

Image field type

Image fields enable you to add images to forms.

Before you begin

Role required: personalize_form

About this task

For example, you can add portraits to the user records in your system. The image type must be `.gif`, `.jpg`/`.jpeg`, or `.png`.

Image resizing is defined by the CSS. Larger images are resized to 250 pixels. For most browsers, the larger of the height and width measurements is reduced to 250 pixels and the proportion of the image is maintained. For example, an image with a size of 1508 x 663 pixels is resized to 250 x 110 pixels. Some browsers resize both the height and width measurements to 250 pixels, resulting in a square image.

You can add a new image field.

Procedure

1. Open the desired form.
2. Create a new field with the **Type** set to **Image**.
   For instructions, see Add and customize a field in a table.
3. Add the new field to the form and save your customization.
   The form displays the new blank image field.
4. Click **Click to add** in the image field, select an image to upload, and click **OK**.
   The selected image is attached to the form and displayed in the image field.

Result

Related topics

Configuring the form layout

IP address field type

The **IP Address (Validated IPV4, IPV6)** [ip_addr] field type stores valid IPv4 and IPv6 addresses.

Support for IPv4 and IPv6

Standard dot-decimal notation is supported for IPv4 and standard hex notation is supported for IPv6. For IPv6, you can use canonical notation to compress zeros. Familiarize yourself with IETF RFC5952 for examples and guidelines on text representation for IPv6 addresses.

The following are examples of valid IP addresses:

- **A standard IPv4 address in dot-decimal notation**:

  `10.34.51.20`

- **An expanded IPv4 address would be zero padded**:

  `10.1.10.210` would be represented as `010.001.010.210`

- **An IPv6 address in fully expanded notation**:

  `1507:f0d0:1002:0051:0000:0000:0000:0004`

- **The same IPv6 address above using the canonical notation**:

  `1507:f0d0:1002:51::4`

How IP addresses are stored in the database

The IP address `ip_addr` field is a Variable Character (VARCHAR) field with a length of 40 characters. The values are normally stored in standard (dot-decimal or canonical) notation for both IPv4 and IPv6 addresses regardless of the notation you use when you input values in the field.

For example, if you enter an IPv6 address in a non-canonical format (expanded or any other valid formats per IETF RFC5952), by default, the system will store it in canonical format unless an IP Data Control attribute value is defined on that column with a different address format. See the [ip_data_control] glide attribute section in Altering tables and fields using dictionary attributes for more details.

Support for data migration

Provide support in migrating IP addresses stored as string fields to the new **IP Address(Validated IPV4,IPV6)** type. The new IP address type accepts only a valid IP address. The existing data is formatted according to the ip_data_control attribute. The ip_data_control attriute can be set to exactly one of the following four values:

- *`canonical`*: Incoming IP addresses are validated and canonicalized before they are entered in the database. Invalid IP addresses are rejected.
- *`canonicalize_when_possible`*: Valid IP addresses are canonicalized before before being entered in to the database. Invalid IP addresses are also entered in to the database without being changed.

  Note:
  `canonicalize_when_possible` is the default value for the ip_data_control attribute, unless it's otherwise specified.

- *`expanded`*: The IP addresses are validated and stored in an expanded form to support efficient database range queries. Invalid IP addresses are rejected.
- *`none`*: No validation or normalization is performed in this value type. It is used only as an emergency mechanism as it essentially reverts this type to a plain string field.

Journal field type

There are three types of journal field: journal, journal_list, and journal_input.

Journal field types

Journal field types	Description
journal	Allow and store input, and display the combined inputs below the input box. Journal fields display in the activity stream in the form and in the list view.
journal_input	Allow and store input, but do not display the combined inputs. Journal input fields only display with the record they are associated with, so they do not display in the activity stream on the list view.
journal_list	Do not allow or store input; they merely display the contents of other Journal fields upon which the journal_list field is dependent. If a journal_list field is dependent on more than one Journal field, it will chronologically interweave those fields' inputs. The journal_list field does not display content within the activity stream, but rather in a separate block.

Journal fields on a form

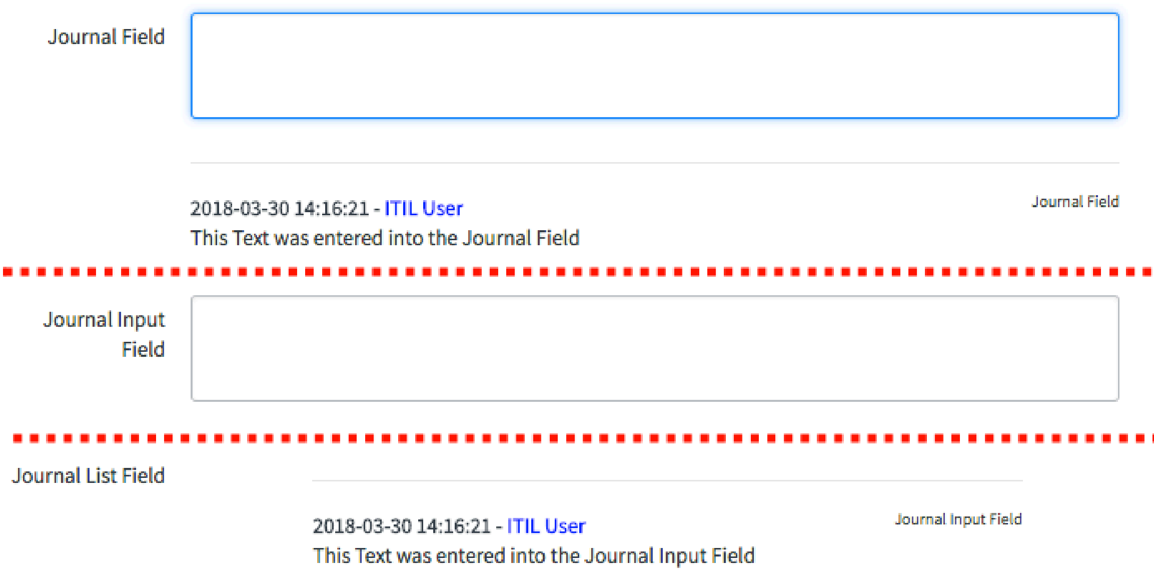

The example image contains the three available journal field types. The first is a journal field, with its inputs displayed below the field. The second is a journal input field, which does not show its previous inputs. The third is Journal list field, which is configured to show the input from the journal input field above it.

Restricting journal entries sent in a notification

Administrators can control the number of journal entries notifications include with the following system property.

Property	Label	Description
glide.email.journal.lines	Number of journal entries (Additional comments, Work notes, etc.) included	Specifies the number of entries from a journal field (such as **Additional comments** and **Work notes**) included in

Property	Label	Description
	in email notifications (-1 means all).	email notifications. A value of -1 includes all journal entries. • Type: integer • Default value: 3 • Location: System Properties > Email

Code for getting the contents of a journal field into an array

To put the contents of a journal field into an array so that you can iterate through each entry, you can use the code in this page.

```
var notes = current.work_notes.getJournalEntry(-1);
//gets all journal entries as a string where each entry is
 delimited by '\n\n'
var na = notes.split("\n\n");

//stores each entry into an array of strings
 for (var i = 0; i < na.length; i++)
  gs.print(na[i]);
```

Journal field script values

The *setValue()* method is not supported for journal fields. Instead, assign values in script as in the following example.

```
var now_GR = new GlideRecord('incident');

//query priority 1 incidents in the state of either 'new' or
 'active'.
gr.addQuery('priority', 1);
var gc = gr.addQuery('state', 1);
gc.addOrCondition('state', 2);
gr.query();

while(gr.next())
{

//print a list of the incident numbers updated
gs.print(gr.number);

//add an entry to the 'work notes' journal field for each
 incident
gr.work_notes = "This is a high-priority incident. Please
 prioritize.";
gr.update();
}
```

Render journal field entries as HTML

Journal fields can render text enclosed within code tags as HTML.

Before you begin

- Role required: any role that grants write access to a journal field
- System property: the glide.ui.security.allow_codetag is set to the default value of **true**

Note:

To learn more about this property, see Allow JavaScript tags in embedded HTML (instance security hardening) in Instance Security Hardening Settings.

About this task

By default, a High Security Setting escapes any HTML code you type in a journal field by replacing it with its equivalent HTML entity value. Escaping causes the system to display HTML code as text rather than forwarding it to the browser as rendering instructions.

Procedure

1. Enter `[code][/code]` tags around any code you want to render as HTML.

   Note:
   A single journal entry can contain multiple code tags as long as each code tag has a beginning and ending tag.

   Example
   For example, enter these lines:

   Results of entering HTML code

Code entered	Result
`[code]<a href="http://www.service-now.com">ServiceNow</a>[/code]`	The system renders a hyperlink to the ServiceNow web site.
`[code]<b>This text will be bold.</b>[/code]`	The system renders the sentence in bold.
`<b>This text will not be bold.</b>`	The system escapes the bold tags and renders them as text.
`[code]<script> gs.info(gs.getUserDisplayName());</script>[/code]`	The system escapes the content of the script tag. **Note:** By default, the HTML Sanitizer prevents the entry of <script> elements.

   For more examples of HTML formatting options, see the blog post Formatting within Journal fields using HTML & [code] by a ServiceNow Technical Support Engineer in the ServiceNow Community.

2. Click **Post**.

   Note:
   You cannot edit previous journal entries.

Result

The system renders the text within code tags as HTML.

Related topics

[General security settings properties]

[HTML sanitizer]

Restrict the CODE tag in journal fields

You can prevent journal fields from rendering HTML code by disabling support for the `[code]` tag.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Clear the check box for **Allow support for embedding HTML code by using the [code] tag** (the `glide.ui.security.allow_codetag` property).

   Note:
   To learn more about this property, see Allow embedded HTML code (instance security hardening) in the Instance Security Hardening Settings.

3. Click **Save**.

Validate HTML in journal fields

Prevent users from saving invalid HTML in a journal field.

Before you begin
Role required: admin

Procedure

1. Add the property `glide.ui.allow_deep_html_validation`.
   For instructions, see Add a system property.
2. Set the **Value** to **true**.
3. Click **Save**.

   Users now see a warning in the activity formatter when they enter invalid HTML code in a journal field.

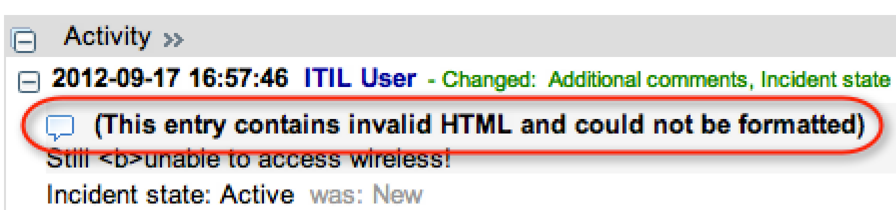

Journal field display limits

Journal fields can greatly increase the size of task records because they allow users to enter very large string values.

Display limits prevent the instance from loading the entire journal field into memory. Administrators have the option to:

- Set the length at which journal fields stop displaying the entire field's contents and instead only display a portion (called a preview) of the field's contents. Users can still access the field's entire contents by clicking a **Show All** button.
- Set the size of the preview text the journal field displays.
- Set the maximum number of journal entries journal fields can display.

In Core UI, journal fields and the activity formatter must be in the same form section. You cannot place another field in between a journal field and the activities formatter. The Activity Stream is built to keep journal fields stacked on top of the activity formatter.

Note:
Journal field display limits are not supported in the legacy **Comments and Worknotes** field.

Set the maximum display size for journal fields

You can set the maximum display size for journal fields by adding a system property.

Before you begin

Role required: admin

About this task

When a journal field exceeds the length set in this property, the instance shows a preview of the journal field instead of the field's entire contents. The preview includes a **Show All** button to display the rest of the field's contents. A separate property determines the number of characters the preview displays.

Procedure

1. Add a system property with the following settings.

Property	Value
Name	glide.max_journal_list_size
Description	Size in megabytes when a journal field should display a preview rather than the field's entire contents.
Type	Integer
Value	10

2. Click **Submit**.

Set the journal preview size

Set the size of the *glide.max_journal_list_size* system property. When a journal field exceeds this size, the instance displays a preview rather than the field's entire contents. Users can click the **Show All** button to see the rest of the field's contents.

Before you begin

Role required: admin

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Select **New**.
3. Set the values for the property as listed in the table.

Property	Value
Name	glide.shortened_journal_length
Description	Number of characters to display as a preview of journal fields.
Type	Integer
Value	512000

4. Select **Submit**.

Set the maximum number of journal entries

To set the maximum number of entries the system shows in the activity formatter, edit the following system property. The activity formatter displays the entries starting with the most recent entry up to the maximum number.

Before you begin

Role required: admin

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. Search for the property *`glide.history.max_entries`*.
3. Edit the **Value** of the property.

Property	Value
Name	glide.history.max_entries
Description	Maximum number of entries the system shows in the activity formatter. The default is 250.
Type	Integer
Value	250

4. Click **Update**.

Related topics

Activity formatter

Enable the text field character counter

By default, multi-line text fields have a 4000 character limit. To help users see how many characters remain before they reach the limit, you can enable the *`glide.ui.textarea.character_counter`* property.

Before you begin

Role required: admin

About this task

This property adds a counter under text fields, such as the **Additional Comments** and **Work notes** fields. The counter is dynamically updated as users enter text.

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Select the check box next to **Character counter for textarea (journal and multi line text fields)**.
3. Click **Save**.

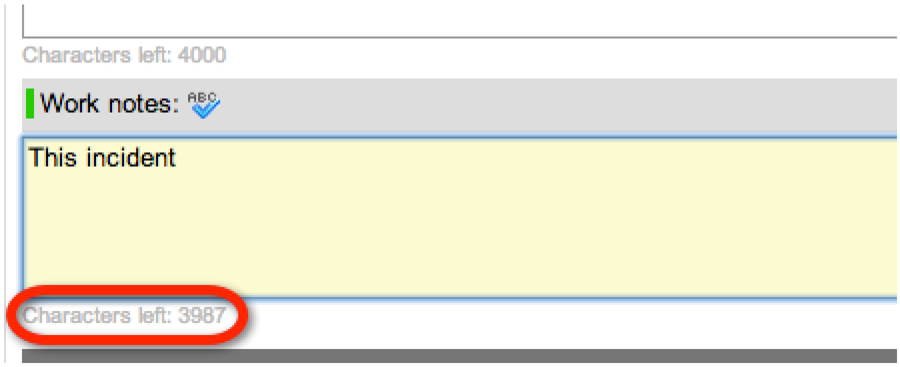

Name-value pairs field type

You can access the values stored in a name-value pairs field in scripts using the name.

Sample script

The following example demonstrates how to add mappings to a name-value pairs field, and how to query existing values using the name.

```
// Script example demonstrating setting and getting values
var gr = new GlideRecord('u_nv_table');
gr.initialize();

gr.nv_field.name1 = "value1"; //add a name-value Pair mapping
 with the name "name1" and value "value1"
gr.nv_field.name2 = "value2"; //add another name-value Pair
 mapping with the name "name2" and value "value2"

// Access by dot notation
gs.print("name1 = " + gr.nv_field.name1); // Expected output:
 name1 = value1

// Iterate over each property and print name and value
for(var name in gr.nv_field) {
    gs.print(name + " = " + gr.nv_field[name]);
}

gs.print(gr.nv_field); // Expected output:
 {"name1":"value1","name2":"value2"}
```

Percent complete field type

Administrators can create percent complete fields, which accept decimal input and appear as progress bars when displayed in lists.

For example, use a percent complete field to set the completion percentage for a task or project in a form and then see that percentage displayed as a progress bar in a list.

Administrators can also configure views that compare actual progress with a target value to determine if goals are being met and then apply color to provide visual alerts where progress does not meet expectations.

Form view of a percent complete field

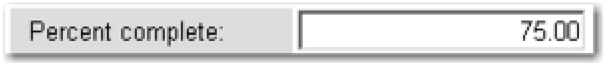

List view of a percent complete field

Target threshold colors attribute

If the *target_field* attribute is configured, a second attribute called *target_threshold_colors* enables an administrator to define additional parameters.

The parameters are:

- Different thresholds at which the colored bar should change color
- A specific color for each threshold

The format of this attribute's value is **number1:color1;number2:color2** and so on. Use this attribute to apply warning colors to completion percentages that are lower than target percentages. These values are defined as the percentage of target accomplished. For example, a value of **0:red;50:yellow;90:green** displays a red bar if the progress to target pecentage is between 0-49. If the percent of target is between 50 and 89, the color is yellow. Percent of target 90 and above displays in green. Completion percentages that exceed target percentages also display in green. Order the color attributes from the smallest percentage to the largest.

If you do not specify a target_field, then a target of 100 is used, allowing you to use the color thresholds with a single field value.

```
target_field=percent_complete_target,target_threshold_colors=0:tomato;50:khaki;90:lightgreen
```

The following table lists examples of percent of target calculation using the colors defined above.

Target threshold colors attribute

Target percent	Percent field value	Percent of target calculation	Color
100	40	40%	tomato
65	59	90.7%	lightgreen
15	10	66.7%	khaki

servicenow.

Percent complete color example

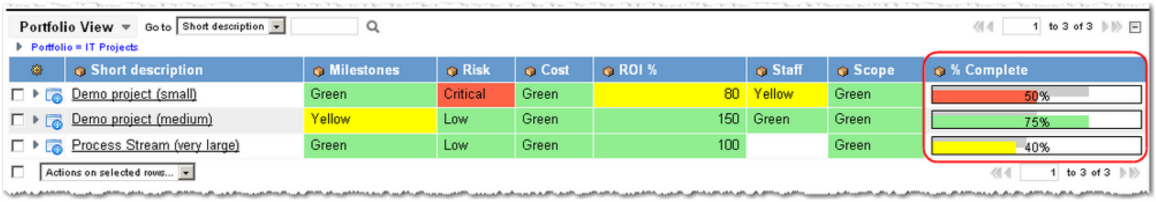

Add a target field attribute

Add an optional attribute (*target_field*) to a percent complete field to compare the actual completion percentage of a task or project with a target percentage in a different decimal field that specifies where the task should be at this point.

Before you begin

Role required: personalize_dictionary

About this task

If a target field is not specified, the target of 100 is assumed.

Procedure

1. Right-click the **% Complete** field in a form.
2. Select **Configure Dictionary** from the pop-up menu.
3. In the Dictionary Entry form, add the following attribute:
   target_field=percent_complete_target
4. Update the dictionary record.

   In the list, a gray bar appears behind the colored bar to indicate the target value. The gray target bar appears only if you defined a target field.

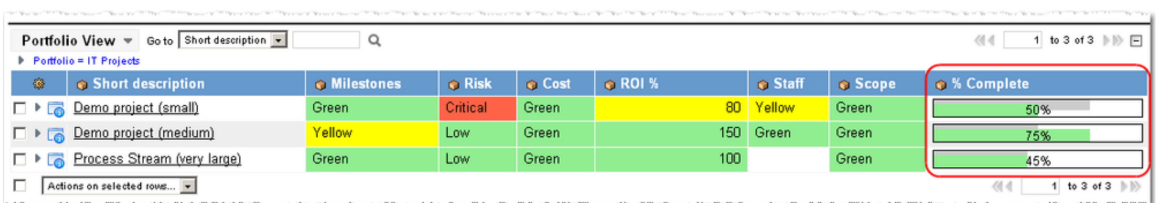

Phone number field type

The E.164 phone number standard ensures that all necessary information for a phone number is included and properly formatted to successfully route an international call over a territory's public telephone network.

When a user enters a phone number, it is received and stored as a string of numbers. An E.164 phone number field automatically formats and validates the numbers so that they are E.164-compliant when displayed as local and international numbers. The E.164 phone number field type does not replace the phone field type.

An E.164 phone number field displays:

- [Optional] A choice list for the phone number territory.
- [Always] An input box for entering phone numbers.

- [By Default] A red underline when a phone number does not match the format for the selected phone territory and cannot be saved.
- [Optional] A green underline when a phone number does not match the format for the selected phone territory but can be saved with **Other / Unknown** as the territory.

Phone e164

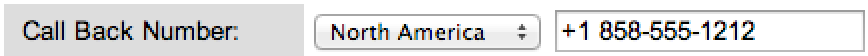

Territories assigned

Territories are assigned to locations, and are not assigned directly to users.

A user's territory, and so the user's E.164-compliant phone functionality, is based on the user's location. For example, if a user has a location of **SHS quadra 5, Bloco E., Brasilia** defined in the User [sys_user] table, the parent record for Brazil in the location table defines the phone territory. The phone territory may be assigned at any level of the Map pages hierarchy, which is searched going up to the next parent until the territory is found or no parents remain.

Dependent fields

In the dictionary, you can specify a dependent field in the **User** or **Location** fields, which displays the appropriate territory in the selector choice list when a user enters a phone number.

For example, if you enter `caller_id` in the dependent field in the Incident table, the appropriate territory is added to the territory selector choice list when a user enters caller information.

Related topics

Make a field dependent

E.164 phone number field configuration

Administrators can use the phone number system properties and dictionary attributes to do certain things.

The system properties apply the configuration option to all phone number fields that do not have a comparable dictionary attribute. The dictionary attributes apply the configuration only to the phone number field it is added to. Since dictionary attributes take precedence over system properties, administrators can set a global configuration with a property and then apply exceptions on a field-by-field basis.

Note:
Alpha numeric characters are not allowed and is not considered as a valid phone number.

Requiring territory format validation

By default, phone number fields require that a phone number match the display format of the selected territory.

By default, phone number fields require that a phone number match the display format of the selected territory. If a phone number does not match this format, the input box displays a red line underneath the phone number or the whole box turns red, and users are prevented from saving it.

Phone e164 strict

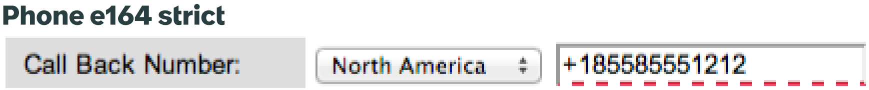

Setting the *glide.phone_number_e164.strict* system property to **false** or adding the *pn_strict* dictionary attribute allows the phone number input box to display a green line underneath numbers that do not match the territory format listed for the selected territory or the whole box turns green. You can use the **Other / Unknown** territory to store otherwise invalid phone numbers.

Phone e164 not strict

> **Warning:**
> Switching from optional territory format validation to required territory format validation may result in some phone numbers failing validation altogether. In such cases, the E.164 phone number field displays an error message.

Requiring entry of international format

By default, users can enter phone numbers in their territory's local format and do not have to format the number for international dialing.

The phone number field automatically formats local phone numbers into E.164-compliant international phone numbers when the user finishes editing the field. As long as the phone number entered matches the territory's format for a local number, users can save the phone number.

Phone e164 entry local

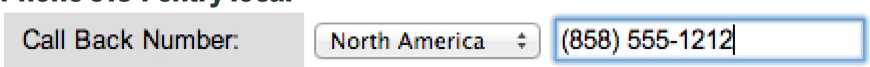

Phone e164 display international

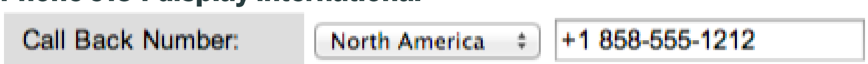

Changing the *glide.phone_number_e164.allow_national_entry* system property or adding the *pn_allow_national_entry* dictionary attribute requires users to enter a phone number in the territory's international format, which starts with the plus (+) character. Users cannot save a locally formatted phone number, and the phone number input box displays a red line underneath phone numbers without the proper international formatting.

Configure the display of the local format

By default, an E.164 phone number field always displays phone numbers in an international format.

Phone e164 display international

Change the E.164 phone number to use the local format in certain circumstances by adding the following system property:

Property	Description
glide.phone_number_e164.display_national	• **Type** : string • **Default value**: false • **Other possible values**:

(continued)

Property	Description
	○ **true or form**: displays phone numbers in a local format on forms, but displays an international format on lists. ○ **all**: always displays phone numbers in a local format. ○ **user**: only displays phone numbers in a local format when the phone number matches the local setting of the current user. ○ **false**: does not display phone numbers in local format.

Select one of the possible values to determine how the system handles the E.164 phone number.

You can also add the following dictionary attribute to a specific field to override the system property:

Dictionary attribute	Description
pn_display_national	Overrides the *glide.phone_number_e164.display_national* property setting for how an E.164 phone number field displays phone numbers. Available values are identical to those described for the *glide.phone_number_e164.display_national* property. • Type: string • Default value: false • Example:pn_display_national=all

Configure the phone territory selector choice list

By default, an E.164 phone number field always displays the phone territory associated with the phone number.

Changing the system property or adding the dictionary attribute hides the territory selector choice list. If the territory selector choice list is hidden, users can only enter a local or national number.

Phone e164 hide selector

Add the following system property to show or hide the territory selector choice list.

Property	Description
glide.phone_number_e164.display_territory_selector	Determines whether to display the territory selector choice list. Hiding the territory selector choice list restricts users to entering only local or national phone numbers. • **Type** : true | false • **Default value**: true

You can also add the following dictionary attribute to a specific field to override the system property:

Dictionary attribute	Description
pn_display_territory_selector	Overrides the *glide.phone_number_e164.display_territory_s* property setting that determines whether to display the territory selector choice list. Available values are identical to those described above for the *glide.phone_number_e164.display_territory_s* property. • Type: true | false • Default value: true • Example: pn_display_territory_selector=false

Configure the display of territory labels

A property controls how territory labels are displayed.

You can display territory labels to the right of the number in an E.164 phone number field by setting the system property or adding the dictionary attribute. This is useful if the territory selector choice list is turned off and you want the user to see the territory for the entered phone number.

Enabling territory labels also displays the phone territory in lists.

Add the following system property to display the territory label to the right of the number in an E.164 phone number.

Property	Description
glide.phone_number_e164.display_territory_text	Determines when an E.164 phone number field displays a territory label. • **Type** : string • **Default value**: read-only • **Other possible values**: ○ **all**: always displays the territory label. ○ **national**: displays the territory label only if the phone number is in local format. ○ **read-only**: displays the territory label in read-only mode, regardless of whether the number is in local or international format. ○ **read-only-national**: displays the territory label in read-only mode only if the number is in local format. ○ **list**: displays the territory label in a list. ○ **list-national**: displays the territory label in a list if the number is in national format. ○ **none**: does not display the territory label.

You can also add the following dictionary attribute to a specific field to override the system property:

Dictionary attribute	Description
pn_display_territory_text	Overrides the *glide.phone_number_e164.display_territory_t* property that defines when a phone number field displays a territory label. Available values are identical to those described above for the *glide.phone_number_e164.display_territory_t* property. • Type: string • Default value: read-only • Example:pn_display_territory_text=all

Configure the international direct dialing prefixes

A property is available to control the display of prefixes.

You can enable the display of the international direct dialing prefix, which appears between the territory selector choice list and the input box for an E.164 phone number field on forms, by setting the system property or adding the dictionary attribute.

E164 phone display idd

Add the following system property to display the international direct dialing prefix.

Property	Description
glide.phone_number_e164.display_users_idd	Determines whether to display the international direct dialing prefix between the territory selector choice list and the input box on forms. • **Type** : true | false • **Default value**: false

You can also add the following dictionary attribute to a specific field to override the system property:

Dictionary attribute	Description
pn_display_users_idd	Overrides the *glide.phone_number_e164.display_users_idd* property that determines whether to display the international direct dialing prefix between the territory selector choice list and the input box on forms. Available values are identical to those described above for the *glide.phone_number_e164.display_users_idd* property. • Type: true | false • Default value: false • Example:pn_display_users_idd=false

E.164 phone number field system properties

Several properties are available to configure E.164 phone number fields.

Note:
To open the System Properties [sys_properties] table, enter `sys_properties.list` in the navigation filter.

e.164 phone number field system properties

Property	Description
glide.phone_number_e164.strict	Determines whether all phone number fields must match the display format for the selected territory.

e.164 phone number field system properties (continued)

Property	Description
	When the value is **true**, the phone number input box displays a red line underneath phone numbers that do not match the format for the selected territory. Users cannot save the phone number. When the value is **false**, the phone number input box displays a green line underneath phone numbers that do not match the format for the selected territory. Users can save the phone number. The territory selector choice list offers the option to select an **Other / Unknown** territory format. • **Type** : true | false • **Default value**: true • **Location**: Add to the System Properties [sys_properties] table.
glide.phone_number_e164.allow_national_entry	Determines whether users can enter phone numbers in the local format or whether they must enter phone numbers in international format. When **true**, users can enter phone numbers in the local format for the selected territory. When **false**, users must enter phone numbers in the international format for the selected territory. • **Type** : true | false • **Default value**: true • **Location**: Add to the System Properties [sys_properties] table.
glide.phone_number_e164.display_national	Determines whether to display E.164 phone numbers in local format. • **Type** : string • **Default value**: false • **Other possible values**: ◦ **true or form**: Displays phone numbers in a local format on forms, but displays an international format on lists.

e.164 phone number field system properties (continued)

Property	Description
	○ **all**: Always displays phone numbers in a local format. ○ **user**: Only displays phone numbers in a local format when the phone number matches the local setting of the current user. The local setting of the current user is determined by the location associated with the user record. ○ **false**: Does not display phone numbers in local format. • **Location**: Add to the System Properties [sys_properties] table.
glide.phone_number_e164.display_territory_selector	Determines whether to display the territory selector choice list. Hiding the territory selector choice list restricts users to entering only local or national phone numbers. • **Type** : true | false • **Default value**: true • **Location**: Add to the System Properties [sys_properties] table.
glide.phone_number_e164.display_territory_text	Determines when an E.164 phone number field displays a territory label. • **Type** : string • **Default value**: read-only • **Other possible values**: ○ **all**: Always displays the territory label. ○ **national**: Displays the territory label only if the phone number is in local format. ○ **read-only**: Displays the territory label in read-only mode, regardless of whether the number is in local or international format. ○ **read-only-national**: Displays the territory label in read-only mode only if the number is in local format. ○ **list**: Displays the territory label in a list.

e.164 phone number field system properties (continued)

Property	Description
	◦ **list-national**: Displays the territory label in a list if the number is in national format. ◦ **none**: Does not display the territory label. • **Location**: Add to the System Properties [sys_properties] table.
glide.phone_number_e164.display_users_idd	Determines whether to display the international direct dialing prefix between the territory selector choice list and the input box on forms. • **Type** : true | false • **Default value**: false • **Location**: Add to the System Properties [sys_properties] table.

E.164 phone number field dictionary attributes

You can override the global system property with certain dictionary attributes.

e.164 phone number field dictionary attributes

Dictionary attribute	Description
pn_strict	Overrides the *glide.phone_number_e164.strict* property setting that requires all phone number entries match the format for the selected territory. Available values are identical to those described for the *glide.phone_number_e164.strict* property. • Type: true | false • Default value: true • Example:pn_strict=false
pn_allow_national_entry	Overrides the *glide.phone_number_e164.allow_national_entry* property setting that determines whether users can enter phone numbers in the local format or whether they must enter phone numbers in international format. Available values are identical to those described for the *glide.phone_number_e164.allow_national_entry* property. • Type: true | false • Default value: true • Example:pn_allow_national_entry=false
pn_display_national	Overrides the *glide.phone_number_e164.display_national* property setting for how an E.164 phone number field displays phone

e.164 phone number field dictionary attributes (continued)

Dictionary attribute	Description
	numbers. Available values are identical to those described for the *glide.phone_number_e164.display_national* property. • Type: string • Default value: false • Example:pn_display_national=all
pn_display_territory_selector	Overrides the *glide.phone_number_e164.display_territory_selector* property setting that determines whether to display the territory selector choice list. Available values are identical to those described for the *glide.phone_number_e164.display_territory_selector* property. • Type: true | false • Default value: true • Example: pn_display_territory_selector=false
pn_display_territory_text	Overrides the *glide.phone_number_e164.display_territory_text* property that defines when a phone number field displays a territory label. Available values are identical to those described for the *glide.phone_number_e164.display_territory_text* property. • Type: string • Default value: read-only • Example:pn_display_territory_text=all
pn_display_users_idd	Overrides the *glide.phone_number_e164.display_users_idd* property that determines whether to display the international direct dialing prefix between the territory selector choice list and the input box on forms. Available values are identical to those described for the *glide.phone_number_e164.display_users_idd* property. • Type: true | false • Default value: false • Example:pn_display_users_idd=false

Configure a territory phone display rule

The string of numbers that make up a phone number is automatically validated and formatted for a specific territory by applying a series of regular expressions.

Before you begin

Role required: admin

About this task

The number is first validated against the phone validations that have been defined for the territory, and in the order specified by the **Order** field. To be valid, the number must match the regular expression defined in the **Condition** field for at least one phone validation.

After a number has been validated, the **Condition** expression for each format defined for the territory is applied to the number in the order determined by the **Order** fields. The **Pattern** and **Format** regular expressions are applied to produce a phone number that is formatted correctly for the territory.

The Sys Phone Territory screen allows administrators to edit the display rules for a given territory. Administrators may want to modify the **Active**, **Display**, or **Order** fields. To edit the display rules for a territory:

Procedure

1. Navigate to **All** > **System Policy** > **Rules** > **Telephone Display Rules**.
2. Click a territory **Name**.
3. Edit the fields, as appropriate (see table).
4. Click **Update**.

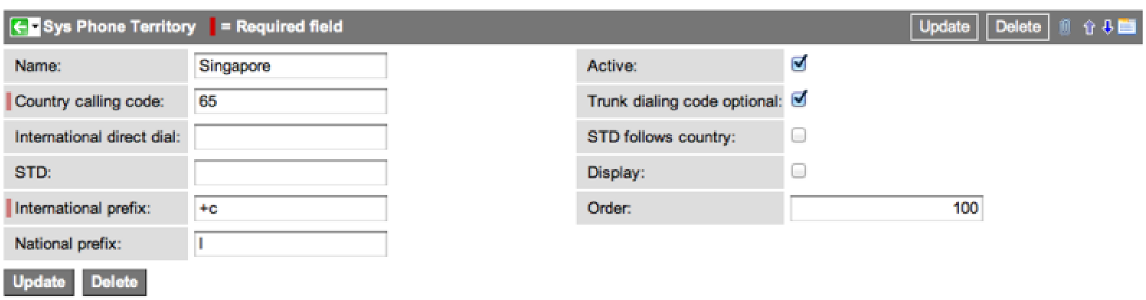

Territory phone display rules

Field	Description
Name	The unique name of the territory.
Country calling code	The country code for dialing numbers from outside the territory.
International direct dial	The prefix for calling internationally from the territory, such as 00 or 001.
STD	The subscriber trunk dialing code, also know as the direct distance dialing code, which is a sequence of numbers before the telephone number that indicate whether the call is to be routed outside of the local calling area.
International prefix	The prefix required to dial an international call, such as a plus sign (+).
National prefix	The prefix required to dial a local call.
Active	An indicator for whether the territory phone definition is active. If deactivated, this territory unavailable to users.
Trunk dialing code optional	An indicator for whether the STD code is optional.

Field	Description
STD follows country	An indicator for whether the STD code should be displayed to the right of the country calling code.
Display	An indicator for whether to display the territory in the choice list. Clearing this check box removes the territory from the choice list. If an international number is entered for a territory that is not displayed in the territory selector choice list, that territory is temporarily added to the selector choice list for that field only. For example, If the United Kingdom **Display** field is not selected, the United Kingdom does not appear in the territory selector choice list. However, if the user enters an international number beginning with +44, the United Kingdom is added to the list and the number is formatted and validated accordingly.
Order	The order in which a territory appears in a choice list. Territories are sorted numerically by the number assigned here. If more than one territory is assigned the same number, they are subsorted alphabetically. All territories are assigned a default value of 100. To display a territory at the top of the list, assign a value that is less than 100. To display a territory at the end of the list, assign a value that is greater than 100. For example, if a territory is assigned an order of 500, it is displayed at the end of the list, and if more than one territory is assigned an order of 500, they are listed alphabetically at the end of the list.

Phone validations

Phone validations are already configured for all territories and are automatically applied to the phone number to ensure that the number is valid for the territory.

E164 phone validations

Phone validations New Go to Order 1 to 7 of 7

Phone territory = Singapore

<table>
<tr><th>Name</th><th>Condition</th><th>Order</th></tr>
<tr><td>SG fixedLine</td><td>[36]\d{7}|[17-9]\d{7,10}</td><td>100</td></tr>
<tr><td>SG fixedLine</td><td>6[1-9]\d{6}</td><td>110</td></tr>
<tr><td>SG mobile</td><td>(?:8[1-7]|9[0-8])\d{6}</td><td>120</td></tr>
<tr><td>SG tollFree</td><td>1?800\d{7}</td><td>130</td></tr>
<tr><td>SG premiumRate</td><td>1900\d{7}</td><td>140</td></tr>
<tr><td>SG voip</td><td>3[12]\d{6}</td><td>150</td></tr>
<tr><td>SG uan</td><td>7000\d{7}</td><td>160</td></tr>
</table>

Actions on selected rows... 1 to 7 of 7

Phone formats

Phone formats are already configured for all territories and are automatically applied to the phone number to ensure that the number is valid for the territory.

E164 phone formats

Phone formats New Go to Order 1 to 4 of 4

Phone territory = Singapore

<table>
<tr><th>Name</th><th>Condition</th><th>Format</th><th>Order</th><th>Pattern</th></tr>
<tr><td>SG availableFormats</td><td>[369]|8[1-9]</td><td>$1 $2</td><td>100</td><td>([3689]\d{3})(\d{4})</td></tr>
<tr><td>SG availableFormats</td><td>1[89]</td><td>$1 $2 $3</td><td>110</td><td>(1[89]00)(\d{3})(\d{4})</td></tr>
<tr><td>SG availableFormats</td><td>70</td><td>$1 $2 $3</td><td>120</td><td>(7000)(\d{4})(\d{3})</td></tr>
<tr><td>SG availableFormats</td><td>80</td><td>$1 $2 $3</td><td>130</td><td>(800)(\d{3})(\d{4})</td></tr>
</table>

Actions on selected rows... 1 to 4 of 4

Reference field type

A reference field stores a reference to a field on another table. For example, the **Caller** field on the Incident table is a reference to the User [sys_user] table.

When you define a reference field, the system creates a relationship between the two tables. Adding a reference field to a form makes the other fields in the referenced table available to the form.

Note:
A reference field can refer only to records from one other table. To add a field that can refer to records on any table, use the Document ID element type.

Administrators can create new reference fields and configure several options for reference fields.

Reference field options

Option	Description
Display values	Each reference field stores a sys_id for each referenced record in the database, but the sys_id is not shown. The reference field shows the specified display value.
Decorations	A reference decoration is an icon that appears next to a reference field.
Reference styles	Reference styles are specialized field styles that control the appearance of reference fields.
Reference qualifiers	Reference qualifiers restrict the records that are available for reference fields.
Cascade delete rules	Cascade delete rules specify what should happen to records that reference a record that is deleted.
Auto-complete	By default, a reference field auto-completes as the user types in the field. Administrators can configure auto-complete settings.
Reference key	A reference key saves a field other than sys_id as the unique identifier for a reference field.
Enable dynamic creation	When dynamic creation is enabled, entering a nonexistent value in a reference field creates a new record on the referenced table instead of returning an error.

Tip:
For troubleshooting information, see the Reference field is not showing the expected display value when selected or it appears blank [KB0693859] article in the Now Support Knowledge Base.

Add a reference field

Add reference fields to a table using the same method as for any other field.

Before you begin

Role required: personalize_form

About this task

The related table also appears in the Available Tables list for future form customizations.

Procedure

1. Open the desired form.
2. Right-click the header and select **Configure** > **Form Layout**.
3. Use dot-walking to locate and select the field in the referenced table that you want to add.

Example
It appears as *`Table name.Field`*. For example, the caller's email address appears as *`Caller.Email`*.

4. Click **Save**.

Enable dynamic creation for reference fields

When dynamic creation is enabled, entering a nonexistent value in a reference field creates a new record on the referenced table instead of returning an error.

Before you begin

Role required: personalize_dictionary

About this task

By default, a user must enter a value in a reference field that matches an existing record in the table that the reference field refers to. For example, the **Caller** field in an Incident must have a value that is an existing user. You can enable dynamic creation to create a new record on the referenced table when a user enter a nonexistent value in a reference field instead of returning an error.

Procedure

1. Right-click the field label in the form and select **Configure Dictionary**.
2. Populate the following fields (you may need to configure the Dictionary form):
   - **dynamic_creation**: Select the check box.
   - **dynamic_creation_script**: Enter a script that dynamically creates the record.
3. Click **Update**.
   Examples:

   You could use the following `dynamic_creation_script` to create a record on the referenced table.

   ```
   current.name = value;
   current.insert();
   ```

   Note:
   The parent object can be used to access anything from the parent record.

   You could create a script include named *`MyUserReferenceCreator`* with the following contents:

   ```
   var MyUserReferenceCreator = Class.create();
   MyUserReferenceCreator.prototype = {
       initialize: function() {
       },

       create: function(current, value) {
           current.name=value;
           return current.insert();
       },

       type: 'MyUserReferenceCreator'
   }
   ```

When the script include is created, the following `dynamic_creation_script` generates a new location for an invalid reference field value:

```
new MyUserReferenceCreator().create(current, value);
```

Related topics

Configuring the form layout

Configure cascade delete rules

When a record is deleted, there are different options for how the deletion will affect records that reference the deleted record. You can configure what happens to records that reference a record when that record is deleted.

Before you begin

Role required: personalize_dictionary

About this task

For example, if you delete a user record that is referenced in the **Caller ID** field on several incident records, you can configure what happens to those incident records. By default, the references are cleared, so the incident records are maintained with an empty **Caller ID** field.

Procedure

1. Navigate to a reference field on a form.
2. Right-click the field label and select **Configure Dictionary**.
3. Under **Related Links**, click **Advanced view**.
4. Locate the Reference Specification section.
5. In the **Reference cascade rule** field, select one of the following options.

   Note:
   The option for this field have no effect for tables with m2m relationships, meaning that in this example, if the tables have an m2m relationship, the user record is deleted.

6. Click **Update**.

Related topics

Modify dictionary entries

Define the reference key

By default, reference fields store the sys_id of the record in the database.

Before you begin

Role required: personalize_dictionary

About this task

By defining a reference key, you can identify a field other than sys_id to use as the unique identifier for the reference field. The value of the reference key field, instead of the sys_id, is stored in the database for that reference field.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the field record (for example, **resolved_by** on the Incident table).
3. In the **Reference key** field, enter a field name on the referenced table (for example, email on the sys_user table).

    Note:
    Always choose a field from the referenced table that is both required and unique.

4. Select **Update**.

Display a reference field as a choice list

You can display a reference field as a choice list instead of opening a lookup window.

Before you begin

Role required: personalize_dictionary

About this task

Note:
A choice reference field with a defined reference key won't function.

Procedure

1. In the form, right-click the label for the reference field and select **Configure Dictionary**.
2. In the Choice List Specification section, select one of the following options in the **Choice** field.
    - **Dropdown with --None--**
    - **Dropdown without --None--**
3. Right-click the form header and click **Save**.
4. Under **Related Links**, click **Advanced view**.
5. In the **Attributes** field, add the following attribute, separated from other attributes by a comma.

    ```
    ref_auto_completer=AJAXReferenceChoice
    ```

6. Click **Update**.
    The form reopens, with the reference field as a choice list.

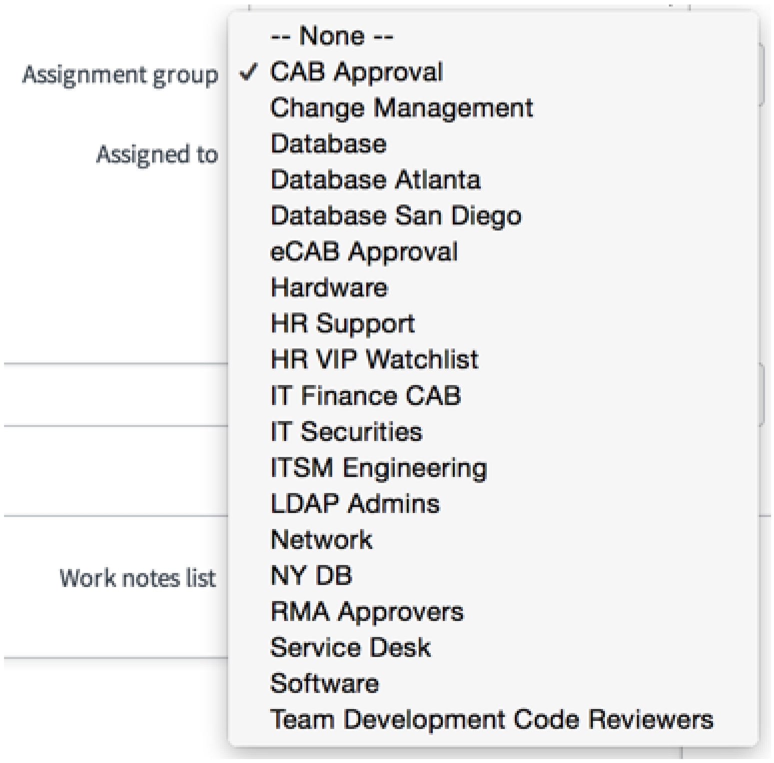

Note:

The *glide.ui.max_ref_dropdown* system property determines whether a reference field is displayed as a choice list or a reference icon. When the number of available choices exceeds the value of this property, a reference icon displays instead of a choice list. The default value of this property is 25. This property affects the entire instance, however you can override this property on an individual field by using the *max_ref_dropdown* dictionary attribute. Modifying dictionary attributes requires the personalize_dictionary role. If the column is not defined on the same table, you should add these attributes as a dictionary attributes override.

The *glide.xmlhttp.max_choices* system property determines how many choices display on a choice list. This property has a value of 15 by default, however, when the property is not present on the instance, the instance uses a hard coded value of 10.

Related topics

System dictionary

Display values

Reference fields store a sys_id for each referenced record in the database, but the sys_id is not shown.

The reference field shows the display value. For example, an incident's **Assigned to** field stores the sys_id of a particular user, but actually displays the user's name. The following example shows how **Charlie Witherspoon**, which is the display value of a user record, is shown in the **Assigned to** field.

Display value xml

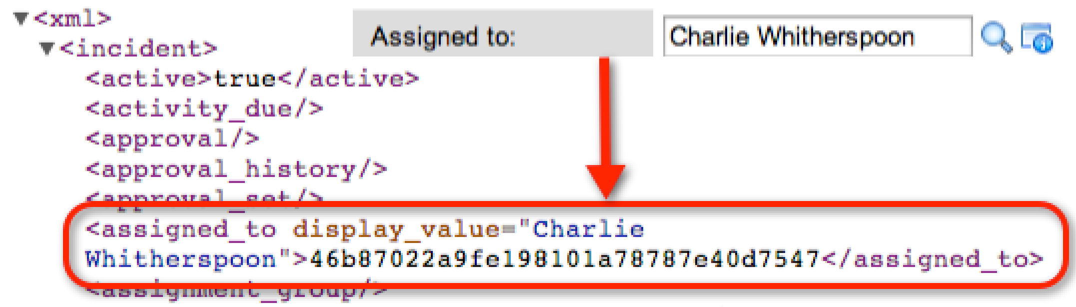

Reference field

Reference field	Value stored in database	Display value field of source table	Value displayed in UI
Assigned to [task.assigned_to]	46b87022a9fe198101a78787e40d7547	Name [sys_user.name]	Charlie Whitherspoon

Reference fields show display values in:

- Lists
- Forms
- Reports
- Auto-complete suggestions
- Slushbuckets

Related topics

Unique record identifier (sys_id)

Select a field as the table display value

Only one field can be defined as the display value for a table.

Before you begin

Role required: personalize_dictionary

About this task

When you set the **Display** value to **true**, a business rule sets the **Display** value to **false** for all other fields on the table. In previous versions, you must manually ensure that no other fields on the table have a value of **true** in the **Display** column.

Note:
Extended tables inherit the display value of the parent table. Setting a separate display value for the extended table overrides the parent table's display value.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Filter on **[Table] [is] [<name of the referenced table>]**.
3. Locate the desired field and set **Display** to **true**.

   For best results, choose a field that is required and unique in each record as the display value field.

> **Note:**
> If you make a field the display field for a table, be sure to translate all values for the field in the Translated Text [sys_translated_text] table into all the languages provided. Display field options left untranslated are not presented by the autocomplete (type ahead) feature.

Reference fields look for the display value in the following order:

a. A field with display=true in the system dictionary on the lowest sub-table for extended tables.
b. A field with display=true in the system dictionary on the parent table.
c. A field named **name** or **u_name** or **x_name**.
d. A field named **number** or **u_number** or **x_number**.
e. The field named in the Glide Property *`glide.record.display_value_default`*
f. The **Created on** field of the referenced record.

Decorations

Reference decorations are icons which appear next to the reference field.

The reference lookup is always visible and is used to select a record to reference. Other reference decorations appear when a record is selected.

Reference lookup

On forms, the reference lookup icon () appears by editable reference fields. Clicking the reference lookup icon displays a list of records on the referenced table.

When List v3 is enabled, the list appears in a popover. Otherwise, the list appears in a pop-up window.

Reference lookup list popover

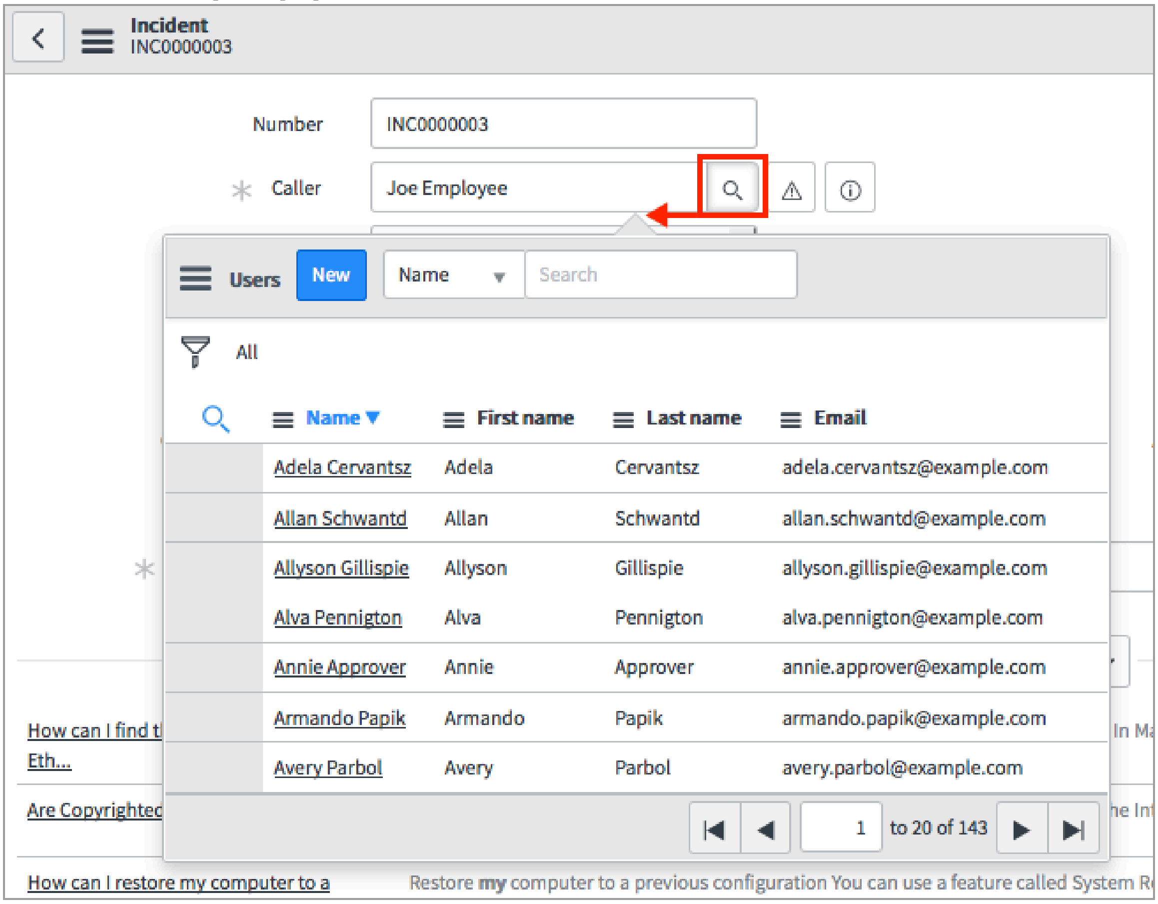

Reference lookup list pop-up

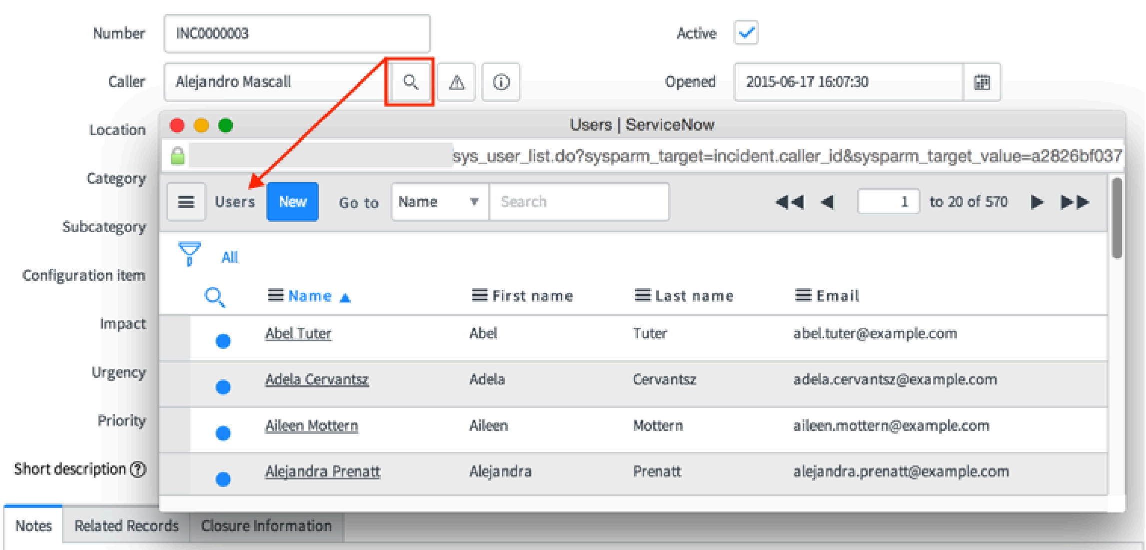

Tree picker lookup

The reference lookup can be rendered in the tree picker format by modifying the dictionary and adding the attribute `tree_picker`.

You cannot customize the label names used in the tree picker. The label names are taken from the values in the table.

Note:
The tree display hierarchy does not appear in reference fields when using mobile applications.

Related topics

[Tree picker](#)

Reference field icon

On forms, the reference icon ( ) appears by populated reference fields. Clicking the icon opens a read-only preview of the referenced record.

Referenced record preview on a form

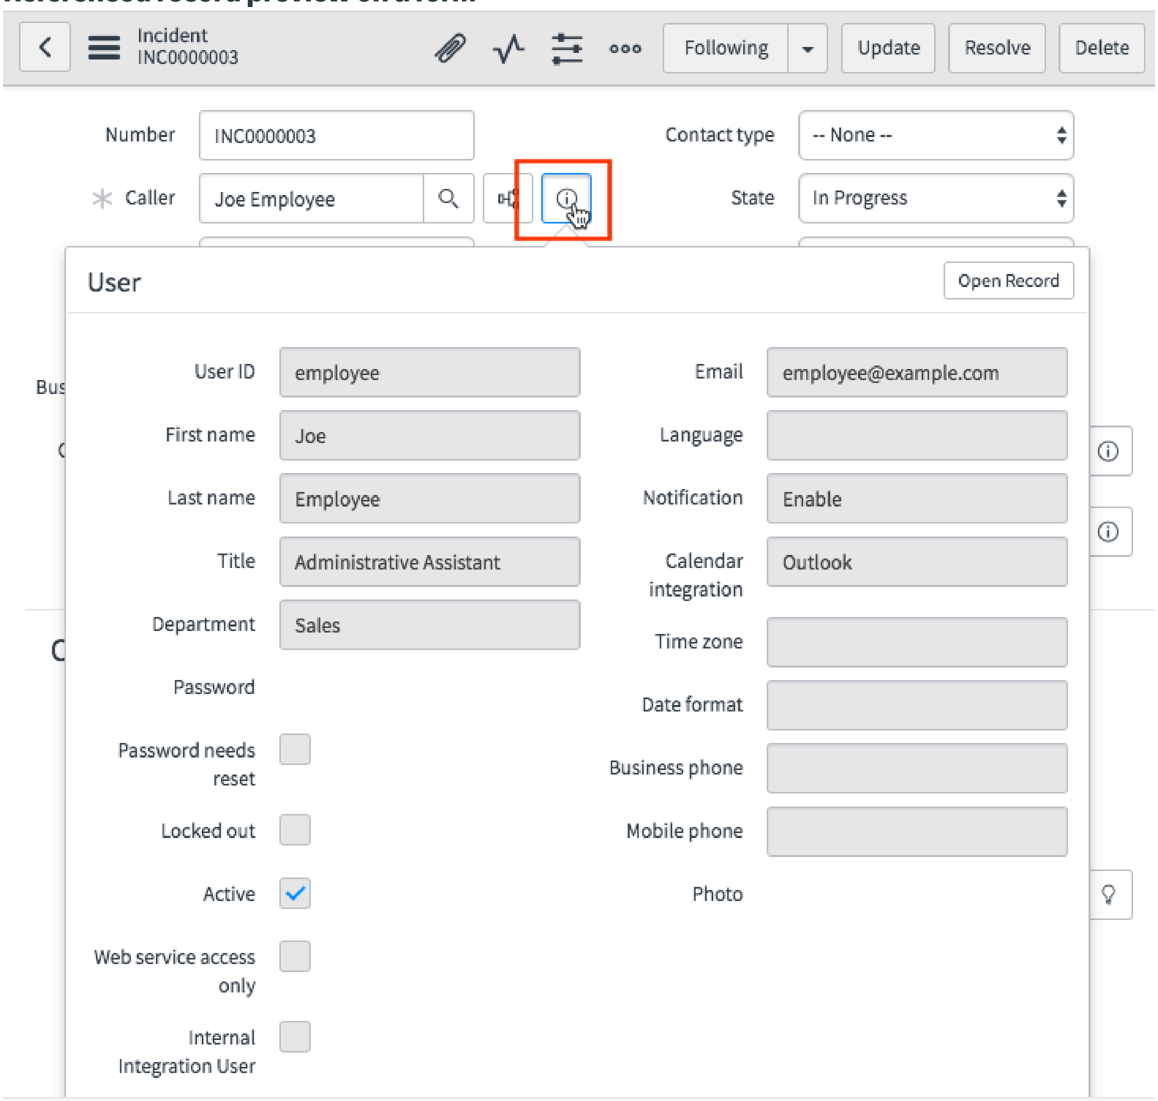

The preview remains open until you click somewhere else on the form.

Click **Open Record** to navigate to the referenced record.

Note:
When using a reference icon in Service Portal, a form opens using the default view of the form rather than a popup view. Write permissions to this form are the same as opening the form in the default view in the standard UI.

Configure the reference icon view of fields

Use a table's sys_popup form view to configure the fields in the pop-up form that appear when pointing to a reference icon. If the table has no sys_popup view, the pop-up uses the default view.

Before you begin
Role required: personalize_form

About this task
The pop-up form displays in the top form section only.

Procedure

1. To configure a reference field popup form for a table using the default sys_popup view, navigate to the following URL format, substituting the instance name and table name:

```
<your instance name>.service-now.com/<table
 name>.do?sysparm_view=sys_popup
```

Note:
This URL format only shows a table and default sys_popup view. It does not work for records that use a different view.

Example
An example of an instance for Acme, the sys_user table, and the sys_popup default view:

```
acme.service-now.com/sys_user.do?sysparm_view=sys_popup
```

2. **Optional:** To configure a reference field popup form for a table using a non-default sys_popup view, navigate to the following URL format, substituting the instance name, table name, and name of view:

```
<your instance name>.service-now.com/<table
 name>.do?sysparm_view=sys_popup,<name_of_view>
```

Example
An example of an instance for Acme, the sys_user table, and the sys_popup ESS view:

```
acme.service-now.com/sys_user.do?sysparm_view=sys_popup,ess
```

3. Configure the form to add or remove fields as appropriate.

Configure pop-ups on read-only fields

Reference pop-ups and click-throughs are hidden by default when a client script, UI policy, variable, or ACL makes the field read-only. The ability to see or click through to the target record does not depend on whether the reference field is writable. You can change the read-only setting.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Change the value of the **Enable click-through of a reference field when the reference field is read-only.** (`glide.ui.reference.readonly.clickthrough`) property.
   If set to true, the pop-up appears for read-only fields and for variables.

What to do next

If this system value is set to **false**, you can override the setting for a specific read-only reference field. Configure the dictionary entry and add the `readonly_clickthrough=true` attribute.

Dictionary entry showing the clickthrough attribute

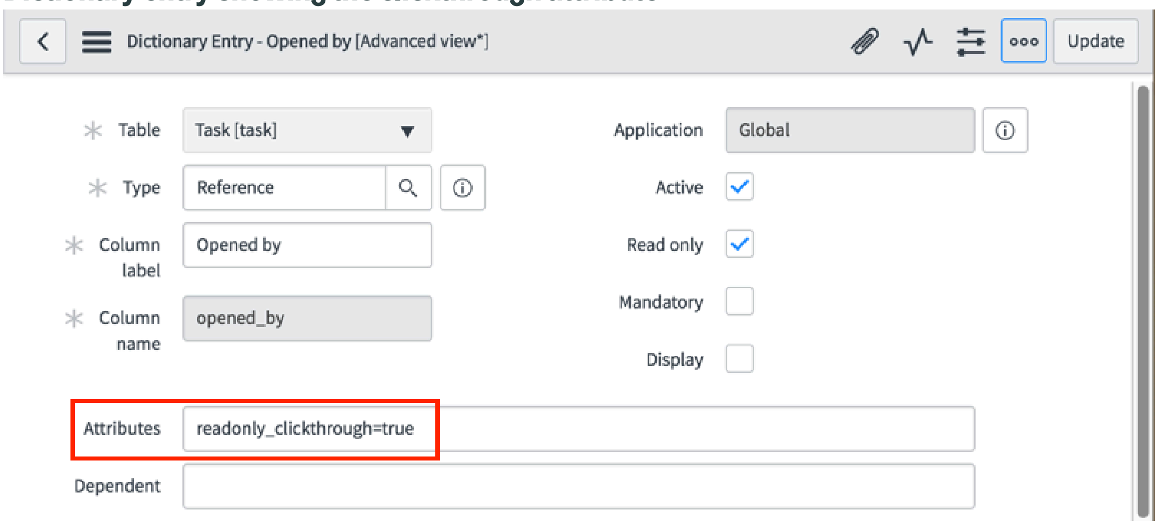

Configure the related incidents icon

You can configure an icon to appear beside a reference field in a form, such as the **Caller** field, to display related incidents.

Before you begin

Role required: admin

About this task

The show related incidents icon () displays other incidents related to the referenced record.

Procedure

1. In the form, right-click the label for the reference field and select **Configure Dictionary**.
2. Add the `ref_contributions=user_show_incidents` dictionary attribute in the **Attributes** field.
3. Click **Update**.
   The form reopens and the related incidents icon appears beside the field on the right.

Related topics

Altering tables and fields using dictionary attributes

Configure the show workflow icon

You can configure an icon to appear beside a workflow field to display the related workflow in the Workflow Editor.

Before you begin

Role required: admin

About this task

The show workflow icon ( ) opens the workflow in the Workflow Editor.

Procedure

1. In the form, right-click the label for the workflow field and select **Configure** > **Configure Dictionary**.
2. Add the *ref_contributions=show_workflow* dictionary attribute in the **Attributes** field.
3. Click **Update**.
   The form reopens and the show workflow icon appears beside the field on the right.

Related topics

Altering tables and fields using dictionary attributes

Reference qualifiers

Use reference qualifiers to filter the data that is returned for a reference field.

A reference field stores a link (reference) to a field on another table, making the records/fields in the referenced table available to the form containing the reference field.

For example, the **Assigned to** field on the Incident table is a reference to the User [sys_user] table. By default, all values for the field that is being referenced appear in the reference lookup and can be directly accessed through the reference field (type ahead). Expanding on the prior example, if a reference qualifier is not defined, all users in the User table appear in the reference lookup. Including those users that are inactive. Sometimes, this might be the desired functionality. In other cases however, only a subset of the available values may be desired. In this case, create a reference qualifier to filter the available data so that only the desired values are returned and made available to the form. Such as only the active users or users that have a specific role. Reference qualifiers are robust and can consist of simple AND/OR conditions, inline JavaScript, or complex script include.

You can modify the reference qualifier for a table, and any table based on that table (parent or extended), by defining a reference qualifier through the Dictionary Entry form. You can also modify the reference qualifier only on an extended table and its children (not the parent table), through a dictionary override. You can only define a single reference qualifier per field, per form/table. Reference qualifiers are not applicable to condition builders. For information on using filtering in condition builders, see Create a dynamic filter option.

Note:

- Creating reference qualifiers requires knowledge of the underlying ServiceNow data model (tables and fields) and knowledge of the Web services and Scripts.
- To restrict what data specific users can access, use ACLs not reference qualifiers.

You can define a reference qualifier using one of the following methods.

Simple reference qualifier

Simple reference qualifiers use AND/OR statements (conditions) to create simple filters. Use simple reference qualifiers when filtering on conditions such as whether a company is active, a user has a specific role, and/or a caller is in a specific time zone. Simple reference qualifiers can

have a maximum of 13 reference qualifier conditions. For additional information on how to use condition builders, see Condition builder.

Simple reference qualifier example

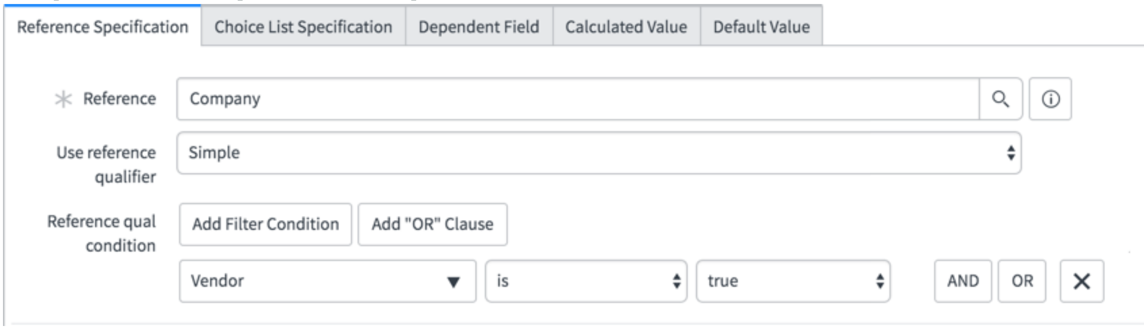

Dynamic reference qualifiers

Dynamic reference qualifiers enable you to use a Create a dynamic filter option to run a query against a reference field to filter the returned data set. Dynamic filter options are stored filters that can contain Encoded query strings, JavaScript, or script includes, and can be used in multiple dynamic reference qualifiers. Changes made to a dynamic filter option automatically apply to all reference qualifiers that use the same dynamic filter option. Use this type of reference qualifier when you want to use the same filter on multiple forms or to provide filter functionality to "non-code savvy" implementers.

The base instance provides several OOB dynamic filter options. If a dynamic filter option that meets your needs does not exist, you can create a new dynamic filter option that is specific to your requirements. An example of an OOB dynamic filter option is the reference qualifier on the **Model ID** field on a configuration item form, such as the Computer form. The reference qualifier calls the **CI Model Qualifier** dynamic filter option, which in turn calls the *`ModelAndCategoryFilters`* script include. This script include filters the data set based on the class of the CI. The only options for the model ID are options that belong to the same class as the current CI. For example, only CIs that belong to the Computer class are available in the **Model ID** field on the Computer form.

To locate the available dynamic filter options, navigate to **System Definition** → **Dynamic Filter Options**. In the right-corner of the Dynamic Filter Options list, click the filter icon and create the filter condition `Available for ref qual is true`. All dynamic filter options that can be used in dynamic reference qualifiers appear.

Dynamic reference qualifier example

Advanced reference qualifier

Advanced reference qualifiers enable you to define an inline encoded query string or JavaScript (actual code or the name of an existing script include or business rule) filter directly in the **Reference qual** field of the reference qualifier. Similar to the other reference qualifier types, when the form loads, the filter is executed, and only the records that match the filter appear in the reference field. Use this type of reference qualifier for implementations that only require

a simple, unique filter, that cannot be handled by a simple reference qualifier, and is not used across multiple reference fields.

Note:
As a good practice, make JavaScript calls to functions in a script include instead of a global business rule.

An example of an encoded query string is `vendor=true`, which returns all companies that are designated as vendors. Entering this string is the same as using the condition builder as shown in the example for the simple reference qualifier. For additional information on valid encoded query string syntax and examples, see Encoded query strings.

Advanced reference qualifier examples

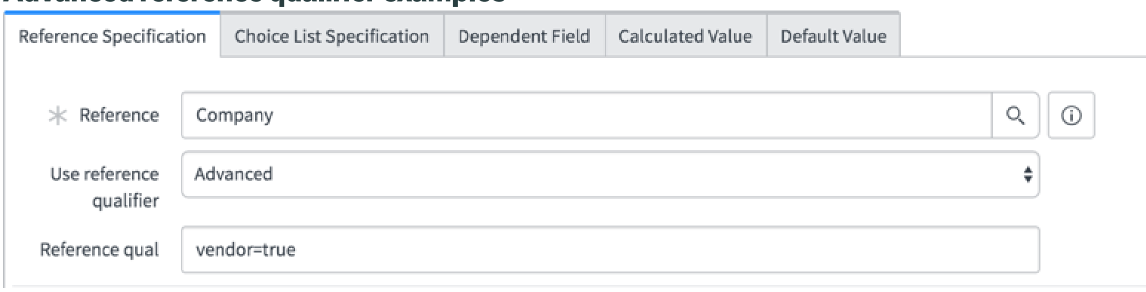

An example of a JavaScript call is `javascript:new mySciptInclude().my_refqual()`. This code calls the function *`my_refqual()`* in the script include *`myScriptInclude()`*. The function must return a query string that can filter the options available on a reference field.

Note:
You can also use encoded JavaScript filters such as `javascript:'u_active=true^' + "u_hr_service="+current.hr_service` in reference qualifiers.

In another example, if you are trying to filter based on the current company, you don't have use a lookup script but can simply add this line to your dynamic qualifier:

```
return "company=" + current.company;
```

You can also simplify this and use an advanced reference qualifier instead of a dynamic one:

```
javascript:"company=" + current.company
```

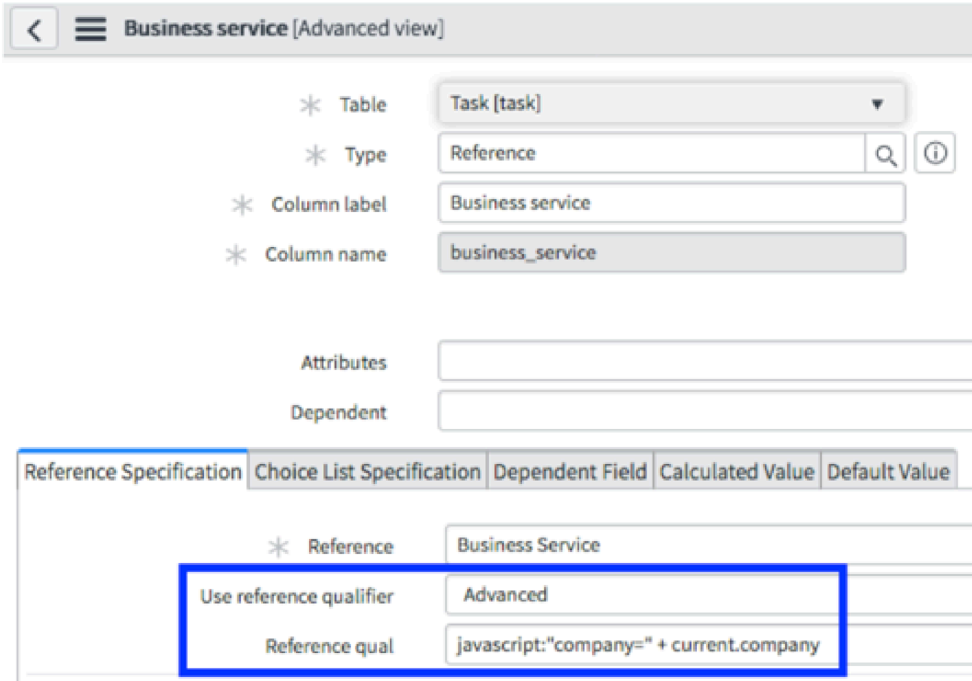

For an additional example, configuring the reference field to show only users of a specific group, see the How to select only users of a specific group into a reference field [KB0831564] article in the Now Support Knowledge Base

.

Related lists and reference qualifiers

When a field appears on multiple Related lists on a single form view, it may be necessary to validate which related list is being referenced to properly build the reference qualifier for the field. In this situation, configure the list control for the related list and enter a unique tag in the **List edit tag** field. This tag value is available to filter scripts as a variable named *listEditRefQualTag*. The following script include code is an example of a function that uses this type of tag.

```
// Advanced reference qualifier on the CI Relationship Child
 field that takes into account
// the related list that we are editing the child field on, if
 the field is being edited
// from a tagged related list.

 cmdb_rel_ci_child_refQual:function(){

  if(listEditRefQualTag =="application") return "sys_class_name
 = cmdb_ci_appl";

  if(listEditRefQualTag =="database") return "sys_class_name =
 cmdb_ci_database"

  }
```

Using Javascript current syntax in reference qualifiers

`current` is a JavaScript object that contains the fields and field values of the active (current) record. For forms, this is the record that is displayed (loaded) in the form. Within advanced and dynamic reference qualifiers, you can use the JavaScript `current` object to define filters such as `javascript:"company=" + current.company`.

This JavaScript, within a reference qualifier, only returns the records from the referenced table that are equal to the company field value of the current record. So, if the value that appears in the **Company** field is Acme, the JavaScript returns all reference field records whose company value is equal to Acme (company="Acme"). If you then bring up a record whose company value is "ViewRite", the JavaScript resolves to company="ViewRite."

All fields within the currently loaded form (tables) are available for use with the `current` object. Use dot-walking to access values in a table, including the referenced table. For example, on the Incident form, the **Assigned To** field references the User table. To access the email address of the user, use the following syntax: `javascript:"emailAddress=" + current.assigned_to.email`.

Related topics

- Create a dynamic filter option
- Create scripted filters
- CMDB classifications

Configure reference qualifiers

System administrators can configure reference qualifiers that enable filtering of the associated reference field.

About this task

This task walks you through creating a reference qualifier in the system dictionary. When creating a reference qualifier in the system dictionary, it not only applies to that table, but also to all the children of the table.

> **Note:**
> You can also create reference qualifiers on extended tables through dictionary overrides. Dictionary overrides only impact the extended table and any of its children, not the parent table.

Procedure

1. Navigate to the reference field on the form or table on which you want to define a reference qualifier.
2. Right-click the reference field label and select **Configure Dictionary**.
3. Under **Related Links**, select **Advanced view**.
   The simple reference qualifier is available in both the default view and the advanced view. The dynamic and advanced reference qualifiers are available only in the advanced view.
4. In the Reference Specification section, verify that the table in the **Reference** field contains the table that you want associated with the reference field, or select another table if necessary.
5. From the **Use reference qualifier** choice list, select the reference qualifier to implement.
6. Configure the qualifier based on the type of reference qualifier.

- **Simple**: Build the condition using the choice lists.
- **Dynamic**: Either select an existing dynamic filter option or Create a dynamic filter option.
- **Advanced**: In the **Reference qual** field, enter an Encoded query strings, JavaScript that returns a query string, or the name of a script include or business rule to run.

7. Complete the **Reference Specification — Additional Customization** section, as appropriate.

Field	Description
Reference key	Identifies a field other than sys_ID to use as the unique identifier for the reference field.
Reference cascade rule	Defines what happens to a record if the record it references is deleted. Possible options include: ◦ None ◦ Cascade ◦ Clear ◦ Delete ◦ Delete no workflow ◦ Restrict
Reference floats	Flag that determines whether to Configure the edit option for one-to-many relationships.
Dynamic creation	Flag that determines if the system should create a new record when a value for the reference field does not match an existing record. If you select this option, enter a script that specifies how to create the record in the **Dynamic creation script** field.

8. Click **Update**.

Constrain the assigned to field by role

This example shows how to use JavaScript and a business rule to restrict the incident **Assigned to** field choices to only the users with the itil_admin role.

Before you begin

Role required: personalize_dictionary or admin

About this task

You can also change itil_admin to any other role on a reference field that refers to the User [sys_user] table.

Procedure

1. Open an incident.
2. In the upper-left corner of the screen, click the form context menu, and then select **Configure** > **Dictionary**.
3. In the **Reference qual** field, enter `javascript:"sys_idIN"+getRoledUsers("itil_admin").join(",")`.
4. Save the record.

5. To see the base system business rule that this JavaScript code calls, navigate to **System Definition** > **Business Rules**.
6. Open **getRoledUsers**.
   The business rule uses the following JavaScript code.

```
// Return an array of sys_ids of the users that have at least
 one role
// optional parameters allow the exclusion (NOT IN) of some
 roles or
// look for specific roles (IN)
//
// optional: queryCondition - 'IN' or 'NOT IN'
// optional: roleList - a comma separated list of role names
//
function getRoledUsers(queryCondition, roleList) {
   var roleListIds;
   if (queryCondition && roleList) {
      roleListIds = getRoleListIds(roleList);
   }

   var users = {};
   var now_GR = new GlideRecord('sys_user_has_role');
   if (roleListIds) {
      now_GR.addQuery('role', queryCondition, roleListIds);
   }
   now_GR.query();
   while (now_GR.next()) {
      users[now_GR.user.toString()] = true;
   }

   var ids = [];
   for (var id in users)
      ids.push(id);

   return ids;
}

// get sys_id's for the named roles
function getRoleListIds(roleList) {
   var ids = [];
   var now_GR = new GlideRecord('sys_user_role');
   now_GR.addQuery('name','IN',roleList);
   now_GR.query();
   while (now_GR.next()) {
      ids.push(now_GR.sys_id.toString());
   }
   return ids;
}
```

Constrain the assignment group field

This example shows how to use an advanced reference qualifier with JavaScript and a script include to restrict the incident **Assignment group** choices to only the groups that contain the user specified in the **Assigned to** field.

Before you begin

Role required: personalize_dictionary or admin

servicenow

Procedure

1. Open an incident.
2. Right-click the **Assignment group** label and select **Configure Dictionary**.
3. If the form appears in **Default view**, under **Related Links**, click **Advanced view**.
4. In the **Use reference qualifier** field, ensure that the **Advanced** option is selected.
5. In the **Reference qual** field, enter `javascript:new ReferenceQualifierHelper().backfillAssignmentGroup()`.
6. Save the record.
7. Navigate to **System Definitions** > **Script Includes**.
8. Click **New**.
9. Name your script includes as `ReferenceQualifierHelper` to match the Javascript in the following step.
   This name ensures the API Name is generated correctly, based on the Name field.
10. Script includes with the following JavaScript code.
    Ensure that the **Accessible form** field is properly set for the environment.

```
var ReferenceQualifierHelper = Class.create();
ReferenceQualifierHelper.prototype = {

 backfillAssignmentGroup:function() {
  var assigneeGroupsQualifier = '';
  var assignee = current.assigned_to;

  //return if the assigned_to value is empty (this causes all
 groups to be returned)
  if(!assignee)
   return;

  //sys_user_grmember has the user to group relationship
  var grAssigneeGroups = new GlideRecord('sys_user_grmember');

  grAssigneeGroups.addQuery('user',assignee);
  grAssigneeGroups.query();

  while(grAssigneeGroups.next()) {
   if (assigneeGroupsQualifier.length == 0) {
    // Create the beginning of the qualifier and add the 1st
 group
    assigneeGroupsQualifier = 'sys_idIN' +
 grAssigneeGroups.getValue('group');
   }
   else {
    // build a comma separated string of groups if there is
 more than one
    assigneeGroupsQualifier += (',' + grAssigneeGroups.group);
   }
  }
  // return Groups where assigned to is in those groups we use
 IN for lists

  gs.log('DP: RefQual = ' + assigneeGroupsQualifier);
```

```
  return assigneeGroupsQualifier;
 },
 type: 'ReferenceQualifierHelper'
};
```

What to do next

To test, create an incident and select a user in the **Assigned to** field. Click the **Assignment group** lookup icon. Only the groups that contain the user you selected appear. For example, you assign an incident to Bob Smith, who belongs to the **Database** group and the **Networking** group. The only options that appear in the assignment group are **Database** and **Networking**.

The INSTANCEOF operator in reference qualifiers

You can use the INSTANCEOF operator in a reference qualifier to shorten or simplify a complex class qualifier.

For example, use the INSTANCEOF operator for a reference field to the cmdb_ci table to specify that all subclasses of a class are included in the results. The following reference qualifier returns all servers, including Linux, UNIX, Windows, and so on, because each of those subclasses extend the cmdb_ci_server class.

```
sys_class_nameINSTANCEOFcmdb_ci_server
```

In another example, you can simplify the following reference qualifier in a similar way.

```
 u_active=true^sys_class_name=cmdb_ci_acc
^ORsys_class_name=cmdb_ci_computer
^ORsys_class_name=cmdb_ci_server
^ORsys_class_name=cmdb_ci_win_server
^ORsys_class_name=cmdb_ci_unix_server
^ORsys_class_name=cmdb_ci_linux_server
^ORsys_class_name=cmdb_ci_appl
^ORsys_class_name=cmdb_ci_netgear
```

Using the INSTANCEOF operator, the reference qualifier is rewritten as follows because the server subclasses extend the cmdb_ci_computer class.

```
 u_active=true^sys_class_name=cmdb_ci_acc
^ORsys_class_nameINSTANCEOFcmdb_ci_computer
^ORsys_class_name=cmdb_ci_appl
^ORsys_class_name=cmdb_ci_netgear
```

Auto-complete for reference fields

By default, a reference field auto-completes as the user types in the field.

Administrators can configure additional auto-complete options. A user must have table-level read permission on the referenced table for auto-complete to display any options.

Auto complete

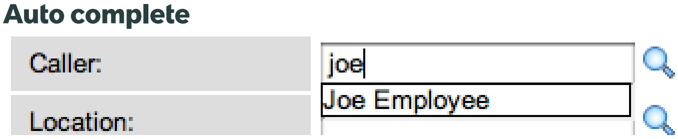

Dictionary attributes for auto-completion of reference fields

There are dictionary attributes that are specific to reference fields and that determine auto-complete behavior.

Reference field auto-completion dictionary attributes

Attribute	Description
ref_auto_completer	Specifies the name of the client-side JavaScript class that creates the drop-down auto completion choices. Valid class values include: • *`AJAXReferenceCompleter`*: Displays matching auto-complete choices as a drop-down choice list. The list only displays the reference table's display value column. Reference fields automatically use this class if there is no other auto-completion class specified. • *`AJAXTableCompleter`*: Displays matching auto-complete choices as rows in a table. The table displays the reference table's display value column and any columns listed in the *`ref_ac_columns`* attribute. • *`AJAXReferenceChoice`*: Displays matching auto-complete choices as a drop-down choice-list. The list only displays the reference table's display value column. Furthermore, the list only displays up to 25 matching choices. If there are more than 25 auto-complete choices, the reference field instead displays the choices with the AJAXTableCompleter class.
ref_ac_columns	Specifies the list of reference table columns to display. Separate column names with a semi-colon. For example, *`ref_ac_columns=user_name;email;sys_created_on`* allows auto-complete to match text from the user_name, email, and sys_created_on columns.
ref_ac_order_by	Specifies the reference table column that sorts the auto-completion choices. For example, *`ref_ac_order_by=name`* sorts the auto-completion choices alphabetically by name.

Auto-complete UI features

The AJAX table completer class has a number of UI improvements.

- The table completer always displays the number of records the auto-complete query finds.
- The table completer highlights the entire selected row by changing the color of the background and text.
- The table completer lists a value for every column.
    - The first time a value appears in a column, the table completer displays it in black text.
    - The table completer displays subsequent duplicate values in grey text. Previously, the table completer displayed an empty cell in a column containing a duplicate value.

Selected row highlight

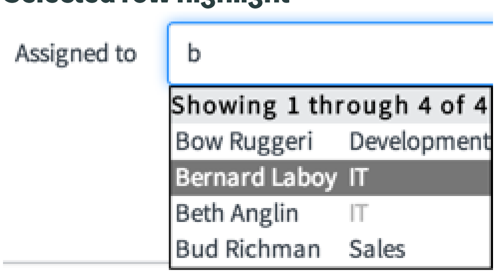

Set the *ref_auto_completer=AJAXTableCompleter* dictionary attribute to use these improvements.

Define auto-complete attributes for all references to a table

A field inherits and uses the reference table's auto-complete attributes unless the field has its own value for the same attributes. You can define the attributes for references to a table, and it affects every form that references that table.

Before you begin

Role required: personalize_dictionary

About this task

A field-level attribute overrides a table-level attribute of the same name. If a field uses different reference attributes from those that are defined for the reference table, then the field uses both sets of attributes.

Use these steps to define auto-complete attributes for all fields in a table that do not already have their own auto-complete attributes. This example describes how to define auto-complete attributes for all references to the User [sys_user] table.

Note:
A field's auto-complete attribute value supersedes a table's auto-complete attribute value. This means that any existing field-level value for an auto-complete attribute supersedes any value the administrator applies to the auto-complete attribute from the reference table.

Procedure

1. Navigate to a list of the target table, such as **All** > **User Administration** > **Users**.
2. Perform the appropriate action for your list version.
3. Select the row that does not list a column name.
   This row is typically the first row in the list. For example, select the first **sys_user** link.
4. Under **Related Links**, click **Advanced view**.
5. In the **Attributes** field, enter a comma-separated list of auto-complete attributes you want to apply to all fields in the table.

   Example
   For example, to display the user's department with all references to the sys_user table, enter:

   ```
   ref_auto_completer=AJAXTableCompleter,ref_ac_columns=department
   ,ref_ac_order_by=department
   ```

6. Click **Update**.

What to do next

To test the new auto-complete attributes, open a form that references the User [sys_user] table, such as an open incident. Enter a single character in the **Assigned to** field. The auto-complete options now include both the user name and department.

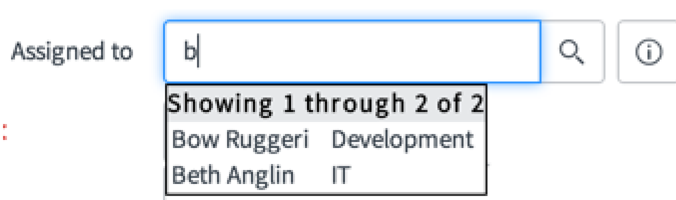

Remove the display value column

You can remove the display value column from a reference field by setting the *ref_ac_display_value* attribute to false.

Before you begin

Role required: personalize_dictionary

About this task

This causes the reference field to remove the display value column and only display the columns listed in the *ref_ac_columns* attribute. This feature requires the use of the *AJAXTableCompleter* class and the *ref_ac_columns*, *ref_ac_columns_search*, and *ref_ac_display_value* attributes.

> **Note:**
> Auto-complete cannot match text from additional columns when the reference field is a product of the *ui_reference* UI macro. This means any auto-complete action against a selector, such as the Impersonate User list, can only match text against the display value.

This example describes how to remove the display value column from references to the User [sys_user] table and replace it with references to the first_name and last_name columns.

Procedure

1. Navigate to a list of the target table, such as **All** > **User Administration** > **Users**.
2. Perform the appropriate action for your list version.
3. Select the row that does not list a column name.
   This row is typically the first row in the list. For example, select the first **sys_user** link.
4. Under **Related Links**, click **Advanced view**.
5. In the **Attributes** field, add the *ref_auto_completer*, *ref_ac_columns*,*ref_ac_columns_search* , and *ref_ac_display_value* attributes.

   Example

   For example, to hide the display value column and only display the user's first and last names enter the following.

   ```
   ref_auto_completer=AJAXTableCompleter,ref_ac_columns=first_name
   ;last_name,ref_ac_columns_search=true,ref_ac_display_value=fa
   lse
   ```
6. Click **Update**.

What to do next

To test the new auto-complete attributes, open a form that references the User [sys_user] table, such as an open incident. Enter a single character in the **Assigned to** field. The auto-complete options now hide the display value column (user_name) and only display the first_name and last_name columns.

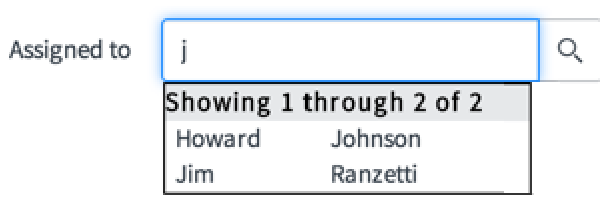

Improve auto-complete queries

By default, all reference fields use a **starts with** query to search for matching text in the reference table. This prevents auto-complete from executing inefficient **contains** queries every time a user searches a reference field. You can require all reference fields to use a **starts with** query.

Before you begin

Role required: admin

About this task

The following example illustrates a **contains** query. Note that the letter "d" appears anywhere in the user's first or last name.

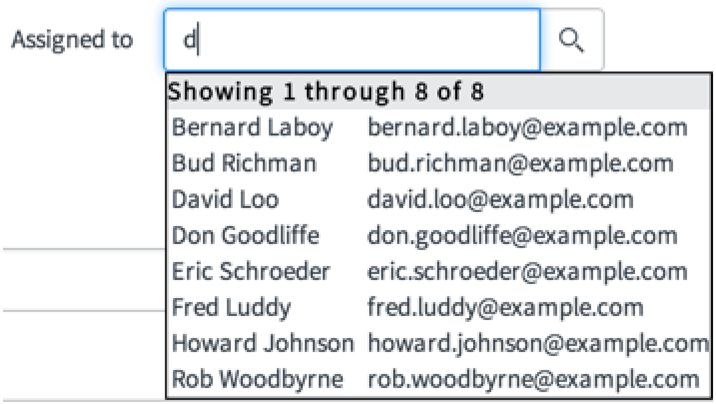

This procedure describes how to change the `glide.ui.ref_ac.startswith` system property to always use a `starts with` query.

Procedure

1. In the navigation filter, enter `sys_properties.list` and press the Enter key.
2. Select the `glide.ui.ref_ac.startswith` property.
   To search for the property, enter `*startswith` in the **Go to** search filter for the **Name** column.
3. In the **Value** field, replace **false** with **true**.

   Note:
   Setting the `glide.ui.ref_ac.startswith` system property to **true** overrides any existing `autocomplete.contains` settings in both user and system level preferences. This property changes the autocomplete query method for all users regardless of preferences.

4. Click **Update**.
5. Test the change by opening a record with a reference field and entering a character in it, as illustrated in the example below.

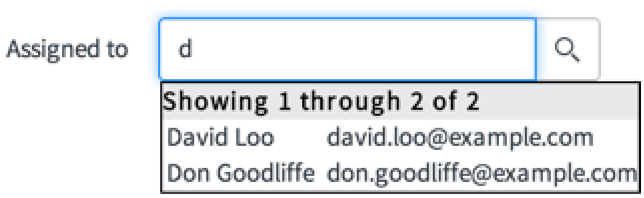

Configure auto-complete to match text from any reference field

By default, auto-complete only matches text in the display value column. You can configure a reference field to match text from any additional column the reference field displays.

Before you begin

Role required: personalize_dictionary

About this task

You can add the `ref_ac_columns_search` attribute to enable auto-complete to match text in any column listed in the `ref_ac_columns` attribute. Set the `ref_ac_columns_search` attribute to **true** to match text from all reference field columns. By default (or when this attribute is **false**) auto-complete only matches text in the display value column.

Procedure

1. Right-click the label of a reference field.
2. Select **Configure Dictionary** from the choice list.
3. Under **Related Links**, click **Advanced view**.
4. In the **Attributes** field, add the desired auto-completion attributes.
   For example, these attributes add the department field to the caller list and sort callers by their department:

```
ref_auto_completer=AJAXTableCompleter,ref_ac_columns=department
,ref_ac_order_by=department,ref_ac_columns_search=true
```

Auto-complete attribute

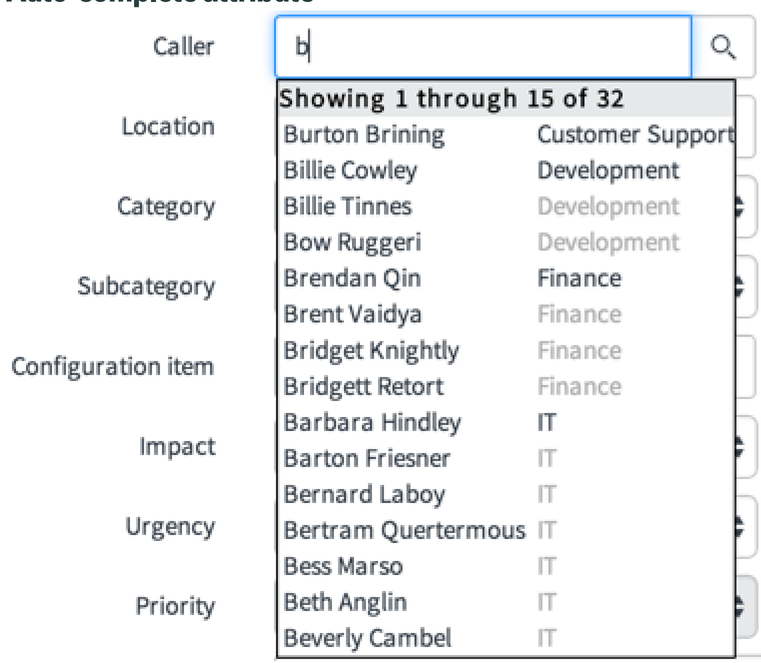

5. Click **Update**.

Example:

The following example describes how to set the **Configuration Item** field display the CI class names from auto-complete choices for the Configuration item [cmdb_ci] table.

```
ref_auto_completer=AJAXTableCompleter ,ref_ac_columns
 =sys_class_name ,ref_ac_order_by
```

```
=sys_class_name ,ref_contributions
=task_show_ci_map ;ci_show_incidents
```

Note:
The *ref_contributions* attribute controls the icons that appear next to the reference field.

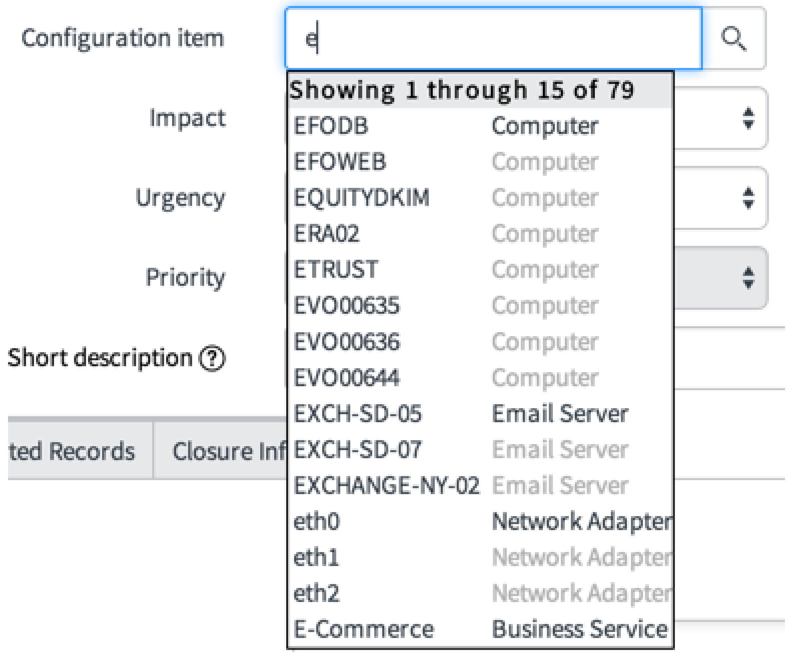

Enable contains auto-complete searches

By default, the reference auto-complete uses a **starts with** search. A user preference can be created to implement a **contains** search.

Before you begin

Role required: admin

Procedure

1. Disable the *glide.ui.ref_ac.startswith* system property.
   For more information, see Improve auto-complete queries.

   Note:
   Setting the *glide.ui.ref_ac.startswith* system property to **true** overrides any existing "autocomplete.contains" settings in both user and system level preferences. This property changes the auto-complete query method for all users regardless of preferences.

2. Navigate to **User Administration** > **User Preferences**.
3. Select the preference **"<referenced table>.autocomplete.contains"**.
4. Set the **value** field to **true**.

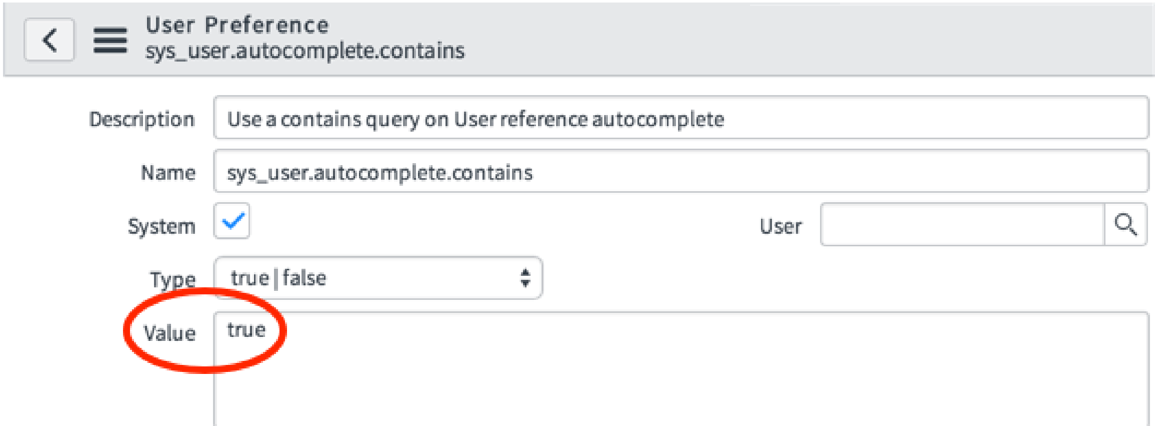

5. Click **Update**.

What to do next

Log out and log back in to display the updated search.

Wildcards in reference auto-completes

Wildcard searches use the auto-complete functionality.

Use an asterisk in the reference field for wildcard searches.

One asterisk wildcard search

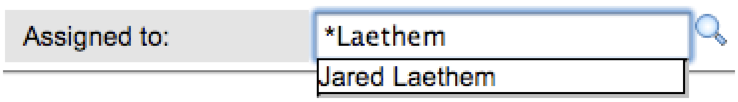

If two asterisks are entered, a list of available records display in the auto-complete suggestions.

Two asterisk wildcard search

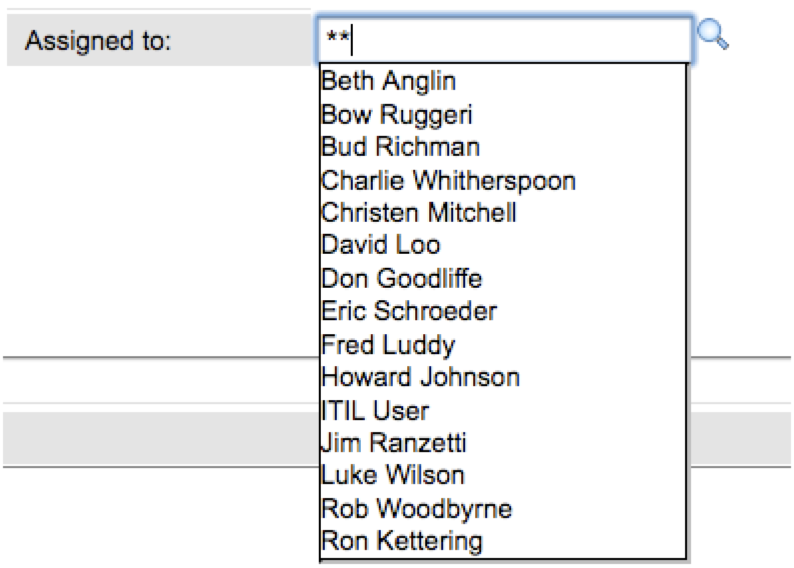

Related topics

Wildcard characters allow searching for patterns and variations

Recent selections

Reference fields store a list of each user's recent selections to enable you to quickly select past values when filling in a reference field.

By default, the system stores up to 15 selections from a reference field for each user in the Recent Selection [sys_ui_recent_selection] table. You can see the recent selections list by selecting an empty reference field.

Note:

- The system doesn't store recent selections for Service Catalog reference variables.
- Recent selections aren't available in the Service Portal or mobile user interfaces.

Recent selections

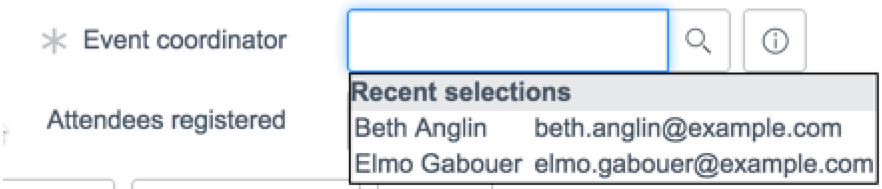

The system uses auto-complete to filter the list of recent selections to match the values that you enter.

Recent selections filtered

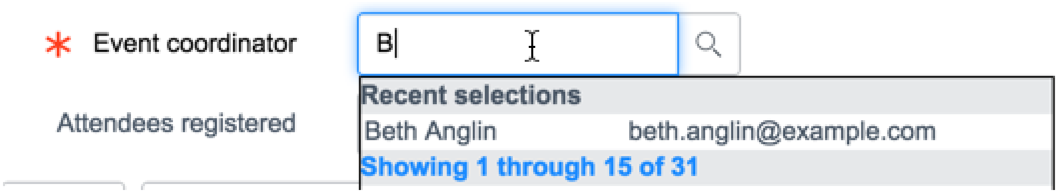

The system adds a Recent Selection record whenever you insert or update a reference field value. Administrators can control the number of recent selections the system displays with the `glide.xmlhttp.max_choices` system property. Setting the property to zero disables recent selections.

Note:

- The `glide.xmlhttp.max_choices` property also controls the number of entries the system displays in lists.
- For information on clearing recent selections, see the How to clear out the "Recent Selections" that appears up on the reference fields. [KB0747334] article in the Now Support Knowledge Base.

Related topics

Available system properties

Suggestion field type

To help users find what they are looking for, you can add a suggestion field to any form.

Before you begin

Role required: personalize_dictionary

Procedure

1. Open the form to which you are adding a suggestion field.
2. Add a new field or locate an existing field.
3. Right-click the field label and select **Configure Dictionary**.
4. In the **Choice** field, select **Suggestion**.
5. Click **Update**.

   Suggestions are presented to the user.

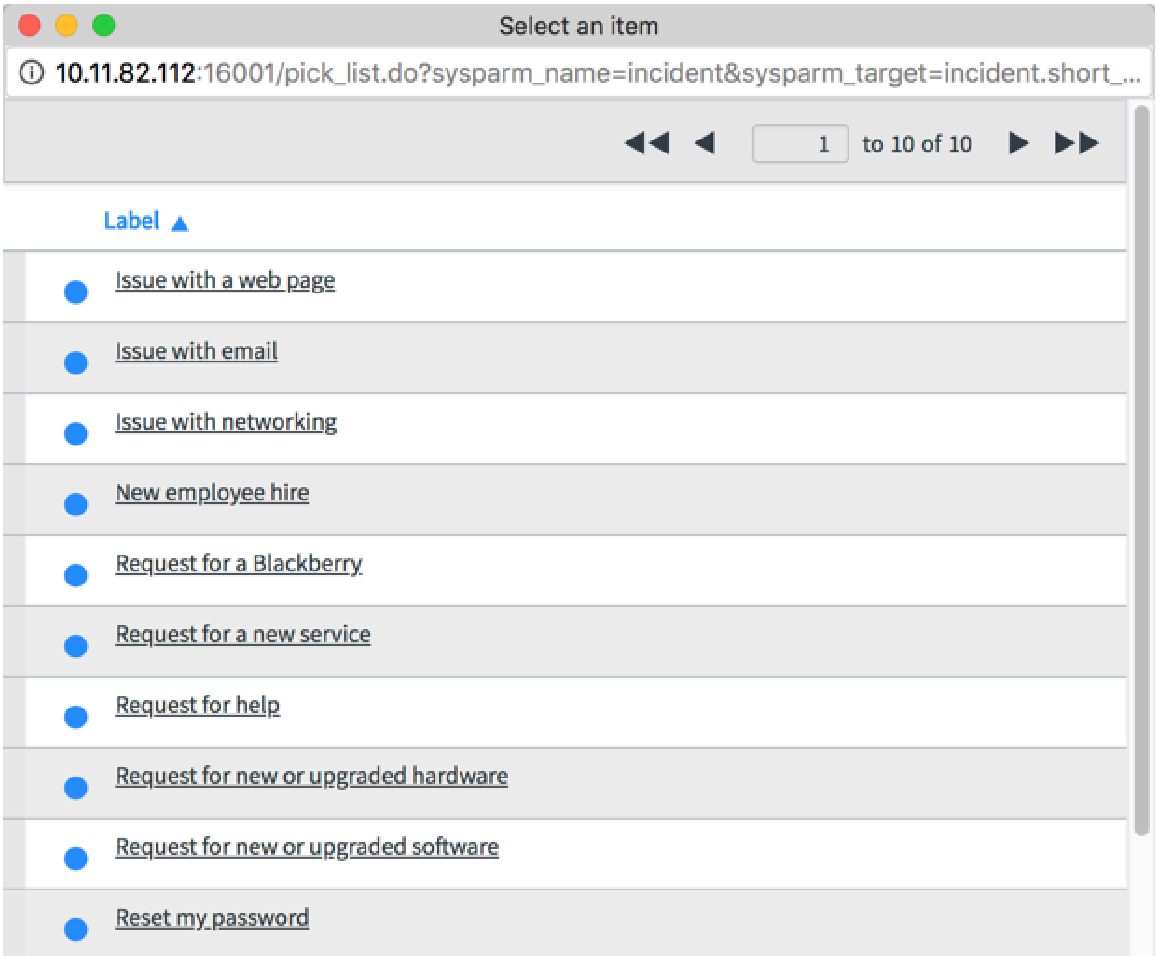

Configure suggested text for string fields

You can configure the suggested text options for string fields.

Before you begin

Role required: personalize_choices

Procedure

1. Right-click the field label and select **Configure Choices**.
2. Using the slushbucket, select options and the order in which you want them to appear.
   - To create a new option, enter the suggested text in the **Enter new item field** and click **Add**.
   - The **Apply to Table** field is available when the current table extends another table (for example, Incident extends Task). This field allows suggested text options to be configured for all tables that extend the parent or for only the current (child) table.

3. Click **Save**.
4. To edit existing options, right-click the field label and select **Show Choice List.**.

Example:

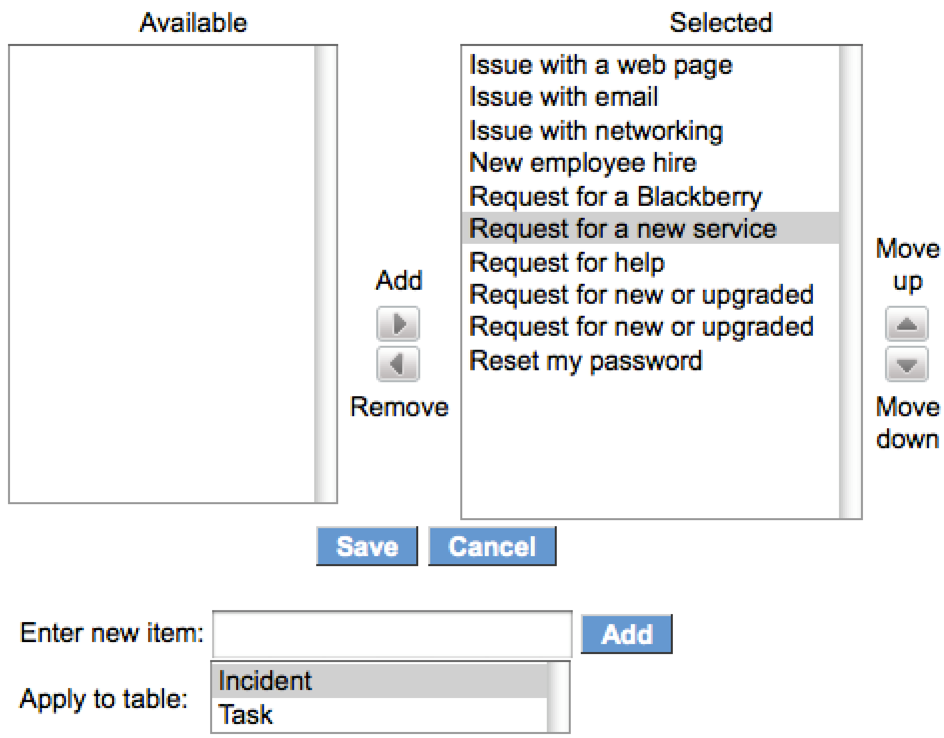

Configure suggested text for journal fields

You can configure the suggested text options for journal fields, such as work notes or the activity stream. Suggested text options for journal fields are unavailable in Core UI.

Before you begin

Role required: personalize_responses

Procedure

1. Right-click the field label and select **Configure Responses**.
2. Click **New**.
3. Enter a label, or brief description, for the option.
4. Enter the complete text in **Response text**.
5. Click **Save**.

Example:

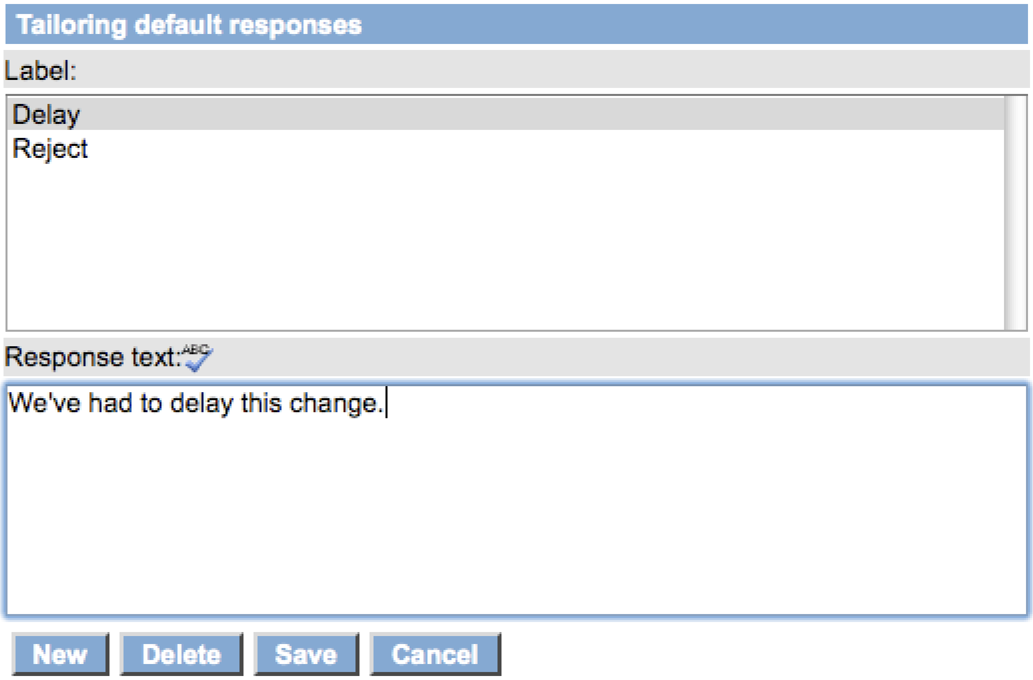

Wiki field type

Wikitext fields use the basic wiki markup language and support links to external URLs and to sources within the system.

Some fields in the base system support Wikitext and administrators can add Wikitext fields on any form in the system.

Supported wiki tags

Basic tags are supported for Wikitext fields in the system.

Supported wiki tags

Format	Wiki tag
Headers	= Header 1 =, == Header 2 ==, etc.
Numbered step	#Step 1
Bullets (multi-level)	*Bullet 1, **Bullet 2
Indentation	:Level 1, ::Level 2
Line break	
Tables	Full table support, including cell spacing, padding, borders, background shading, and width measurements. For example: `{| class="wikitable"` `|-` `! Header 1`

Supported wiki tags (continued)

<table>
<tr><th>Format</th><th>Wiki tag</th></tr>
<tr><td></td><td><pre>! Header 2
|-
| row 1, cell 1
| row 1, cell 2
|-
| row 2, cell 1
| row 2, cell 2
|}</pre>
Note:
Sortable tables are not supported (class="sortable").</td></tr>
<tr><td>Code blocks</td><td><pre> Generic block </pre>

<javascript> Javascript Syntax </javascript>

<xml> XML Syntax </xml></td></tr>
<tr><td>Bold</td><td>'''text'''</td></tr>
<tr><td>Italics</td><td>''text''</td></tr>
<tr><td>Ignore Wiki and HTML formatting</td><td><nowiki> wikitext </nowiki></td></tr>
<tr><td>Web link</td><td>[http://community.service-now.com/ <link text>]</td></tr>
<tr><td>Placed images</td><td>[[Image:<image name>]] For more information on uploading images or files to the database, see Storing images in the database.</td></tr>
<tr><td>HTML</td><td>Most common HTML tags are supported

bold, strong, <i>italics</i>, <h1>header1</h1>, <h2>header2</h2>, <h3>header3</h3>, <p>paragraph</p>, _{sub}, ^{sup}, <center>center</center></td></tr>
</table>

Create a Wikitext field

You can create a Wikitext field on any form.

Before you begin

Role required: personalize_form

Procedure

1. Create a new field with the **Type** set to **Wiki**.
2. Click **Wikitext** to begin editing.

The Wikitext field appears on the form.

Related topics

Add and customize a field in a table

Extend the functionality of a Wikitext field

Configure the Wikitext field to link to other sources within the system.

Before you begin

Role required: personalize_dictionary

Procedure

1. Right-click in the header of the new Wikitext field.
2. Select **Configure Dictionary**
3. In the **Dependent** field of the dictionary form, enter the field you want to use for linking to other pages.

   For example, on the Incident table, you might choose the **number** field.

   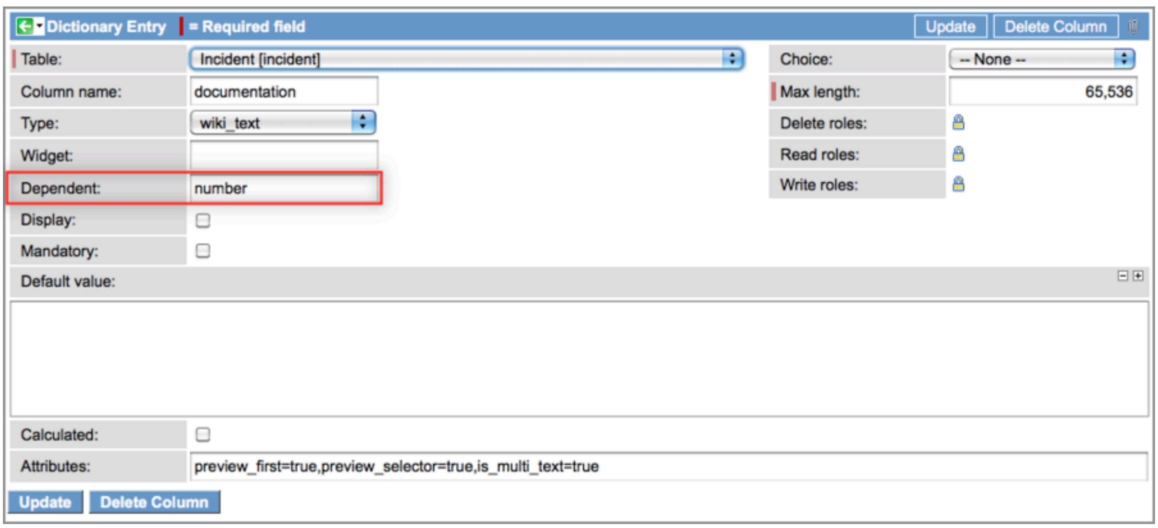

   This configuration enables you to link to any incident by using that incident's number. For example, to open INC0000002, you enter `[[INC0000002]]`.

   The following dictionary attributes are available for wiki_text fields.

   - *`preview_first=true`* - sets the preview mode to display on page load, otherwise editor will be displayed
   - *`preview_selector=true`* - enables the toggle button to switch between the editor and preview
   - *`dual_mode=true`* - enables displaying both the editor and the preview simultaneously during edit mode.

Add and customize a field in a table

Administrators can add new fields to a table to store and display data.

Before you begin

Role required: admin

About this task

Note:
Certain ServiceNow AI Platform subscriptions include custom table entitlements. You can create custom tables for any purpose, up to the entitlement limit in the subscription. To learn more about how your usage administrator maps the custom tables that you create to subscriptions, see Map custom tables to a product subscription in Subscription Management.

Be aware of these database limitations:

- The system can only have a maximum of 1000 columns per table. Although 1000 columns is a specified limit, this limit doesn't mean you can physically have 1000 columns within a table. The number of columns within a table is defined by the database used in the ServiceNow datacenter, not by the ServiceNow AI Platform.
- Every table, regardless of storage engine, has a maximum row size of 65,535 bytes. Storage engines may place additional constraints on this limit, reducing the effective maximum row size.
- The system can't have more than 10 medium-length or longer **String** fields to a single table. Attempting to save a large number of characters in 11 or more String fields can result in the following error: Syntax Error or Access Rule Violation detected by database (Row size too large (> `8126`).
- When you create fields, the u_ prefix is automatically added to the column name. If the column label that you enter contains leading numeric characters, they are replaced by the u_ prefix.

For more information on database limitations and general questions on tables in your ServiceNow instance, see KB0749585.

Procedure

1. Navigate to any form.
2. Right-click the form header and select **Configure** > **Form Layout**.
3. In the **Create new field** section, fill in the following fields

Option	Description
Name	Enter the name of the field as you want it to appear on forms and lists.
Type	Select a field type.
Field length	Select a field length. This field is visible only for certain field types.

4. Click **Add**.
5. Use the slushbucket to place the field in the desired location on the form.
6. Click **Save**.
   The field now appears on the form in the designated location.

Make a field mandatory

Fields can be marked as mandatory, meaning they must contain a value before the record can be saved. Mandatory fields are marked with a field status indicator before the label.

Before you begin

Role required: personalize_dictionary

About this task

A mandatory field that is pre-populated by the platform with default data, such as a value from a client script, does not display the indicator. If you delete this value from the field, however, the indicator appears. The color of the indicator depends on the field state.

Using a dictionary entry to make a field mandatory does not work when using WebServices. You can use a data policy instead. For more information on data policies, see Data policy.

Procedure

1. Right-click the field's label in the form and select **Configure Dictionary**.
2. In the **Dictionary** form, select the **Mandatory** check box.
3. Click **Update**.

   The next time the form is opened, a field status indicator appears next to the field label, indicating that a value is mandatory. If the mandatory field is pre-populated for any reason, the mandatory field indicator does not appear.

   Mandatory fields are global. The field is marked as mandatory everywhere it appears in a form.

   Making a **True/False** mandatory will have no effect. **True/False** fields are always considered to have a value. A selected check box has a value of true and an unselected check box has a value of false. Either of these values satisfies the requirement of a mandatory field.

   A form can be saved with an empty mandatory field if that field is a reference field, and if the parent field is also blank. However, if the mandatory reference field shows a value from the parent field, then the form cannot be saved if this value is deleted. If the value in the referenced field is changed, the value for that field is changed everywhere it appears.

4. **Optional:** Use *`dictionary override>override mandatory`* to make a field that is mandatory only in the child table and not the parent table.

Related topics

Set a CI field to be mandatory

Define a dictionary override

Change the field label or hint

You can change a field's label or the text that appears as a hint when you point your cursor to the field.

Before you begin

Role required: personalize_dictionary

Procedure

1. Navigate to the form the field appears on.
2. Right-click the field label and select **Configure Label**.
3. Update the form.

Changing the Field Label or Hint

Field	Description
Table	The table the label appears on.
Label	The label that the field displays. You can enter up to 80 characters for the label, but not all databases support this many characters. The recommended maximum is 30 or fewer characters. HTML in field labels is not supported.
Plural	The plural version of the field label.
Element	The dictionary name of the field.
Help	Helpful text stored in the record.
Hint	A short description of the field that displays when the user hovers over it.
URL	A URL link that displays on the label if this field is not blank.
URL target	A target attribute that determines where the URL will open. For information on the target attributes, see this W3Schools article.

4. Click **Update**.

Related topics

Add help to a field label on a form

Delete a field from a table

You can delete custom fields that you created. Custom fields begin with **u_**. It is recommended that you remove the field from forms and lists instead of deleting it.

Before you begin
Role required: admin

About this task
You cannot delete base system fields. In addition, any missing base system fields are recreated when the instance is upgraded.

Procedure

1. Navigate to a form that contains a custom field to delete.
2. Right-click the field label and select **Configure Dictionary**.
3. Click **Delete Column** in the form header, and then click **OK**.
4. To delete multiple custom fields, complete the following steps.
   a. Navigate to **System Definition** > **Dictionary**.
   b. Locate the custom fields to delete.

      Example
      For example, search for column names that start with `u_`.

   c. Check the boxes next to the fields to delete and select **Delete** from the action list below the list.

A confirmation dialog opens and reminds you that this may result in deletion of related records. If there are dependencies for the selected fields, they are listed.

d. To proceed, click **Delete**.

Add users to a watch list

Watch lists enable you and others to subscribe to notifications of a task.

Before you begin

Role required: none

Procedure

1. Navigate to a record that you want to add watchers to.

   Note:
   Not all records have watch lists. If you're unable to see a watch list, you may need to adjust your form view or ask your administrator to update the form to add a watch list.

2. Expand the watch list by selecting the lock icon ( ).
3. Select users with the list controls (see table).

Glide controls

Icon	Function
	Add yourself to the watchlist
	Remove the highlighted user
	Open the highlighted user's record (active only when a user record exists)
	Open the Edit Members window to add or remove multiple users (not available until a record has initially been saved)
	Indicates that the watch list is unlocked; select the icon to lock the watch list
	Open the Users list in a new window to select a single user
	Enter an email address for users who aren't in the User [sys_user] table or don't have an email address defined in their user record

Configure email notifications for watch lists

Watch lists (glide_list field type) allow multiple users to subscribe to notifications of a task. You can specify conditions in an email notification to send email notifications to the members when the conditions are met.

Before you begin

Role required: admin

About this task

To receive these notifications, users must define an email address in their user record or enter an email address into the watch list email field.

Important:
Administrators configure email notifications for watch lists (see Configure email notifications for watch lists).

An advanced configuration using watch lists involves placing two watch lists on a form, one for the general comments on a task and another for work notes or non-public comments. By configuring separate email notifications, separate users on each watch list can be notified about different information.

If users on a watch list are getting more than one email for each update to an incident, it can be because other recipients are replying all to an email notification. Recipients may be receiving email through their email system (Outlook, Groupwise, and so on) and through the base system. To stop this duplication, remove the names of other users from the email or the watch list.

Procedure

1. Open the notification to configure.
2. In the Who will receive section, select the icon beside **Users/groups in fields**.
3. Double-click **Watch list** in the **Available** column to move it to the **Selected** column.
4. Click **Update**.

Related topics

Configure email notifications for watch lists

Email and SMS notifications

Hide email addresses in a watch list

You can remove the email address text entry element from a watch list by modifying the dictionary.

Before you begin

Role required: admin

Procedure

1. Open a task record that displays the **Watch list** field.
2. Right-click the label and select **Configure Dictionary**.
3. In the **Attributes** related list, click **New**.
4. Enter or select the **no_email** attribute and enter `true` in the **Value** field.
5. Click **Submit**.

Result

The email entry field is hidden. Users can select users from the reference field to add to the watchlist.

Related topics

Altering tables and fields using dictionary attributes

Configure order buttons on the watch list slushbucket

When you add multiple users to a glide_list, such as a watch list, the slushbucket does not display the order buttons for the list of selected members. You can set a dictionary attribute to display the order buttons.

Before you begin
Role required: admin

About this task
The slushbucket opens when the user clicks the add/remove multiple users icon.

By default the order buttons on the right are not displayed. Follow the procedure to display the order buttons.

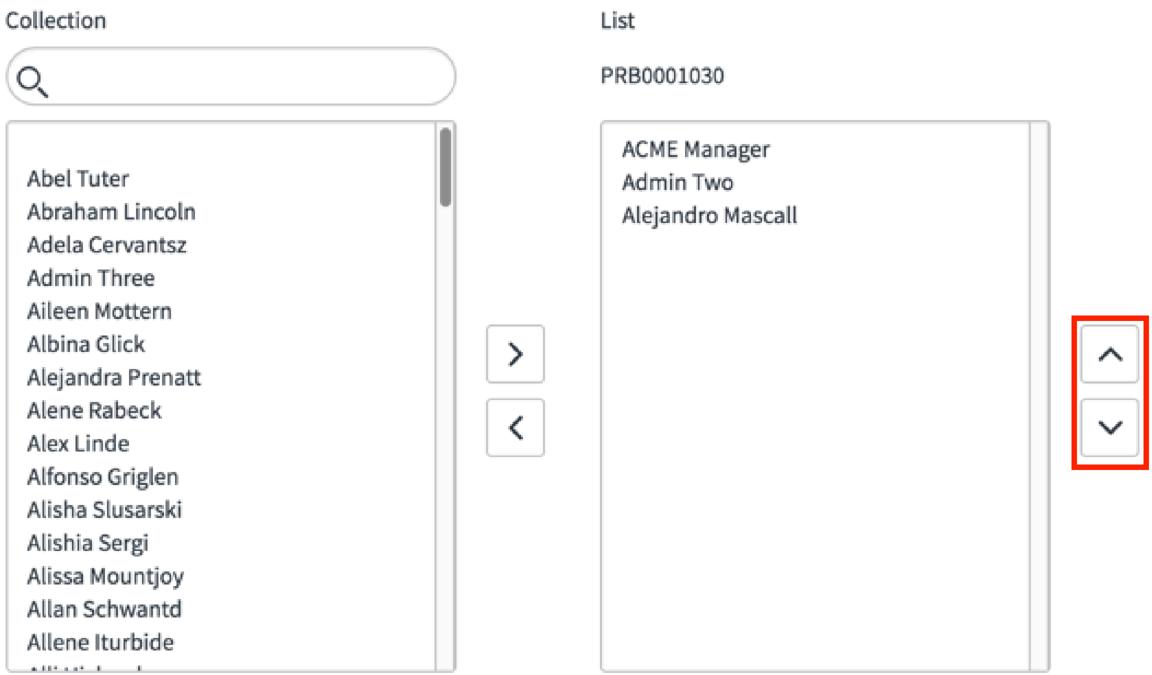

Procedure

1. Open a task record that displays the **Watch list** field.
2. Right-click the label and select **Configure Dictionary**.
3. Under **Related Links**, click **Advanced view**.
4. In the **Attributes** field, enter `maintain_order=true` separated by a comma if necessary.
5. Click **Update**.

Related topics

Altering tables and fields using dictionary attributes

Highlight list fields

Color fields in lists to call an agent's attention to them.

Before you begin
Role required: admin

About this task
Highlighted fields in lists appear colored and optionally, have a leading bullet.

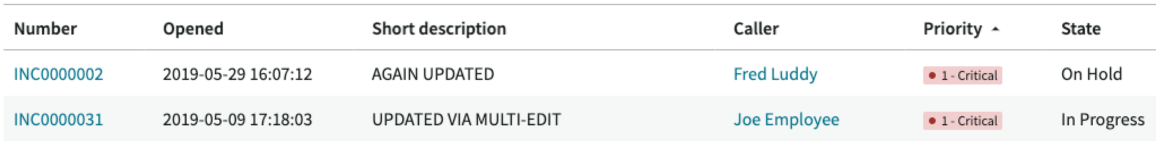

Number	Opened	Short description	Caller	Priority	State
INC0000002	2019-05-29 16:07:12	AGAIN UPDATED	Fred Luddy	1 - Critical	On Hold
INC0000031	2019-05-09 17:18:03	UPDATED VIA MULTI-EDIT	Joe Employee	1 - Critical	In Progress

The set up involves specifying the fields that gets highlighted and the conditions that trigger the highlighting.

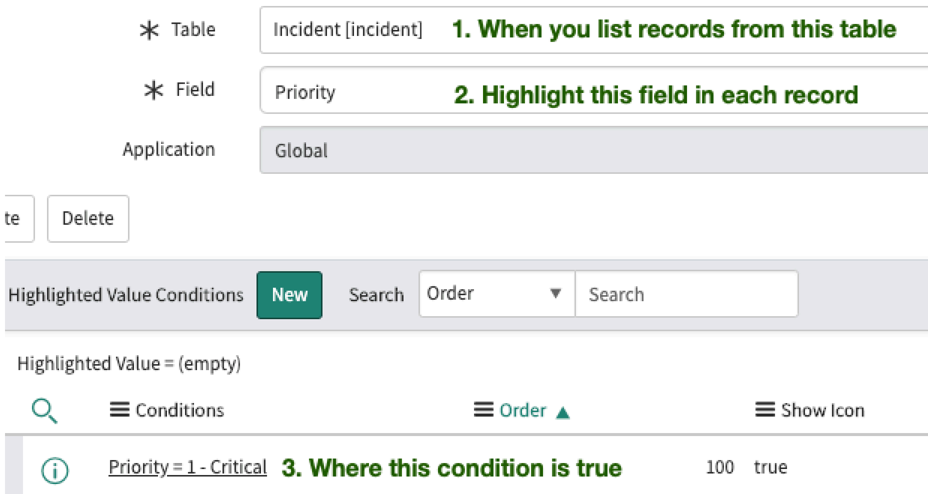

Available highlighting colors are:

You cannot highlight the following field types:

- Reference
- URL
- Document IDs

Procedure

1. Navigate to **All** > **Workspace Experience** > **Administration** > **Highlighted Values** and click **New**.
2. On the form, fill in the fields.

Highlighted Value

Field	Definition
Table	Name of the table that contains the fields you want to highlight.

Field	Definition
Field	Name of the field you want to highlight if the triggering conditions are met.
Application	Scope of applications this highlighting pertains to.
Workspace	The workspace this highlighting applies to.

The **Field** value can be unrelated to the condition that triggers the highlighting. In the following example, if the **Caller** is Fred Luddy, Workspace highlights the **Priority** field, not the **Caller** field.

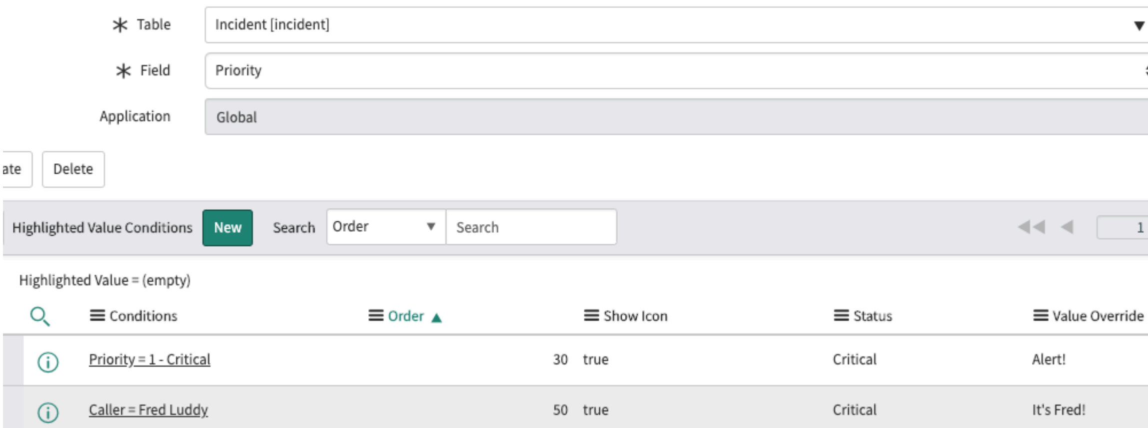

3. Right-click the header and click **Save**.
   You've specified the field where the highlights appear. Now, it's time to specify the conditions that trigger the highlighting.
4. In **Highlighted Value Conditions**, click **New**.
5. In the Highlighted Value Condition form, fill in the fields.

Field	Description
Conditions	Conditions that trigger the highlighting. The condition can be different from the field that gets highlighted. For example, you might highlight the **Priority** field in a record if the **Caller** in that record is Fred Luddy. A simpler example is highlighting the **Priority** field if the priority value is **critical**.
Status	The highlighting color.
Show Icon	Toggle to display (selected) or hide the leading bullet when highlighting.
Value Override	Text that replaces the field values from the table. If you don't supply a **Value Override**, Workspace highlights the field value from the table. If you supply a value, Workspace replaces the field value with your text and highlights that.
Application	Scope of applications this highlighting pertains to.
Order	The order in which to evaluate the conditions. When you have multiple conditions for highlighting the same field, they are evaluated in order from lowest to highest. If any of the conditions are True, Workspace highlights the field.

6. Click **Submit**.

   In the following example, the highlighting was triggered on the **Priority** field because the priority value was critical.

Number	Opened	Short description	Caller	Priority ▲	State
INC0000002	2019-05-29 16:07:12	AGAIN UPDATED	Fred Luddy	1 - Critical	On Hold
INC0000031	2019-05-09 17:18:03	UPDATED VIA MULTI-EDIT	Joe Employee	1 - Critical	In Progress

7. **Optional:** Click **New** on the **Highlighted Value Conditions** Related List and create another condition for highlighting the same field.
   For example, highlight the **Priority** field associated with the Incident table when the **Priority** value is "critical," or when the **Caller** value is "Fred Luddy." The conditions are ranked according to the **Order** value. Workspace evaluates the lowest value first.

Conditions	Order ▲	Show Icon	Status	Value Override
Priority = 1 - Critical	30	true	Critical	Alert!
Caller = Fred Luddy	50	true	Critical	It's Fred!

8. **Optional:** Remove all highlighting from a filtered list that would otherwise have highlighting.
   When setting up highlighting for a table, all filtered lists associated with the table display highlighted fields unless you disable highlighting for specific, filtered lists.

   a. Navigate to `sys_aw_list.list`.

   b. In the **Name** column, enter the name of the table with highlighted fields, for example, Incident.
      All filtered lists associated with the table appear.

   c. Click the filtered list you want to remove highlighting from.

   d. On the Workspace List form, if Global is the current application, click **here** in **To edit this record click here.**

   e. Select the **Ignore Highlight in List**check box and click **Update**.
      Now, regardless of the conditions, Workspace will not highlight any fields in the specified filtered list. This option does not work, however, for filtered lists in **My Lists**, which are the filtered lists agents create.

Modify string field length

You can modify the maximum character limit for a string field.

Before you begin

Role required: personalize_dictionary

About this task

The URL type field hides the Max Length field on the Dictionary form by default. Browsers support URL length up to 2,083 characters. The URL field is a String field.

Procedure

1. Navigate to a form that contains the field that you want to modify.
2. Right-click the field label in the form and select **Configure Dictionary**.
3. Change the **Max length** field to the desired length.

   Note:
   You can change between string-based data types as long as length changes do not cause any data loss from truncation. For example, you can change from a MEDIUM database type to a VARCHAR(100) database type if none of the existing data is greater than 100.

4. Click **Update**.
   The system cancels any length change that results in data loss due to truncation.

Related topics

Field types

Specify a default field value

A default value populates a value in a field when a new record is created.

Before you begin

Role required: personalize_dictionary

About this task

The default value populates the field on the blank form for a new record, and also subsequently when the new record is submitted if the field is empty. Default values can be specified as either a constant or generated through script.

Procedure

1. Right-click the field's label in the form and select **Configure Dictionary**.
2. Enter the default value in the **Default value** field.

   To set a constant value, type it into the **Default value** text box field. To assign a default value using a more complex formula, use JavaScript to output a default value.

   Note:
   The default value should be the underlying value that would be present in the field, not the label. For example, in a choice list field, use the **value** of the choice as the default value, not the choice's name.

Default field value examples

Review the following examples about specifying default field values.

Example: Constant default values

Here is an example that sets a default value for the **Priority** field.

Default value constant

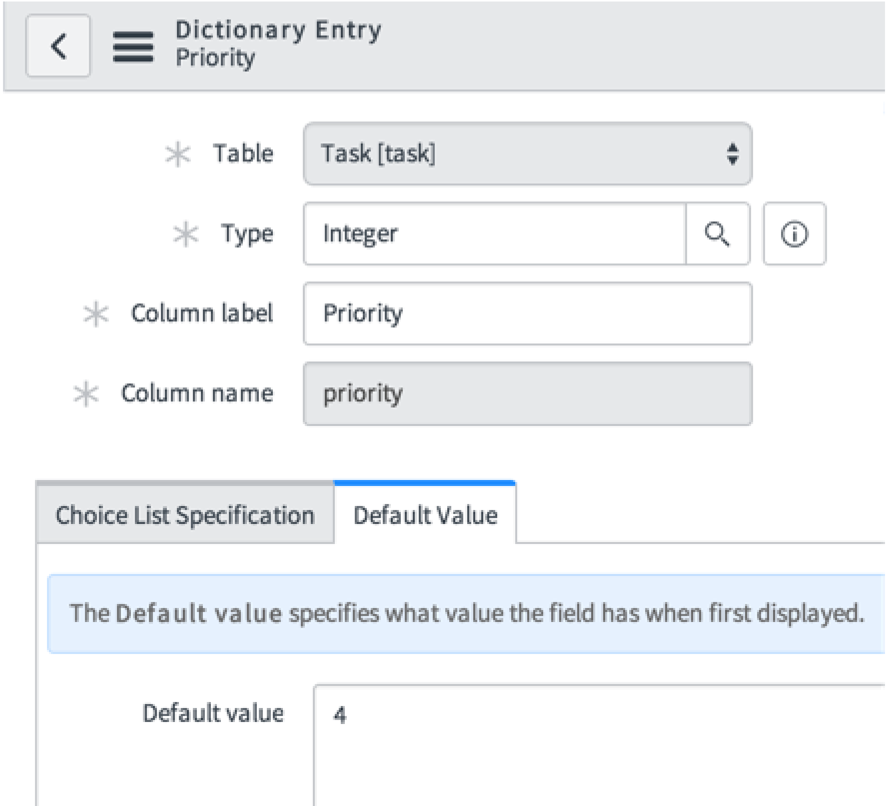

Example: Javascript default values

To view out-of-box examples of JavaScript default values, navigate to **System Definition** > **Dictionary** and enter this filter: **[Default value] [starts with] [javascript]**. Open some of the records and view the default value javascript entries.

Example: Set a default value for assignment_group

The following example sets a default value in a [sys_user_group] reference field by getting the ID from the name of a group.

```
javascript:GetIDValue('sys_user_group', 'Development');
```

Example: Set a default value for assigned_to if user has the itil role

This example for the Task [task] table describes how to configure the default value for the [assigned_to] user equal to the current user id if the user has a role of itil.

- Column label: **Assigned to**
- Column name: **[assigned_to]**
- Reference Specification section
  - Reference: **User** [sys_user]
  - Reference qual condition: **[Roles] [is] [itil]**
- Default value script:

```
javascript:if (gs.hasRole("itil"))current.assigned_to =
 gs.getUserID();
```

Example: Set a default value for a duration field

To set a default value for a duration field, use the following in the **Default value** field of the duration field's dictionary entry:

```
javascript:current.duration_field.setDisplayValue('3 04:30:14');
```

Avoid hard-coding a particular date-time. If the system date-time format changes, the value becomes invalid. Instead, select the **Use dynamic default** check box and assign a dynamic filter option. For more information on dynamic filters, see Create a dynamic filter option .

Make a field dependent

A choice or reference field can be declared dependent on another field on the same table to limit the values available to select based on the value of the dependent field.

Before you begin
Role required: personalize_dictionary or admin

About this task
If a required dependency does not function as expected, as might happen if there is a many-to-many relationship between the fields, consider using reference qualifiers to accomplish the goal.

Note:

- Fields cannot be made dependent on derived fields.
- When you create a dependency between reference fields (similar to assignment group and assigned to), you must add `_group` to the name of the field on which the other is dependent on to get the expected behavior. See KB0998861: Dependent field does not work as expected on task "Assigned to" field for more information.

Procedure

1. Navigate to a form that contains the field that you want to make dependent on another field.

   Example
   For example, to configure a field in the Incident [incident] table, navigate to **All** > **Incident** > **All** and select **New**.
2. Right-click the field label in the form and select **Configure** > **Dictionary**.
3. From the Dependent field tab, select **Use dependent field** and select the name of the field that this field will depend on.

   Note:
   For some field types, you may only need to enter the name of a field in the same table on which the field depends.
4. Select **Update**.

Example: Dependent field

In the example below, the Subcategory field is made dependent upon the Category field. On an incident form, the value selected for the Category field determines which options appear for the Subcategory field.

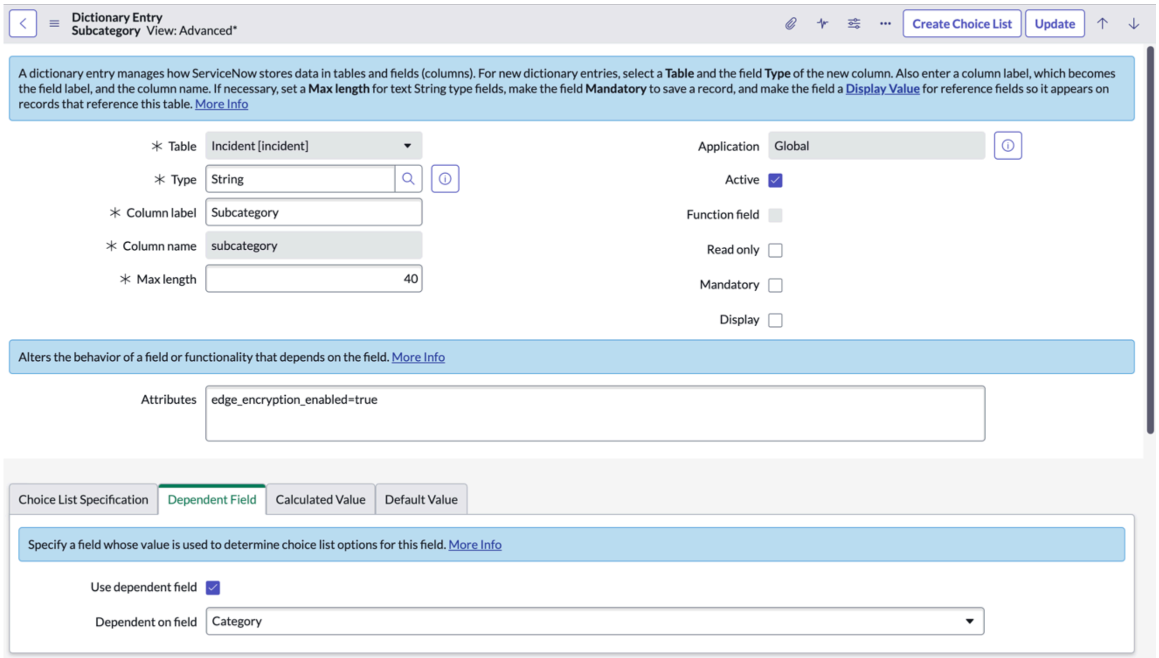

Related topics

Reference qualifiers

Require unique values for a field

The system allows you to require that a field's values be unique. When this is done, the system will not let two records have the same value for that field.

Before you begin

Role required: personalize_dictionary

About this task

By default, fields are created without this constraint. A field can have unique values only if there aren't already duplicate values in the database for that field. The system doesn't allow you to make a field unique while there are duplicate values in the table.

The system also doesn't allow you to add a unique index if a non-unique index already exists. In this scenario, you can't select the **Unique** check box on the form; instead, the system directs you to the **Tables & Columns** module, so you can drop the non-unique index, and then create the unique index.

Procedure

1. Verify that no records in the table for the field have values, or that they all have the same value.
2. Right-click the field label in the form and select **Configure Dictionary**.
3. Configure the form to add the **Unique** field if it does not already appear.
4. Select the **Unique** check box.
5. **Update**.

Related topics

Configuring the form layout

Define field styles

Field styles enable you to declare individual CSS styles for a field in a list or form.

Before you begin

Role required: personalize_styles or admin

About this task

The CSS lets you change the following attributes of a field.

- Change the color.
- Change the font attributes (bold, italics, underline).
- Change the padding and alignment of text.

You can define field styles for tables and database views that are in the same scope as the field style and for other tables that have at least one field in the same scope as the field style. Field styles defined for a table don't apply to database views that include the table. Create separate field styles for database views.

> **Note:**
> Field styles aren't applied to comments and work notes fields used with the activity formatter. Background color for these fields can be set using the `glide.ui.activity_stream.style.comments` and `glide.ui.activity_stream.style.work_notes` system properties located in the sys_properties table.

Procedure

1. Navigate to the list of styles for the field in one of the following ways.
   - Right-click the field label on the form and select **Configure Styles**.
   - Navigate to **System UI** > **Field Styles**, and locate the field to apply a style (admin only).
2. Select **New**, or select the style to modify.
3. Complete the form.

Field	Description
Table	Table that contains the field. **Note:** The list shows only tables and database views that meet the scope protections for field styles.
Field name	The field to which the style applies.
Value	The exact value or script-based-condition required to apply the style. **Note:** The value only affects list field styles when the field isn't read-only. The value affects both list and forms if the field is read-only. To apply field styles to forms only, leave the value field empty.
Style	The CSS style to apply. For example, to make the background of the Incident Category fields red in record lists, with a 24-px font size and white text color, enter the following code.

Field	Description
	`background-color:red; font-size:24px; color:white;`

You can use a script-based-condition with the syntax `javascript:<script>` in the **Value** field to determine whether the system should apply a style to a field in the list. If the evaluated script returns true, the system applies the style to the list item.

The script has access to the *current* script object, which enables you to build conditions based on the field values of the current record. For example, this script checks for overdue items based on a field value:

```
javascript:gs.dateDiff(gs.now(),
current.u_datefield.getDisplayValue(), true) < 0
```

Note:
In this example, `current.u_datefield` refers to the name of a custom date field.

The system only supports one `javascript:<script>` entry in a **Value** field. If you want to enter multiple JavaScript conditions, consolidate the conditions into one statement. For example, instead of this invalid entry:

```
javascript: current.state == "Completed" javascript:
 current.error_tables > 0
```

Consolidate the conditions into one statement such as:

```
javascript: current.state == "Completed" &&
 current.error_tables > 0
```

4. **Optional:** To add alternative text for a style, configure the form to add the alternative text field. For styles like the VIP style icon, you can add alternative text so that screen readers can differentiate between the style and the text.
5. Select **Submit** or **Update**.

Record numbering

In the base system, several tables are numbered, including Incident, Problem, Change Request, and Knowledge. You can also use these numbers anywhere that script is present, for example to generate watermarks for emails. Records in tables can be numbered automatically.

Administrators can manage record numbering by navigating to **System Definition** > **Number Maintenance**. The current number format for a table, including the prefix (such as **INC** for incidents or **CHG** for changes), is stored in a record on the Number [sys_number] table.

Number table

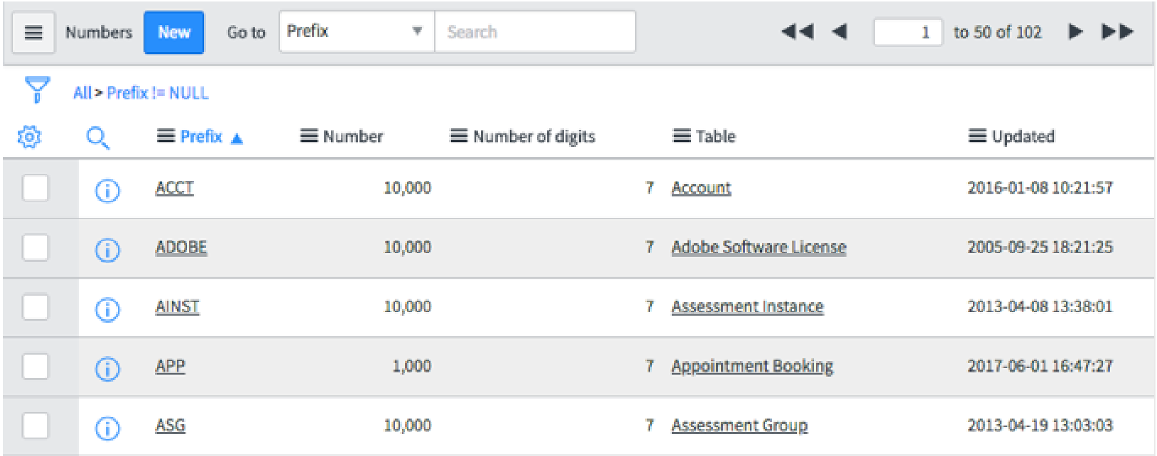

Prefix ▲	Number	Number of digits	Table	Updated
ACCT	10,000	7	Account	2016-01-08 10:21:57
ADOBE	10,000	7	Adobe Software License	2005-09-25 18:21:25
AINST	10,000	7	Assessment Instance	2013-04-08 13:38:01
APP	1,000	7	Appointment Booking	2017-06-01 16:47:27
ASG	10,000	7	Assessment Group	2013-04-19 13:03:03

You can renumber auto-incremented tables that extend the task table or manage numbering with a database field named **Number this field**.

Note:
If you reset numbering in Number Maintenance, it only renumbers new records. It leaves the numbering as-is for existing records in system tables.

Add auto-numbering records in a table

You can define one number format per table in the system.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Number Maintenance**.
2. Select a table name to view the number record for that table, or click **New**.
3. Define the number format by completing the fields (see table).
4. Select **Submit** or **Update**.

   If an auto-numbered field does not already exist, a new field is automatically created on the table with the following values:

   - Label: Number
   - Name: u_number
   - Default value:

   ```
   javascript:getNextObjNumberPadded();
   ```

   Note:
   This script renumbers records when the **Number of digits** is updated. To use a script that does not renumber records when the **Number of digits** is updated, open the dictionary entry for the **Number** field and enter the following script in the **Default value** field.

   ```
   javascript:getNextObjNumber();
   ```

Auto-numbering records in a table

Field	Description
Table	Select a table.
Prefix	Enter a prefix for every number in the table (for example, **INC** for Incident).
Number	Enter the base number for this table (default value is 1000). Record numbers are automatically incremented, and the next number is maintained in the Counter [sys_number_counter] table. If you set the base number to a value higher than the current counter, the next record number uses the new base number. Otherwise the next record number uses the current counter. The counter does not reset to a base number lower than itself. To see the current counter, click the **Show Counter** related link.
Number of digits	Enter the minimum number of digits to use after the prefix (default value is 7). ◦ Leading zeros are added to auto-numbers, if necessary. For example, **INC0001001** contains three leading zeros. ◦ The number of digits can exceed the minimum length. For example, if **Number of digits** is **2** and more than 99 records are created on the table, the numbers continue past 100 (such as **INC101**). **Warning:** Changing this field may update all number values for existing records on a table. Take care when changing this field on a production instance.

When moving some applications these records may not be automatically included. Take care to ensure that all necessary records are present when moving applications.

Note:
To change the default values for new number formats, change the **Default value** field on the system dictionary record for the **Number** or **Number of digits** field. These fields are on the Number table.

Related topics

System dictionary

Prepare to left-pad number fields in custom tables

Before you configure left padding of number fields on a custom table or a table that does not extend the task table, you must prepare business rules and script includes.

Before you begin

Role required: admin

Procedure

1. Duplicate the base system business rule named *Pad Numbers*: Click **Insert and Stay**.
   If **Insert and Stay** is not available on your instance, create a new business rule and manually copy all field values from Pad Numbers.
2. Give the business rule a new name and change the script field to read as follows:

```
padCurrentCategory();
function padCurrentCategory() {
```

```
var target = new GlideRecord(current.category+"");
if (!target.isValidField("u_number") ||
 target.isValidField("number")) return;
var nm = new UNumberManager();
nm.padTableNumbers(current.category, current.maximum_digits);
};
```

3. Save the changes.
4. Duplicate the base system script include named *NumberManager*.
5. Change the following two lines as indicated:

```
105: currentNumber = records.u_number.toString();
117: records.u_number = prefix + currentNumber;
```

6. Name the new script include `UNumberManager` and save your changes.
   You can now continue with the process of configuring left-padding and renumbering records.

Configure left padding of a system number in a table

You can configure the left padding of the system numbers on a table. For example, pad the **Number** field on an Incident, Problem, or Change Request.

Before you begin

If you are configuring numbers on a custom table or a table that does not extend the task table, then, before performing the following procedure, you must prepare business rules and script includes. For more information, see Prepare to left-pad number fields in custom tables.

Role required: admin

Procedure

1. Navigate to the form, then right-click the **Number** field and select **Configure Dictionary**.
2. Enter the following script in the **Default value** field and click **Update**.

```
javascript:getNextObjNumberPadded();
```

3. Navigate to **System Definition** > **Number Maintenance**.
4. Open the table record.
5. Enter a value in the **Number of digits** field.
6. Click **Update**.

   Number padding is applied to both existing and new records.

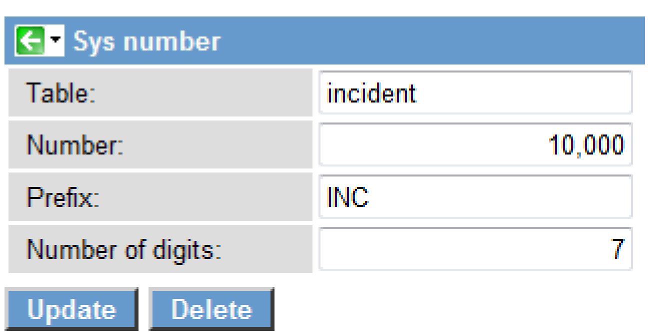

The result of the configuration in the image is an Incident number that is left padded.

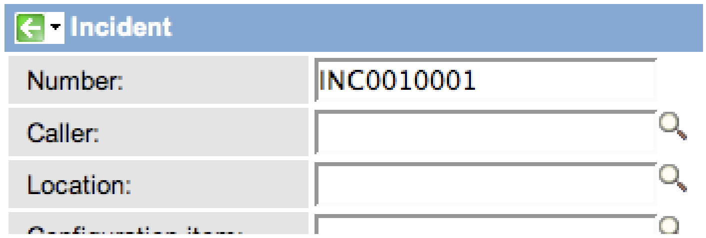

Prevent numbering gaps

By default, numbers are generated every time a new record is created.

Before you begin

Role required: admin

About this task

When records are created but not saved, a gap in the numbering is created. You can prevent these numbering gaps by generating numbers only when records are saved.

Procedure

1. Navigate to **All** > **System Properties** > **System**.
2. Set the property **Assign a task number only upon insert (prevents unused numbers)**, *`glide.itil.assign.number.on.insert`*, to **true**.
3. Click **Save**.

Enforcing unique numbering

Although duplicate numbers are rare, numbering does not enforce uniqueness, by default.

To enforce uniqueness, you can:

- Create a before business rule on insert only to check for duplicate values and replace duplicates with the next available number.
- Enable a unique index on the table. For more information, see Create a table index.

> **Note:**
> While unique indexes ensure data integrity they also prevent any insert involving a duplicate number. This may cause unexpected errors during data entry.

Sample business rule

This sample script can be used as part of a before business rule on insert only to check for duplicate numbers and replace them with the next available number. The following script references a script created in Configure left padding of a system number in a table.

```
var curNum = current.number + '';

if(curNum) {

  var recordClass = current.getRecordClassName();
  var now_GR = new GlideRecord(recordClass);
  now_GR.addQuery('number', curNum);
  now_GR.setLimit(1);
  now_GR.query();
```

```
  if(now_GR.getRowCount() > 0) {
    var newNum = getNextObjNumberPadded();
    gs.addInfoMessage("The number " + current.number + " was already used
by another " +
      recordClass + ". The " + recordClass + " number has been changed to " +
 newNum);
    current.number = newNum;
  }
}
```

Related topics

Business rules

Field normalization and transformation

Field Normalization includes normalization and transformation, which are two different ways to alter field values for increased data integrity and reduced duplication.

Normalization

Normalization searches for variations of the same field value and converts them into a single preferred value. By consolidating multiple variations of the same value into a single simple recognizable value, the system eliminates duplicate records and provides better search results. When a process or a user enters a value in a normalized field, the system determines whether to replace it with a normal value. Normalization also automatically adjusts queries to return normalized results and normalizes values in scripts.

While normalization is available for every field in the platform, it works best for descriptive values such as names or standard units of measurement. For example, you might create normalization rules to:

- Set the CPU type of a computer CI to a standard model name such as Xeon.
- Set the suffix used for the names of corporations to a standard format such as ServiceNow, Inc..

Transformation

Transformation converts raw field input values into standardized values that are more meaningful to an organization. Administrators control when transformation happens by defining rules and conditions for specific fields. For example, you might create transformation rules to:

- Remove suffixes from user names such as Jr. and II.
- Round computer CI RAM sizes to the nearest whole number such as rounding 4112 MB to 4000 MB.

Transform records make up the rules that define how a field transformation is executed. Order values determine the order in which each rule is evaluated. A check box on each transform enables an administrator to determine where processing stops when a rule evaluates to true.

Enabling normalization and transformation by field type

Field type records specify which data types are available for normalization and transformation. By default, the system supports normalization and transformation for these field types.

Field type	Use to normalize	Use to transform
Decimal	false	true
Float	false	true
Integer	false	true
Numeric	false	true
String	true	true
URL	true	true

A field type entry applies to all fields whose dictionary entry data type matches the field type entry. Administrators can create additional field type records for other data types as needed.

Warning:
Avoid creating field type records for fields that store a Sys ID value such as a reference field, field name field, or a table name field. Directly altering a Sys ID value is more likely to produce data corruption and broken references than to produce meaningful standard values. Instead, normalize or transform the display value associated with the Sys ID. For example, normalize the user name instead of the Sys ID of a specific user.

Identifying normalized fields

The system displays the normalization icon on fields with an associated normalization or transformation record. Users with the normalizer role can click the icon to access the associated normalization or transformation record. Users without the normalizer role instead see a help page. Administrators can configure who sees the normalization icon with a preference called `Restrict to roles`.

Activate Field Normalization

Activate the Field Normalization [com.snc.field_normalization] plugin to install demo data and activate related plugins if they are not already active.

Before you begin

Role required: admin.

Procedure

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

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

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

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

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

Related topics

List of plugins (Washington DC)

Installed with field normalization

Several components are installed with Field Normalization.

The following preferences (**Field Normalization** > **Administration** > **Preferences**) control features for normalization and transformation:

Field Normalization preferences

Preference	Description
Enable field normalization	This preference enables or disables the Field Normalization functionality, including transformations.
Enable field normalization auditing	This preference enables or disables auditing of field normalization tables.
Logging	Select the maximum level of logging detail desired. For example, select Error to log only errors and Information to log errors, warnings, and information.
Decoration URL	Enter the URL for the help/information link associated with the Field Normalization decoration. This link opens the page presented to users who do not have rights to see the normalization or transformation record. The default link opens a ServiceNow Wiki page with a general explanation of Field Normalization. The purpose of this page is to explain to normal users why a field value they entered was changed automatically.
Restrict decorations to roles	Use this preference to define the user roles that can view the reference icon (decoration) adjacent to a normalized field. This decoration displays transformation and normalization records to users with the normalizer role. For all other users, it opens the URL specified in the Decoration URL preference. If this preference does not specify any roles, then all roles can see the decoration.

Normalization preferences

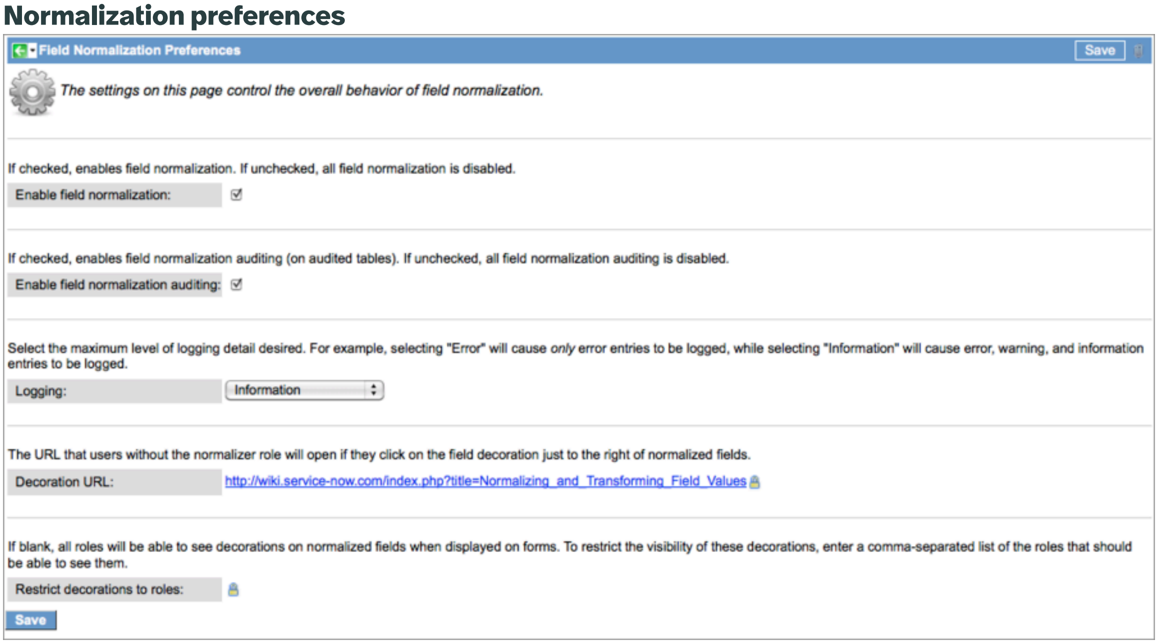

Field Normalization business rules

The following business rules were created for this plugin:

- Ensure Rules Application Job
- Ensure Transform Application Job
- Ensure Rules and Alias Jobs
- Ensure Transform Application Job
- Ensure Decoration Attribute
- Cleanup on field or table change
- Prevent duplicates in class hierarchy
- Handle potential duplicates
- Run job
- Ensure Pending Value Collection Job
- Ensure name changer job
- Flush forms when activating

Field Normalization script includes

Name	Description
FieldNormalizationAjax	Provides AJAX services for the Field Normalization plugin.
FNEnsureJob	Ensures that a job record exists for a particular extant data job target.
FNExtantDataJobChoices	Generates a list of tables for use in document_id field of the fn_extant_data_normalization table.
FNExtantDataJobUtil	Determines whether an extant data job is allowed to run.
FNFields	Returns valid fields for normalization.
FNTransformChooserUtil	Provides support for the fn_transform_chooser UI page.
Position	Finds a position within a given string.
Round	Rounds numbers with various rounding modes and intervals.

Field Normalization UI pages

Name	Description
aliases	Creates the slushbucket with pending values for choosing aliases for a normal value.
fn_transform_chooser	Displays available transforms. Invoked by the New button on the fn_transform related list.

Field Normalization roles

Role	Description
normalizer	Person authorized to manage field normalization and transformation.
normalization_tester	Person authorized to create test records for normalizations and transformations. Only records opened by users with this role are available for testing normalization and transformation rules in Test mode.

Related topics

Enable auditing for a table

Normalizer [normalizer]

Manage field normalization and transformation.

Contains Roles

The Normalizer [normalizer] role contains the Normalization Tester [normalization_tester] role.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Normalization Tester [normalization_tester]

Create test records for normalizations and transformations. Only records opened by users with this role are available for testing normalization and transformation rules in Test mode.

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.

Enable a field type for normalization or transformation

Create or modify a normalization field type record to enable a specific field data type for use with normalization or transformation.

About this task

To normalize or transform a value in a reference field, apply the processing to the field in the target table.

Procedure

1. Navigate to **All** > **Field Normalization** > **Administration** > **Normalization Field Types**.
2. Click **New** in the record list.
3. Enter a **Name** for the field type that clearly describes the type in the dictionary.
   This value is for reference only and is not used in any processing. For example, you might enter **IP Address** for the field type of `ip_address`.
4. Enter the **Type** from the dictionary.
5. Select the appropriate check box to use this field type to normalize or transform fields.
   In the base system, only the **String** field type is used for normalization.
6. Right-click in the header bar and select **Save** from the choice list.
   The **Transform Categories** Related List appears.
7. If this field type is being used for transforms, click **Edit** to associate an existing Transform Category with this field type.

   Note:
   If you create a custom field type that is used for normalizations only, a link to a transform category is not necessary.

   The relationship of a field type to a category, and the category to a list of transformation definitions, is completely configurable.

   The String field type is associated with the Text Transform Category, which contains these transform definitions. Any of these associations are configurable.

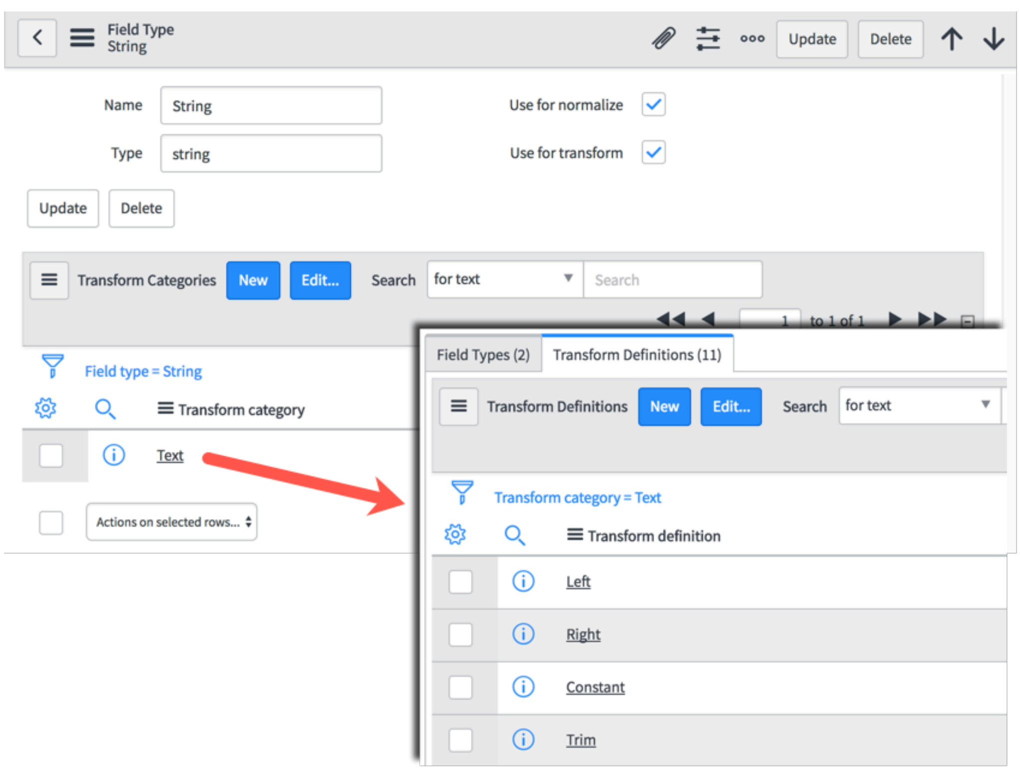

Create a raw field

A raw field is a custom field created by an administrator to show the original (raw) input in a field on a form after it has been normalized or transformed.

About this task

An administrator might add a custom field to a form to show the original, or raw, value of a normalized field.
This is a read-only field that might be called something like **Raw CPU type** or **Original Name**. In the following example, the **CPU type** field was normalized to **Xeon** from an original, raw value of **Xeon L3350**.

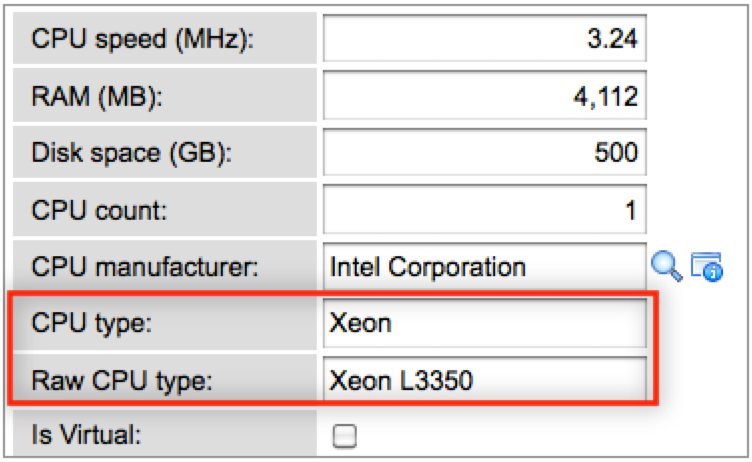

Procedure

1. In the form containing the field that is being normalized or transformed, right-click in the header bar.
2. Select **Personalize** > **Form Layout**.
3. Complete the Create new field form at the bottom of the page, and then click **Add**.
   - **Name**: Type the field label. In this case, use Raw + <field label>.
   - **Type**: Select a data type from the list for this field.
   - **Field length**: Select the character limit for this field. The default is 40.
4. Move the new field adjacent to the normalized field using the direction arrows in the slushbucket.
5. Click **Save**.

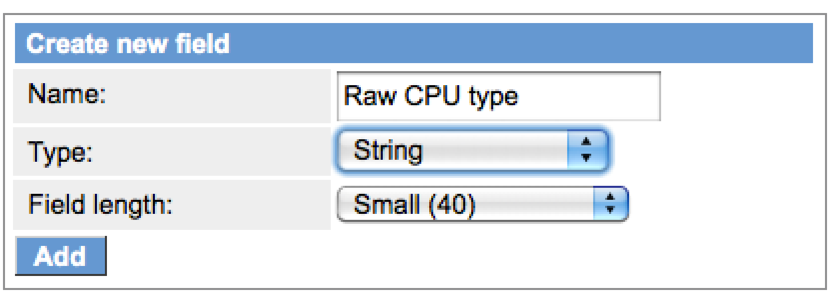

Normalization and transformation data jobs

The system uses data jobs to change field values during normalization and transformation.

Administrators must manually start data jobs that apply field changes. All transformation records open in Test mode by default, which blocks administrators from running manual data jobs that

apply changes to existing data. Most data jobs have a roll back option to revert changes after the job completes. By default, the system provides these data jobs.

Note:

If the Field Normalization plugin is active, the normalization engine executes on insert of records into a table according to normalization rules. To learn more about this plugin, see Activate Field Normalization.

If the Field Normalization plugin is not active, a data job runs to normalize the values. The data job saves the previous values in the rolled back fields to revert them if necessary.

Name	Used for	Description
Pending value collection	Normalization	The system runs this job to collect field values that will change during the next field normalization run. The system runs this job whenever you insert or update a field normalization record. Since this data job does not change any field values, there is no roll back option.
Normal value change	Normalization	Run this job to rerun field normalization and update field values when you change a normal value. The system creates this job whenever you change a normal value. Roll back the data job to revert the normalized fields to the previous normal value.
Alias application	Normalization	Run this job to normalize all field values that match an alias. The system creates this job whenever you create or update an alias. Roll back the data job to revert normalized fields to their previous values.
Rule application	Normalization	Run this job to normalize all field values that match a rule. The system creates this job whenever you create or update a rule. Roll back the data job to revert normalized fields to their previous values.
Coalesce too normal	Normalization	Run this job to normalize the reference field values by replacing references to multiple duplicate records with one reference to a normalized record. The system creates this job whenever you change a normal value that applies to a reference field. Roll back the data job to restore the duplicate records as valid options. The rollback does not restore references to the previous records.
Transform application	Transformation	Run this job to transform all field values that match the transformation

Name	Used for	Description
		conditions. You cannot start this job while the Transformation record is in **Test** mode. The system creates this job whenever you create or update a transformation record. Roll back the data job to revert transformed fields to their previous values.

Run a single data job

Running the data job/jobs is the fourth step in transforming a field. This topic explains how to run a single data job.

Procedure

1. Start the Transform application data job to apply the transform to all the appropriate records in the database.
2. Open the Transformation record and switch the **Mode** to **Active**.
3. Select the **Data Jobs** Related List.
4. Run the data job using one of the following methods.

   - In the list of data jobs, select the check box of the jobs you want to run, and then select **Start** from the **Actions** menu.

   - Click the link in the Created column to open the data job, and then click the **Start** Related Link.

   The **State** of the jobs turns to **Queued** as it runs and to **Completed** when it finishes successfully.

Run multiple data jobs

Running the data job/jobs is the fourth step in transforming a field. This topic explains how to run data jobs for multiple field transformations.

Procedure

1. Start the Transform application data job to apply the transform to all the appropriate records in the database.
2. Navigate to **Field Normalization** > **Data Jobs** > **All**.
3. Select the check boxes for the jobs you want to start.

   Note:
   These jobs must have the **Mode** set to **Active**.

4. Select **Start** from the **Actions** menu.

   Note:
   The platform only runs data jobs from Active transformations. The Action menu indicates the number of Active data jobs that can run. For example, the menu might display Start (3 of 4).

Rollback a data job

Use rollback to revert changes made by a normalization or transformation data job.

Before you begin
Role required: admin or normalizer

About this task
You can only rollback completed data jobs.

Procedure

1. Navigate to **Field Normalization** > **Data Jobs** > **All**.
2. Select the completed data job to rollback.
   The system displays the data normalization job record.
3. From **Related Links**, select **Rollback**.
   The values in the **Pending Values** Related List are removed and the data job **State** indicates that it was rolled back.

Normal values

A normal value replaces similar but ambiguous field values with one standard value.

Field value variations

Records values can come from multiple sources such as:

- Automated entries made by Discovery.
- Automated entries made by importing records from external systems or files.
- Manual entries made by users.

Each of these sources may describe the same field value in several different forms. For example, the **CPU Type** field on a computer CI form might display any of the following similar values:

- **E3350 (Intel) 4.5.2234**
- **Intel Xeon 5.4.554**
- **Xeon L3350**
- **L3350**

Without normalization, these variant field values results in:

- Duplicate CPU types
- Poor search results
- Complex queries and conditions to apply business logic

Creating a normal value record solves these issues by consolidating on one standard value such as `Xeon`.

Identifying variations with aliases and rules

Each Normal value record specifies how to identify variations of a normal value using a combination of aliases and rules.

Aliases

Aliases are known variations of an input value that normalization converts to the normal value. Use aliases when there is a short list of variant values.

For example, you could create a normal value `Xeon` that has these aliases.

- **E3350 (Intel) 4.5.2234**
- **Intel Xeon 5.4.554**
- **Xeon L3350**
- **L3350**

Whenever a normalization data job or normalized query sees a field value matching an alias, it automatically replaces the field value with the normal value. Normalization data jobs and queries process aliases before rules.

Note:
Aliases are logically equivalent to rules using the **[is]** operator in a condition where **[Field name][is][Alias value]**. For example, the sample aliases are equivalent to these rules: **[CPU Type][is][E3350 (Intel) 4.5.2234] OR [CPU Type][is][Intel Xeon 5.4.554] OR [CPU Type][is][Xeon L3350] OR [CPU Type][is][L3350]**

Rules

Rules specify the conditions under which normalization replaces an input value with the normal value. Use rules when there are a large number of possible variant values, or when you must create complex conditions.

For example, the normal value `Xeon` could have this rule.

[CPU Type][matches regex][.*\bxeon\b.*]

Whenever a normalization data job or normalized query sees a field value matching a rule, it automatically replaces the field value with the normal value. Normalization data jobs and queries process rules after aliases.

Rules and aliases can be combined to normalize a field. Make sure to test your normalization methods before applying them to all the existing records in the database.

Normalized queries

An administrator can configure normalization to apply to queries issued against normalized fields in lists. Select the **Normalize query** check box on the Normalization form to enable this functionality. In a list containing normalized values, Filters and breadcrumbs using the original (raw) value for the normalized field in the query condition.

Normalized query example

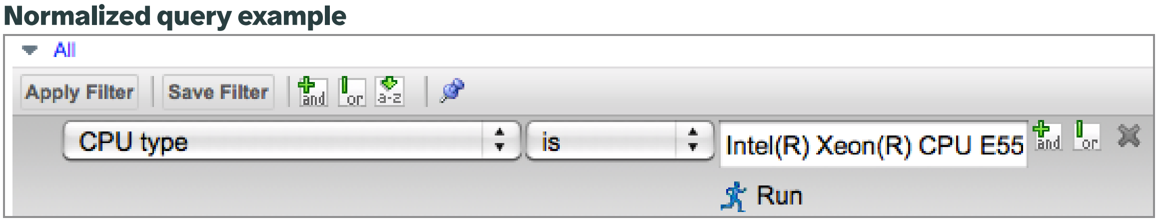

The filtered list returns records with the normal value substituted for the raw value. However, the breadcrumbs for the filter display the original query conditions.

Normalized query results

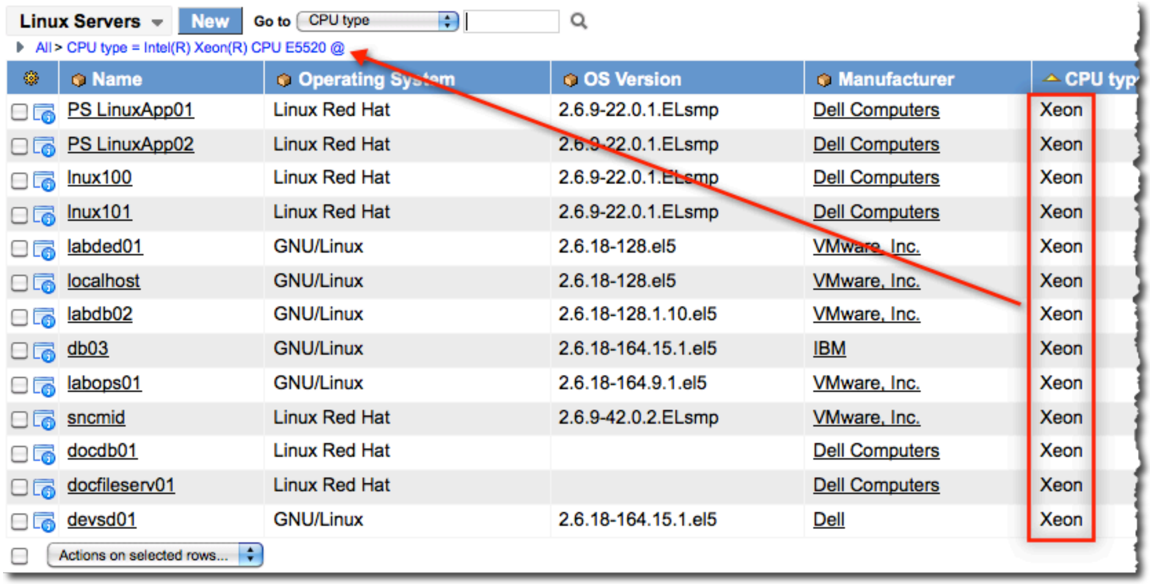

Linux Servers | New | Go to CPU type

All > CPU type = Intel(R) Xeon(R) CPU E5520

Name	Operating System	OS Version	Manufacturer	CPU typ
PS LinuxApp01	Linux Red Hat	2.6.9-22.0.1.ELsmp	Dell Computers	Xeon
PS LinuxApp02	Linux Red Hat	2.6.9-22.0.1.ELsmp	Dell Computers	Xeon
lnux100	Linux Red Hat	2.6.9-22.0.1.ELsmp	Dell Computers	Xeon
lnux101	Linux Red Hat	2.6.9-22.0.1.ELsmp	Dell Computers	Xeon
labded01	GNU/Linux	2.6.18-128.el5	VMware, Inc.	Xeon
localhost	GNU/Linux	2.6.18-128.el5	VMware, Inc.	Xeon
labdb02	GNU/Linux	2.6.18-128.1.10.el5	VMware, Inc.	Xeon
db03	GNU/Linux	2.6.18-164.15.1.el5	IBM	Xeon
labops01	GNU/Linux	2.6.18-164.9.1.el5	VMware, Inc.	Xeon
sncmid	Linux Red Hat	2.6.9-42.0.2.ELsmp	VMware, Inc.	Xeon
docdb01	Linux Red Hat		Dell Computers	Xeon
docfileserv01	Linux Red Hat		Dell Computers	Xeon
devsd01	GNU/Linux	2.6.18-164.15.1.el5	Dell	Xeon

Actions on selected rows...

Scripting and normalization

Scripts that update or insert records into the database (GlideRecord) are normalized automatically when field normalization is applied. For example, if a script to insert a CI record contains a CPU type of **Xeon L3350**, the script is normalized to insert the CI with a CPU type of **Xeon** instead. Scripts that query the database for normalized field values (using the conditions of equals or not equals) can be configured to return the normal value (such as **Xeon**) rather than the original (raw) value.

Related topics

GlideRecord

Create the normalization record

Regardless of the normalization method selected, all field normalization requires a list of existing variants and a normal value that is configured to replace these variants in forms and in queries.

About this task

Normalize a field by selecting aliases for a normal field value or by creating rules that use condition statements to match field variants with a normal value.

Note:
Users must have the normalizer role to create and manage normalization records.

Special Cases

- Reference fields cannot be normalized. To normalize values appearing in reference fields, normalize the field in the target table. Examples of this would be normalized values for the **Name** fields in the Company [core_company] and Location [cmn_location] tables, which are referenced by other tables in the platform. The normalized names are available to all fields that reference these tables.
- Fields in a choice list can be normalized if they are of a type **string**.

Procedure

1. Activate the Field Normalization plugin.
2. Navigate to **Field Normalization** > **Configurations** > **Normalizations**.
3. Click **New**.
4. Create a normalization record.

Normalization record fields

Field	Input value
Name	Unique name for this normalization record. This value is for reference only and is not used in any processing.
Table	Select the ServiceNow table for the field being normalized.
Field	Select the field to normalize.
Mode	The three available modes are Off, Test, and Active. All normalization records are created in the test mode by default. If you are planning to select aliases for your normal values, change the Mode to Active. If you intend to normalize the field using rules, be sure to leave this record in the Test mode. To disable this normalization, switch the mode to Off.
Normalize query	Select this check box to apply the field value normalized by this record to all queries involving this field. Queries formed with the raw (original) field value return records displaying the normal value. Queries issued by a script using the conditions equals or not equals return normalized values. See Normalized Queries for details.
Coalesce each normal	Select this check box to normalize reference field values. This option resets all references to records containing an alias field value to the record using the normal value. See Coalesce Normal Values for details.
Raw field	Select the field to use to display the original (non-normalized) input values on a form in which a field value has been normalized. For the selection to appear in the drop-down list, add a custom field to the form for the table selected. For instructions on adding a field for raw data, see Creating a Raw Field.

After the normalization record is submitted, the platform runs the Pending value collection job in the **Data Jobs** related list automatically to gather all the current values (Pending Values) for the field being normalized.

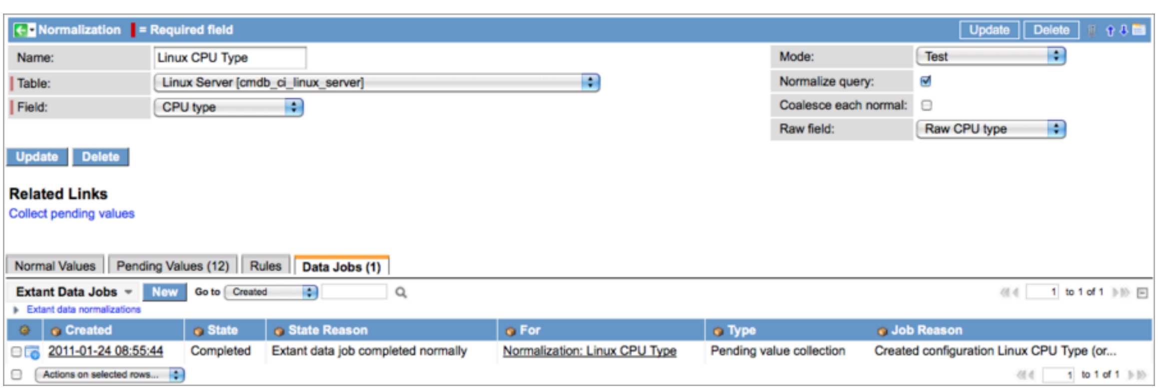

Click the **Pending Values** related list to view values for the normalized field that have been entered manually, imported into the platform, or created by Discovery.

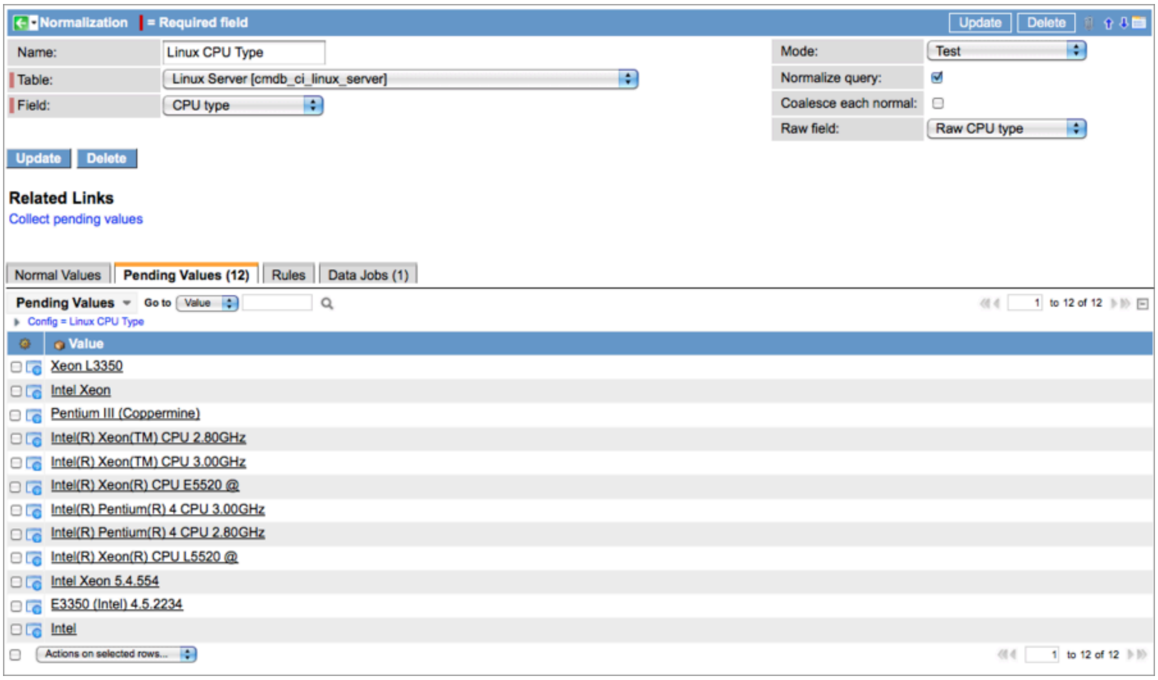

Create a normal value

A normal value is a simplified, generic value for a field that replaces all the possible variants of that value that exist in the database.

Before you begin

Role required: admin or normalizer

About this task

Normal values should be clear and unambiguous.

After the platform runs the data job, the **Pending Values** related list on the Data normalization jobs form is populated with all the unique values for the field in the database. Examine the values in the list and decide which normalizing method is best for the existing data. For example, define an alias for a small pool of values and a rule for a large pool of values. The following screenshot shows the pending values for CPU types in Linux servers in a network. The list contains several choices for Intel Xeon CPUs, which might be normalized as **Xeon**.

Normalization pending values 2

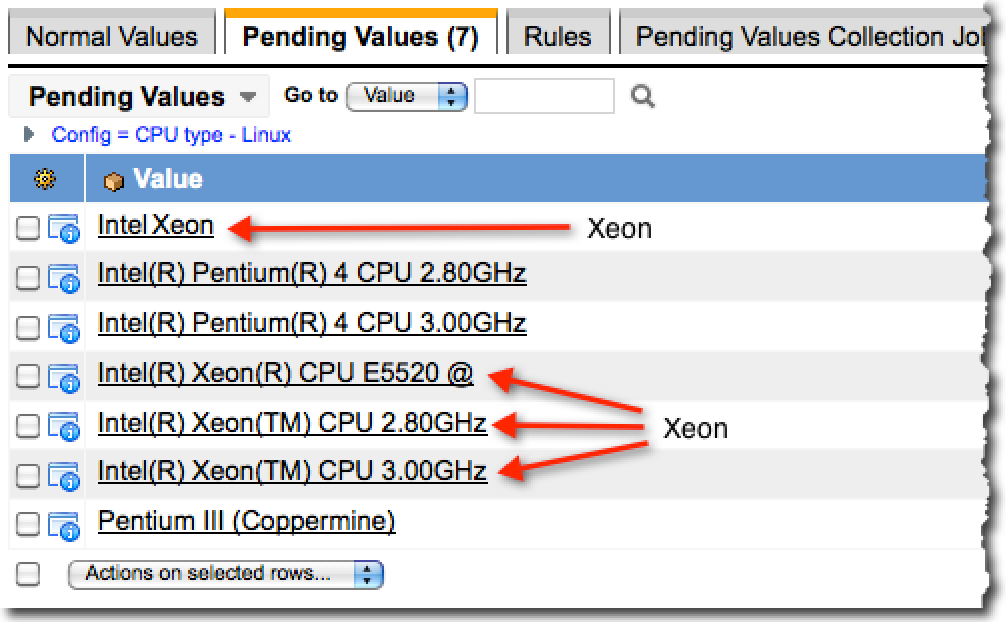

Procedure

1. Navigate to **All** > **Field Normalization** > **Normalizations**.
2. Open the appropriate normalization record.
3. Click the **Normal Values** related list.
4. Click **New**.
5. In the Normal Value form, create normal values for the variants in the **Pending Values** related list.
   These are the values the platform uses to replace the variants configured as aliases.

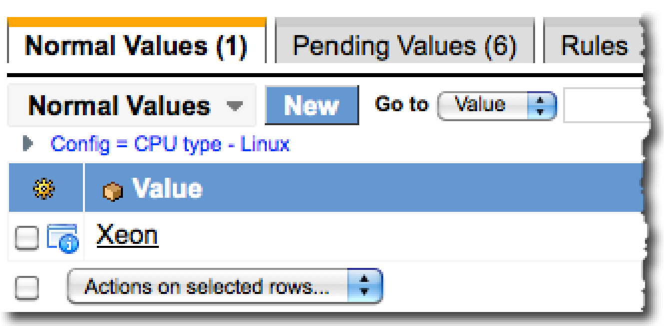

Create aliases

Aliases are the variants of a field value in the instance that will be replaced by the normal value.

Before you begin

Role required: admin or normalizer

About this task

The list of potential aliases is the contents of the **Pending Values** related list. After creating a normal value, assign aliases to this value if the pool of pending values is small. A normalized field can have a combination of aliases and rules.

Procedure

1. Navigate to **All** > **Field Normalization** > **Configurations** > **Normalizations**.
2. Open a normalization record.
3. Click the **Normal Values** related list.
4. Select one of the values.
5. In the normal value record, click the **Aliases** related link.
6. Select aliases for this normal value from the available (pending) values that appear in the slushbucket, and then click **OK**.

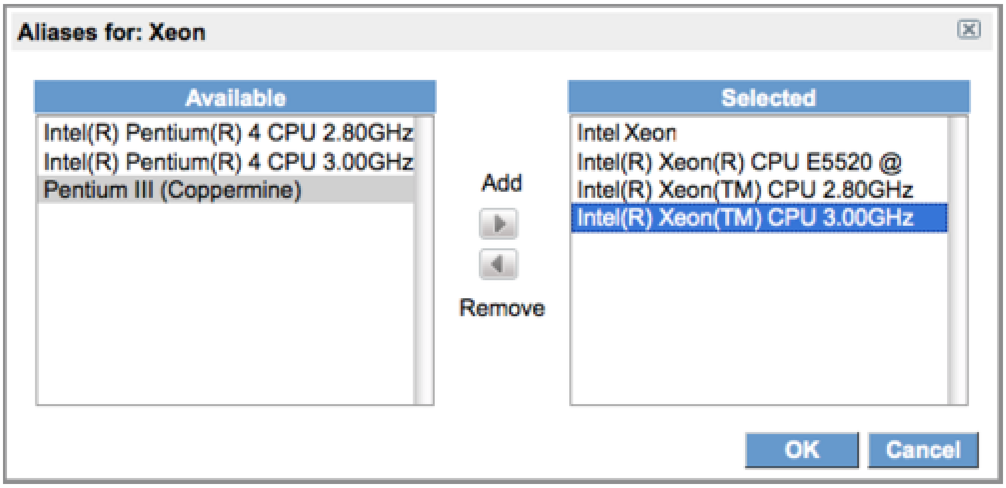

The aliases for this normal value now appear in the **Aliases** related list.

What to do next

Apply the aliases by running the associated data jobs.

Apply aliases

After testing, aliases can be normalized in all new records or in existing records when they are updated.

About this task

Each time an alias is created for a normal value, a data job is created. The alias is not applied to values in the entire database until its data job is started manually. Run each job separately or run the jobs together to apply all aliases at once.

Procedure

1. In a normalization record, ensure that the Mode is set to **Active**.

   Data jobs cannot run in the **Test** mode.

2. Click the **Normal value** related list.
3. Select a value from list.
4. In the Normal Value record, select the **Data Jobs** related list.

   A data job is listed for each alias configured for this normal value.

5. Run the extant data jobs to replace the aliases with the normal value in all existing records in the database.

a. Select the check box next to a job, and then select **Start** from the Actions menu.

b. To run all data jobs at once, select all the check boxes, and then select the **Start** action.

c. Refresh the list to check the progress of the data jobs to ensure that they complete normally.

Create rules

The use of rules to normalize a field is intended for large lists of variant field values.

About this task

Always test your rules before applying them to all the existing records in the database. Prior to creating the rule, make sure to generate the list of Pending Values and create a normal value for the field. A normalized field can have a combination of aliases and rules.

The rules in this example are based on the following Pending Values:

Normalization rules 3

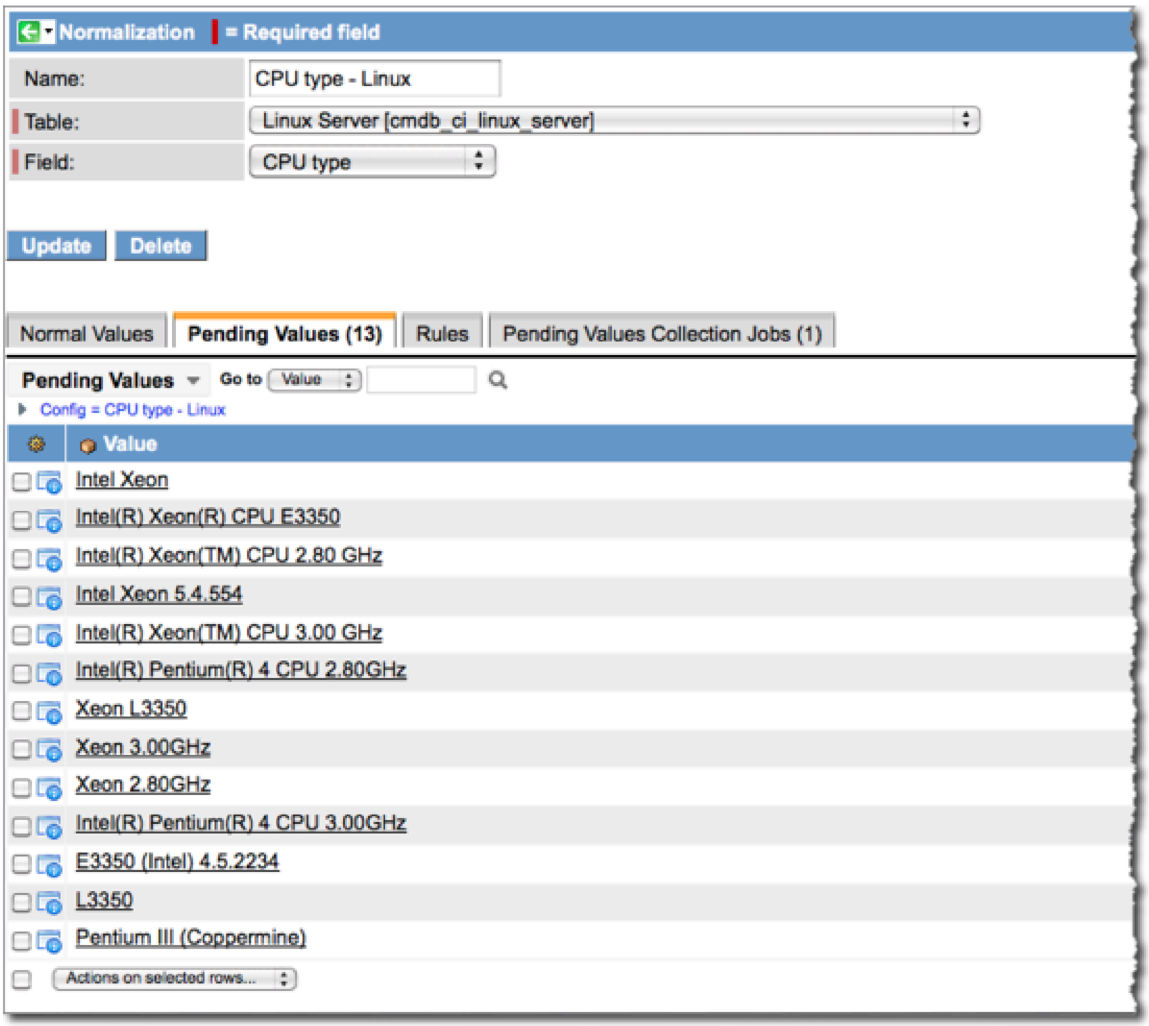

Procedure

1. In a Normalization record, open the **Normal Values** related list.
2. Open a Normal Value record.
3. Open the **Rules** related list in the Normal Value form, and then click **New**.

   The Field Normalization Rule form provides the following fields:

Field normalization rule fields

Field	Description
Name	Unique name for this rule. For example, this name might be Xeon CPU Type.
Order	Order in which this rule should be evaluated. The platform parses the rules in the order configured until it finds one that evaluates to true.
Active	Enables or disables this rule. By default, new rules are active (true)
Make alias	If this check box is selected, and the rule evaluates to true, the rule makes an alias automatically from a pending value. If this check box is not selected, and the rule evaluates to true, the platform automatically changes the value for the named field in the record to the normal value, but does not create an alias.
Case sensitive	All pending values for the named field must match the case of the same value in a rule to be normalized.
Rule	Use the ServiceNow rule builder to construct the rules for normalizing fields. Rules automate the transformation of large numbers of pending field values into aliases.

4. Create a set of rules to incorporate all the possible variants of the pending values.

   For this example, the rules might be:

   - **Intel Xeon**: CPU type matches pattern *Intel*Xeon*. This rule normalizes all variants in which **Intel** precedes **Xeon**, including **Intel Xeon**, **Intel(R) Xeon(TM) CPU 2.80GHz**, and **Intel(R) Xeon(TM) CPU 3.00GHz**.
   - **Xeon**: CPU type contains Xeon
   - **L3350**: CPU type contains L3350
   - **E3350**: CPU type contains E3350

   For every rule that is created, the platform generates a **Rule applier** data job. In the testing mode, the **Start** controls are not available, and the job cannot be run until the mode is changed to **Active**.

5. Test all the rules before making the normalization record **Active**.

   Testing Rules

   > **Note:**
   > Users must have the normalization_tester role to create test records.

   Field normalization records are created in the **Test** mode by default, enabling administrators to test normalization rules thoroughly before applying them to the existing records in the database. In the testing mode, the **Start** controls are not available for the **Rule applier** data job. The job can be run only when testing is complete and the **Mode** has been changed to **Active**. In the testing mode, only records that have been created or updated by a user with the normalization_tester role are normalized. The normalizer and normalization_tester roles can be combined for a single user or granted separately.

Coalesce records on a normal value

Coalescence enables an administrator to redirect references to multiple records containing variants of the same field value to point to a single record, based on a normal value.

Before you begin
Role required: admin

About this task
An example of this is the Company table that might have multiple variants of a company name, such as Hewlett-Packard, Hewlett-Packard, Inc., Hewlett-Packard Incorporated, HP, and so on. Potentially, thousands of records might reference each of these duplicate company records. Using the variants of the Hewlett-Packard name as aliases, coalescence unifies all these references into a single record that normalizes the **Name** field in the Company record to a normal value such as **HP**.

Note:
Coalescing normal values changes the record values permanently. If a rollback is performed, records will be returned to the table, but the normalized values will not be rolled back to the original variants.

When the references are fixed, all table fields directly corresponding to sys_dictionary get fixed. The secondary references are not fixed. If you have filter conditions that has the old, pre coelesced company names, they aren't fixed either.

Procedure

1. Navigate to **All** > **Field Normalization** > **Configuration** > **Normalizations**
2. Select a normalization record.
3. Enable **Coalesce each normal**.
   The system adds the **Coalesce to** field to the Normal Value form.
4. Create one or more normal value records for this normalization record.
   Create related aliases and rules as needed.
5. For each normal value record, set **Coalesce to** by selecting the record that contains the normal value.

   Example
   For example, suppose your Company table contains several variations of the name ServiceNow. When you create the Normal Value record, you select the Company record for ServiceNow, Inc.. During normalization, the system updates any references to variant records to instead refer to the normal record.
   The system updates any references to records that match aliases and rules to instead point to the normal record. The system also deletes the duplicate records from the table.
6. Start all the **Alias application** data jobs to replace the aliases with the normal value in existing records in the database.
   The system starts the **Coalesce to normal** data jobs for each alias.

Transforms

Transforms update field values by applying a transformation definition to the field contents such as replace text, change case, or round value.

Transform definitions

Transform definitions define the transformation actions available for a given field type. Administrators select a definition when they transform a field, and then provide the definition with the specific parameters that are applied to the transformation. For example, a definition can round up an integer or insert a value at a defined position in a string.

The system provides a number of definitions that are designed to meet the needs of most organizations, but administrators can create new definitions as needed. Transform definitions can be associated with existing Transform Categories or to new categories.

Transform categories

Transform categories are used to group the transform definitions together appropriately to present to users when creating new field transforms. You can create new categories for existing definitions or change the default associations of categories to definitions.

Transform categories are associated with normalization field types to present the correct definition option list for the field being transformed. The base system provides two transform categories: Text and Numeric.

Text

Definitions in the Text category transform string type field values. Included in the Text category are:

- Left
- Right
- Constant
- Trim
- Prefix
- Suffix
- Change Case
- Delete
- Insert
- Substring
- Replace

Numeric

Definitions in the Numeric category arithmetically manipulate integer type field values. Included in the Numeric category are:

- Limit
- Round
- Constant

Transformation testing mode

Transformation is a powerful tool that can cause data issues if used incautiously. Test all transformations before committing the changes to the CMDB.

All transformation records open in **Test** mode by default, which blocks administrators from running manual data jobs that apply changes to existing data. Only users with the normalization_tester role can transform data when a record is in the **Test** mode. This is limited to new records opened by the tester or records in which the transformed field is updated by the tester.

To apply transformations to the CMDB after testing, change the mode to **Active** in the transformation record and run the appropriate data job. For information on testing transforms, see Test a transform.

Transform a field

Transform the contents of field using a set of rules and conditions.

Before you begin

Role required: admin or normalizer

Procedure

1. Create a transformation record.
2. Create one or more related transform records.
3. Test the transform.
4. Runs data jobs.

What to do next

If you want to also show what the original (raw) input value was prior to transformation, create a raw field to store this value.

Create a transformation record

Creating a transformation record is the first step in transforming a field.

Procedure

1. Activate the Field Normalization plugin.
2. Navigate to **Field Normalization** > **Configurations** > **Transformations**.
3. Click **New**.
4. Create a transformation record.

Transformation record fields

Field	Description
Name	Unique name for this transformation record. This value is for reference only and is not used in any processing.
Table	Select the ServiceNow table containing the field being transformed. It is important to understand the table hierarchy when setting up a field transform. For example, if you configure transformation for a field in the Computer [cmdb_ci_computer] table, that field will be transformed for all workstation machines, Windows servers, Linux servers, and UNIX servers.
Field	Select the field to transform. The list presented contains only those field types (integer and string) from the table selected that can be transformed.

Field	Description
	Note: The sys_user record that initiates the transform process must have its date format set to the default format of "yyyy-MM-dd." Any other date format causes an error during transformation. This problem is only specific to transforming **TO TARGET** fields of type Date/Time. This problem is not an issue if the target field type is of type String or if the field mapping for the date field is changed to the same date format as the transformation process.
Mode	The three available modes are **Off**, **Test**, and **Active**. All transformation records are created in the test mode by default. Do not change the mode until you have thoroughly tested the transformation. When testing is complete, change the mode to **Active**. To disable this transformation, switch the mode to **Off**.
Normalize query	Select this check box to apply the field value transformed by this record to all queries involving this field. Queries issued with the raw (original) field value will be edited to use the transformation value.
Raw field	Select the field to use to display the original input (non-normalized) values on a form in which a field value has been normalized. For the selection to appear in the drop-down list, add a custom field to the form for the table selected. For instructions on adding a field for raw data, see Create a raw field.

5. Click **Submit**.

   The **Transforms** and **Data Jobs** Related Lists appear on the form.

Create one or more related transform records

Each related transform record performs a specific transformation type such as adding characters to the beginning of the value or replacing one string for another. You may need to create multiple related transform records to generate a preferred output field value.

Procedure

1. In the Transformation record, select the **Transforms** Related List.
2. Click **New**.

   A selection list of transform types appears, displaying only those transformations appropriate for the field type selected.

Transform types

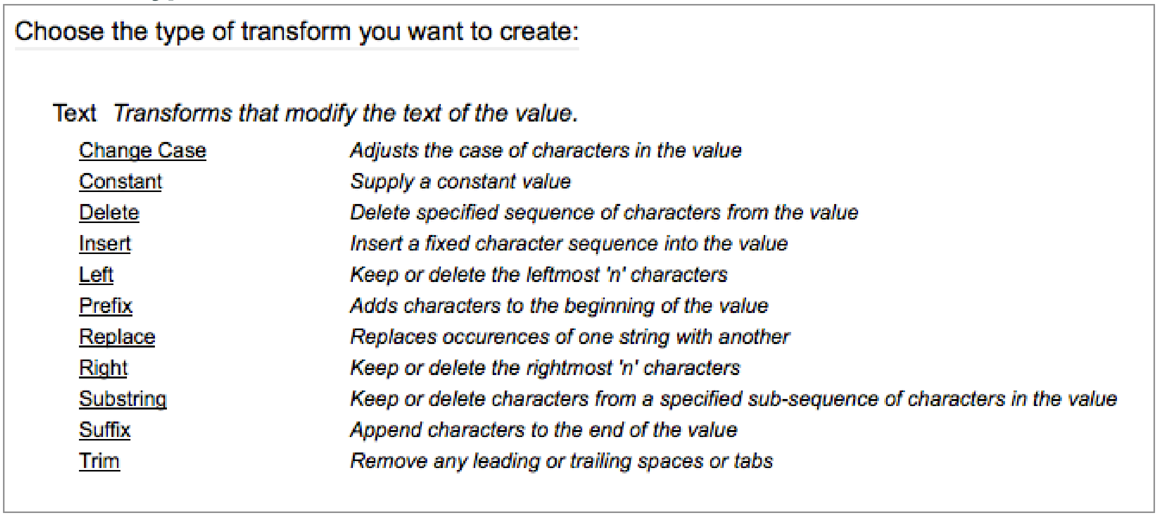

3. Select a transform type and provide the appropriate parameters.
4. Select an **Order** number for this transform.

   Note:
   The conditions for the transforms are executed according to the order numbers assigned.

5. Select the **Final** check box to stop processing with this transform if the condition evaluates to true.
6. Select the **Case sensitive** check box to force case sensitivity in the condition statement.

   The following transform example replaces the *INC* at the beginning of an incident number with the string *ENG* if the assignment group is *ITSM Engineering*.

Transformation record

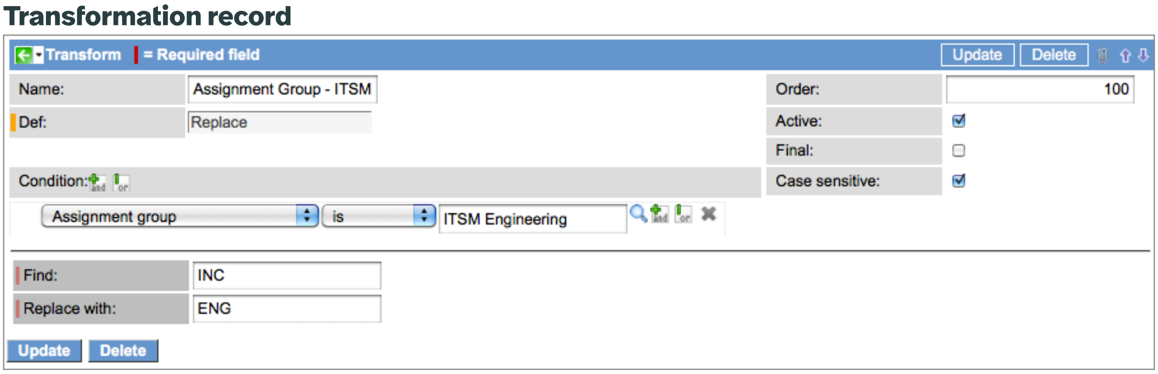

7. Click **Submit**.

   The new Transform appears in the Related List of the Transformation record.

   When the Transform is created, a Transformation application data job is also created. This data job applies this transform to appropriate records in the entire database and should not be run until testing is complete.

8. **Optional:** Repeat steps 2 through 8 until the output value meets your desired criteria.

Test a transform

Verify the transform changes the field value as desired before applying them to existing records in the database.

About this task

Note:
Users must have the normalization_tester role to create test records.

New transformation records open in the **Test** mode by default, enabling administrators to test transforms thoroughly before applying them to the existing records in the database. In the test mode, the **Start** controls are not available for the **Transform application** data job. There are two methods, listed below, for testing transforms before committing the transformations to existing data.

Procedure

- Manually create or update test records.

  In the test mode, only records that have been created or updated by a user with the normalization_tester role are transformed. Grant the normalizer and normalization_tester roles to the same user or grant them to separate users.

- Use the Test transforms utility to enter a raw value and see the resulting transformed value.

  This feature enables a normalization tester to transform field values on the fly without opening or updating records. This utility tests all the transforms configured for this field.

1. Open a Transformation record.
2. Click the **Test transforms** Related Link.

   A dialog box appears for testing field values.
3. Enter a value to transform in the **Raw data** field.

   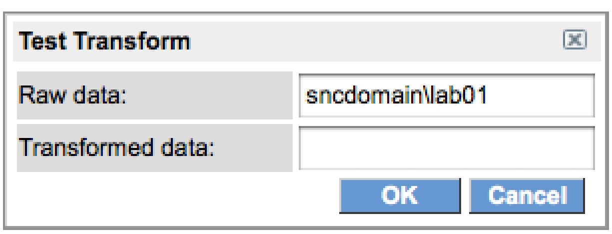

4. Click **OK**.

   The platform transforms the raw value in the **Transformed data** field.

   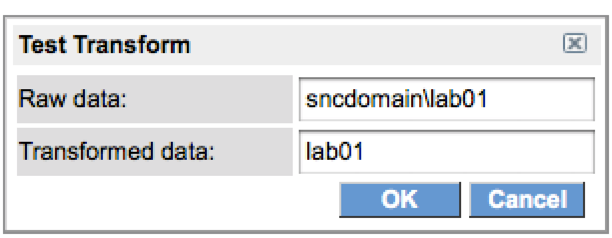

5. Enter new raw data to test other transforms.

6. Click **Cancel** to end the test.
7. When testing is complete, change the **Mode** to **Active** and run the data job.

Default Transform Definitions

The system offers default transform definitions for fields containing text, text numeric, and numeric values.

Transform Definitions

Transform Type	Category	Description	Parameters
Change case	Text	Changes the case of the characters in the field value.	*Mode*: Select one of the following modes: • *Upper*: Converts the value to all upper case characters • *Lower*: Converts the value to all lower case characters • *Proper*: Converts the value to title case, with the first character in each string in upper case, and the remaining characters of the string in lower case. • *Formal*: Converts the value to a string in which only the first letter of the first word is in upper case.
Constant	Text Numeric	Converts the value in this field to a constant.	*Constant*: The constant with which to replace the value in this field.
Delete	Text	Delete a specified sequence of characters from a field value.	• *Starting position*: Specifies the first character in a sequence of characters to delete from a string. See the discussion of position modes at the beginning of this section for details. • *Ending position*: Specifies the final character in a sequence of characters to delete from a string. See the discussion of position modes at the beginning of this section for details.
Insert	Text	Insert a fixed character sequence into a field value.	• *Position*: The character position at which to insert the new value. See the discussion of

Transform Definitions (continued)

Transform Type	Category	Description	Parameters
			position modes at the beginning of this section for details. • *Insert*: The value to insert into this field.
Left	Text	Deletes or keeps a specified number of characters from the left side of this field value.	• *Position*: Specifies the number of characters to keep or delete from the left side of the value. See the discussion of position modes at the beginning of this section for details. • *Mode*: Select the mode for this transform: *Keep* or *Delete*.
Prefix	Text	Adds characters to the beginning of a field value.	*Prefix*: Defines the characters to add to the beginning of the transformed field value.
Replace	Text	Replaces occurrences of one string with another string. The special characters backslash (\) and dollar sign ($) in the replacement string can cause the transform to be different than if the replacement string were being treated as a literal replacement string. Use a regular expression to replace a string or parts of a string.	• *Find*: Enter the string or regular expression to replace. • *Replace with*: Enter the replacement string.
Right	Text	Retains or deletes a specified number of characters from the right	• *Position*: The number of characters to delete or keep from the right side of this transformed field. See the discussion of

Transform Definitions (continued)

Transform Type	Category	Description	Parameters
		side of a field value.	position modes at the beginning of this section for details. • *Mode*: Select the mode for this transform: *Keep* or *Delete*.
Round	Numeric	Rounds integers to a configured rounding interval using specific criteria. The interval must be appropriate to the value being transformed, such as an interval of 12 for a value expressed in dozens or 0.01 for decimal values expressed in hundredths.	• *Interval*: Select the rounding interval that is appropriate to the units of the field value. For example. an interval of 256 is appropriate for expressing RAM values in megabytes, but does not work for Disk space expressed in gigabytes. The rounding interval for the examples below is 1 • *Mode*: Criteria for applying the rounding interval. ○ *Half up*: Always round up a value that is exactly half way between two intervals. For example, 3.5 is always rounded up to 4, and -3.5 is always rounded up to -3. ○ *Half down*: Always round down a value that is exactly half way between two intervals. For example, 3.5 is always rounded down to 3, and -3.5 is always rounded down to -4. ○ *Half away from zero*: Always round an integer that is half way between the specified interval away from zero. For example, 3.5 is always rounded to 4, and -3.5 is always rounded to -4. ○ *Half toward zero*: Always round an integer that is half way between the specified interval toward zero. For example, 3.5 is always rounded to 3, and -3.5 is always rounded to -3. ○ *Half to even*: Always round an integer that is half way between the specified interval to the nearest interval whose least significant digit is even. For example, 3.5 is always rounded

Transform Definitions (continued)

Transform Type	Category	Description	Parameters
			to 4, and 4.5 is always rounded to 4. ◦ *Half to odd*: Always round an integer that is half way between the specified interval to the nearest interval whose least significant digit is odd. For example, 3.5 is always rounded to 3, and 4.5 is always rounded to 5. ◦ *Up*: Always round an integer up by the specified rounding interval. For example, 3.4 is always rounded to 4 by a rounding interval of 1.0. ◦ *Down*: Always round an integer down by the specified rounding interval. For example, 4.6 is always rounded to 4 by a rounding interval of 1.0. ◦ *Away from zero*: Always round an integer away from zero by the specified rounding interval. For example, 3.3 is always rounded to 4, and -3.3 is always rounded to -4 by a rounding interval of 1.0. ◦ *Toward zero*: Always round an integer toward zero by the specified rounding interval. For example, 3.3 is always rounded to 3, and -3.3 is always rounded to -3 by a rounding interval of 1.0.
Substring	Text	Keep or delete characters from a specified sub-sequence of characters in the field value.	• *Starting position*: Specifies the first character in a sub-sequence of characters within the value. See the discussion of position modes at the beginning of this section for details. • *Ending position*: Specifies the final character in a sub-sequence of characters within the value. See the discussion of position modes at the beginning of this section for details.

Transform Definitions (continued)

Transform Type	Category	Description	Parameters
			• *Mode*: Select whether to `Delete` the sub-sequence selected or `Keep` only those characters defined.
Suffix	Text	Appends characters to the end of a field value.	*Suffix*: Defines the suffix to add to the end of the field value.
Trim	Text Numeric	Removes blank spaces from the field value.	No parameters

Create a transform definition

The following example describes the procedure for creating a new transform definition. In this example, we create a definition that transforms a number field to an odd or even integer. The transform category is **Numeric** and the normalization field type is **Integer**.

Procedure

1. Navigate to **All** > **Field Normalization** > **Administration** > **Transform Definitions**.
2. Click **New** in the record list.
3. Enter a name for this definition.

   Note:
   In this example, we enter `Odd/Even`.

4. Enter a brief description of the action, such as, `Transforms an integer to an odd or even value`.

   This information appears in the definition choice list when a user selects a new transform.

5. Right-click in the header bar and select **Save** in the context menu.

   Two Related Lists appear in the form.

   - Transform Categories: Click **Edit** and select **Numeric** as the category to which this definition belongs. Currently, field transformation supports two categories: **Numeric** and **Text**. The **Integer** normalization field type is already associated with this category.
   - Transform Variables: Define any variables required by this transform definition to perform an action on a field value. Variables are not necessary if a script can perform the action alone.

Create a transform variable for a transform definition

Transform variables enable an administrator to apply the same definition to different fields in different ways.

About this task

Transform variables contain values used by a script to perform a field transformation. Scripts and variables can be created in either order, but the script must use the transform variables. Transform variables are populated with values when a user configures a transform type.

Procedure

1. In the Transform Definition record, click **New** in the **Transform Variables** Related List.
2. Complete the form.
   Important considerations for completing a form:
   - The Column name is an entry in the fn_transform_var table for this variable. This becomes the variable in the script, in the form of variables.<variable name>. For example, `odd_even`.
   - The value in the Label field appears as the variable field label in the Transform form. For example, `Odd/Even`.
   - The field Type defines the field type of the variable value. Because the values for the variables used are **even** and **odd**, this is a type of string.
   - The Order of the variables controls the order in which they are displayed in lists and records.
   - This variable has a choice list with two options: **Even** and **Odd**. Select **Dropdown without - None** as the format for the list in the Choice field and define a Default value of **even** when the list is displayed.
   - Create a Hint that becomes a tooltip for the variable in the Transform record.
3. Right-click in the header bar and select **Save** from the context menu.

   The **Variables Choice List** Related List appears.
4. Click **New** in the Variables Choice List and define the list options.
5. Create records for **Even** and **Odd**.

   Note:
   The **Element** value is the same as the **Column** name in both selections for the choice list.
6. Save the choice list variables and return to the transform definition form to create the script.

Create a script for a transform definition

Create the script at any time during the configuration of a definition.

About this task

The script can perform a transform action without using a variable, but the action of the definition will be the same for all fields. Variables create more flexibility for the definition, enabling an administrator to use the same definition in different ways in different places. If a variable is defined, the script must reference the variable using the correct format.

There are three arguments in the script:

- Variables: Contains the variables using the format variables.<variable name>.
- Value: Contains the un-transformed value
- Parameters: Special objects that set debug messages.

All position parameters (such as Starting position and Ending position) have three modes that apply to all the transform types that use this variable.

Position parameters

Positive positions	If the position is expressed as a positive integer, the platform calculates the starting position beginning from the left side of the field value. For example, in the string ABCDE, a position of 3 places the starting point of the action after C.

Position parameters (continued)

Negative positions	If the position is expressed as a negative integer, the platform calculates the position beginning from the right side of the field value. For example, in the string ABCDE, a position of -3 places the starting point of the action before C.
Regex	If the position value starts with /regex/, everything after that is a regular expression that is used to calculate the starting position. For example, in the string ABCDE, a position of /regex/B.*D places the starting point of the action after C (B and all the characters between B and D).

Procedure

1. Open the **Odd/Even** record in the Transform Definitions module.
2. Enter the following script to pass values with the odd_even variable.

```
function(variables, value, parameters) {
 var odd = ('odd' == variables.odd_even);
 var val = value - 0;
 var val_odd = ((val & 1) == 1);
 if (odd != val_odd)
  val++;
 return '' + val;
 }
```

The script references the variable in the form variables.odd_even.

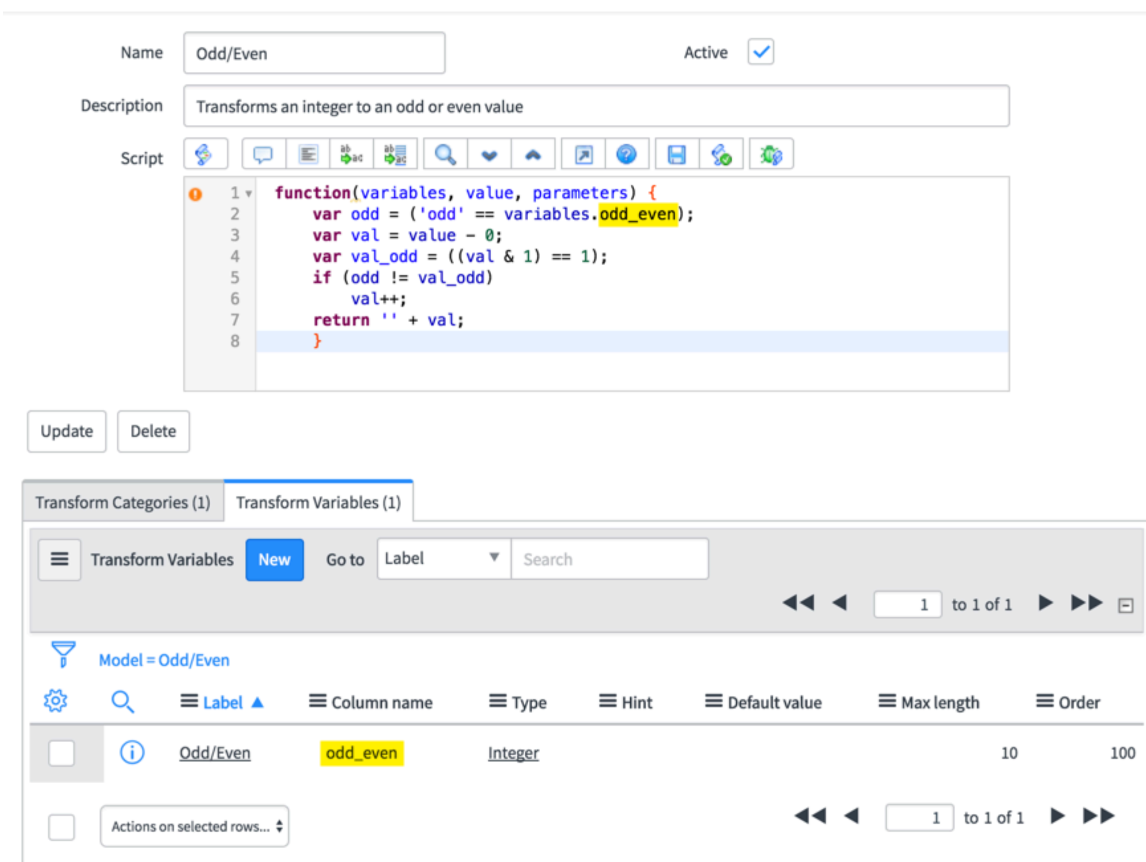

3. Update the record to complete the configuration.

   The Odd/Even transform definition is now ready to use in a field transformation.

Create a transform category

Create a transform category to group the transform definitions together.

Procedure

1. Navigate to **All** > **Field Normalization** > **Administration** > **Transform Categories**.
2. Click **New** in the **Transform Categories Related List**.
3. Enter the **Name** of this category and a description.
4. Select an **Order** for this category and save the record.
   The order determines the display order of categories in lists and forms. Two Related lists appear:

   - **Field Types**: Click **Edit** to select an existing field type for this category or click **New** to create a new field type. The normalization field types provided are:

     - Decimal
     - Float
     - Integer
     - Numeric
     - String
     - URL

   - **Transform Definitions**: Click **Edit** to select the transform definitions that are included in this category.

Regular expressions and patterns in field normalization rules

Field Transformation definitions support the use of regular expressions (referred to in the platform as **regex**) and pattern matching for determining the position of characters in a string.

After identifying the target characters, field transformation can replace or delete the identified characters or insert other characters at that position.

Regex

Regular expressions can be used in transform parameters and in condition statements to determine which characters in a field value are transformed.

Regular expressions used as parameters to locate characters in transformed field values must begin with `/regex/`. Everything after that is a regular expression that is used to calculate character position.

Example

The computer names in an organization's Windows network are expressed as domain\machine name, such as **development\devlab01**. The network administrator wants to simplify these names by removing the domain name and backslash. He creates a transformation record for the Computer [cmdb_ci_computer] table and selects the **Name** field to transform.

Transformation regex example

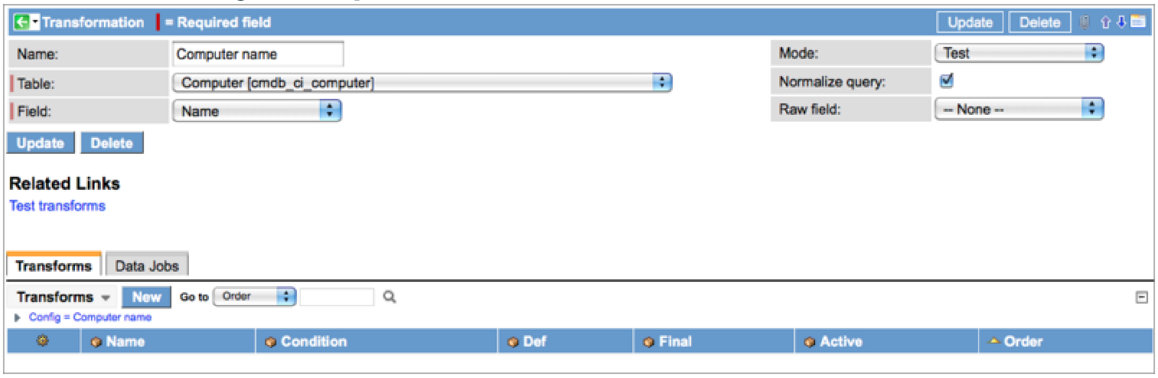

The network contains several domains, and each domain contains numerous computers. The only character common to each name is the backslash. To delete the domain name, the administrator decides to use a regular expression to replace the entire raw value in the field with the characters that appear after the backslash (the actual machine name). He creates a new Transform using **Replace** as the Transform Type and enters the following values:

- Find:**/regex/.*\\(.*)**
- Replace with:**$1**

Transformation regex example

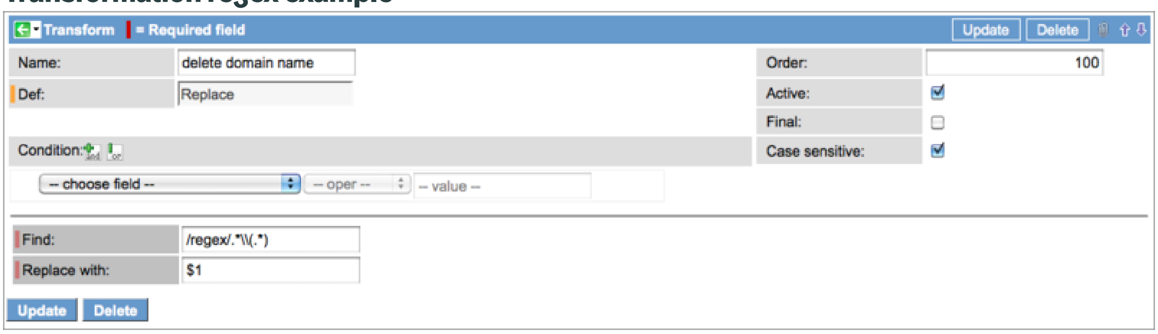

The regular expression **.*\\(.*)** represents the entire raw value in the **Name** field - in this example **development\devlab01**. The first part of the expression, **.***, represents everything before the backslash (the **development** domain name). The backslash by itself is the escape character in regular expressions and requires special syntax to retain its function in the computer name. The administrator must escape it by using another backslash (\\ means \). The part of the expression after the backslash, **(.*)**, represents the computer name (**devlab01**) and is grouped within parentheses for reference. The value in the **Replace with** field, **$1**, references this group and replaces the entire raw value of the field with the contents of the group, **devlab01**.

The administrator clicks **Test transforms** in the transformation record and enters `development\devlab01` in the **Raw data** field. He then clicks **OK** to apply the transform to the test value. The transform replaces **development\devlab01** with **devlab01**.

Transformation Regex 3

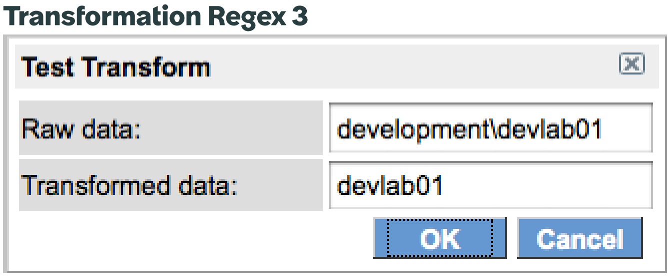

When the transforms for this field are tested successfully, the administrator changes the **Mode** in the transformation record to **Active** and runs the Transformation application data job to apply this transformation to existing records in the database.

Pattern matching

Pattern matching in Field Normalization uses special characters differently from regular expressions to create patterns that the platform recognizes when transforming field values.

Pattern matching can be used only in condition statements. When using pattern matching characters in a condition statement, make sure to select the **matches pattern** operator.

Use the following special characters to create patterns for searches.

- The asterisk in a search string (*) matches any number (including zero) of any character.
- The question mark (?) in a search string matches one of any character.
- Everything else in a search string matches itself.

Examples

- **the story** matches **the story** but not **that story**.
- ***story** matches **the story** and **that story**, but not **that story is the best**.
- **st?ry** matches **story** and **stxry**, but not **my story** or **stairy**.
- ***b?gus*** matches **bogus**, **my bogus story**, and **His bagus machine**, but not **my bgus story** or **my baigus story**.

Domain separation and Field Normalization

Domain separation is unsupported in Field Normalization. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

Related topics

Domain separation for service providers

Data policy

Data policies enable you to enforce data consistency by setting mandatory and read-only states for fields.

Data policies are similar to UI policies, but UI policies only apply to data entered on a form through the standard browser. Data policies can apply rules to all data entered into the system, including data brought in through import sets or web services and data entered through the mobile UI.

For example, suppose that you are configuring a web service that allows users from outside the platform to update problems on the ServiceNow instance. Since these problems are not updated through the instance UI, they are not subject to the UI policies on the problem form. To ensure that the **Close notes** field is completed before a problem is marked **Closed/Resolved**, you can create a data policy that applies to server-side imports. Data that does not comply with this data

policy produces an error. You can also apply the policy on the browser by selecting the **Use as UI Policy on client** check box in the data policy record.

Since UI policies can also manage the visibility of fields on a form, you may want to augment UI policies with data policies rather than replace them.

By default, data policies are applied to all GlideRecord operations including those used in Scripted REST APIs, and the REST Table API. You can opt out of applying the data policy to:

- Target records of SOAP web services
- Import sets
- Client-side UI policies

The admin role is required to edit data policies.

Note:
Defining a data policy enforces the policy when a record is submitted from the UI. This behavior cannot be changed.

Installed with data policy

Data policy includes several components that are installed on the instance.

Tables

Table label	Table name
Data Policy	sys_data_policy2
Data Policy Rule	sys_data_policy_rule

Data Policy Module

The **System Policy** > **Data Policies** module displays a list of all data polices and where they apply.

Roles

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

Data Policy Administrator [data_policy_admin]

Access and maintain data policies to specific users. Grant this role access to the System Policy application

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Create a data policy

You can create a new data policy to define data rules for a table.

About this task

Create data policies to enforce consistency. You can create data policies only for tables and database views that are in the same scope as the data policy and for other tables that have at least one field in the same scope as the data policy. For tables that are in a different scope from the data policy record, you can create data policy rules only for fields in the same scope as the data policy and you cannot make a field mandatory.

Procedure

1. Navigate to **Data Policies** by completing one of the following actions.
   - Navigate to **System Policy** > **Rules** > **Data Policies**.
   - From any form header, right-click the header bar and select **Configure** > **Data Policies**.
   - In List v2, open any column context menu and select **Configure** > **Data Policies**.
   - In List v3, open the list title menu and select **Configure** > **Data Policies**.
2. Click **New**.
3. Select any options for the data policy.
4. Create the condition that must exist for the platform to apply this policy.
   For example, your conditions might include **[Problem state] [is] [Closed/Resolved]**
5. Right-click the header and select **Save**.
   The **Data Policy Rules** related list appears.
6. Click **New** in the related list and create the record that identifies the field and the policy to apply.

   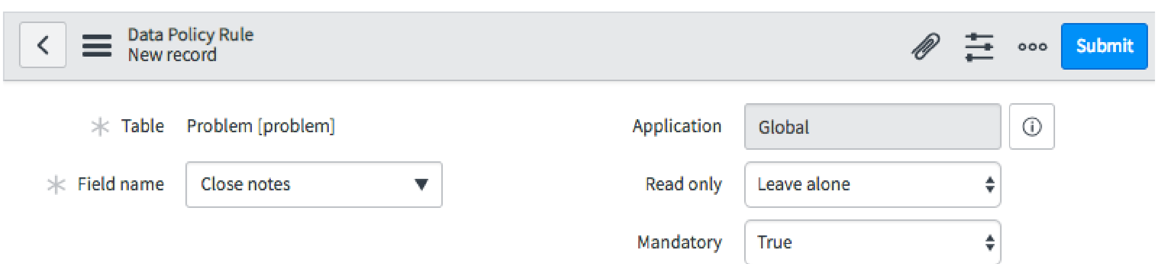

   It is possible to have multiple rules on a single field, but it is not recommended.
7. Click **Submit**.

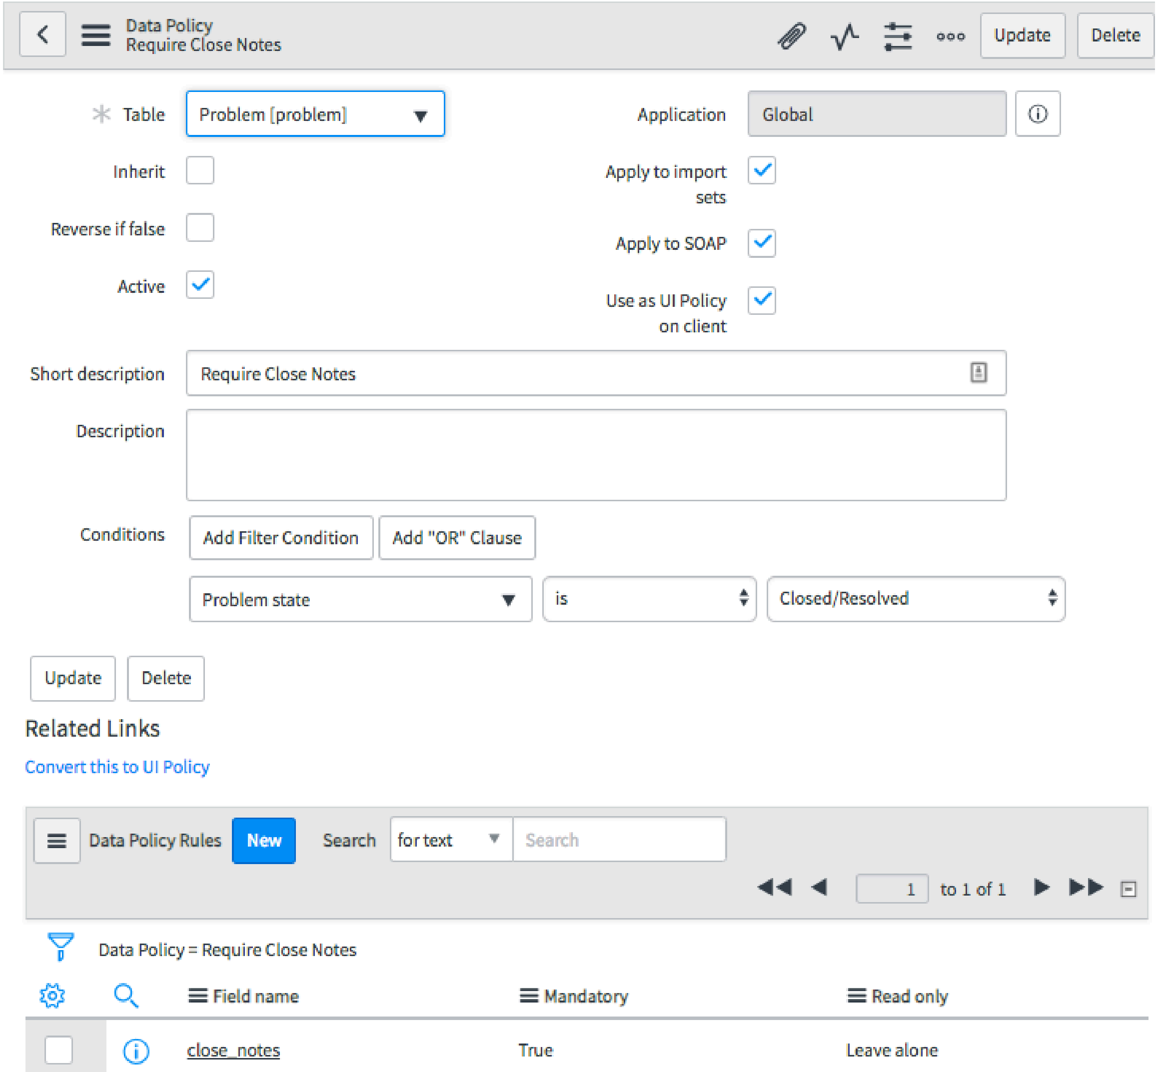

8. Optional: Add more rules by repeating steps 6 and 7.

Data policy fields

These fields appear on the Data Policy form and related forms.

Data policy fields

Field	Description
Table	The table to which this policy applies. **Note:** The list shows only tables and database views that are in the same scope as the data policy.
Application	Application that contains this data policy.
Inherit	If selected, applies this data policy to tables that extend the specified table. For example, incident, problem, and change tables all extend the task table, therefore selecting Inherit on a data policy defined for task would apply the data policy to them as well.

Data policy fields (continued)

Field	Description
Reverse if false	If selected, the data policy action is reversed when the conditions evaluate to false. For example, when the conditions are true, then actions are taken and when they change to false, the actions are reversed.
Active	If selected, the data policy is used.
Short description	A short description that identifies the policy.
Description	A detailed description of the policy.
Apply to import sets	If selected, the data policy applies to data brought into the system from import sets. This option also applies to web service import sets.
Apply to SOAP	If selected, the data policy applies to data brought into the system from a SOAP web service. Scripted SOAP web services are not affected. This field does not affect data policy interaction with REST web services.
Use as UI Policy on client	If selected, enforces the data policy on the UI using the UI policy engine.

Data policy rule fields

Field	Description
Table	The table on which the data policy action applies.
Field name	The field from the specified table to which the data policy will apply.
Read Only	How the data policy affects the read only state of the field. Choices are: • Leave alone • True • False
Mandatory	How the data policy affects the mandatory state of the field. Choices are: • Leave alone • True • False **Note:** For tables that are in a different scope than the data policy record, you cannot make a field mandatory.

Convert a UI policy to a data policy

To make a UI policy the default setting, convert the UI policy to a data policy.

Before you begin

Role required: ui_policy_admin

About this task

You can also apply a UI policy to import sets or to data imported by SOAP web services when you convert it to a data policy. Converting a UI policy to a data policy deactivates the UI policy. To retain the policy in the UI, ensure that the **Use as UI Policy on client** check box is selected on the data policy record.

For a UI policy to be eligible for conversion to a data policy, the following conditions must be met on the UI Policy form.

- The **Run scripts** check box must be cleared.
- The **Global** check box must be selected.
- None of the UI policy actions can have **Visible** set to **True** or set to **False**. **Visible** must be set to **Leave Alone**.

Procedure

1. Navigate to **All** > **System UI** > **UI Policies**.
2. Open an existing UI policy.
3. Under **Related Links**, click **Convert this to Data Policy**.
   A new data policy record is created.
4. Edit the fields on the data policy record as necessary.

Related topics

Data policy

Convert a data policy to a UI policy

Convert a data policy to a UI policy if a data policy already exists but needs to apply only to records created or updated in the browser.

Before you begin

Role required: admin

About this task

Converting deactivates the data policy. The new UI policy is applied only at the UI layer only and not to import sets or data imported from SOAP web services.

Note:
An alternative to converting from a data policy to a UI policy is to select the **Use as UI Policy on client** checkbox on the data policy record. This field extends the data policy to the UI. The main difference between converting and using the **Use as UI Policy on client** checkbox is that converting provides the Visible field on the UI policy record. Use the **Visible** field to select how the UI policy affects the visible state of the field.

Procedure

1. Navigate to **All** > **System Policy** > **Rules** > **Data Policies** and click an existing data policy.
2. Under **Related Links**, click **Convert this to UI Policy**.

   A new UI policy record appears.
3. Edit the fields on the UI policy record as necessary.
   For details about the fields, see Create a UI Policy.

Data policy debugging

Debug messages can help administrators identify and resolve data policy problems.

Debug messages can help you identify and resolve data policy problems. To view data policy debugging messages at the bottom of the screen, navigate to **System Diagnostics** > **Session Debug** > **Debug Data Policies**.

In the example, a data policy is in place to prevent the short description on an incident from being changed when the incident state is set to Open. A user edited the short description while the incident was open and tried to save the changes, but the data policy was enforced.

Data policy debug messages

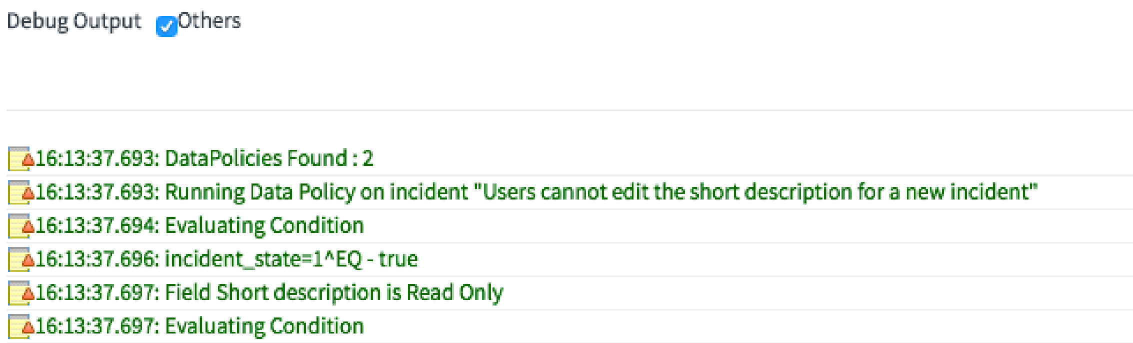

Data lookup and record matching support

The data lookup and record matching feature enables administrators to define rules that automatically set one or more field values when certain conditions are met.

Data lookup rules allow administrators to specify the conditions and fields where they want data lookups to occur. For example, on Incident forms, there are priority lookup rules for the sample data that automatically set the incident **Priority** based on the incident **Impact** and **Urgency** values.

> **Note:**
> Activating the Data Lookup and Record Matching Support plugin replaces the `calculatePriority` business rule with a priority data lookup definition, but does not transfer any custom logic. If you manually activate the plugin, you must recreate any custom business logic that uses the priority lookup rules.

The following codes will be removed by manually activating the `com.glide.data_lookup` plugin:

- sys_script_18105ab0c61122750127a49c9055d29f
- sys_script_client_2bbac34da9fe1561000d89c04779a30c
- sys_script_client_2bc0bf4fa9fe15610019c5c8a8461647
- sys_script_client_ac0081210f030000b12e6903cfe012ac
- ys_script_d1b7d4af4655e7c2001be4c18b99215a
- sys_script_include_9e1904840a0a0b7d00d5ee8d0e2d89ef

Create custom data lookups

Creating custom data lookups involves adding a lookup value, creating the definition, and activating data lookup.

Create a custom data lookup table

Create a custom table to store lookup data.

The custom table must extend the Data Lookup Matcher Rules `[dl_matcher]` table.
For example, this custom lookup table stores information about VIP callers and incident assignments.

Custom data lookup table

Field	Sample value
Label	VIP Caller Lookup
Table name	u_vip_caller_lookup
Extends base table	dl_matcher
Create new module	True
Add module to menu	System Policy

Add a data lookup value to the data lookup table

The columns of a data lookup table contain both matcher and setter field data.

About this task

Each data lookup is a query that searches for a row containing values that match the matcher fields. The data lookup then returns the value listed in the setter fields. For example, this Priority Data Lookup [dl_u_priority] table lists the combinations of impact and urgency (matcher fields) that produce a particular priority value (setter field).

Lookup table

Matcher fields		Setter field
Impact	Urgency	Priority
1 - High	1 - High	1 - Critical
1 - High	2 - Medium	2 - High
1 - High	3 - Low	3 - Moderate
2 - Medium	1 - High	2 - High
2 - Medium	2 - Medium	3 - Moderate
2 - Medium	3 - Low	4 - Low
3 - Low	1 - High	3 - Moderate
3 - Low	2 - Medium	4 - Low
3 - Low	3 - Low	5 - Planning

Procedure

1. In the navigation filter, enter the name of the new custom lookup table.
   For example, u_vip_caller_lookup.list.
2. Configure the list and create new fields.

   For this example, create the new fields in the following table.

New fields

Name	Type	Field length or Table to reference
Caller	Reference	User [sys_user]
Priority	Integer	
Assignment Group	Reference	Group [sys_user_group]

For more information, see List configuration.

3. From the table list, click **New** and enter appropriate matcher and setter field values.

   For example:

New matcher and setter field values

Matcher field	Setter fields	
Caller	Priority	Assignment Group
Beth Anglin	2	VIP Issues
Fred Luddy	1	VIP Issues

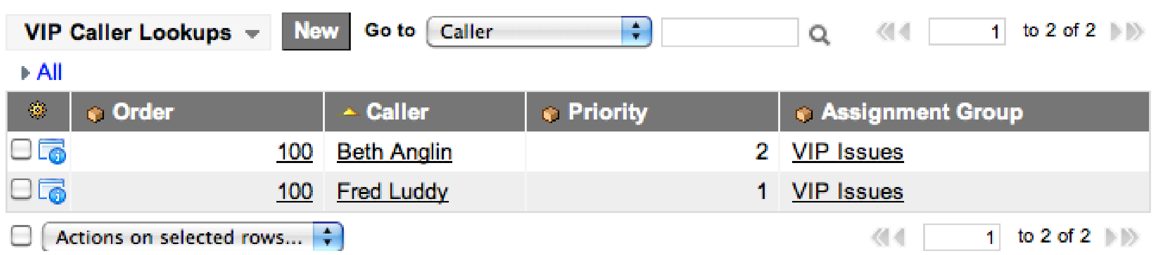

Note:
Each row in a data lookup table must be unique.

Create a data lookup definition record

Data lookup requires a definition record that specifies how to set one or more field values when specified conditions are met.

Procedure

1. Navigate to **All** > **System Policy** > **Rules** > **Data Lookup Definitions**.
2. Select **New**.
3. Select **Data Lookup Rule**.
4. Complete the Data Lookup Rule form using the following table.

Data lookup definitions fields

Field	Description
Name	Enter a unique name to identify the definition record.
Source Table	Select the table containing the fields you want to automatically update with lookup values. Data Lookup Definitions are not inherited by extension tables. For example,

Field	Description
	a Data Lookup Definition on the Task table cannot match values on the Incident incident table.
Matcher Table	Select the table containing the lookup values. This table should always start with a u_ prefix. This table can be any table within the current scope containing the lookup values. **Note:** The list shows only tables and database views that are in the same scope as the catalog data lookup definition.
Active	Select this check box to run this data lookup rule. Clear the check box to ignore this data lookup rule.
Run on form change	Select this check box to automatically look up values whenever a user or onChange client script changes a field value on a source table form. **Note:** This does not include changes automatically made by other data lookup rules, such as the Priority Lookup Rules.
Run on insert	Select this check box to automatically look up values whenever a user creates a new record.
Run on update	Select this check box to automatically look up values whenever a user saves or updates a record.

5. Right-click the form header and select **Save**.
6. From the **Matcher Field Definitions** related list, click **New**.
7. Complete the Matcher Field Definitions using the following table.

   A data lookup only occurs on fields with matcher field definitions. The data lookup uses the values of the source table fields to look up one or more values from the matcher table. Note that data lookup does not work with **Journal** type fields.

Matcher field definitions fields

Field	Description
Data Lookup	Displays the name of the parent data lookup definition record.
Source table field	Select the field from the source table that contains the data to match.
Matcher table field	Select the field from the matcher table that contains the data to match. This table can be any table within the current scope containing the lookup values.
Exact lookup match	Select this check box to require the matcher table to contain a matching row for every possible combination of values (including blank values). Clearing this check box means that any blank values in the matcher table match any value. For example, suppose the Priority field is blank in the matcher table. When this check box is selected, there is a match only when the Priority value is blank in the source

Field	Description
	table row. When this check box is cleared, the blank matcher field value matches any value in the source table field. **Note:** If the lookup does not require an exact match, matcher table rows containing blank values are treated as wild cards, matching all values.

8. Select **Submit**.
9. From the **Setter Field Definitions** related list, click **New**.
10. Complete the Setter Field Definitions form using the following table.

Setter field definitions fields

Field	Description
Data Lookup	Displays the name of the parent data lookup definitions record.
Source table field	Select the field from the source table that the data lookup updates.
Matcher table field	Select the field from the matcher table that provides the new value for the update.
Always replace	Select this check box to replace any existing value with a value from the data lookup. Clear this check box to ignore the update if the field has an existing value.

11. Select **Submit**.
12. Select **Update**.

For example, the following data lookup definition assigns incidents to the VIP Issues group based on the **Caller** field. In addition, the incidents are set to critical or high priority based on the caller.

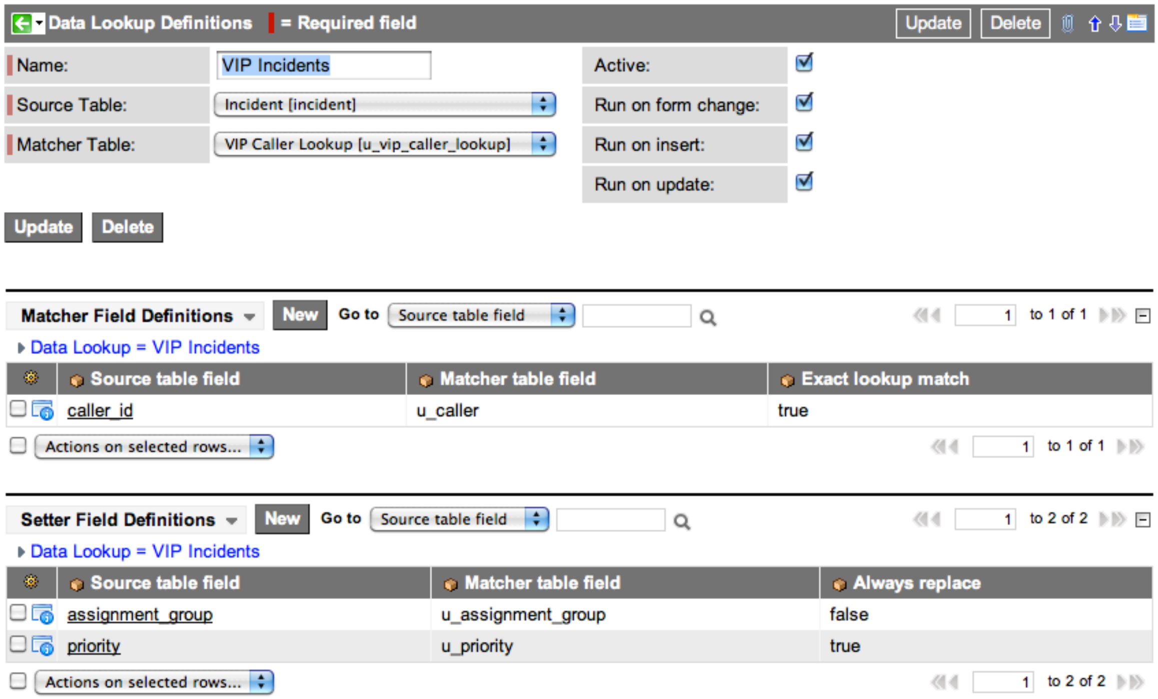

Create a data lookup module

You can create a module for data lookup so it appears in the instance application navigator.

Procedure

1. Navigate to **All** > **System Definition** > **Application Menus**, then select an application to add the module to.
2. In the Modules related list, click **New**.
3. Configure a module for the data lookup table you created using the following properties.

   Module properties

Property	Required value
Table	Select the data lookup table you created in Step 2. For example u_vip_caller_lookup.
Link type	List of Records

4. Click **Submit**.

Troubleshooting data lookup

If the custom data lookup definition rules are not behaving as expected, check for certain conditions.

- Verify that the data lookup definition is set to run on the appropriate events.
- Verify that the matcher field is not read-only. Since users cannot change read-only fields, user interactions cannot trigger an on form change event for read-only fields.
- Verify a client script is not changing a field value. Client scripts can trigger `Run on form change` events even on read-only fields.
- Verify that the data in the matcher table is correct.
- If the lookup requires an exact match, verify that there is a matcher table row for each possible combination (including blank values). The lookup fails if cannot find a matching value.
- Verify that you have not created a recursive rule, such as:

  If Field A = 1, then Field B =2. If Field B = 2, then Field A = 2

List administration

Administrators control the data presented to end users in a list, the controls that appear, and the behavior of the controls.

> **Note:**
> For information about using lists, see Lists.

List configuration

Users with the appropriate roles can configure various aspects of lists. Configuration changes apply to all users.

With list configuration, you can add, remove, and reorder list columns. You can configure calculations to appear under columns. You can also hide controls and define access conditions by role for existing list controls. Configuring lists is called personalizing lists in versions prior to the Fuji release.

> **Note:**
> Configuring a list in this way modifies the list for all users. To make changes to a list that are visible to you only, see Personal lists.

Configure the list layout

You can configure a list to choose which columns appear in a list, create list views, and create fields on the table.

Before you begin
Role required: personalize_list

About this task
Configure the list layout

Procedure

1. Navigate to the list you want to configure.
2. Open the **List controls** menu and select **View** > **(view name)**to select the view you want to

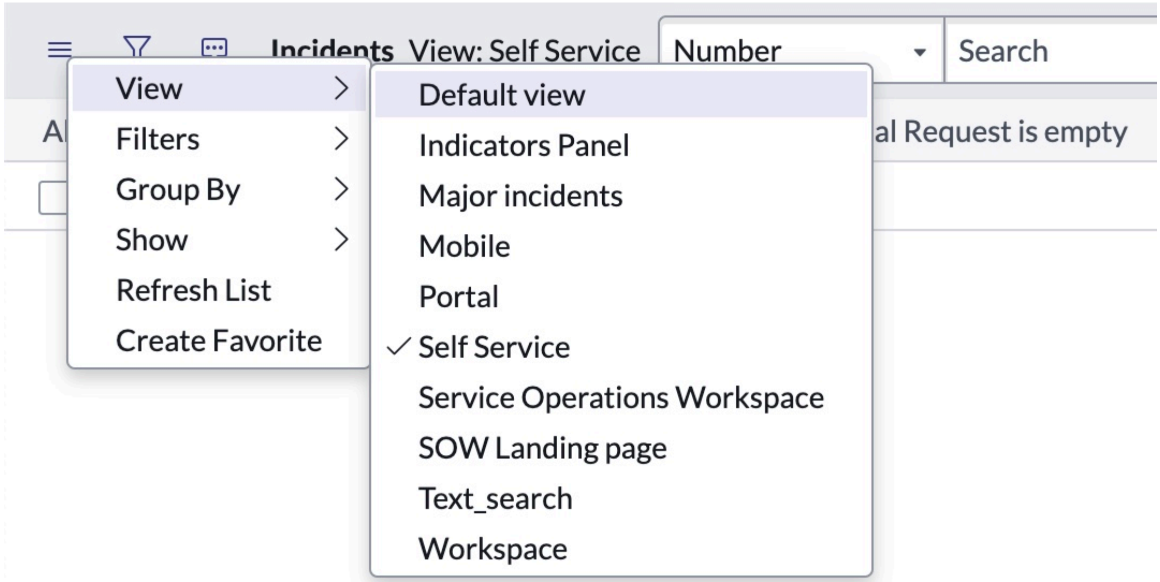

   configure.
3. Right-click any column heading and select **Configure** > **List Layout**.
4. Add, remove, or reorder fields to configure the columns displayed and their order of appearance.
   The first non-reference field automatically links to the form view of the record. For this reason, consider using the record number as the first column in the list layout.
5. Select **Save**.

   Note:

   - If you use Personal lists and then configure the list view differently, the changes do not appear until you reset your personal list to column defaults.
   - Inactive fields display in admin-configured lists until an admin updates the list layout to remove them.

Trouble?

If a reference field in a list displays as (empty) instead of the expected value, it is due to another field in the list containing an incorrect reference to either an orphaned record or a sys_id that does not exist in the referenced table. The solution is to remove the incorrect reference data from any reference fields in that table.

Add an extended field to a base table list

Administrators can configure a property that determines whether fields on extended tables can be added to a parent table list. For example, when this feature is enabled, you can view and filter on the **Caller** field, from the Incident table, on a Task table list.

Before you begin

Role required: admin

Note:
Enabling this property does not show the extended table fields in the personalize list.

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Locate the property **Allow base table lists (task, cmdb_ci, etc.) to include extended table fields (incident_state, os_version, etc.), and allow filtering on extended table fields** (`glide.ui.list.allow_extended_fields`).
3. Select the check box to enable extended fields on parent table lists, or clear the check box to disable the feature.
4. Select **Save**.

Configure list calculations

You can configure list calculations for columns, which calculate column totals, minimums, maximums, and averages.

Before you begin

Role required: personalize_list.

About this task

You can configure list calculations for multiple columns in the same list. List calculations apply only to the view of the list in which they are configured. All users can see configured list calculations.

Note:

- List calculations are also available in list reports.
- List calculations for a calculated field can only operate on the *current values stored in the database for each record and not the values as calculated fresh when the list is rendered.

Procedure

1. In a list of records, right-click the header of a numeric column you want to evaluate with list calculations.
2. Select **Configure** > **List Calculations**.
3. Select one or more of the following list calculation options:
   - Total value (not available for string, date, or date/time fields)
   - Minimum value
   - Maximum value
   - Average value (not available for string, date, or date/time fields)
4. Select **OK**.
   The calculations appear below the last record in the column. If the list is grouped, in addition to the overall calculations at the bottom of the list, group calculations appear below the last record in each group.

Omit record count in a list

Speed up how fast lists load by removing the calculation of the total number of records in a list.

Before you begin

Role required: workspace_admin

About this task

Calculating the total number of records included in a list that's filtered from large tables can take a long time. Platform displays the total number of records in the list.

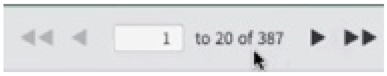

To load lists faster, you can remove the page calculation for all or only specified filtered lists. When you remove the calculation, the total record count is not displayed and the double arrow icon ( ) is grayed out.

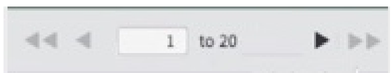

Procedure

1. Open a list, for example navigate to **Incident** > **All**.
2. Right-click a header in the list and navigate to **Configure** > **List Control**.
3. Select **Remove pagination count**.
4. **Optional:** To omit the record count only in specific view:
   a. Select **Remove pagination count for specific views**.
   b. Beside Specific list of views prevented from getting full record count, click the lock icon ( ).
   c. Select the search icon ( ) and select the view where platform omits the page count.
   d. Repeat the previous step for as many views as necessary.
5. Select **Update**.

Configure list controls

You can configure list controls, such as buttons and filters.

Before you begin

Role required: personalize_control.

Procedure

1. Open a list.
2. Right-click any column heading and select **Configure** > **List Control**.
3. Complete the form.

   Note:
   The list control fields that are available for embedded lists are more limited. Unless otherwise noted, the list control fields in the following table are available for both standard and embedded lists.

List Control form

Field	Description
Table	Name of the table for the list or related list. For example, **Change Request [change_request]**. The system automatically sets this value.

Field	Description
Related list	Name of the table and field that define the related list. For example,**sysapproval_approver.sysapproval**. The system automatically sets this value.
Label	Label to display for this list. Allows an admin to customize the label for a related list or list. If not supplied, the default plural label for the file is used. For example, the label for the Incident table would be **Incidents**.
Omit new button	Option to prevent the **New** button from displaying on this list. Clear the check box to display the **New** button or to control the **New** button with roles (**New roles** field). This field is available for standard lists only.
Omit edit button	Option to prevent the **Edit** button from displaying on this list. Clear this check box to display the **Edit** button or to control the **Edit** button with roles (**New roles** field). The **Edit** button does not apply to all lists. This field is available for standard lists only.
Omit if empty	Option to omit the **Related List** from the form entirely (no header) if there are no entries for the **Related List**.
Omit columns if empty	Option for a top-level list to omit the column headers AND filters/breadcrumbs for an empty column.
Omit filters	Option to hide filters or breadcrumbs for this list. Clear this check box if you always want filters or breadcrumbs or to control filters/breadcrumbs with roles (**Filter roles** field). This field is available for standard lists only.
Omit links	Option to hide links for fields that reference other files in this list. Leave this button unchecked to generate links or to control the use of links with roles (**Link roles** field).
Omit drill-down link	Option to disable the link to the record from the first column in list view. Users can still select the reference icon to access the record. This field is available for standard lists only.
Hierarchical lists	See Enable a hierarchical list. Hierarchical list inserted into a record list. Hierarchical lists enable a user to view the contents of a record's related lists without leaving the record list form. This field is available for standard lists only.
List edit type	Ability for the user to edit values directly in individual cells in a list. The options are: ◦ **Save immediately (cell edit mode):** enables cell editing. The entire row is saved when the user enters a new value. ◦ **Save data by rows:** enables cell editing. The row is saved only when the user navigates away from the row or selects the **Save** icon ( ). This mode allows the user to modify multiple values before saving a record. ◦ **Disable list editing:** prevents users from editing cells in the list.

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td></td><td>This field is available for standard lists only.</td></tr>
<tr><td>List edit tag</td><td>Unique tag that is sent to a reference qualifier as the script variable <code>listEditRefQualTag</code>. Enter an arbitrary string of letters or numbers.</td></tr>
<tr><td>List edit insert row</td><td>

Ability for a user to create records in list view. When it is enabled, an empty row appears at the bottom of the list.

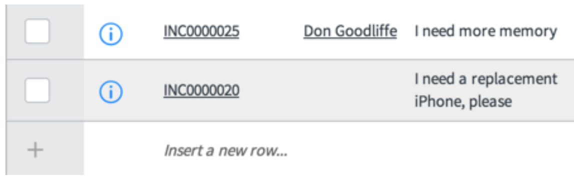

This field is available for standard lists only.

</td></tr>
<tr><td>Remove pagination count</td><td>Option to remove pagination count from the list.</td></tr>
<tr><td>Disable Natural Language Query (NLQ)</td><td>

Option to disable the NLQ filter in specific lists when NLQ on lists is enabled in the instance. The NLQ filter enables you to use natural language to filter lists instead of using the condition builder. For more information on NLQ filters, see Natural Language Query.

Note:
When the system property (<code>com.snc.listv2.nlq.lists.enabled</code>(com.snc.listv2.nlq.lists.enabled) is set to **false**, the query destructs. Text entered in the search replaces existing conditions.

</td></tr>
<tr><td>Application</td><td>Scope of this setting's applicability. Global means that all workspaces can use this setting.</td></tr>
<tr><td>New roles</td><td>User roles that can access the New button on this list. Leave the field blank to enable all users to access the New button. This option does not apply to embedded lists, which do not contain New buttons.</td></tr>
<tr><td>Filter roles</td><td>User roles required to have the filter appear in the list.

This field is available for standard lists only.</td></tr>
<tr><td>Link roles</td><td>User roles required to have links generated for fields that reference other files.</td></tr>
</table>

Example of configuring list controls

The following example shows how to configure a list of related incidents in a problem record to suit specific business needs.

Before you begin

Role required: personalize_control

About this task

In this example, a related list label is configured to say `Child Incidents`, not just `Incidents`. The related list is also configured to remove the **New** button to prevent users from creating incidents from the Problem form.

Procedure

1. Navigate to a problem record.
2. Open the list context menu in the **Incidents** related list and select **Configure** > **List Control**.
3. On the List Control form, change the **Label** field to `Child Incidents`.
4. Select the **Omit new button** check box.
5. Select **Submit** to save your changes and return to the previous record, in this case, the problem record.
   In the problem record, the label for the embedded incident list has changed and the **New** button is no longer available.

Advanced list control with scripts

You can write scripts to specify which list elements, such as links, filters, and buttons, are present.

These scripts can react to the record that is being displayed. For related list controls, the global variable **parent** identifies the parent record. For primary lists, there is no parent record.

Configure the List Control form and add the script field that you need:

- **Omit Columns Condition:** If there are no conditions, omit column headings.
- **Omit Edit Condition:** Omit the **Edit** button.
- **Omit Empty Condition:** If there are no conditions, omit the list header.
- **Omit Filter Condition:** Omit the list filter.
- **Omit Links Condition:** Omit related links.
- **Omit New Condition:** Omit the **New** button.

In these script fields, if the script evaluates to **true**, the item is omitted. The following script is an example that you can use on the **Affected CIs** related list to show only the **Edit** button if the parent task is active. This script is placed in the **Omit Edit Condition** script field.

```
var answer;
 if (parent.active == true ) {
 //Do not remove the 'Edit' button
  answer = false;
}
else {
 //Remove the 'Edit' button
  answer = true;
}
answer;
```

Configure a list UI action

You can create a UI action that opens the slushbucket interface to customize lists in the list view.

Before you begin

Role required: ui_action_admin

Procedure

1. Navigate to **All** > **System Definition** > **UI Actions**.
2. Click **New**.

3. In the **Table** field, select the **List [sys_ui_list]** table.
4. Select the **List context menu** check box.
5. Ensure the **Active** check box is selected.
6. Paste the following script in the **Script** field:

```
var url  = "slushbucket.do?sysparm_form=list&sysparm_list=" +
 current. name + "&sysparm_view=" + current. view. name ;
url  += "&sysparm_collection=" + current. parent +
 "&sysparm_collection_relationship=" + current. relationship ;
url  += "&sysparm_referring_url=sys_ui_view.do?sys_id=" +
 current. view. toString ( ) ;
action. setRedirectURL (url ) ;
```

7. Click **Submit**.

Show/hide filter controls

The Show/hide filter used in the list configuration, lists the fields to configure in the filter conditions.

The fields mentioned in the Show/hide filter are sorted in an alphabetical order.

> **Note:**
> If the name of the fields start with an accented character, those fields are placed at the end of the list.

You can enable the *guide.ui.condition_builder.sort_labels_by_locale* sys property to sort the fields of the Show/hide filter as per the locale-based language. The sys property is set to **False**, by default. If you have a customized script that's based on alphabetical sorting of the fields, the locale-based language setting might break your script. So, you can enable the sys property only when you need the filter to be locale-based language.

> **Note:**
> If the sys property is not enabled, the fields are sorted in an alphabetical order. The fields starting with an accented character are placed at the end of the list.

Controlling the sort sequence used to display lists

All lists have a default sort sequence based on the type of fields present in the list.

When a user displays a list for the first time, it is sorted by one of the following items.

- The **ORDERBY** arguments found the URL. (See the following section on sort order control.)
- The **<table>.db.order** and **<table>.db.order.direction** user preferences.
- The **isOrder** dictionary attribute.
- The **Order** field if one is present in the table.
- The **Number** field if one is present in the table.
- The **Name** field if one is present in the table.
- The field specified as the display field for the table.

> **Note:**
> The task table is an exception to the preceding sequence. Task and tables extended from task use the **Number** field when no URL arguments or user preferences are found.

Setting default order with user preferences

The **<table>.db.order** and **<table>.db.order.direction** user preferences determine the field and sort direction seen by the user. A default sort can be created by creating user preference records with no value in the **User** field and checking the **System** check box. This preference is applied to any users who do not have their own preference.

Setting the default sort order in the system dictionary

An administrator can set the sort order of records displayed for a table in the system dictionary when there is no other sort specified. This is done by adding an attribute called **isOrder** with a value of **true** to the dictionary entry of the desired field. This sort is the sort order presented to all users initially. Once a user sorts a list, that user preference is saved, and the list is always sorted for that user based on the previous sorting preference.

Note:
isOrder can define which field is used for list sorting, but cannot control the sort direction.

Related topics

- System dictionary
- User preferences
- Altering tables and fields using dictionary attributes

Sort order control

For any list, if the user has clicked a column header to sort by that column, then that action is remembered. The next time that list is displayed, the same field is used to order the list.

This remembered sort field can be overridden by adding a `sysparm_order` specification to the definition of the module. For example, if each time **News** is displayed, you want the new items listed in order of importance, then the **News** module can be updated accordingly.

Set module sort order to the Importance field value

The actual order (ascending or descending) can be specified by adding a `sysparm_order_direction` specification. Here are two examples of sorting a list by the **sys_created_on** field, one ascending and one descending:

```
&sysparm_order=sys_created_on &sysparm_order_direction=desc
&sysparm_order=sys_created_on &sysparm_order_direction=asc
```

Sort by multiple fields

To sort a list by multiple fields, remove everything from the filter field and place it in the arguments field. For example, if the filter specifies **[Active] [is] [true]** and you want to sort by category and subcategory, remove the condition and put the following in the module arguments:

```
active=true^ORDERBYcategory^ORDERBYsubcategory
```

This argument first orders by the category and then orders by the subcategory. If you want any field to be a descending sort, then change *ORDERBY* to be *ORDERBYDESC*.

List editor administration

The list editor allows users to edit field values directly from a list without navigating to a form.

Administrators can manage this feature by using the following options. Some of the options are not supported in List v3.

- Configure global properties
- Configure list control settings for a table
- Configure contextual security rules
- Manage user preferences (List v2 only)

Warning:
In List v2, the list editor does not enforce client scripts or UI policies. In List v3, the list editor enforces UI policies and mandatory dictionary attributes, but not client scripts. Allowing list editing with client scripts running on fields in a form can result in incorrect data being saved to the record. For systems in which client scripts or UI policies apply to forms, consider disabling list editing or creating appropriate business rules or access control to control the setting of values in the list editor. See Use business rules and client scripts to control field values for information on managing form and list editing.

Warning:
You cannot edit state fields that are extended from the Task [task] table from a list of task records. Examples include:

- problem_state
- incident_state
- phase_state
- request_state

While these fields are visible from the task list, they must be edited from the extended table.

Related topics

Use the list editor

Configure list editor properties

Configure list editor properties that control whether lists can be edited and, in List v2, which field types cannot be edited.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. To disable list editing, set the **Enable list editing** (`glide.ui.list_edit`) property to **No** by clearing the check box.
   This property is enabled by default, and it globally enables list editing. When you disable it, the list editor is disabled globally.
3. To configure the field types that cannot be edited for v2 lists, complete the following steps.

a. Locate the **List of element types (comma-separated) that cannot be edited in the list editor** (*glide.ui.list_edit_ignore_ types*) property, which contains several element types that cannot be edited by default.

Note:
This property does not impact v3 lists. There is no equivalent property for List v3.

The following field types are not editable from the list editor by default.

- Conditions [conditions]
- Currency [currency]
- Document ID [document_id]
- Field List [field_list]
- HTML [html]
- Image [user_image]
- List [glide_list]
- Price [price]
- Template Value [template_value]
- Time [glide_time]
- User Roles [user_roles]
- Video [video]
- Array [array]

b. Add any other field types you want to disable to the end of the list, separated by a comma.

4. Select **Save**.

Configure list control settings for the list editor

You can configure the list control settings that affect the list editor.

Before you begin
Role required: personalize_control

About this task
List control settings customize the behavior of list functions for a table.

Procedure

1. Navigate to a list view for the desired table.
2. Right-click any column heading and select **Configure** > **List Control**.
3. On the List Control form, select the desired settings.

List control settings for the list editor

Field	Description
List edit type	Ability for the user to edit values directly in individual cells in a list. The options are:

Field	Description
	○ **Save immediately (cell edit mode):** enables cell editing. The entire row is saved when the user enters a new value. ○ **Save data by rows:** enables cell editing. The row is saved only when the user navigates away from the row or selects the **Save** icon ( ). This mode allows the user to modify multiple values before saving a record. ○ **Disable list editing:** prevents users from editing cells in the list. This field is available for standard lists only.
List edit insert row	Ability for a user to create records in list view. When it is enabled, an empty row appears at the bottom of the list. 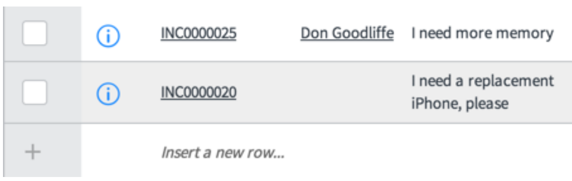 This field is available for standard lists only.

4. Click **Update**.

Configuring contextual security for the list editor

The list editor enforces existing access control rules (ACLs) and additional security controls to restrict editing from a list.

The list_edit security operation specifically controls the ability to edit information in a list. Apply this operation in the same manner as the write operation to limit list editing for fields that require the user to navigate to the form. Reasons you may require a user to edit a field in a form include complex UI policy constraints or other policies you have in place.

- write and list_edit access to the field
- write and list_edit access to any dependent fields of the field
- write and list_edit access to any fields that depend on the field being edited

To configure access controls, navigate to **System Security** > **Access Controls**. The following examples use the list_edit security operation to restrict list editing in certain contexts.

Example - Restrict a table

This access control prevents everyone from editing all fields in the Incident table in a list.

Restrict the Incident Table

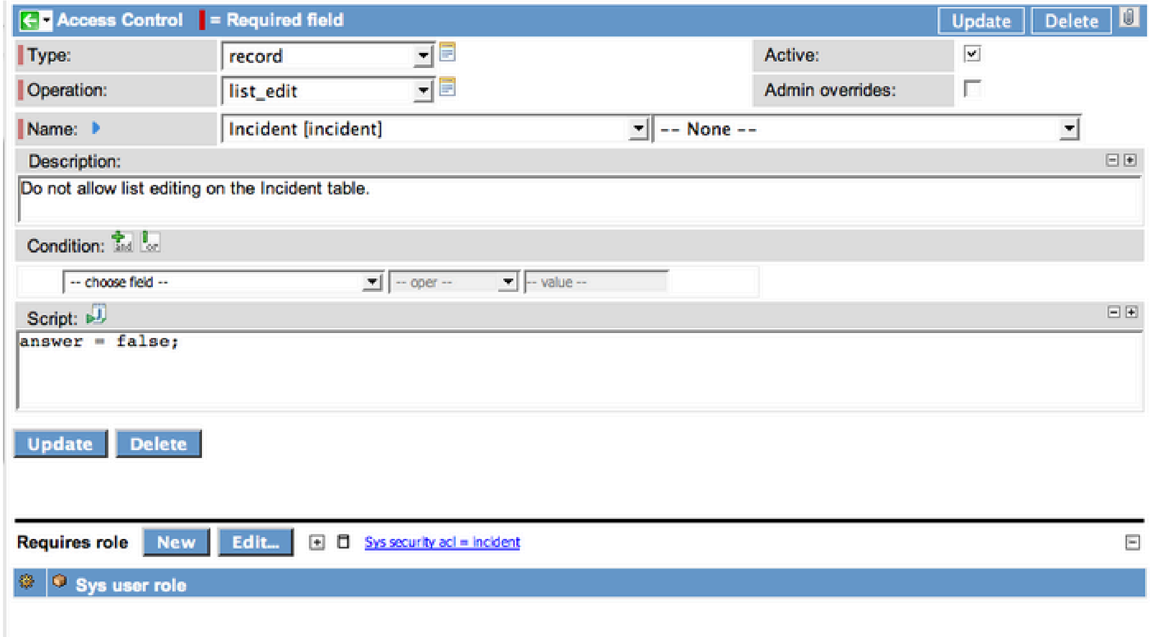

- **Type:** record
- **Operation:** list_edit
- **Name Incident:**[incident]
- **Admin overrides:** Clear the check box.
- **Script:** `answer = false;`

Example - Restrict a field

This access control prevents everyone except an administrator from editing the Short Description field of an incident record in a list.

Restrict the Incident Short Description

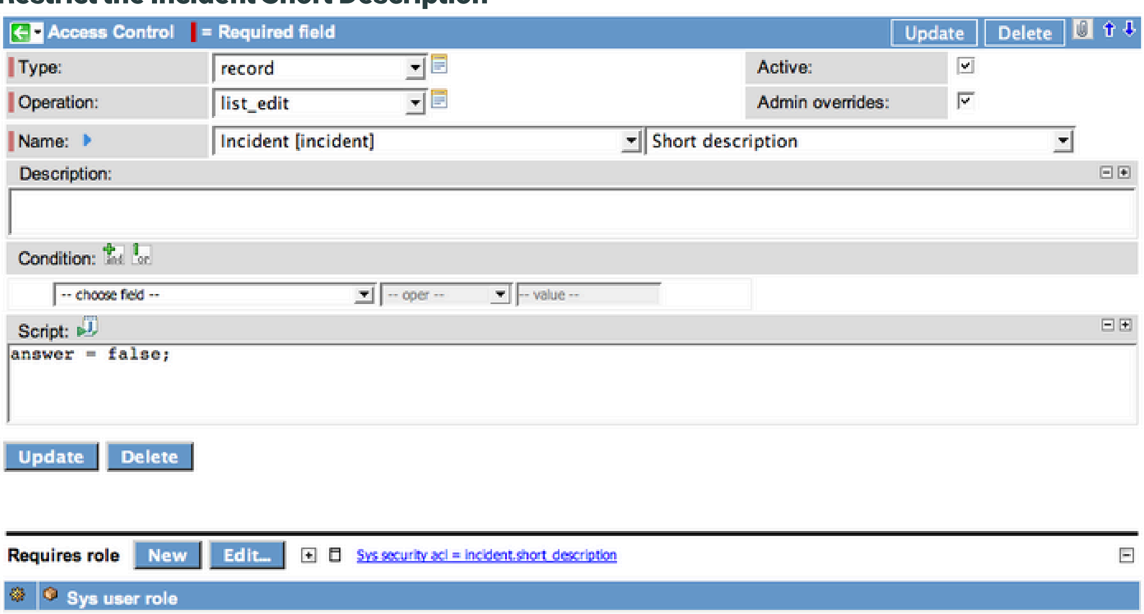

- **Type:** record
- **Operation:** list_edit
- **Name:** Incident [incident], Short Description
- **Admin overrides:** Select the check box.
- **Script:** `answer = false;`

Example - Restrict a field with a script

This access control prevents everyone from editing an incident with a category of Software in a list. It is defined by a script.

Restrict Software Incidents

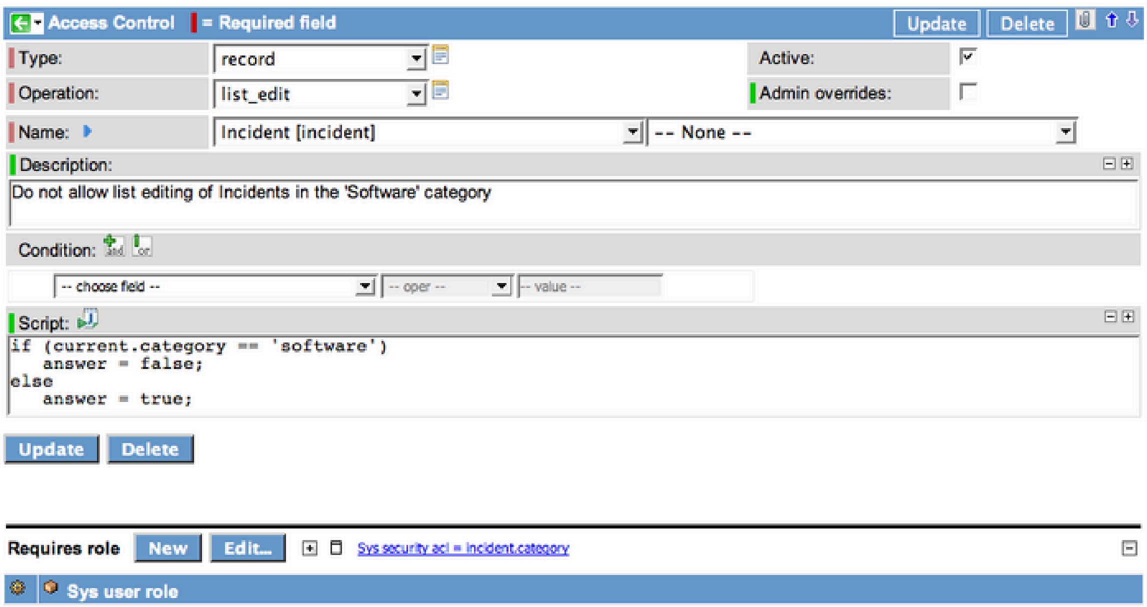

- **Type:** record
- **Operation:** list_edit
- **Name Incident:**[incident]
- **Admin overrides:** Clear the check box.
- **Script:**

```
if (current.category == 'software')
answer = false;
else
answer = true;
```

Example - Restrict a field with a condition

This access control prevents everyone from editing a Critical Incident in a list. It is defined by a condition.

Restrict Critical Incidents

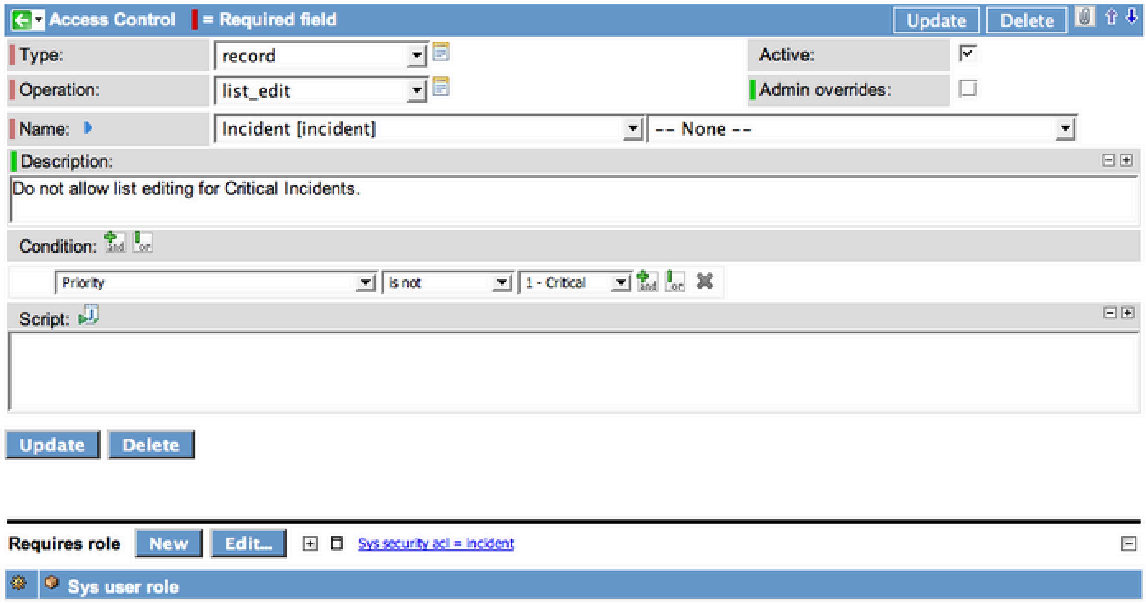

- **Type:** record
- **Operation:** list_edit
- **Name Incident:**[incident]
- **Admin overrides:** Clear the check box.
- **Condition:** Priority is not 1 - Critical

User preferences for list editing

Users can set user preferences for list editing by personalizing a list (available when the list mechanic is activated).

To view and manage list editor user preferences, navigate to **User Administration** > **User Preferences**.

User preferences for list editing

User preference name	Description
list_edit_enable	Set the value to **true** to enable or **false** to disable list editing. Default value is **true**.
list_edit_double	Set the value to **true** for double-click or **false** for single-click. Default value is **true**.

Related topics

Personal list administration

Personal list administration

Users can customize the layout for any list view by personalizing a list. Administrators can configure options related to list personalization.

In List v2, administrators can also disable or restrict access to list personalization. This functionality is not supported in List v3.

Enable or disable personal lists

You can enable or disable personal lists through the UI Macros module.

Before you begin

Role required: ui_macro_admin

Procedure

1. Navigate to **All** > **System UI** > **UI Macros**.
2. Activate (or deactivate) the *list_mechanic2* UI macro.

Control which roles can personalize lists

You can control who can create personal lists.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Locate the property labeled **List of roles (comma-separated) that can use Personal Lists. If blank, all can use** (*glide.ui.list_mechanic.roles*).
3. Enter the roles for which personal lists are available, separated by commas, or leave the field blank to allow access for all users (default).

Manage personal lists

Administrators can manage personal lists.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System UI** > **Lists**.
2. Personal lists have a user value in the **User** field.

Related topics

Personal lists

Administer detail rows

Detail rows provide additional information about records in a list.

Before you begin

Role required: admin

About this task

To display detail rows, an administrator must enable them globally and add them to lists.

> **Note:**
> Detail rows are not supported in List v3.

Procedure

1. To enable detail rows, add a new system property with the following field values.

Name	glide.ui.list.detail_row
Type	true | false
Value	true

2. To add detail rows to a list, complete the following steps.

   a. Add the *detail_row* dictionary attribute to the table.
   Specify the name of the field to display in the following format:
   `detail_row=short_description`

   b. Use the *all_tables* element descriptor with the *detail_row* attribute to display detail rows for all child tables that contain the designated field.

   Example
   For example, *all_tables.detail_row=short_description* is on the Task [task] table by default, so the short description detail row appears in lists for Incident [incident], Problem [problem], and other child tables unless you set *detail_row* to a different field for each one.

Related topics

Lists

Altering tables and fields using dictionary attributes

Restrict filters and breadcrumbs with fixed queries

The record list view allows users to navigate to different subsets of a table using breadcrumbs and filters. You can limit access to parts of the table by restricting active links in breadcrumbs or by suppressing breadcrumbs and filters for specific roles.

Before you begin

Role required: admin

About this task

A breadcrumb option enables an administrator to control the base view of a record list presented to users. By adding a fixed query to the argument for a module, an administrator can prevent users from expanding their view past a specified starting point. The argument for this fixed query is written as *&sysparm_fixed_query=active=true*. A use case for this query is to prevent users from using the breadcrumbs to switch a list of open incidents to a list of all incidents. When users select **Incident** > **Open**, they are limited to viewing and filtering a list of open (active=true) incidents.

Note:
A new Create ACL allows all users to save filters by default. This overrides any custom ACLs in place if administrators are restricting filter access. The new ACL gives all users access to the User field by default, and access to the Group field only if users have the filter_group role and are in the currently selected group.

Procedure

1. Point to the application menu that contains the module to edit and click the edit application (pencil) icon.
   To open the module directly, point to the module and click the edit module (pencil) icon.
2. Select the module to edit.
   For example, select **Open**.
3. In the **Link Type** section of the Module form, select **List of Records** for the **Link type**.
4. Delete the **Active is true** filter, if present.
5. Add *&sysparm_fixed_query=active=true* to the **Arguments** field and update the record.

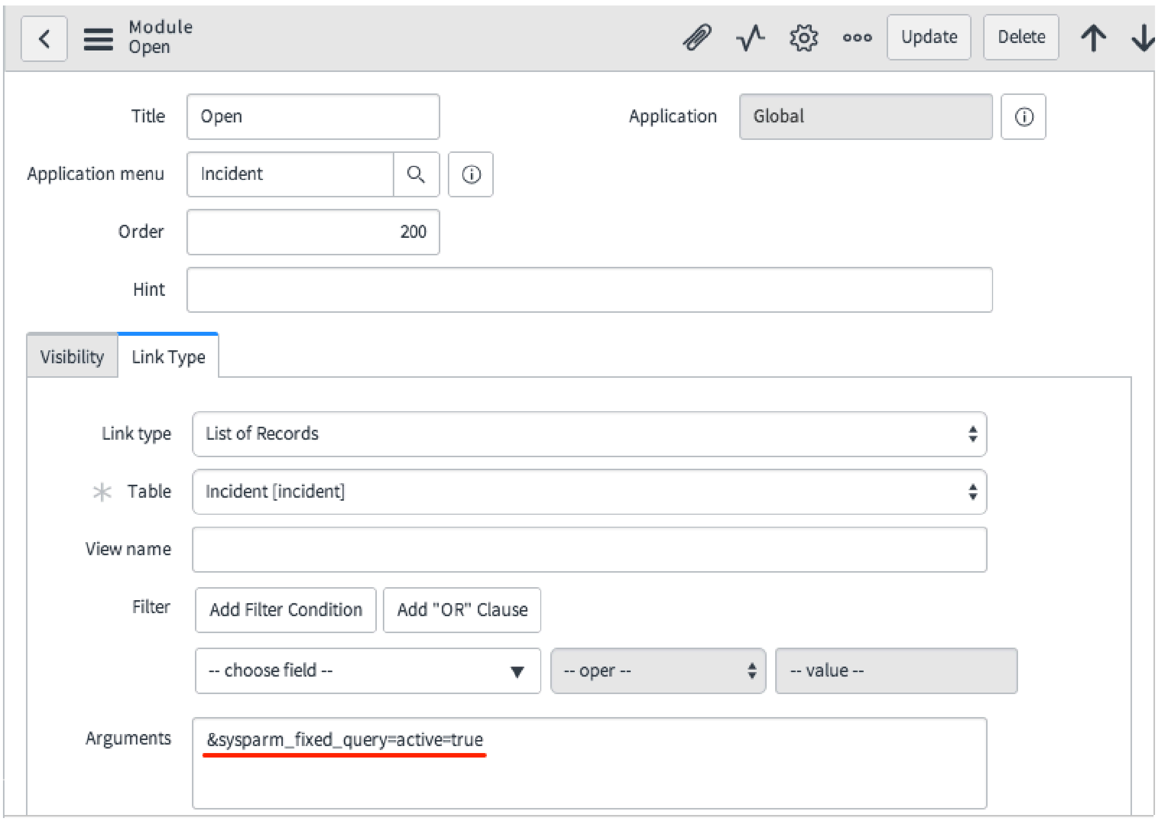

6. Navigate to **Incidents** > **Open** and examine the breadcrumbs.
   The **Active=true** link is the widest view permitted in the list of Open incidents shown. The breadcrumb for **All** is not a link because of the fixed query.

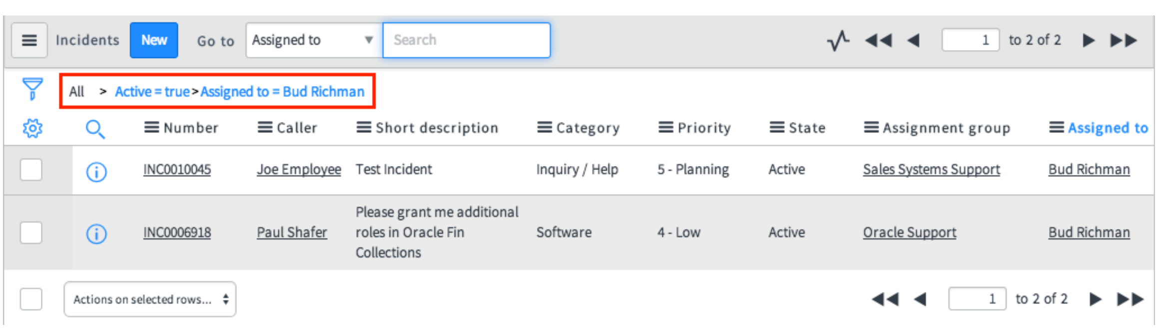

Suppress filters and breadcrumbs with list controls

Use list controls to suppress filters and breadcrumbs on defined tables for users with specific roles.

Before you begin

Role required: personalize_control

Procedure

1. Navigate to the list view of the table to restrict breadcrumbs.
2. Open list control by performing the appropriate action for the list version.
   You can also navigate to **System UI** > **List Control**.
3. On the List Control form, select the **Omit Filters** check box, or clear the check box to include filters.
   The **Filter Roles** field allows an administrator to specify certain roles that can see the filters.

Use script includes to suppress filters and breadcrumbs

You can use a script to restrict filters and breadcrumbs to specific roles, either on a per-table or global basis. Using a script is an advanced option that offers additional flexibility compared to using list control.

Before you begin

Using a script include requires knowledge of JavaScript.
Role required: script_include_admin

About this task

The examples shown must be modified for your environment.

Procedure

1. Create a script include with the name `<tablename>DisplayFilter.`
   The script section contains one function with the same name as the script include.
2. Use your function to set the global variable *answer* to either true (show the filters and breadcrumbs) or false (hide them.)

   Example
   The following example restricts filters and breadcrumbs on the Incident table to users with any role. Be sure that the name of the script matches the function name exactly, including case.

   ```
   function incidentDisplayFilter() {
       if (gs.getUser().hasRoles()) {
           answer = "true";
       } else {
           answer = "false";
       }

       return answer;
   }
   ```

3. **Optional:** To exclude a specific role from having access to filters and breadcrumbs, make the following change.

   ```
   function incidentDisplayFilter() {
       if (gs.getUser().hasRoles()
    && !gs.getUser().hasRoles('newrole')) {
   ```

```
            answer = "true";
        } else {
            answer = "false";
        }

        return answer;
}
```

Users with the role `newRole` do not have access to filters and breadcrumbs.

4. **Optional:** To let all users use filters and breadcrumbs on the Incident table, make the following change to your script.

```
function incidentDisplayFilter() {
    var answer = true;

    return answer;
}
```

5. **Optional:** To modify filter and breadcrumb access for another table, create a script include using the name of that table instead of Incident.

Increase the allowed number of breadcrumb entries

You can add a property to allow for a larger number of breadcrumb entries in the filter.

Before you begin
Role required: admin

Procedure

1. Enter `sys_properties.list` in the Navigation filter.
   The entire list of properties in the System Properties [sys_properties] table appears.
2. Verify that the property does not exist by searching for *`glide.ui.breadcrumb_max_entries`*.
3. Click **New**.
4. Complete the form.

5. To verify this property, go to any table and use the filter modifier **is one of** to search for any number of items.
   The number of entries you entered in the Value field displays before ending in a [...].

Breadcrumb entries

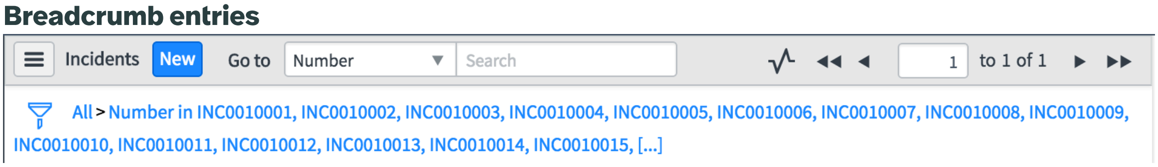

Enable a hierarchical list

You can enable hierarchical lists to make data from related lists directly accessible from within a v2 list.

Before you begin
Role required: personalize_control

About this task

Hierarchical lists enable a user to view the contents of the related lists of a record without leaving the record list form.

> **Note:**
> You cannot enable a hierarchical list view for an embedded list. Hierarchical lists are not supported in List v3, even if they are enabled.

Procedure

1. Right-click any column header, and select **Configure** > **List Control** from the context menu.
2. Select the **Hierarchical lists** check box in the List Control dialog box and then click **Submit** or **Update**.
3. Click the arrow for a record number to display the related lists for the record.

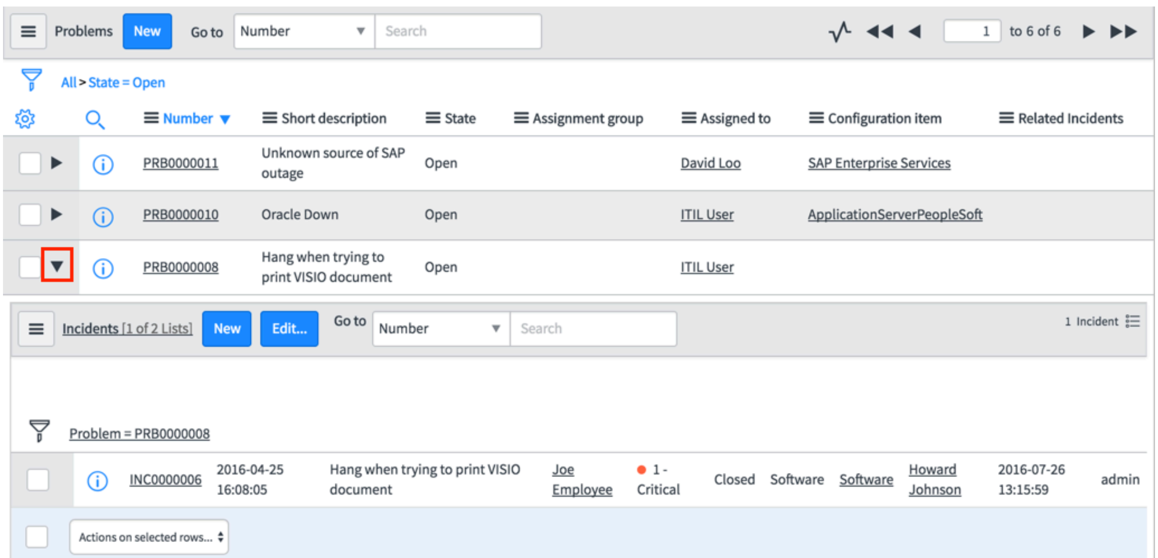

Use list controls in hierarchical lists

By default, no related list tabs are visible in a hierarchical list and only a single related list is displayed at a time. The column headers are hidden by default.

Before you begin

Role required: none

About this task

> **Note:**
> Hierarchical lists are not supported in List v3.

Procedure

1. To select a different related list, click the link in the list title.
2. Select a list from the menu.

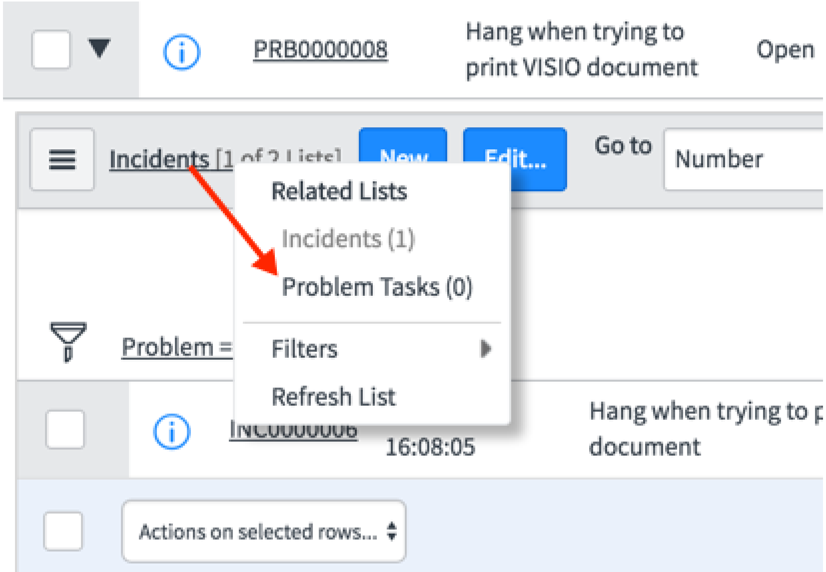

If the related list is empty, the hierarchical list indicates that there are no records to display.

The platform remembers this preference and displays the last selected related list for each table when you open another hierarchical list within that table.

3. To display column headers in the related list, click the show/hide headers icon in the upper right corner of the hierarchical list.

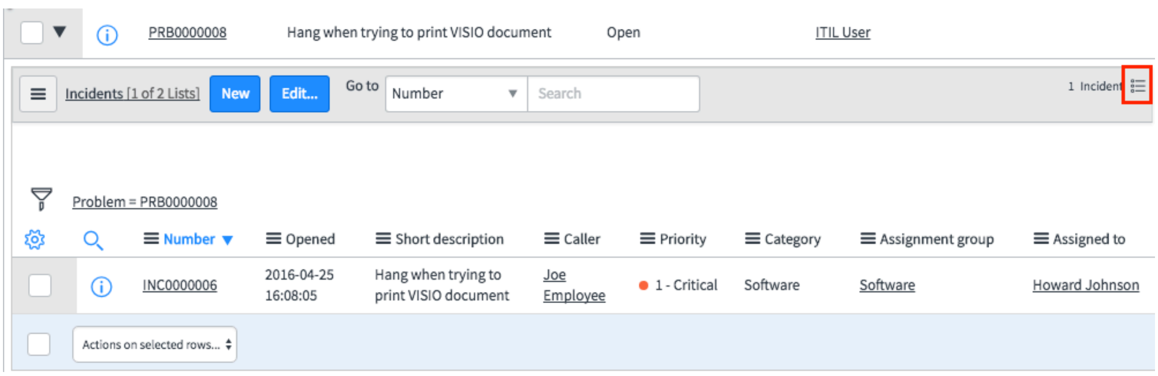

The platform remembers this preference and opens all hierarchical lists with column headers enabled until the headers are hidden.

Context ranking

Context ranking allows a user to sort a collection of records preferentially, that is, independently of the attributes of those records.

For example, a Customer Support manager can organize a list of incidents in the order in which a technician should work on the tasks. Creating such an arbitrary list with a list filter is not possible.

Activate context ranking

If you have the admin role, you can activate the Context Ranking plugin (com.glide.sorting). This plugin activates related plugins if they are not already active.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Context Ranking plugin (com.glide.sorting) using the filter criteria and search bar.

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

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

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

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

Create a ranking definition

Enable context ranking for a related list by creating a ranking definition.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Ranking Definitions**.
2. Click **New**.
3. Enter a **Name** for the ranking definition.
4. In the **Record table** field, select the table for which ranking is enabled.
   This is the table of the records in the ranked related list. For example, to enable ranking for a related list of incidents in the Problem form, select the Incident [incident] table.
5. In the **Context column** field, select the reference column providing the context in which ranking is to be performed.
   The context in this case is the form in which the related list appears and not a table name. For example, a related list with a **Record table** of Incident [incident] can have several contexts, including the **Problem** form or a **Configuration item** form.
6. Complete the relationship-based fields if the ranking is to be performed on a related list that is defined by a relationship.
   - **Relationship table:** select the table forming the other side of the relationship. In the case of the ranking definition for Product stories in scrum, you might use the Product Model [cmdb_model] table.
   - **Relationship column:** select the column from the relationship table to compare with the context column. To avoid adding a related list of user stories to the Product Model [cmdb_model] table, you might use the **Sys ID** column, which defines a subclass of the Product Model table (cmdb_model.sys_id).
7. In the **Attributes** field, enter attributes to change and restrict the contents of the Rank dialog box, as appropriate.

- *`visible_columns`*: a semi-colon separated list of columns to be displayed in the Rank dialog box. If not specified, the Rank dialog box uses the default view of the related list you selected in the **Record table** field.
- *`extra_conditions`*: an encoded query to restrict the records shown in the Rank dialog box. For example, the query `extra_conditions=^state!=7` displays all incidents whose state is not **Closed**. If not specified, the Rank dialog box shows all records for the relationship.

8. Click **Submit**.
   The index is generated automatically. The system completes these read-only fields:
   - Index column
   - Rank ui action
   - Context menu item

   Note:
   After you submit a ranking definition, only the **Attributes** field can be edited. If additional changes are required, delete the record and create a new one with the revised settings.

What to do next

The system does not generate indexes for ranking definitions inserted into the database by an update set. To use a ranking definition inserted by an update set, you must generate the indexes manually. To generate indexes, open the Ranking Definition form, and click **Generate Indexes**.

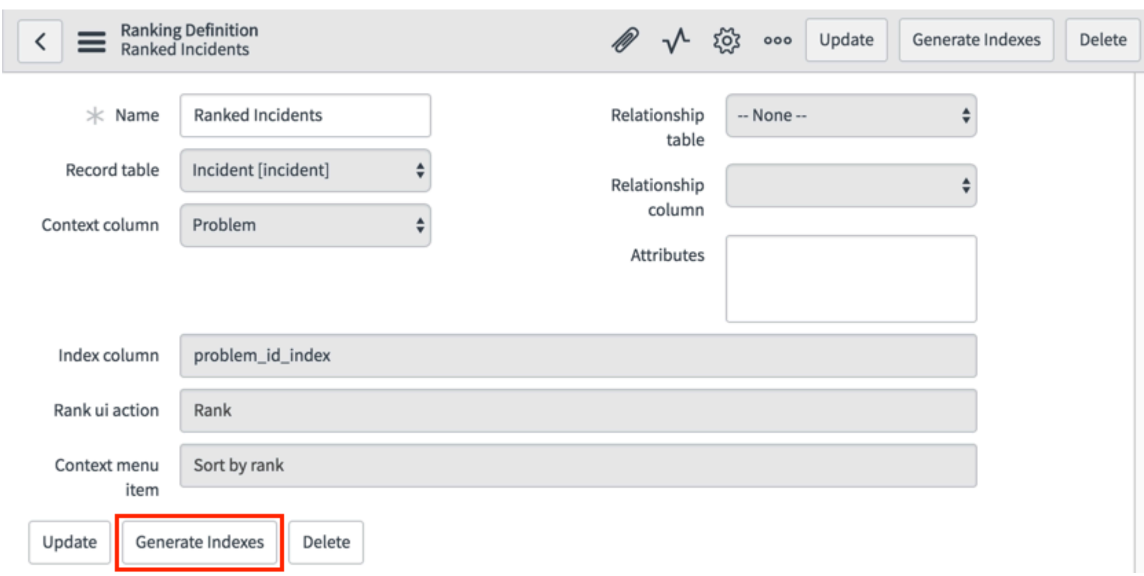

Rank a list

After you create a ranking definition for a table, related lists based on that table and context include options for ranking the list and displaying the list by rank.

Before you begin

Role required: none

About this task

You can set the preferential order of records.

Procedure

1. Click the **Rank** button in the related list to reorder the records.

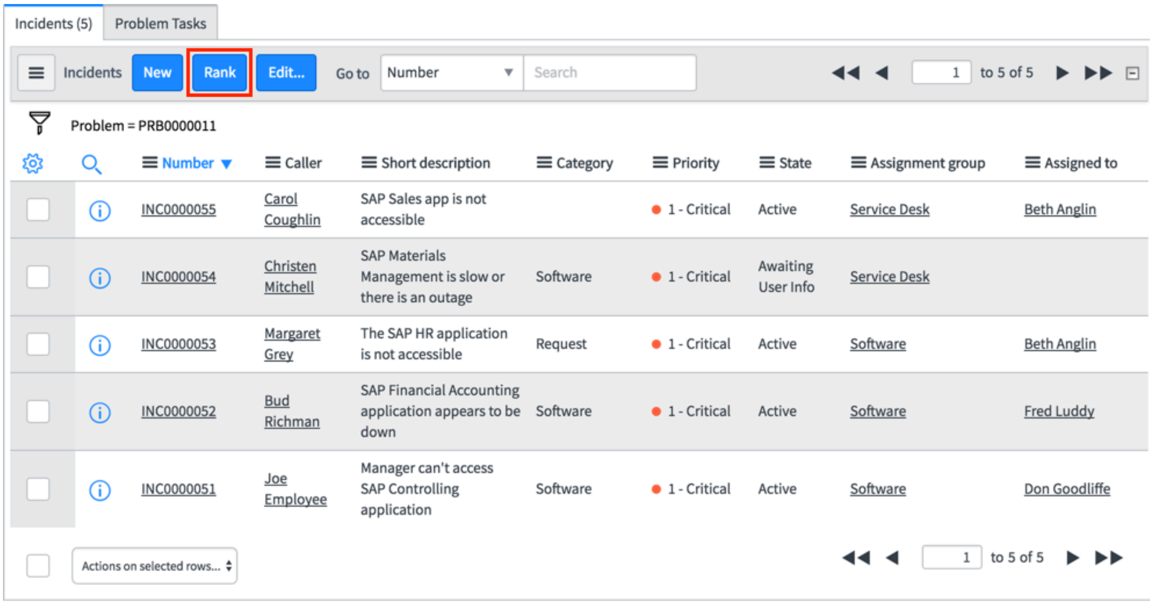

The Rank dialog box appears, allowing you to sort the records in any order.

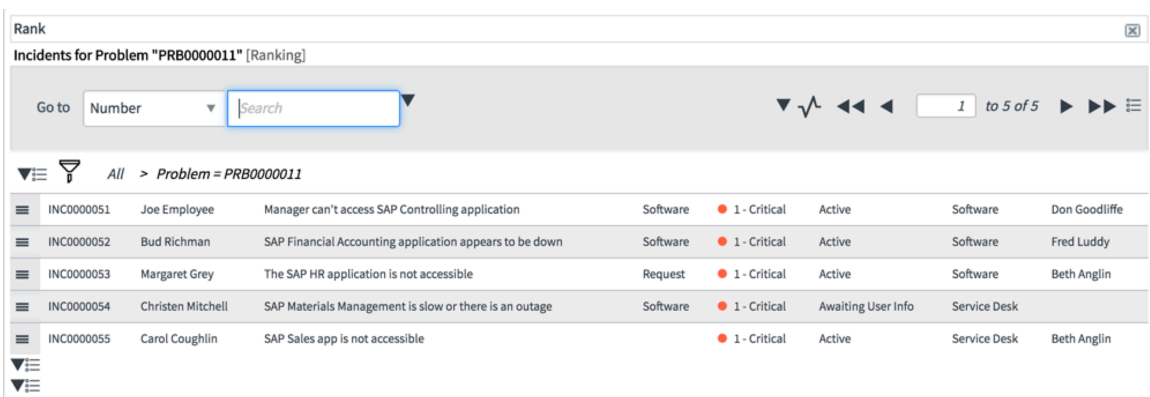

2. Click and drag the move icon ( ) to change the position of a record.

> **Note:**
> Administrators can define which columns appear in the Rank dialog box in the **Attributes** field of the Ranking Definition form.

Apply a new sort order to a list

You can apply a new sort order to the records in a list.

Before you begin

Role required: none

Procedure

1. Open the context menu from the list title and select **Sort by rank**.

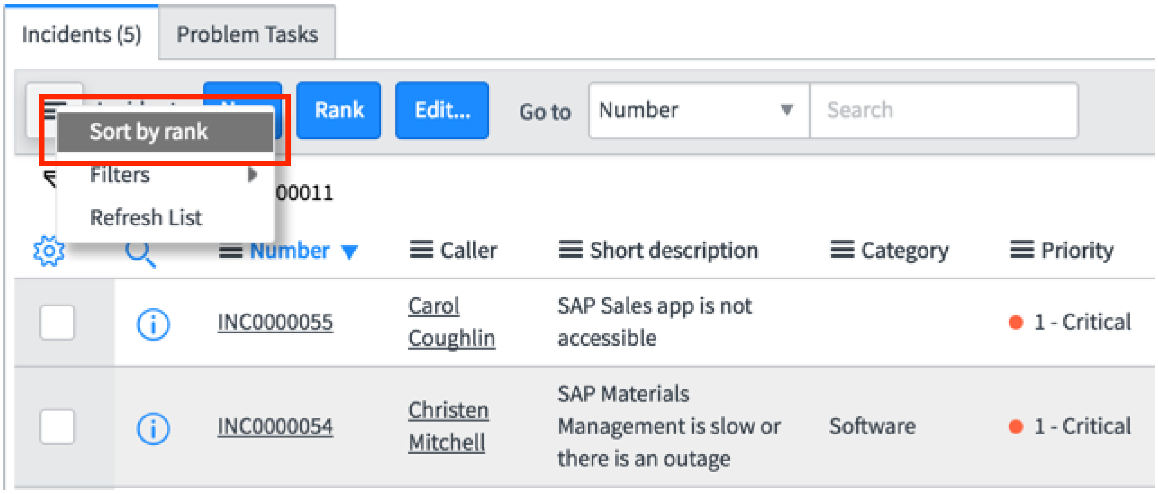

2. To remove the ranking and return to the original order of the list, click the label in any column that contains data.

Scrum ranking definitions

Ranking definitions for user stories are part of the SDLC (scrum process) application.

These ranking definitions enable scrum users with the proper roles to rank related lists of stories in these tables:

- Product stories: Application Model [cmdb_application_product_model] table
- Release stories: Release [rm_release] table
- Sprint stories: Sprint [rm_sprint] table

Users can rank stories in the scrum planning board by priority. Story lists ranked in the planning board appear in the new order in product, release, and sprint forms. Stories ranked in one of these scrum forms changes the order in the planning board.

Story points for each story are listed in a ranked view, which is useful for prioritizing stories.

Ranking story points

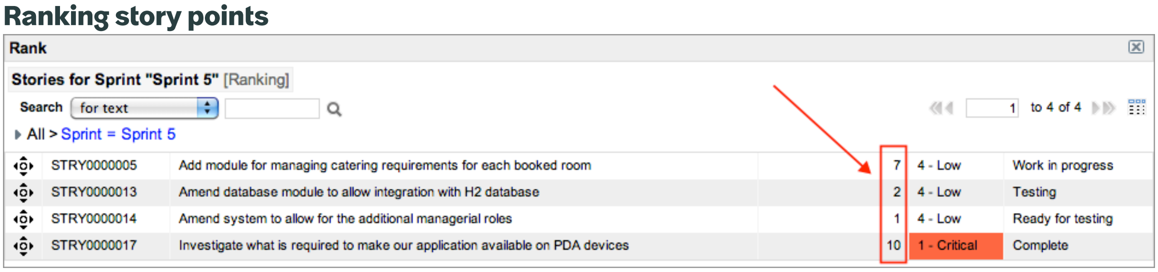

Related topics

Agile Development

Rank stories in a related list

The Context Ranking feature enables all scrum users to manually sort a related list of stories by priority.

About this task

A product owner or release planner uses this tool to establish the order in which to work the stories. Stories ranked in a related list (in the Product, Release, and Sprint forms) appear in the same order in the planning board when the appropriate backlog is displayed. Conversely, stories ranked in the planning board appear in that order when viewed in the related list in the form. Users can switch the view in a Stories related list from ranked to any sort order without changing the ranking in the planning board.

Procedure

1. In the **Stories** related list, click **Rank**.
   A dialog box appears, allowing you to arrange the stories in any order, such as by priority.
2. Click and drag each story into position using the move icon to the left of the story number.

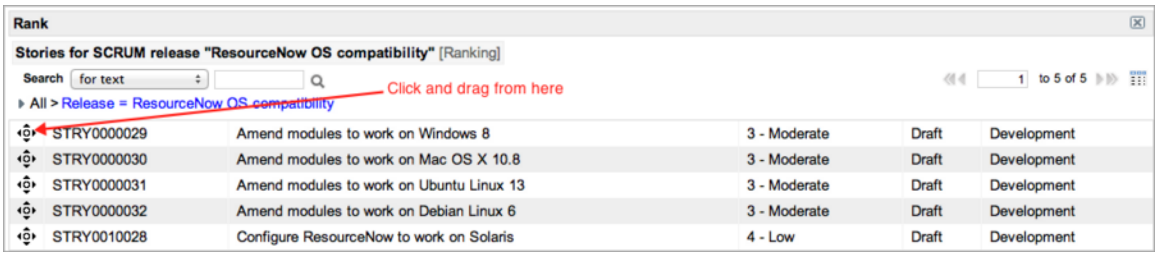

3. When you are done, click the **X** in the upper right corner to close the list.
   The **Stories** related list is not sorted by your ranking initially.
4. To view the related list in its ranked order, open the context menu from the related list and select **Sort by rank**.

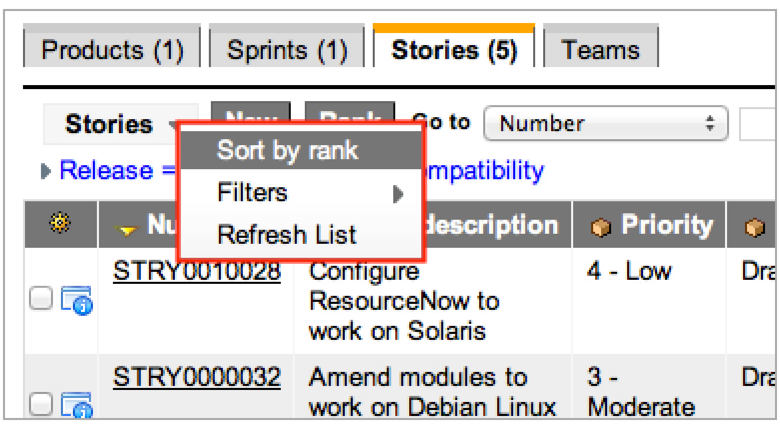

5. To return the sort order of the list to an unranked state, click once in the heading of any column that contains data.
   The system uses this ranked list to display the appropriate backlog in the planning board. For example, if you rank stories in the **Stories** related list in the Release form, the release backlog in the planning board uses the same ranking to display the stories. Conversely, scrum masters, product owners, and release planners can create a new ranking order for the **Stories** related list by rearranging the list of stories in the planning board.

Related topics

Agile Development

Customize list context menus

You can customize the content and behavior of list context menus.

Before you begin

Role required: admin

About this task

For example, you might add an option to the list column heading menu for a specific table. You might also want to change the roles required to access a global context menu option.

Procedure

1. Navigate to **All** > **System UI** > **UI Context Menus**.
2. Click **New**.
3. Complete the form.

Context menu form fields

Field	Description
Table	Select the table to which this context menu option is attached. The base system menu items are attached to the Global [global] table, which applies the context menu option to all lists for all tables. If you specify a particular table, the option is available only on context menus in lists from that table. **Note:** The list shows only tables and database views that are in the same scope as the context menu.
Menu	Select the list context menu in which this option appears. ◦ List title ◦ List header ◦ List row
Type	Select the type of menu option to create: ◦ **Action**: A menu option that performs an immediate action. ◦ **Menu**: Creates a parent menu that can display a submenu. ◦ **Separator**: Draws a line between groups of options on a menu. Menus do not display separator lines next to one another or at the bottom of a menu. If a condition that removes options for a role forces two separators together, one of the separators is removed from the view. ◦ **Label**: Create an unlinked label for a menu or section of a menu. ◦ **Dynamic actions**: Menu options dynamically created, such as the available views or user filters that can only be generated at the time the list is displayed.
Name	Enter the label for the action as it will appear in the menu.
Parent	If this action is part of a submenu, type the name of the parent menu item. For example, in the base system, **Configure** is a parent.
Order	Assign an order number to this item, menu, or separator to specify where in the menu it appears.
Active	Enable or disable this context menu item. Only active items are shown in the context menu.

Field	Description
Run onShow script	Select this check box to display the **onShow script** field.
Condition	Create the conditions under which this menu option appears. For example, define the role that has permission to see this item.
Action script	The action script runs when someone selects the context menu option. For more information, see Action script for list context menus.
Dynamic actions script	The dynamic actions script populates the context menu with dynamic options, such as filters or views. For more information, see Dynamic actions script for list context menus.
onShow script	The onShow script runs before the context menu is displayed, to determine which options appear in the context menu. For more information, see onShow script for list context menus.

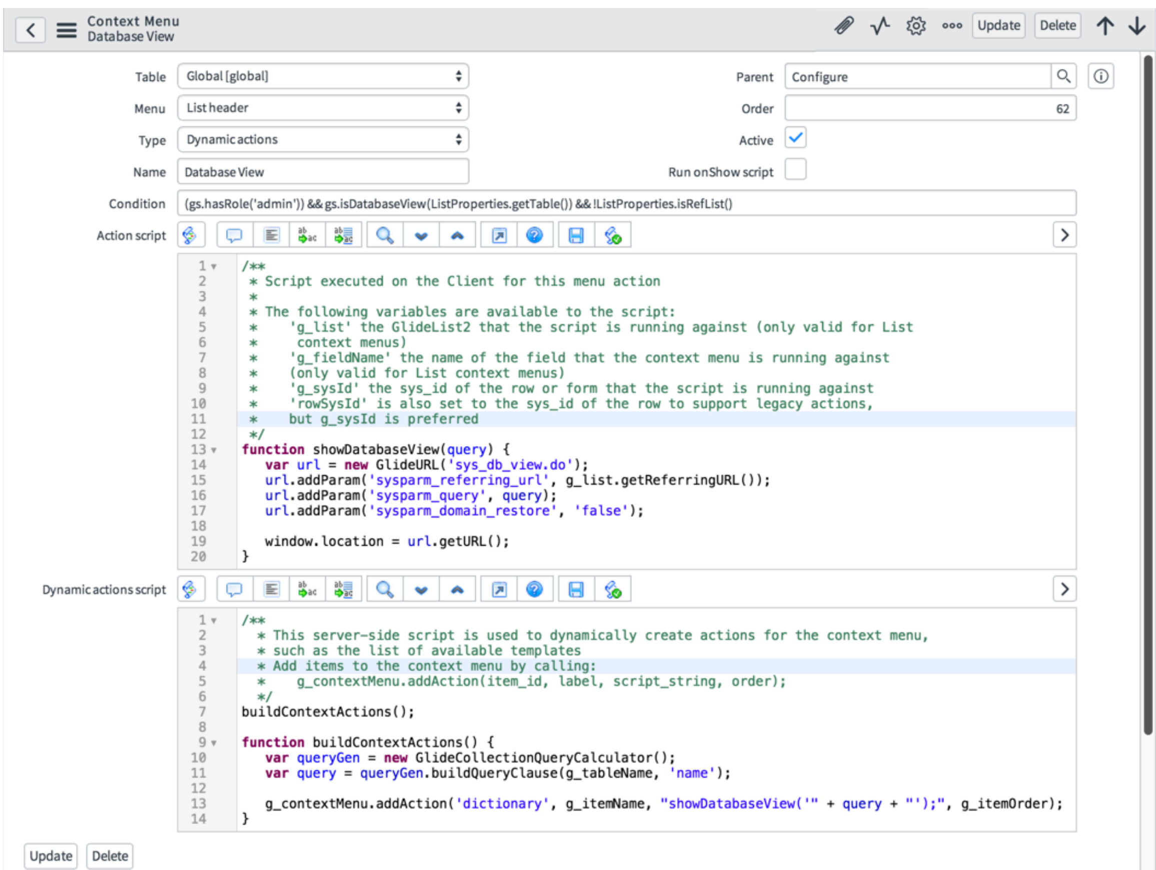

4. Click **Submit**.

Action script for list context menus

The **Action script** field, on the Context Menu form, defines a script. The script runs when someone selects the context menu option.

This script is client-side and runs in the user's browser. The following JavaScript variables are available to the Action script when it is executed.

Action script variables

Variable	Description
g_list	GlideList2 against which the script runs.
g_fieldName	Name of the field against which the context menu runs.
g_fieldLabel	Label of the field against which the context menu runs.
g_sysId	The sys_id of the row or form against which the script runs.

The base system uses the following code in an action script to refresh the platform view.

```
g_list.refresh(1);
```

Another example is the use of these variables in a list header menu to sort a list by the selected field in descending order (z to a).

```
g_list.sortDescending(g_fieldName);
```

Dynamic actions script for list context menus

The **Dynamic actions script** field, on the Context Menu form, defines a script. The script populates a list context menu with dynamic options, such as filters or views.

The following JavaScript variables are available to the dynamic actions script when it is executed.

Dynamic actions script variables

Variable	Description
g_tableName	Name of the current table.
g_listId	ID of the list for which the context menu is built.
g_itemName	Name defined in the UI context menu record.
g_itemOrder	Order defined in the UI context menu record. Use this variable to pass the value of the **Order** field to the dynamic actions script.
g_contextMenu.addAction(item_id, label, script_string, order)	Add options to the context menu and select the order in which they appear.

The following example displays a list title menu item that controls the number of records per page in the list view.

```
g_contextMenu.addAction('50', g_itemName,
 'showRowsPerPage("50");', g_itemOrder);
```

Note:
The action script for this item must define the `showRowsPerPage` function so that when selecting this menu item, that function is called with an argument of 50.

onShow script for list context menus

The **onShow script** field defines a script that runs before the context menu is displayed to determine which options appear in the context menu.

Use this script to change the menu items on the list header menu based on the current field column. The following JavaScript variables are available to the *onShow* script when it is executed:

onShow script variables

Variable	Description
g_menu	Context menu to be displayed.
g_item	Current context menu item.
g_list	GlideList2 against which the script runs.
g_fieldName	Name of the field against which the context menu runs.
g_fieldLabel	Label of the field against which the context menu runs.
g_sysId	The sys_id of the row or form against which the script runs.

An example of an *onShow* script is one that determines when to enable or disable the **Ungroup** option in a list column heading menu based on whether the list is grouped or not.

```
if (g_list.getGroupBy()) {
   // list is grouped so enable to Ungroup menu item
   g_menu.setEnabled(g_item);
} else {
   // list is not grouped, so disable the Ungroup menu item
   g_menu.setDisabled(g_item);
}
```

Search administration

The ServiceNow AI Platform® includes multiple search engines and interfaces to help users easily and quickly find the information they need.

Search engines

The ServiceNow AI Platform includes two search engines:

AI Search

AI Search is the default search engine for the ServiceNow AI Platform. It provides consumer-grade search capabilities for Customer Service Management, Now Mobile, Service Portal, and Virtual Agent .

Zing text indexing and search engine

Zing is a legacy search engine which provides fast results and emphasizes relevancy scoring.

Search plugins

These search plugins offer specialized search experiences for particular search contexts.

Contextual Search

The Contextual Search plugin enables search that focuses on query context and user intent to fetch the most relevant set of results. Contextual search displays related results within a form or record producer based on the text you enter in a field.

Intelligent Search for CMDB

Intelligent Search for CMDB enables users to query for a set of Configuration Items in the CMDB. Complex queries open in CMDB Query Builder, where you can refine them or run them.

Search interfaces and widgets

These are search fields and widgets offered in the ServiceNow AI Platform UI or in ServiceNow AI Platform applications.

Global search

Global search in the Next Experience UI allows you to search multiple tables from a single search field.

Service Portal

The Service Portal framework displays search data within a widget on the search page.

Service Portal search widgets

Service Portal includes widgets for multiple forms of search.

Now Mobile app

The Now® Mobile app provides global search, voice search, and photo search.

Mobile Agent

The Mobile Agent app enables service desk agents to promptly manage and resolve end user issues from their mobile devices.

Virtual Agent

The Virtual Agent platform uses search to help parse user statements and intents in chats.

Agent assist

The Agent assist feature in Workspace automatically displays possible solutions for records that agents open.

Machine learning solutions for Search administration

Machine learning in Search administration helps to enhance business scalability and improve business operations for organizations.

Businesses face a challenge with growing volumes of data, in extracting meaningful information from a huge set of raw data, and in deriving meaningful business insights. Machine learning can serve as a solution to a variety of business complexities as its algorithm is built using historical data. It helps avoid duplicate and inaccurate data being entered into the database and enables businesses to compute and process information much faster.

Predictive Intelligence for Contextual Search

The Predictive Intelligence for Contextual Search capability uses machine-learning algorithms to search for Similar Incidents, Similar Knowledge Articles, Similar Open Incidents, and Similar Resolved Incidents.

Predictive Intelligence suggests related results to help you solve an issue or resolve a question.

The sources such as Similar Incidents, Similar Knowledge Articles, Similar Open Incidents, and Similar Resolved Incidents are available only when the Predictive Intelligence plugin (com.glide.platform_ml) and the Predictive Intelligence for Contextual Search plugin (com.snc.contextual_search_ml) are active.

For more information, see Predictive Intelligence.

Features of Search administration

Enable and configure Search administration features.

Feature	Description	Top tasks	Status
Zing text indexing and search engine	Index and search record data by table.	• Configure a table for indexing and searching • Zing can expand search results with synonyms • Zing can include attachments in search results	Active
Contextual search	Display search results on forms and record producers when users enter text in a field.	Define a search context	Active

Search Suggestions

The Search Suggestions application uses saved user search data to create auto-complete suggestions for AI Search applications and search suggestions for Zing. These suggestions appear as you enter terms into an application's search field. You can refine your search or jump to a relevant search result by selecting an applicable suggestion.

Get started

Explore	Configure
Learn about Search Suggestions concepts, workflow, and benefits.	Configure settings for creation of auto-complete suggestions and search suggestions from saved search results.

Use	Reference
View and select suggestions in AI Search applications and Zing. Analyze saved search data to learn how users interact with search.	Component and domain separation information for Search Suggestions.

Troubleshoot and get help

- https://www.servicenow.com/community/now-platform/ct-p/now-platform
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Search Suggestions

Learn how Search Suggestions displays suggestions to users in AI Search applications and the Zing text indexing and search engine. Review the rules that dictate which suggestions are offered in each search application and context.

Search Suggestions overview

The Search Suggestions application makes it possible for search applications to display suggested search completions or results to users as they enter terms into a search field.

Search Suggestions saves query terms and results from user searches. A Build Search Suggestions script runs against this saved data periodically to generate auto-complete suggestions for AI Search and search suggestions for Zing text search. The Prune Search Suggestions script runs periodically to remove low-relevancy suggestions.

Administrators can limit the age of saved search data used to generate suggestions. They can also configure specific suggestions to prevent them from being displayed to search application users.

Search Suggestions workflow

In this infographic, see a sample workflow of how different users in an organization configure and use auto-complete suggestions and search suggestions with the Search Suggestions application.

Configuring and using Search Suggestions

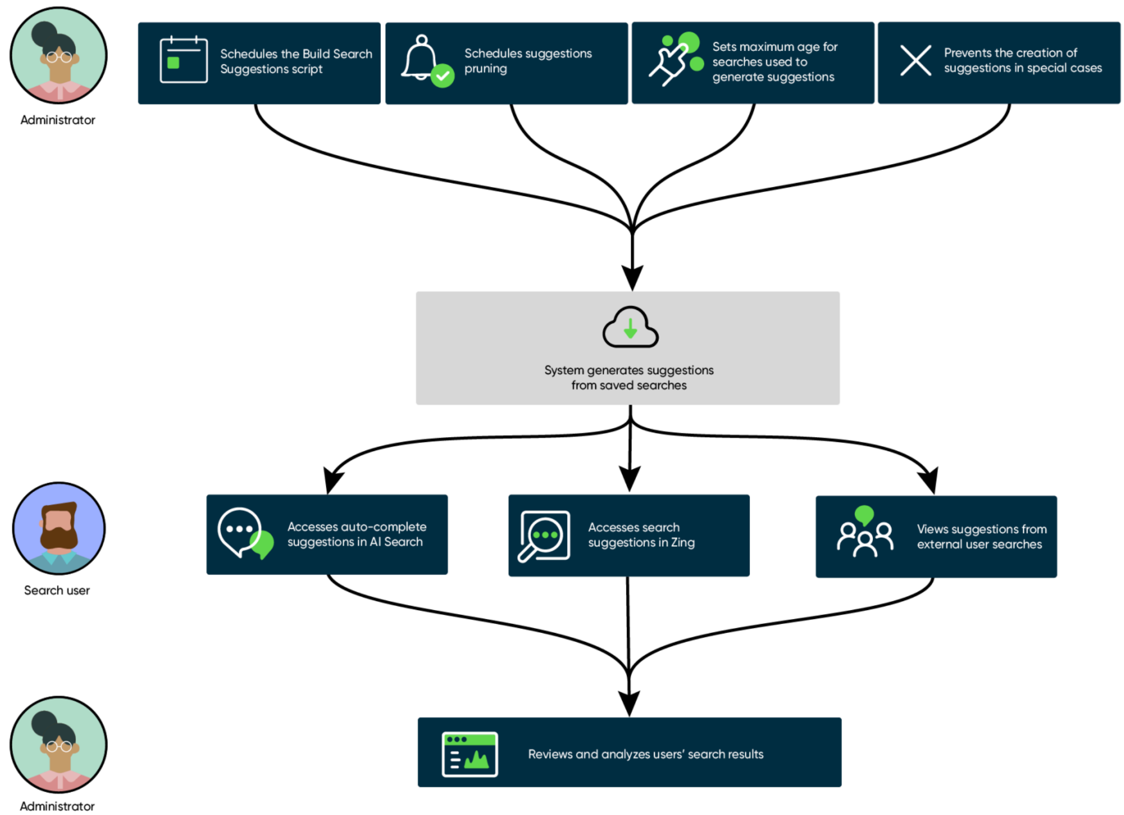

In this Search Suggestions workflow:

1.
   - Administrators configure Search Suggestions settings, including scheduling of scripts and restrictions on which saved search results the system uses to create auto-complete suggestions and search suggestions.
   - The system saves users' search results and uses them to generate new auto-complete suggestions and search suggestions.
2.
   - Search users view and select generated auto-complete suggestions in AI Search applications.
   - Search users view and select generated search suggestions in Zing.
   - Non-authenticated guest users access and review the list of generated suggestions created from searches by external users.
3. Administrators review and analyze saved search results to improve understanding of how users interact with generated suggestions, allowing for more effective tuning of Search Suggestions settings.

Search Suggestions benefits

Benefit	Feature	Users
View and select relevant auto-complete suggestions when searching in AI Search applications.	Auto-complete suggestions in AI Search applications	AI Search users

Benefit	Feature	Users
View and select relevant search suggestions when searching with the Zing text search engine.	Zing displays search suggestions as users enter search terms	Zing search users
See the list of auto-complete suggestions and search suggestions created from searches by external users.	Use the GraphQL REST API to view suggestions created from external user searches	Non-authenticated guest user
Control suggestion generation time by configuring how often the system builds auto-complete suggestions and search suggestions from user search strings.	Schedule the Build Search Suggestions script	Administrator
Improve suggestion relevancy by configuring how often the system prunes the least relevant auto-complete suggestions and search suggestions from the Search Suggestion [sys_search_suggestion] table.	Schedule suggestion pruning	Administrator
Increase suggestion freshness by using only the most recent search event records to create auto-complete suggestions and search suggestions.	Set maximum age for searches used in suggestion generation	Administrator
Prevent creation of unwanted auto-complete suggestions and search suggestions.	Prevent the creation of suggestions in special cases	Administrator
Understand how users interact with search by analyzing data from the Search Suggestions tables.	Analyze search relevancy	Administrator

Configuring Search Suggestions

You can configure Search Suggestions settings to control the creation of auto-complete suggestions and search suggestions on your instance.

Configuration overview

Task	Description
Schedule the Build Search Suggestions script	Specify how often and at what time the system generates auto-complete suggestions and search suggestions from stored search strings. By default, this task runs daily.
Schedule suggestion pruning	Specify how often and at what time the system prunes the lowest-rated suggestions to limit the Search Suggestion [sys_search_suggestion] table to 500,000 suggestions. By default, this task runs weekly.
Set maximum age for searches used in suggestion generation	Specify the maximum age of search strings the system considers when creating auto-complete suggestions and search suggestions. By default, the system only

Task	Description
	generates suggestions using search strings from the last 180 days.
Prevent the creation of suggestions in special cases	Disable generation of suggestions from specific search strings to keep unhelpful terms from appearing in suggestions or to prevent disclosure of personal or secure information. As part of this task, you may want to test regular expression patterns in your Search Suggestion Exclusion List Rule entries.

Schedule the Build Search Suggestions script

Specify when to run the script that builds auto-complete suggestions and search suggestions from user search strings.

Before you begin

Role required: admin

About this task

Search Suggestions stores search strings entered by users in the Search Event [sys_search_event] table. The `Build Search Suggestions` script turns those search strings into auto-complete suggestions and search suggestions and stores them in the Search Suggestion [sys_search_suggestion] table.

By default, this script runs daily. You can customize how often and at what time the script runs. How often you run the script depends on how frequently you want to update suggestions for users. When you run the script might depend on performance considerations.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs** and select **Build Search Suggestions**.
2. In the Scheduled Search Execution form, edit how often the script runs in the **Run** field.
3. Edit the time of day the script runs in the **Time** field.
4. Select **Update**.
5. **Optional:** Select **Execute Now** to run the script immediately.

Schedule suggestion pruning

Search Suggestions limits the number of auto-complete suggestions and search suggestions stored in the Search Suggestion [sys_search_suggestion] table to 500,000. A periodic pruning job removes the least relevant suggestions and increases the overall relevancy of suggestions.

Before you begin

Role required: admin

About this task

By default, the pruning job runs once a week and limits the Search Suggestion [sys_search_suggestion] table to 500,000 suggestions. Once the table reaches that limit, the pruning job removes the lowest-rated auto-complete suggestions and search suggestions to keep the maximum number of suggestions at 500,000. Pruning the least relevant suggestions improves the relevancy of the suggestions.

How often you run the pruning job might depend on how frequently you generate suggestions from the Search Suggestion table, and how quickly the table exceeds 500,000 suggestions.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Open the **Prune Search Suggestions** job.
3. On the Scheduled Script Execution form, edit the following fields.

Field	Description
Run	Frequency of execution for the pruning job. To execute the job manually, select **On Demand** and select **Execute Now**.
Day	Day of the week to run the pruning job.
Time zone	Time zone to use when interpreting the **Time** field value.
Time	Time of day to run the pruning job. Specify hours (00–23), minutes (00–59), and seconds (00–59).
Repeat interval	Interval between executions of the pruning job when **Run** is set to `Periodically`. Specify days (01–07, where 01 is Sunday), hours (00–23), minutes (00–59), and seconds (00–59).
Conditional	Script to run when executing the pruning job.

Note:
The fields on the form change with the choice of the **Run** value.

4. Select **Update**.

Set maximum age for searches used in suggestion generation

Set the Auto Flush parameter to limit the age of search strings used to create auto-complete suggestions and search suggestions.

Before you begin

Role required: admin

About this task

The Search Event [sys_search_event] table contains all of the strings used in searches. Scripts use this table to generate search analytics, search suggestions, and auto-complete suggestions. A system property sets the maximum age for records preserved in the table. The default is 180 days, but you can change this number to preserve search events for a shorter or longer duration. Longer durations increase the size of the table and its disk consumption.

Procedure

1. Navigate to the Auto Flush [sys_auto_flush] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_auto_flush.list`.
   c. Press Enter.
2. Search on `*search` and open the **sys_search_event** record.
3. On the Auto Flush form, edit the **Age in seconds** value.
   The default is 15,552,000 seconds, which is 180 days.

Prevent the creation of suggestions in special cases

Disable generation of suggestions from specific search strings to keep unhelpful terms from appearing in suggestions or to prevent disclosure of personal or secure information.

Before you begin

Make sure you're familiar with Java's regular expression pattern syntax. To learn about regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class.

Role required: admin

About this task

Search Suggestions generates auto-complete suggestions and search suggestions from the search strings that users enter. You may not want to create suggestions from all search strings. You can prevent the generation of suggestions using one of the following approaches:

- Exclude words or regular expression patterns from auto-complete suggestions and search suggestions.

  For example, a user might search on `KB2938272`. This search string is too opaque to provide a meaningful suggestion. Or, someone might enter a social security number or phone number as search strings which shouldn't be surfaced as suggestions to anyone else. So, you might exclude all search strings that contain multiple numbers. Likewise, if someone uses an expletive in a search string, you may not want to surface that word in suggestions. The Search Suggestion Exclusion List [sys_search_suggestion_blacklist] table contains the list of excluded words and regular expression patterns.

  Note:
  In the base system, the Search Suggestion Exclusion List table contains excluded words for English and language-independent regular expression patterns. Search Suggestions inserts additional language-specific excluded words into this table when you activate any of the following languages: Chinese, Czech, Dutch, Finnish, French, German, Hebrew, Hungarian, Italian, Japanese, Norwegian, Polish, Portuguese, Russian, Spanish, Swedish, Thai, or Turkish.

- Assign roles to users that prevent the generation of suggestions from their search strings.

  There may be individuals whose searches shouldn't appear in search or auto-complete suggestions for privacy or security reasons. To prevent their searches from generating suggestions, assign them the suggestion_exempt role. This role is unnecessary in most situations. Search Suggestions never shows who performed the initial search. To prevent people from seeing auto-complete suggestions and search suggestions, assign them the cannot_read_suggestions role.

Procedure

1. Navigate to the Search Suggestion Exclusion List [sys_search_suggestion_blacklist] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_search_suggestion_blacklist.list`.
   c. Press Enter.
2. To make a word or regular expression pattern active or inactive:
   a. Select the term in the list that you want to activate or deactivate.
   b. In the Search Suggestion Exclusion List form that appears, select or clear the **Active** option.
3. To add a new excluded word or regular expression pattern:
   a. Select **New**.
   b. On the Search Suggestion Exclusion List form, fill in the fields.

Search Suggestion Exclusion List form

Field	Description
Phrase	Word or regular expression pattern to exclude.
Language	The language that the excluded word is in. Use the ISO 639-1 code for the language, such as "en" for English. Because regular expression patterns work for all languages, search suggestions ignores this field for regular expression patterns.
Description	Explanation of why you excluded the word or regular expression pattern if it's not obvious. This field is especially important for regular expression patterns.
Active	Option to remove the excluded word or regular expression pattern from suggestions.
Type	Option that specifies whether the **Phrase** field represents one word or a regular expression pattern.

Note:
If you use regular expression patterns, you should verify that they exclude the words you want. For more information, see Test regular expression patterns in Search Suggestion Exclusion List Rule entries.

4. To assign someone a role that prevents them from seeing auto-complete suggestions and search suggestions, or to prevent their search strings from becoming suggestions:
   a. Navigate to **All** > **User Administration** > **Users** and select the name of a user.
   b. In the Roles related list, select **Edit**.
   c. To prevent a user from seeing auto-complete suggestions and search suggestions, move **cannot_read_suggestions** from the Collections column to the Roles List column and select **Save**.
   d. To prevent a user's search strings from generating auto-complete suggestions and search suggestions, move **suggestion_exempt** from the Collections column to the Roles List column and select **Save**.

Test regular expression patterns in Search Suggestion Exclusion List Rule entries

Regular expression patterns are powerful and often require editing to get the correct behavior. When using regular expression patterns to exclude search strings, test the patterns thoroughly so as not to have unintended results.

Before you begin

Make sure you're familiar with Java's regular expression pattern syntax. To learn about regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class.

Role required: admin

About this task

The Exclusion List Rule and Search Suggestion Relations [m2m_blacklist_search_suggestion] table lists the suggestions eliminated from the Search Suggestion [sys_search_suggestion] table by entries in the Search Suggestion Exclusion List [sys_search_suggestion_blacklist] table.

Procedure

1. In your browser, navigate to `https://<instance name>.service-now.com/m2m_blacklist_search_suggestion_list`.
   A list of words excluded from search suggestions appears in the Exclusion List Rule and Search Suggestion Relations table. The Suggestion column shows the search string that was eliminated in the creation of suggestions.
2. Add a regular expression pattern to the exclusion list table.
   For more information, see Preventing suggestions in special cases.
3. In a search field, for example, on the ServiceNow® Service Portal, enter words that satisfy the regular expression pattern.
4. Run the script that builds the search suggestions.
   For more information, see Schedule the Build Search Suggestions script. Select **Execute Now** to run the script immediately.
5. In the m2m_blacklist_search_suggestion_list table, select the menu icon (≡) for the Exclusion List column heading, then select **Group By Exclusion List** to see the regular expression patterns and what they eliminated.
6. Review the suggestions in the table to see if the regular expression patterns removed the suggestions they should have and no more.
7. In the sys_search_suggestion_blacklist table, set an excluded regular expression pattern's **Active** field to **false**, then run the Build Search Suggestions script again to make sure that only excluded suggestions were removed.
8. Revise the regular expression patterns if necessary and repeat the procedure.

Using Search Suggestions

Search applications display auto-complete suggestions and search suggestions as you type in their search fields. Administrators can analyze data from Search Suggestions tables to learn how users interact with search. Guest users can view suggestions created from self-registered external user searches.

Using Search Suggestions in AI Search

AI Search applications display auto-complete suggestions created by the Build Search Suggestions script. Auto-complete suggestions appear as you enter terms into the application's search field. You can select an auto-complete suggestion to accept it as your search query.

To learn more about auto-complete suggestion usage and configuration in AI Search applications, see Auto-complete suggestions in AI Search applications.

Using Search Suggestions in Zing text indexing and search engine

Zing displays search suggestions created by the Build Search Suggestions script. Search suggestions appear as you enter terms into the Zing search field. You can select a search suggestion to accept it as your search.

For more details on how search suggestions appear in Zing, see Zing displays search suggestions as users enter search terms.

Analyze search relevancy

Analyze data from the Search Suggestions tables to understand how your users interact with search.

Before you begin

Role required: admin

About this task

Search Suggestions uses the following tables:

- Search Event [sys_search_event]: Stores all search strings and search results.
- Search Source Event [sys_search_source_event]: Stores search strings and search results and ties them to a data source, for example, the suggestion came from the knowledge base table rather than the catalog table.
- Search Suggestion [sys_search_suggestion]: Stores suggestions for searches.

For details on the Search Suggestions tables, see Search Suggestions tables.

Procedure

1. Display the Search Event [sys_search_event] table by navigating in your browser to `https://<instance name>.service-now.com/sys_search_event_list.do`.
2. Find search queries with **Has results** column values of **false**.
   A value of false means that someone searched but didn't get any search results. You can choose to create information, such as a Knowledge article, to provide search results in the future, create synonyms for the search words that would surface a search result, or not provide information if the question is irrelevant.
3. Average the numbers in the **Click rank** column.
   Click rank shows which search result you selected. If you click the first result in the list of search results, the click rank value is 1. If you click the 6th result listed, the click rank value is 6. If you click multiple search results, the click rank is the highest listing value. For example, if you click the first result, return to the search page and click the 6th result, return to the search page and click the 4th result and find an answer, the click rank is 6 (not 4). A click rank of 0 means that you didn't select any search result, or that you selected a Genius Result answer card in AI Search. Lower positive click rank values indicate that the top search results are the most relevant, so the goal is to reduce the average value for this field.
4. Determine the click-through rate by dividing the number of records where the click rank isn't zero by the total number of records.

   Higher click-through rate numbers are better. Click-through rate is a good indicator of how often users find relevant results, but it doesn't reveal how much effort it took to find those results. That's where click rank comes in. Click rank describes how much effort it took the user to find what they perceived to be a relevant result.

5. Average the numbers in the **Refinements** column.
   Refinement is the number of actions a user took to reduce the number of search results. Actions include sorting and filtering. More refinements means more effort to find relevant search results, so lower numbers are better.
6. Display the Search Source Event [sys_search_source_event] table by navigating in your browser to `https://<instance name>.service-now.com/sys_search_source_event_list.do`.
7. Find search queries with **Has results** column values of **false**.
   False means that someone searched but didn't get any search results from a particular source table, for example, the Knowledge Base table. You can choose to create information, such as a Knowledge article, on the source to provide search results in the future, create synonyms for the search words that would surface a search result in that source table, or not provide information if the question is irrelevant.

Use the GraphQL REST API to view suggestions created from external user searches

Submit a request to the GraphQL REST API endpoint to retrieve search suggestions created from searches by self-registered external users. This endpoint allows access by non-authenticated guest users.

Before you begin

Ensure that the following conditions are met in your instance:

- The Domain Separation plugin isn't activated. For details on this plugin, see Domain separation plugin.
- The Explicit Roles plugin (com.glide.explicit_roles) is activated. To learn more about this plugin, see Explicit Roles.
- External users with the snc_external role have submitted searches in search applications that generate search suggestions, such as ServiceNow® Service Portal, Now Mobile, or ServiceNow® Virtual Agent.
- The *`glide.search.suggestions.allow_guest_user`* system property is set to **true**. For details on setting a system property value, see Add a system property.

Obtain the sys_id of the search application that you want to retrieve external user search suggestions for.

Note:
Users with the search_application_admin, agent_workspace_user, or workspace_user roles can obtain this sys_id by copying it from the search application record in the Search Application Configuration [sys_search_context_config] table.

Download a third-party GraphQL client, such as the GraphiQL, Insomnia, or Postman client.

Role required: none

About this task

Search Suggestions creates suggestions for searches performed by self-registered external users. When the *`glide.search.suggestions.allow_guest_user`* system property is set to **true**, non-authenticated guest users can access the GraphQL REST API endpoint for your instance to retrieve external user search suggestions for a search application.

No authentication is required or enforced for this GraphQL REST API endpoint.

servicenow

Note:
Non-authenticated guest users can only retrieve search suggestions from the Global domain. If domain separation is enabled in your instance, guest users can't retrieve any search suggestions.

Procedure

1. In your third-party GraphQL client, configure a new request that uses the HTTP POST method.
2. Set the endpoint for the request to the GraphQL REST API endpoint for your instance, `https://<instance name>.service-now.com/api/now/graphql`.
3. Set the request body to the following GraphQL query, replacing `<search_application_sys_id>` with the sys_id of the search application that you want to retrieve external user search suggestions for:

```
{
  GlideSearch_Query {
    suggestions(searchContextConfigId:
 "<search_application_sys_id>", searchTerm: "") {
      term
      data {
        name
        records {
          type
          columns {
            label
            fieldName
            value
            displayValue
          }
        }
      }
    }
  }
}
```

Result

The system returns search suggestions from application searches by self-registered external users as a JSON object, as shown in this example:


```
{
  "data": {
    "GlideSearch_Query": {
      "suggestions": {
        "term": "",
        "data": [
          {
            "records": [
              {
                "type": "POPULAR_QUERY",
                "columns": [
                  {
                    "label": "name",
                    "displayValue": "automatic replies",
                    "value": "automatic replies",
                    "fieldName": "name"
                  }
```

```
                ]
              },
              {
                "type": "POPULAR_QUERY",
                "columns": [
                  {
                    "label": "name",
                    "displayValue": "contact group",
                    "value": "contact group",
                    "fieldName": "name"
                  }
                ]
              },
              {
                "type": "POPULAR_QUERY",
                "columns": [
                  {
                    "label": "name",
                    "displayValue": "tablet",
                    "value": "tablet",
                    "fieldName": "name"
                  }
                ]
              },
              {
                "type": "POPULAR_QUERY",
                "columns": [
                  {
                    "label": "name",
                    "displayValue": "laptop",
                    "value": "laptop",
                    "fieldName": "name"
                  }
                ]
              },
              {
                "type": "POPULAR_QUERY",
                "columns": [
                  {
                    "label": "name",
                    "displayValue": "new email account",
                    "value": "new email account",
                    "fieldName": "name"
                  }
                ]
              }
            ],
            "name": "Popular searches"
          }
        ]
      }
    }
  }
}
```

Search Suggestions reference

Reference topics provide additional information about components and domain separation behavior for Search Suggestions.

Installed with Search Suggestions

Several types of components are installed with activation of the Search Suggestions plugin, including tables, user roles, and scheduled jobs.

Roles installed

Role title [name]	Description	Contains roles
Search Suggestions blocked searcher [cannot_read_suggestions]	Doesn't see search or auto-complete suggestions displayed when searching.	None
Search Suggestions exempt searcher [suggestion_exempt]	Submits searches that the system ignores when generating search or auto-complete suggestions.	None

For more details on restricting generation of search suggestions, see Prevent the creation of suggestions in special cases.

Scheduled jobs installed

Scheduled job	Description
Build Search Suggestions	Converts searches stored in the Search Event [sys_search_event] table into auto-complete suggestions and search suggestions stored in the Search Suggestion [sys_search_suggestion] table.
Prune Search Suggestions	Removes the lowest-rated auto-complete suggestions and search suggestions from the Search Suggestion [sys_search_suggestion] table to keep the maximum number of suggestions at 500,000.

Tables installed

Table	Description
Search Event [sys_search_event]	Search strings and details about their associated results, including the search user, application, UI source, session, and language, and any user-specified click rank and refinements.
Search Event per source table	Correlations between search strings with search results from a specific table (data source). A search that queries multiple tables produces multiple Search Event per source records.

Table	Description
[sys_search_source_event]	
Search Suggestion [sys_search_suggestion]	auto-complete suggestions and search suggestions generated from user searches by the Build Search Suggestions scheduled job. The Prune Search Suggestions scheduled job removes the lowest-rated auto-complete suggestions and search suggestions from this table to keep its maximum size at 500,000 suggestion records.

To learn about individual columns on these tables, see Search Suggestions tables.

Search Suggestions tables

Search Suggestions uses tables to generate and track relevant search suggestions. You can review data from these tables to gain insight into search suggestions offered in your system.

Note:
Because the Search Suggestions tables are system tables, they are excluded from the Reporting module by default. To learn about enabling reporting on system tables, see Reporting on system tables.

Search Event table

The Search Event [sys_search_event] table stores search strings and their search results.

sys_search_event table

Column	Description
Search query	Search string that the user entered.
Has results	Boolean value that specifies whether the search did (true) or didn't (false) return results. For all false values, consider adding information so that the same search would return suggestions in the future.
Click rank	Click rank shows which search result you selected. If you click the first result in the list of search results, the click rank value is 1. If you click the 6th result listed, the click rank value is 6. If you click multiple search results, the click rank is the highest listing value. For example, if you click the first result, return to the search page and click the 6th result, return to the search page and click the 4th result and find an answer, the click rank is 6 (not 4). A click rank of 0 means that you didn't select any search result, or that you selected a Genius Result answer card in AI Search. Lower positive click rank values indicate that the top search results are the most relevant, so the goal is to reduce the average value for this field.
User	User who performed the search.
Application ID	Application where the search feature was used, for example, you did a search in the Service Portal. See **UI source**.
Language	Language of the query. Search only returns results that are in the same language as the search query.
Refinements	Number of actions that you took to reduce the number of search results. Actions include sorting and filtering. The goal is to reduce this number.
UI source	Specific page (mobile) or page (portal) where the search query came from. See **Application ID**.

sys_search_event table (continued)

Column	Description
Session	ID of the session, which specifies when the search happened.
AI Search Profile	Name of the search profile used for the search. Only applicable for searches made with the AI Search engine.
GR Configurations Executed	Shows the list of Genius Result configurations that the system evaluated for the search. Genius Result configurations in this list may or may not have displayed answer cards for the search. Only applicable for searches made with the AI Search engine.
Genius Result Displayed	Shows whether or not any Genius Result answer cards were displayed for the search. Only applicable for searches made with the AI Search engine.
GR Configurations with Result Displayed	Shows the list of evaluated Genius Result configurations that displayed answer cards for the search. A Genius Result configuration can appear multiple times in this list if it displayed multiple answer cards for the search. Only applicable for searches made with the AI Search engine.
Last Clicked Genius Result Position	Shows which Genius Result answer card you selected an action from. If you select an action from the first displayed Genius Result answer card, the position value is 1. If you select an action from the third answer card, the value is 3. A value of 0 means that you didn't select actions from any Genius Result answer card. Only applicable for searches made with the AI Search engine.
Genius Result Events	JSON document showing details of actions you selected from Genius Result answer cards displayed for the search. Only applicable for searches made with the AI Search engine.

Search Event per source table

The Search Event per source [sys_search_source_event] table stores records correlating search strings with results from a specific table, here called a data source. A search can query multiple tables. The sys_search_event table lists one record for the entire search, while the sys_search_source_event table lists one record for each data source (table) queried.

sys_search_source_event

Column	Description
Search query	Search string that you entered.
Source	System ID of the data source. For example, the system ID of the Knowledge Base might be c45067sample3c82eef3e. Suggestions from one source are used as suggestions for that source only. For example, a suggestion derived from the Service Portal wouldn't be used as a suggestion in Now Mobile.
Has results	Boolean value that specifies whether the source did (true) or didn't (false) supply search results. For all false values, consider adding information so that the same search query would return suggestions in the future.
Created	Time stamp of the search query.

Search Suggestion table

The Search Suggestion [sys_search_suggestion] table stores auto-complete suggestions and search suggestions. A scheduled job uses the sys_search_event table to build the content of the suggestion table. This job segregates generated suggestions by source. Suggestions generated from searches on the Service Portal, for example, are only returned for other Service Portal searches. These suggestions aren't returned for Now Mobile queries.

sys_search_suggestion

Column	Description
Application ID	Application where the search feature was used, for example, you did a search in the Service Portal.
User group	Specifies whether an internal or external user generated the search query. Suggestions for internal users aren't shared with external users and vice versa.
Language	Language of the query. Search only returns suggestions that are in the same language as the search query.
Score	Suggestion rank. The higher the score, the more relevant the suggestion. Once all suggestions are returned for a search query, the suggestions are sorted by **Score**, the highest listed first. Scores reflect how often you selected a suggestion. A popular suggestion implies that it answers the query.
Source	System ID of the data source. For example, the system ID of the Knowledge Base might be c45067sample3c82eef3e. Suggestions from one source are used as suggestions for that source only. For example, a suggestion derived from the Service Portal wouldn't be used as a suggestion in Now Mobile.

Domain separation and Search Suggestions

Domain separation is unsupported in the Search Suggestions 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: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

Search Suggestions domain separation overview

Domain separation provides complete data isolation for domain-specific users.

How domain separation works with Search Suggestions

By default, Search Suggestions results are limited to the specific domain of the logged-in user. Search and auto-complete suggestions aren't shared in parent or child domains.

Related topics

Domain separation

Search signals

Search stores signal data from user searches. The aggregated signal data helps improve search result relevancy in your system.

Signals record details of user searches, search refinements, and interactions with search results.

AI Search machine learning relevancy analyzes aggregated signal data and automatically tunes search relevancy scoring to reflect user behavior patterns. This tuning helps improve result relevancy for search users.

To learn more about AI Search machine learning relevancy, see Machine learning relevancy in AI Search.

Search signal tables

Search uses tables to store details from user searches, search refinements, and interactions with search results. These details are called search signals. Data from the search signal tables helps improve the relevancy of search results in your system.

Search signals refer to user, query, and similar data stored in columns on search suggestion tables. For a guide to search suggestion tables and their columns, see Search Suggestions tables.

Note:
Because the search signal tables are system tables, they are excluded from the Reporting module by default. To learn about enabling reporting on system tables, see Reporting on system tables.

Search Signal Event table

The Search Signal Event [sys_search_signal_event] table stores signal data produced when users submit a search.

sys_search_signal_event

Column	Description
Search event	Reference to the record for the search event on the Search Event [sys_search_event] table.
Browser	User browser agent string.
Latitude	Latitude coordinate of the user location in decimal degrees.
Longitude	Longitude coordinate of the user location in decimal degrees.
Matched Query Intents	List of query intents matched in the user search query.
First page results	JSON array of objects in which each node represents a single result returned in the first page of results for the search. The result order is preserved. Each object includes the unique record identifier (sys_id) of a result and the table where the result is located.

sys_search_signal_event (continued)

Column	Description
	Note: The first page results aren't updated when you refine a search.
Search Response Time(ms)	Time taken by the search, in milliseconds.
User Query Language	Language of the user search query.

Search Signal Result Event table

The Search Signal Result Event [sys_search_signal_result_event] table stores signal data for search results that users select.

sys_search_signal_result_event

Column	Description
Search query	Search query that led to the search result selection.
User	User that selected the search result.
Result DocID	Source table and record names for the selected search result.
Description	Short description for the selected search result.
Source Table	Source table for the selected search result.
AIS DocID	Source table name and record sys_id for the selected search result.
Result Type	Type of search result selected. For example, *REGULAR* for a regular search result or *GENIUS* for a Genius Result answer card.
Search Signal Event	Reference to the record for the search on the Search Signal Event [sys_search_signal_event] table.
Number	Record number for the selected search result.
Domain	Domain to which the selected search result belongs.
Use for Suggestions	Option indicating whether AI Search uses this result selection when computing auto-complete suggestions. For more details, see Auto-complete suggestions in AI Search applications
Source URL	URL for the selected search result.
Action Data	JSON array of objects representing the search result selection action.

Search Signal Result Event Action table

The Search Signal Result Event Action [sys_search_signal_result_event_action] table stores signal data for user interactions with search results.

sys_search_signal_result_event_action

Column	Description
Signal Type	The type of user interaction signal recorded for the search result. For example, `CLICK` indicates that the user selected a search result.
Signal Value	The value for the user interaction signal. For example, when the signal type is `CLICK`, the signal value is the numeric index for the search result that the user selected.
Search Signal Result Event	Reference to the record for the search result on the Search Signal Result Event [sys_search_signal_result_event] table.

Search Autocomplete Query Event table

The Search Autocomplete Query Event [sys_search_autocomplete_query_event] table stores signal data for auto-complete suggestions that users select from the search input field.

sys_search_autocomplete_query_event

Column	Description
Selected Query	The auto-complete suggested query that the user selected in the search input field.
Click Index	The numeric index of the auto-complete suggestion that the user selected in the search input field.
Latitude	Latitude coordinate of the user location in decimal degrees.
Longitude	Longitude coordinate of the user location in decimal degrees.
Search Profile	Reference to the search profile used for the search query.

Genius Result Event Action table

The Genius Result Event Action [sys_search_genius_result_event_action] table stores signal data for user interactions with Genius Result answer cards.

sys_search_genius_result_event_action

Column	Description
Action Id	The identifier for this Genius Result event action signal.

sys_search_genius_result_event_action (continued)

Column	Description
Action Type	The type of action taken from the Genius Result answer card.
Action Value	The value, if any, for the action taken from the Genius Result answer card.
Matched Genius Result Intents	List of NLU model intents matched by the action taken from the Genius Result answer card. **Note:** This field is only populated for Genius Result configurations with NLU triggers.
Card Type	The type of Genius Result answer card that the user interacted with.
Search Signal Result Event	Reference to the record for the Genius Result answer card on the Search Signal Result Event [sys_search_signal_result_event] table.

Zing text indexing and search engine

Index and search record data by table.

The ServiceNow® Zing text indexing and search engine is the default search engine used to search ServiceNow AI Platform® record data. Its engine is created entirely with ServiceNow code and implemented within the relational model.

The Zing text search index is split into multiple database shards for parallel querying. Its search engine logic yields fast results and emphasizes relevancy scoring. Relevancy scores include frequency and proximity weighting. Scores also include weighting for specific indexed fields such as title, short description, and metadata.

Zing provides type-ahead search suggestions and **Did you mean:** global search suggestions.

Administrators and users with the ts_admin role can configure the following Zing text search settings:

- match relevance
- relevancy field weights for each indexed table
- global stop words
- stop words for each index
- synonym dictionaries for each table
- tables with attachment indexing

Explore	Set up	Administer

• Washington DC • Features of Search administration • Features of Zing text indexing and search engine	• Set global text search properties • Zing can include attachments in search results • Zing can expand search results with synonyms	• Zing indexes words • Global search finds records from multiple tables • List search finds records from the current table
Use	Develop	Troubleshoot and get help
• Boolean operators allow conditional search results • Quotation marks allow exact phrase searches • Wildcard characters allow searching for patterns and variations	• Developer training • Developer documentation • Create synonym dictionaries	• Ask or answer questions in the Platform forum • Search the Known Error Portal for known error articles • Contact Customer Service and Support

Features of Zing text indexing and search engine

Enable and configure Zing text indexing and search engine features.

Feature	Description	Top tasks	State
Zing indexes words	Index documents by dividing them into words. Depending on the languages your instance supports, a word may be a single character such as a Chinese or Japanese pictogram or a sequence of characters separated by spaces such as with Latin, Arabic, and Pinyin languages.	• Configure a table for indexing and searching • Regenerate a text index for a table • Configure tables to use the Japanese tokenizer	Active
Zing computes document scores using three components	Compute document scores based on the frequency, sequence, and weight of search terms in the document.	• Score search terms by inverse document frequency (IDF) • Set the relative weight of a field	Active

Feature	Description	Top tasks	State
Global search finds records from multiple tables	Search multiple record types from a single search field.	• Add a search group for Core UI • Configure a table for indexing and searching	Active
Search results page in Core UI and UI15	Display global text search results for each table as Zing generates them.	• Add a search group for Core UI • Revert to the legacy global search UI	Active
List search finds records from the current table	Search records from a table list view.	• Configure a table for indexing and searching • Regenerate a text index for a table	Active
Zing can include attachments in search results	Search content from attachments on indexed tables. Display attachments for search results from the Knowledge [kb_knowledge] table.	• Index attachments on a table • Configure a table for indexing and searching	Active
Zing removes stop words from queries	Remove common words from search queries that don't produce meaningful results.	• Configure a global stop word • Configure a table-specific stop word	Active

Feature	Description	Top tasks	State
Zing matches derived words with stemming	Convert any multiple-character search keyword to its stem form to find derived versions of the word.	• System Localization properties • Activate a language	Active
Zing can expand search results with synonyms	Expand search results to include additional search terms. Requires configuration before use	• Enable search synonyms • Create synonym dictionaries	Inactive

Available search options

Zing offers users several search interfaces and search operators to conduct text searches.

Available search interfaces

Search interface	Description
Global search	Find records in multiple tables using the global search field in the Next Experience Unified Navigation.
Search a list	Find records in a list; search in a specific field (Go to), all fields (Search), or in a specific column.
Knowledge Base	Find knowledge articles.
Live feed	Filter, search, or sort messages in live feed.
UI pages	Create a custom UI page to search for records in a table.

Available search operators

Search operator	Description
Boolean operators	Construct complex searches using Boolean AND, OR, and NOT operators in global text searches.
Quotation marks	Wrap search terms in quotation marks to search for an exact phrase consisting of one or more words.
Wildcard characters	Expand search results to match any non-space character. Zing supports both single-character and multiple-character wildcards.

List search finds records from the current table

Search records from a table list view.

Indexed tables display the *for text* option in the list title bar, which searches all records for matching field values.

List search

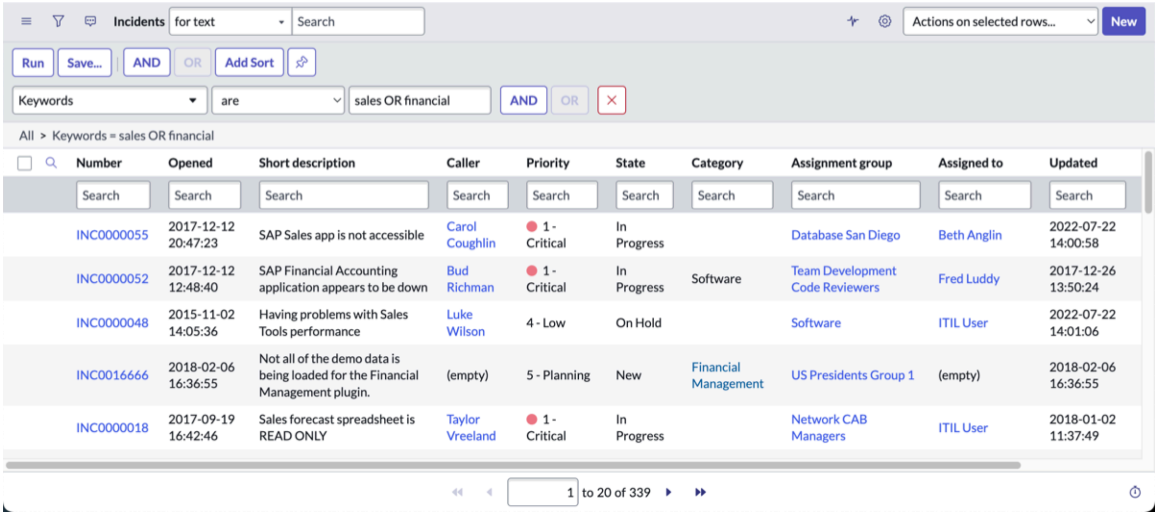

The list search field accepts Boolean operators (AND, OR, and NOT) in search queries. When a user adds a Boolean operator to a search query, the system only returns records that match all search conditions of the query.

The system also converts any search query into an equivalent keyword condition in the list breadcrumbs and filter. For example, searching for the text "Oracle OR SAP" produces the condition **[Keywords] [are] [Oracle OR SAP]**. The standard list controls can modify or remove these breadcrumbs and conditions.

Boolean operators allow conditional search results

Construct complex searches using Boolean AND, OR, and NOT operators in global text searches.

You can use Boolean operators in all global text search interfaces:

- Global text search field
- Knowledge Base search
- List search for text

Note:
Boolean operators are case-sensitive. As an example, the system recognizes OR as a Boolean operator, but treats `Or` and `or` as regular search terms.

Boolean operators

Operator	Description
OR or vertical bar symbol (|)	Finds a match if any of the search terms exist in a document (a union of sets). For example, to find documents that contain either `email password` or just `email`, you can search for `"email password" OR email` or `"email password"|email`. Using the `.or.` operator, which the system uses when you create OR condition groups for fields other than Keywords with the condition builder, can produce different results. For example, this filter searches for records with short

Boolean operators (continued)

Operator	Description
	descriptions containing SSO and then searches separately for records with short descriptions containing SAML: **Searching with the filter choice lists** 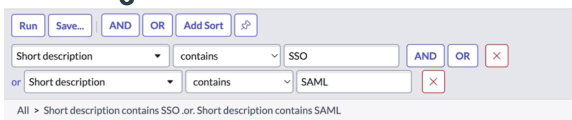 **Note:** By design, you can only have one Keywords filter condition per condition set. If you select the Keywords field in a filter choice list, the OR button for that filter condition becomes unavailable. Similarly, if you construct an OR filter condition, the Keywords field becomes unavailable in the filter choice lists.
AND	Finds a match if all search terms exist in a document (an intersection of sets). For example, to find documents that contain both "CPU load" and "10 minutes", search for `"CPU load" AND "10 minutes"` By default, Zing applies an AND operator between all search terms not separated by OR or NOT operators. For example, if you enter `email server down` as your search terms, the system searches for `email AND server AND down`. **Note:** For Knowledge Base searches, if the default AND search returns poor results, the search is automatically re-run with the OR operator. Administrators can configure the knowledge base search to always use OR by modifying the *`glide.knowman.search.operator`* system property.
NOT, minus (-), or exclamation point (!)	Excludes documents that contain the search term after the NOT operator (a difference of sets). For example, to find documents that contain "CPU load" but not "10 minutes", you can search for `"CPU load" NOT "10 minutes"` or `"CPU load" -"10 minutes"` or `"CPU load" !"10 minutes"`. NOT must be a stand-alone term in your search. For example, `atom NOT ion` excludes the term `ion`, but NOTION searches for the term `notion`. The minus and exclamation point symbols must immediately precede the search term that you want to exclude. For example, a search for `email !Joe` excludes `Joe` but a search for `email ! Joe` includes `Joe`. You can't use the NOT operator to exclude all terms in your search. For example, the following search returns no results: NOT `"10 minutes"`.

Related topics

Knowledge properties

Quotation marks allow exact phrase searches

Wrap search terms in quotation marks to search for an exact phrase consisting of one or more words.

Search terms in quotation marks are treated as exact phrase searches. When you search for an exact phrase, Zing only matches documents that contain the quoted words in the exact order you specify.

As an example, searching for `email password` returns search results that include `email` and `password` in any order, including records with these text strings:

- `enter your email password`
- `your email account password`

If you search for the quoted phrase `"email password"`, Zing returns search results with text `enter your email password` because they contain `email password` as an exact phrase. It doesn't return results with text `your email account password` because the search terms don't occur in the exact order specified in your search phrase.

Exact phrase searches ignore:

- stop words
- punctuation marks
- wildcard characters

Note:
Wildcard characters are ignored as punctuation.

For example, searching for the phrase `"email password"` returns the same search results as these phrases:

- `"email the password"` because the stop word `the` is ignored.
- `"email password?"` because the punctuation mark ? is ignored.
- `"email password*"` because the wildcard character `*` is ignored as punctuation.

Wildcard characters allow searching for patterns and variations

Expand search results to match any non-space character. Zing supports both single-character and multiple-character wildcards.

The following searches support Zing wildcard characters.

- Global text search
- Knowledge Base

  Note:
  Users with the admin role can enable wildcard search in the Knowledge context by navigating to **All** > **Contextual Search** > **Search Contexts**, editing the **Knowledge Base Search** record, and selecting the **Enable wildcard searches** check box.

- Search a list

  Note:
  To perform a wildcard text search of a list, you must select **for text** as the search field. If you select a list field instead of **for text**, wildcard characters in search terms will be interpreted as list search wildcards (for non-reference fields) or as auto-complete wildcards (for reference fields).

To perform a single-character wildcard search, use the percent sign (%) character. This wildcard finds words that contain any one character in place of the percent-sign-character. For example, to find words such as **text** or **test**, search for `te%t`.

To perform a multiple-character wildcard search, use the asterisk (`*`) character. This wildcard finds words that contain zero or more characters in place of the asterisk-character. For example, to find words such as **planned** or **placed**, search for `pl*d`.

You can use wildcard characters anywhere in a search string, and can include multiple wildcards in the same word. For example, to find words including **affordable**, **effort**, and **offering**, search for `%ff%r*`.

If a wildcard search returns too many hits, the system displays a message asking you to refine the search. Administrators can set a global text search property to change the limit for wildcard search results.

Note:
Wildcard text searches are only effective with multiple character words. A wildcard within a sequence of pictogram characters returns too many options for an efficient search. Wildcards are best used with letter-based words.

Wildcard characters and synonym matching

Zing doesn't expand the `%` and `*` wildcard characters when looking for synonyms that match your search term. These characters are treated as literal characters when finding matching synonyms. As an example, if you search for `issu*`, Zing doesn't find synonyms for `issue`, but it does find synonyms for `issu*`.

Enable or disable the Zing junk filter

When its junk filter is enabled, Zing doesn't index or search for 2-digit numbers or single-character words (unless they are Chinese or Japanese characters). You can enable or disable the junk filter for individual tables.

Before you begin

Role required: admin

About this task

For example, to disable matches for search terms like `a` or `40` in the Knowledge Base, turn on junk filtering for the kb_knowledge table.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the dictionary entry for the table.
3. In the Attributes related list, set the value of the **Text index filter junk** attribute:
   a. If you want Zing to include 2-digit numbers and single-character words when indexing and searching the table, set the attribute value to `false`.
   b. If you want Zing to ignore 2-digit numbers and single-character words when indexing and searching the table, set the attribute value to `true`.
4. Select the additional actions icon (≡) in the Dictionary Entry form header, then select **Save**.

5. Select the **Generate Text Index** related link and select **OK**.

Note:

- To have the desired effect, you may also need to reconfigure the automatic stop words. For example, if a specific 2-digit number appears in all of your problem numbers (92-0001, 92-0002, and so on), the automatic stop word threshold is quickly exceeded and the word is no longer found in searches. To find that 2-digit number in the problem table, you must disable both junk filtering and automatic stop words.
- Disabling the junk filter results in a larger table index. For optimal performance, don't apply this attribute unless it's required.

Debug Zing

Debugging messages create a log of search processes as they occur. These messages may help administrators configure optimal performance settings.

Before you begin

Role required: admin

About this task

- To view debugging messages, navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug Text Search**. The search log messages appear at the bottom of the results page for the remainder of your session.

Debugging Messages

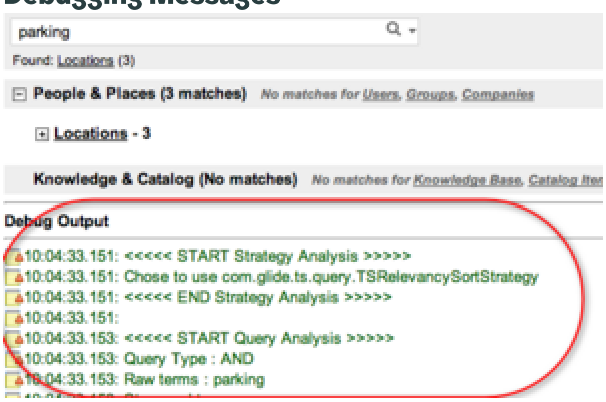

- To debug text indexing problems, you may want to view text_index events in the system log.
- In rare cases, you may need to reset the text search caches after regenerating a text index. If text search fails for known words, diagnose and fix this problem:

Procedure

1. Enable text search debugging messages.
2. In any table with search problems, search for the problem text.
3. In the debug output, note the numbers used for the Stemmed terms.
4. Navigate to the Ts Word [ts_word] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `ts_word.list`.
   c. Press Enter.

5. Locate the stemmed terms and compare the numbers in the list to the debug output from step 3.
   If the numbers are different, the UI node's cache is stale.
6. Navigate to **All** > **System Definition** > **Text Indexes**.
7. Select the **Reset Text Search Caches** related link and select **OK**.

Global search finds records from multiple tables

Search multiple record types from a single search field.

To perform a global search, perform one of the following actions based on your UI version.

UI version	Action
Next Experience UI	Enter your search terms in the Unified Navigation search field, then select **View results** or press Enter.
Core UI	Select the search icon in the banner on the upper right of the System settings for the user interface (UI) , then enter your search terms and press Enter.

All UI versions also have a global search keyboard shortcut. The `ui_global_text_search` UI Macro provides the global search field.

For a list of search source tables that the system searches by default, along with steps for adding new search source tables, see Default global search sources for Next Experience UI and Default global text search groups for Core UI.

Global search supports Boolean operators, quoted phrases, and wildcard characters. For details on these search features, see Boolean operators allow conditional search results, Quotation marks allow exact phrase searches, and Wildcard characters allow searching for patterns and variations.

Related topics

Keyboard shortcuts

Global search displays a page of matching results

Global search displays a page of results that match your search, grouped by table.

The elements included on the search results page differ by UI, as described in the following sections.

Search results page in Next Experience UI

Global search displays a Search Results page containing results that match your search terms, grouped by search source. Results appear when the search is complete.

Search Results page in Next Experience UI

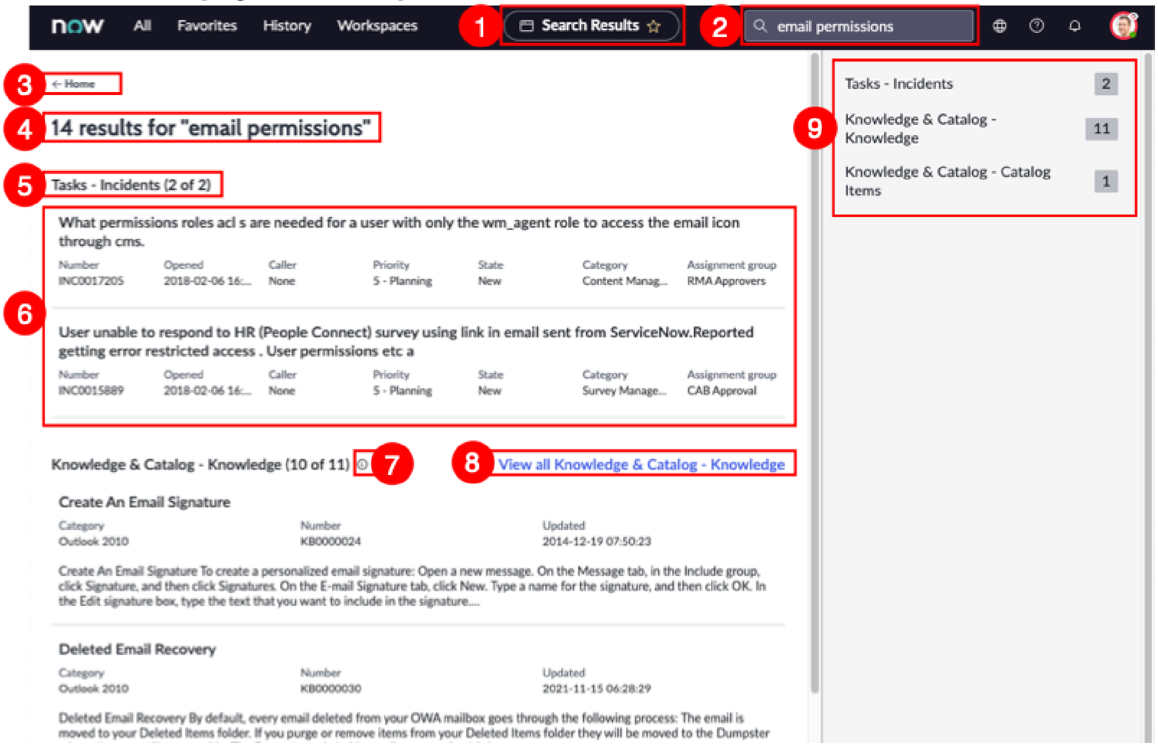

Zing displays global search results in a dedicated page consisting of these elements.

Elements of the global search results page in Next Experience UI

UI element	Description
1. Search Results page title	Contextual app pill for the search results page.
2. Global search field	Field displaying your current search terms. If you edit these terms, refresh the results page by selecting the **View results** link that appears in a panel below the search field.
3. Contextual return link	Link to return to the page you were on when you performed the search. The title of this link depends on which page you searched from.
4. Total result count	Total number of records matching your search.
5. Search source table name and result count	Name of a search source table and the number of previewed and total matching records in that table. Users with the admin or ui_builder_admin roles can hide the per-table counts of matching records. For details on making this change, see Hide per-table search result counts in Next Experience UI.
6. Preview of matching records for table	Partial list of matching records for the specified table, with maximum count controlled by the preview limit. By default, the system displays up to 10 matching records per table, sorted from highest to lowest document score. The system displays a separate list for each table in the search source.

Elements of the global search results page in Next Experience UI (continued)

UI element	Description
	By default, the system displays fields from the table's text_search list view for each search result. If this view doesn't exist for the table, it instead displays fields from the default list view. For more details on the text_search list view and the format of individual result listings, see Text search views control format of global search results for tables. Users with the admin role can override the preview results limit for a search application by modifying the application's record in the Workspace Global Search Configuration [sys_aw_global_search_config] table. For more details on this procedure, see Set the preview limit for global or workspace search in Next Experience UI. Administrators can also modify the search results table list view by modifying the value of the `glide.ui.text_search.view` system property. For details on changing this value, see Set global text search properties.
7. Table filter	Icon indicating that a filter was applied to results from this search source table. Pointing to the information icon displays a pop-up containing the search filter applied.
8. Link to all table results	Link to all matching records from the searched table. Select the link to launch a filtered list view based on the query. This link only displays when the total number of matching records in the table exceeds the preview limit.
9. Results summary	List of search source tables, showing the number of matching search results per table. Select the search source table name to display the associated preview list results. Search source tables with no matching results don't appear in this list. Users with the admin or ui_builder_admin roles can hide the per-table counts of matching records. For details on making this change, see Hide per-table search result counts in Next Experience UI.

View and open search results using workspace applications in Next Experience Unified Navigation

Use the search context menu to view and open results for your current search in workspace applications.

Before you begin

You must be using Next Experience UI.

Role required: none

About this task

In addition to global search, you may have access to workspace search applications, such as CSM/FSM Configurable Workspace. If you have access to multiple search applications, the search field's drop-down search context menu displays all available search applications, with the current application preselected. To view and open search results in another application, select it from the search context menu.

Users with the administrator role can add workspace applications to the search context menu. For details on this procedure, see Add a workspace application to the Unified Navigation search context menu.

Procedure

1. Use the search field to perform a search and view search results.
2. To open the search context menu, select the **Choose search context** icon ( ) in the search field.

   Note:
   If the **Choose search context** icon isn't displayed in the search field, the current search applications is the only one available. You can't switch search applications.

   A drop-down menu of available search applications displays, with the current search applications preselected.

   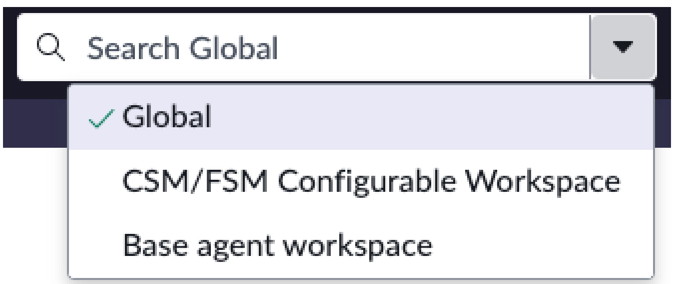

3. Select an alternate search application from the drop-down menu.
   The drop-down menu disappears and the **View results** link appears.
4. To view search results in the selected application, select **View results** or press Enter.

Result

The page reloads and displays results for your current search from the selected search application. If you select a previewed matching record, it opens in the selected search application.

Search results page in Core UI and UI15

Global text search displays a page of results that match your search terms. Search results appear dynamically on this page as Zing generates them.

Global text search results page in Core UI

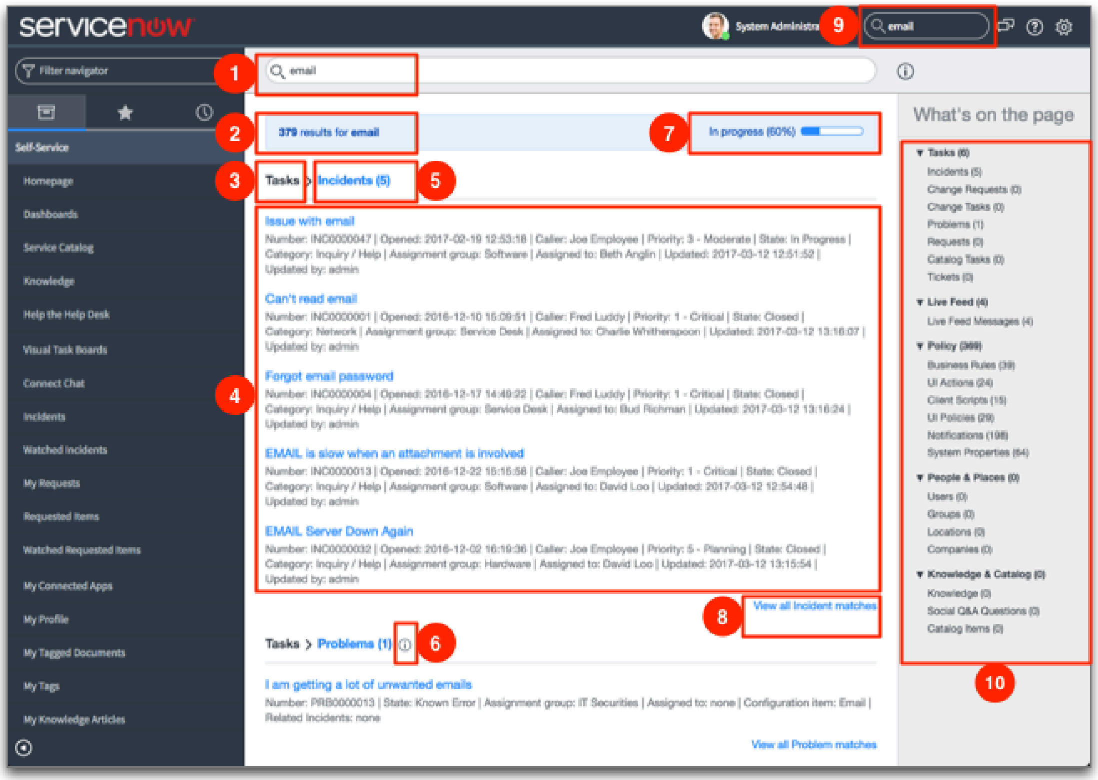

Zing displays global text search results in a dedicated page consisting of these elements.

Elements of the global text search results page in Core UI and UI15

UI element	Description
1. Search terms	Field displaying your current search terms. If you edit these terms, refresh the results page by selecting the search icon ( ) or pressing Enter. **Note:** If you delete your search terms and select the search icon or press Enter, the search results page doesn't refresh. To refresh the page, enter search terms, then select the search icon or press Enter.
2. Total results	Total number of records matching your search.
3. Search group name	Name of the search group that includes the table with matching records.
4. Preview of matching records	Partial list of matching records for the specified table, with maximum count controlled by the preview limit. By default, the system displays up to 10 matching records per table, sorted from highest to lowest document score. The system displays a separate list for each table in the search group.

Elements of the global text search results page in Core UI and UI15 (continued)

UI element	Description
	By default, the system uses the table's text_search list view if available. Otherwise, it displays fields with the default list view. Users with the admin role can modify the default preview results limit and search results table list view by changing the values of the *glide.ui.text_search.rowcount* and *glide.ui.text_search.view* system properties. For more details on these properties, see Set global text search properties.
5. Table name	Name of the search group table and the number of matching records in the table. Click the link to launch a filtered list view based on the query. Users with the admin role can hide matching record counts by setting the *sn_global_searchui.hide_results_count* property to **true** in the System Property table.
6. Table filter	Current filter applied to search results from this table. Pointing to the information icon displays a pop-up containing the search filter applied. For information on applying a table filter to a search group, see Add a search group for Core UI.
7. Progress bar	Current percentage of tables searched. The system hides this element after the search is complete.
8. Link to all table results	Link to see all matching records from the search group table. Click the link to launch a filtered list view based on the query.
9. Global search field	Global search field. **Note:** If you delete your search terms and select the search icon or press Enter, the search results page refreshes, displaying zero search results.
10. Results summary	List of search groups and tables with the number of search results per group and table. Select the table name to display the associated preview list results.

Text search views control format of global search results for tables

The text_search view for a table determines how global search displays results from that table.

Results displayed by global search include multiple elements. For each search result, the system populates these elements with values of fields included in the table's text_search list view.

Sample search result listing in Next Experience UI

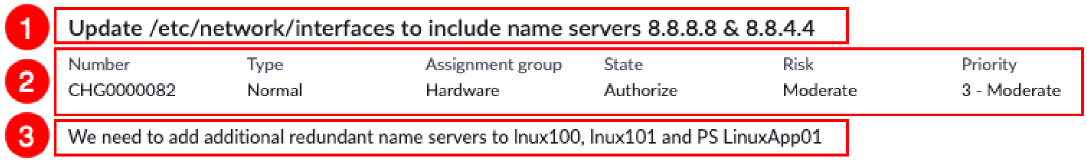

Sample search result listing in Core UI

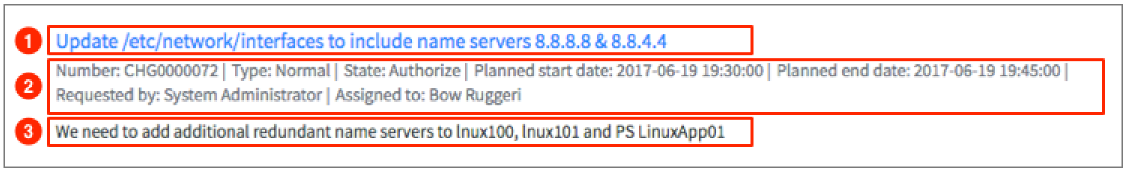

Search result elements

Element	Description	Matching criteria
1. Title	The title is the first line of the search result block. It links to the matching record of the search result.	Global search uses the first string field in the table's text_search list view that is not the number field. If there is no matching string field or the contents of the field is empty, the system displays the text "No title" in the session language. For example, in the default text_search view for the Incident table, the short description field satisfies the title criteria. Global search uses this field as the search result title.
2. Field list	The field list is the second line of the search results. It displays field name-value pairs, separated by the pipe character, for up to 10 fields.	Global search uses the first 10 fields in the table's text_search list view that don't match the title or description criteria. For example, in the default text_search view for the Incident table, the number, opened, caller, priority, state, category, and assignment group fields satisfy the field list criteria. Global search displays the names and values for these fields in the search result field list.
3. Description	The description is the third line of the search results. This line is optional and may not appear for all search results.	Global search uses the first string field in the table's text_search list view that is not the title and is over 100 characters in length. If no field matches these criteria, the search results omit this line. For example, in the default text_search view for the Incident table, the description field satisfies the description criteria. Global search uses this field as the search result description.

By default, the system uses the text_search list view for a table to format global text search results for that table. If a text_search view doesn't exist for the table, the system instead formats search results using the table's default form view. Implementers can

specify which view the system uses to format search results by modifying the value of the *glide.ui.text_search.view* system property.

Related topics

Create and delete views

Default display fields for global search tables

This table lists the tables enabled for global search in the base system. For each table, it shows the fields included in the table's default text_search list view. The system uses fields from this view to format global search results from these tables.

Table	Display fields
Catalog Item [sc_cat_item]	Short description, Description, Category, Price
Catalog Task [sc_task]	Short description, Description, Number, Priority, State, Assignment group, Request item
Change [change_request]	Short description, Description, Number, Type, Assignment group, State, Risk, Priority
Change Task [change_task]	Short description, Description, Number, Type, State, Assignment group
Company [core_company]	Name, Street, City, Zip / Postal code, Phone
Group [sys_user_group]	Name, Description, Type
Incident [incident]	Short description, Description, Number, Opened, Caller, Priority, State, Category, Assignment group
Knowledge [kb_knowledge]	Short description, Article body, Category, Number, Updated
Problem [problem]	Problem statement, Description, Number, State, Resolution code, Assignment group, Configuration item, Related Incidents
Request [sc_request]	Short description, Number, Requested for, State, Opened by
Requested Item [sc_req_item]	Number, Item, Stage, State, Requested for, Opened by, Opened
User [sys_user]	Name, Email, Company, Department, Title

Search settings filter and group global search results by table

Global search filters and groups matching results by table. You can configure search sources (in Next Experience UI) or search groups to define tables and groups for global search.

The configuration settings for table-based filtering and grouping of global search results differ by UI, as described in the following sections.

Search sources filter global search results by table in Next Experience UI

Global search in Next Experience UI displays search results broken out by search source. A search source is a collection of global search settings for a ServiceNow AI Platform® table.

A user must have read access to a table to see search results for it.

Default global search sources for Next Experience UI

The base system includes these default search sources for global search in Next Experience UI.

When Next Experience UI is activated, the following search sources are defined for global search. These search sources appear in the Application Search Sources related list for the **Now Experience Search Configuration** record in the Search Application Configuration [sys_search_context_config] table.

Search source	Table searched
Tasks - Incidents	Incident [incident]
Tasks - Change Requests	Change Request [change_request]
Tasks - Change Tasks	Change Task [change_task]
Tasks - Problems	Problem [problem]
Tasks - Requests	Request [sc_request]
Tasks - Catalog Tasks	Catalog Task [sc_task]
People & Places - Users	User [sys_user]
Tasks - Requested Items	Requested Item [sc_req_item]
People & Places - Groups	Group [sys_user_group]
Knowledge & Catalog - Knowledge	Knowledge [kb_knowledge]
People & Places - Companies	Company [core_company]
Knowledge & Catalog - Catalog Items	Catalog Item [sc_cat_item]

Note:
Search sources only affect filtering of global search results in Next Experience UI. In Core UI, global text search filters results based on settings defined in text search groups. For details on these settings, see Search groups filter global text search results by table in Core UI.

Add a search source for Next Experience UI

Add a new search source to make a table searchable by global search or workspace search in Next Experience UI.

Before you begin

You must be using Next Experience UI.

Note:
You can configure tables for global text search in Core UI by adding new text search groups. For more details, see Add a search group for Core UI.

The tables you want to include in the search group must have text indexing enabled. To enable text indexing for a table, see Configure a table for indexing and searching.

Role required: search_application_admin or workspace_admin

About this task

A search source defines a table that is searchable by global or workspace search in Next Experience UI. Define new search sources to add to the default set of tables configured for global search or your workspace search. For the list of global search sources included in the base system, see Default global search sources for Next Experience UI.

Procedure

1. Define the new search source in the Search Source [sys_search_source] table.

   a. Select **All**.

   b. In the **Filter** field, enter `sys_search_source.list`.

   c. Press Enter.

   d. Select **New**.

   e. On the Search Source form, fill in the fields.

Field	Description
Name	Name to display for the search source in the global search results page.
Table	Table to make searchable for global search.
Conditions	Filter conditions that table records must satisfy to be searchable.
Application	Application scope for the search source. This field is automatically set.

   f. Select **Submit**.

2. Link the new search source to the record for the global or workspace search application in the Search Application Configurations [m2m_search_context_config_search_source] table.

   a. Select **All**.

   b. In the **Filter** field, enter `m2m_search_context_config_search_source.list`.

   c. Press Enter.

   d. Select **New**.

   e. Select the additional actions icon (≡), then select **Configure** > **Form Layout**.

   f. Use the slushbucket to add the Search context config field to the selected field list, then select **Save**.

   g. On the Application Search Sources form, fill in the fields.

Field	Description
Source	Name to display for the search source in the global search results page.

Field	Description
Order	Order of appearance for the search source in the search application configuration context.
Application	Application scope for the search source. This field is automatically set.
Search context config	Record for the global or workspace search application in the Search Application Configuration [sys_search_context_config] table. To add the search source for global search, select the **Now Experience Search Configuration** record. To add the search source for a workspace application, select the record for that application. For example, to add a search source for Agent Workspace, select the **Agent Workspace Search Config** record.

h. Select **Submit**.

3. **Optional:** Reorder search sources for the global or workspace search application.

   a. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.

   b. Open the record for the search application.

   Note:
   For global search, open the **Now Experience Search Configuration** record. For a workspace application, open the record for that application, such as the **Agent Workspace Search Config** record for Agent Workspace.

   c. In the Application Search Sources related list, edit the **Order** values for search source records to alter the display order for table filters on the application's search results page.

Result
The new search source appears in global or workspace search results.

Search groups filter global text search results by table in Core UI

Global text search in Core UI displays search results broken out by search group. A search group is a collection of related searchable tables.

Within each search group, the system divides search results by table. A user must have read access to a table to see search results for it.

Displaying search results by group

The system displays knowledge and service catalog search results within their respective search results pages. The system displays other results in list format.

For exact match searches, the system only returns records from the Task (task) and Knowledge (kb_knowledge) tables.

Note:
The system filters global search results to display only knowledge articles that are in the same language as the preferred language setting.

Example global text search results

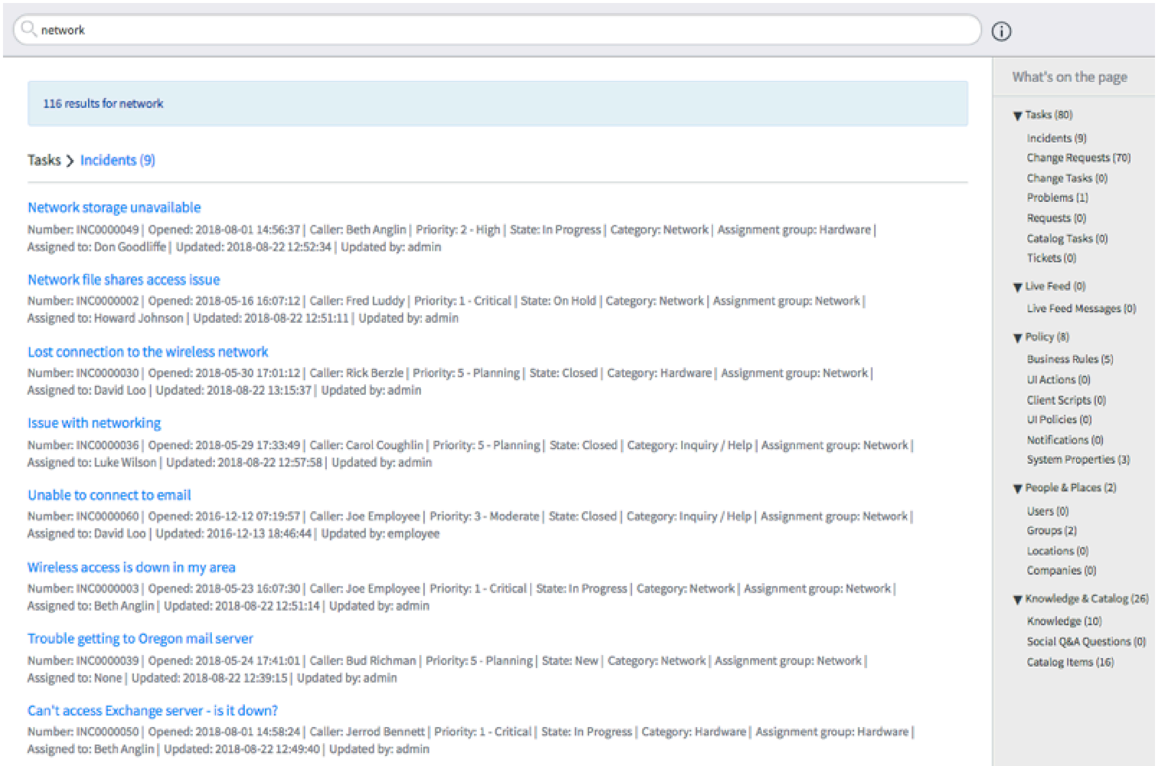

Options for administrators

Administrators can customize these search group settings for all users:

- Specifying which tables are part of the group.
- Specifying whether the search group is available for use.
- Specifying whether searches include results from this search group by default.
- Creating new search groups.

Enabling a search group allows users to see records from the search group's tables. Disabling a search group hides records from the search group's tables.

Default global text search groups for Core UI

The base system includes these default text search groups for global text search in Core UI.

Search group	Tables searched
Tasks	• Incident [incident] • Change Request [change_request] • Change Task [change_task] • Problem [problem]

Search group	Tables searched
	• Request [sc_request] • Catalog Task [sc_task] • Requested Item [sc_req_item]
People & Places	• User [sys_user] • Group [sys_user_group] • Company [core_company]
Knowledge & Catalog	• Knowledge [kb_knowledge] • Catalog Item [sc_cat_item]

Note:
Text search groups only affect filtering of global text search results in Core UI and UI15. In Next Experience UI, global search filters results based on search sources linked to the **Now Experience Search Configuration** record in the Search Application Configuration [sys_search_context_config] table. For details on the search sources included in the base system, see Default global search sources for Next Experience UI.

Add a search group for Core UI

Add a new text search group to filter global text search results by table in Core UI.

Before you begin

You must be using Core UI.

Note:
You can configure tables for global search in Next Experience UI by adding new search sources. For more details, see Add a search source for Next Experience UI.

The tables you want to include in the search group must have text indexing enabled. To enable text indexing for a table, see Configure a table for indexing and searching.

Role required: text_search_admin

About this task

A text search group defines a set of tables to use as filters for global text search results in Core UI. Define new search groups to add to the default set of text search groups configured for global text search.

For the list of text search groups included in the base system, see Default global text search groups for Core UI.

Procedure

1. Navigate to **All** > **System Definition** > **Search Groups**.
2. Select **New**.
3. On the Text Search Groups form, fill in the fields.

Field	Description
Name	Enter a name for the search group.
Roles	Select one or more user roles if you want to restrict access to the search group to those roles.
User group	Select a user group if you want to restrict access to the search group to users in that user group.
Description	Enter text to display as a tooltip for the search group.
Active	Ensure the check box is selected to make the search group available.
Searched	Clear the check box if you want to exclude results from this search group.
Order	Enter a number to determine the order in which this search group is listed relative to other search groups.

4. Select the additional actions icon (≡), then select **Save**.
5. For each table that you want to add to the search group, perform the following steps.

   a. In the search group record's Text Search Tables related list, select **New**.

   b. On the Text Search Table form, fill in the fields.

   Text search table

Field	Description
Search group	The search group that you want to add a table to.
Table	Select the table to include in the search group. **Note:** You can select any table not excluded by the `glide.ui.no_text_search` system property, but global search only displays search results for tables that have text indexing enabled.
Application	Application scope for the text search table record. This field is automatically set.
Active	Option to search the table as part of the search group.
Searched	If you want to exclude results from this table by default, clear this option. A user can manually choose to show results from this table on the search results page.
Order	Order in which this table is listed relative to other tables in the search group.
Conditions	Select one or more conditions to determine which records on the selected **Table** are searched.

   c. Select **Submit**.

Global search displays exact matching records

When your search exactly matches the **Number** field value for a record, global search takes you directly to that record, bypassing the search results page.

As an example, if you search for `INC0038278`, global search shows Incident record INC0038278 in form view instead of showing the search results page.

> **Note:**
> Exact match respects security settings. If you don't have permission to view the record that matches your search, global search won't show you that record.

If more than one record has a **Number** field value that exactly matches your search, global search displays the first exact match. An informational message reports the total number of exact matches and provides links to the other exact matches.

Exact match configuration settings

Exact match only returns matches for Number values that contain prefixes defined in the Numbers [sys_number] table. To view prefixes and their matching tables, navigate to **All** > **System Definition** > **Number Maintenance**.

Search behavior for exact match is controlled by the *`com.snc.agent_workspace.global_search.typeahead.exact_match_request_criterio`* system property. This property's value is a Java regular expression pattern. Only search queries that match the regular expression pattern are evaluated as possible exact record matches. For details on Java regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class.

Exact search matches in Next Experience UI

When your search in Next Experience UI exactly matches a searchable record, the Unified Navigation search field displays a preview of that record. Selecting the preview takes you directly to that record and bypasses the search results page.

Exact search match in Next Experience

The preview of the matching record appears below the Unified Navigation search field with an **Exact Match** tag. Select this tag or press Enter to navigate directly to the record.

To see the full search results page for your search, select **View results**. For details on the elements that make up the search results page, see Search results page in Next Experience UI.

Note:
You may need to wait a few seconds after entering your record number into the search field for the system to process your search request and display the preview. Once the preview appears, you can press Enter to go directly to the matched record. If you press Enter before the preview appears, the system shows you the full search results page for your search instead.

Exact search matches in Core UI

When your search in Core UI exactly matches a searchable record, global search takes you directly to that record, bypassing the search results page.

Exact search match in Core UI

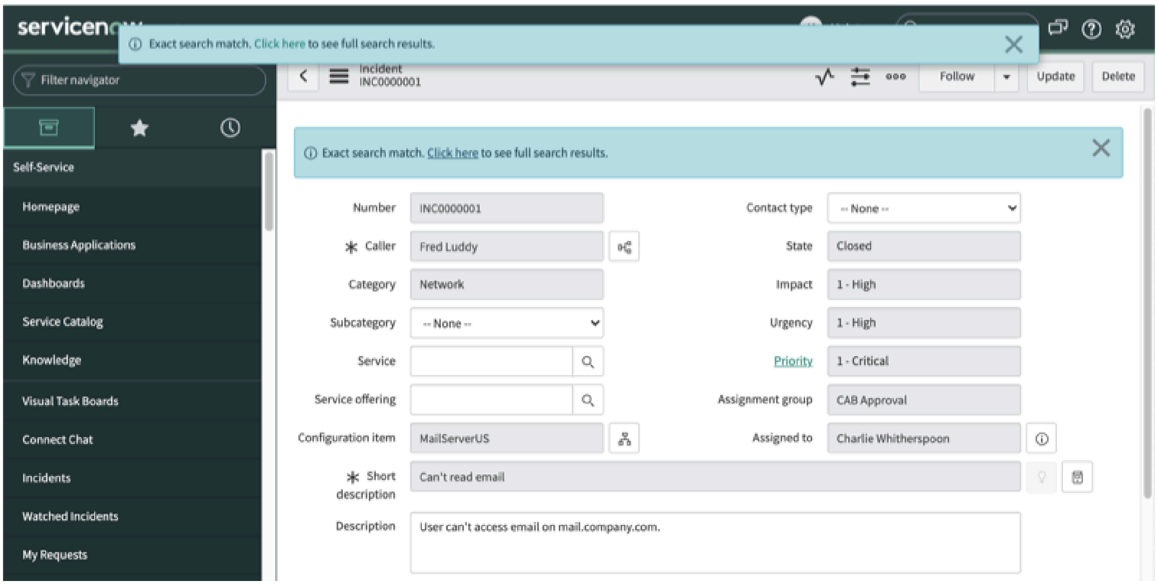

The system displays the matching record with an informational message indicating that your search produced an exact search match.

To see the full search results page for your search, select the link in the informational message. For details on the elements that make up the search results page, see Search results page in Core UI and UI15.

Global search displays your most recent search queries and results in Next Experience UI

When you select the empty Unified Navigation search field, the system displays lists showing your most recent search queries and your most recently viewed search results. Select a **Recently Searched** query to repeat it, or select a **Recently Viewed** search result record to navigate to it.

The **Recently Searched** and **Recently Viewed** lists appear below the Unified Navigation search field.

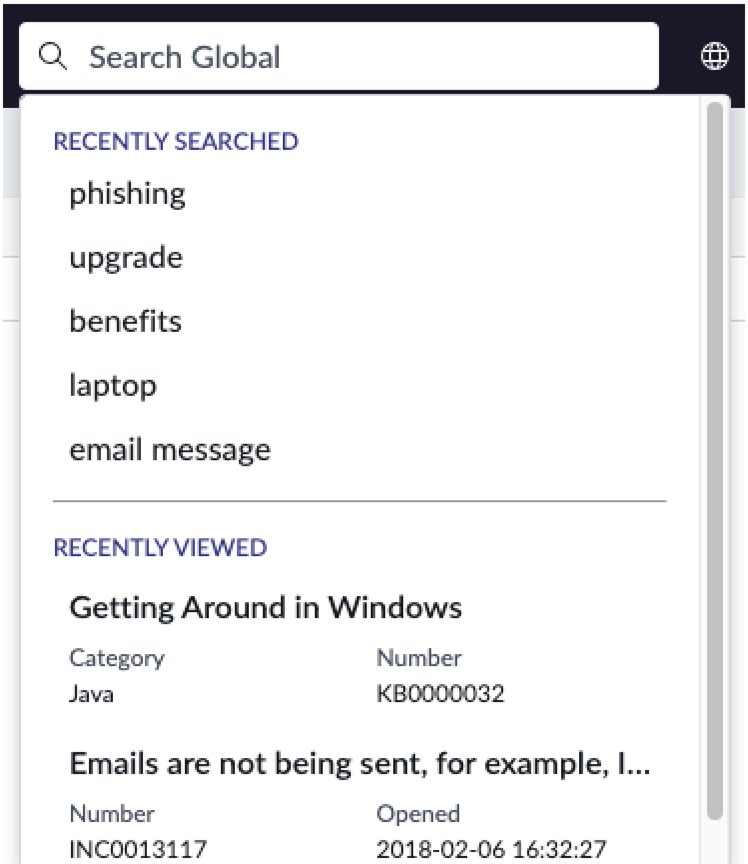

The **Recently Searched** list shows the most recent search queries from your personal search history. This list only appears when your search history includes previous search queries.

The **Recently Viewed** list shows the most recently viewed search results from your personal search history. This list only appears when your search history includes previously viewed search results.

Set global text search properties

Administrators can control how global text search behaves with system properties.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **Global Text Search**.
2. Review the values for the following global text search properties and change them as needed.

UI properties for global text search

Property	Description
Enable Global Text Search (`glide.ui.text_search.enable`)	Option to enable global search in Core UI and UI15. ◦ Type: Boolean ◦ Default value: `true` ◦ Supported values:

servicenow

Property	Description
	▪ *true*: Global search is enabled in Core UI. ▪ *false*: Global search is disabled in Core UI. **Note:** When Next Experience is enabled, this system property has no effect. To enable or disable the Unified Navigation search field in Next Experience, set the *glide.ui.polaris.global_search* system property described in Next Experience system properties.
List of roles (comma-separated) that can use the Global Text Search capability (*glide.ui.can_search*)	Comma-separated list of user roles that allow access to global search. To enable global search for users with a role, add the role to this list. ○ Type: String ○ Default value: *itil, text_search_admin, admin*
List of tables (comma-separated) that won't appear as options when setting up Global Text Search tables (*glide.ui.no_text_search*)	Comma-separated list of tables that don't appear as options when defining text search groups for Core UI. Examples are system or maintenance tables. ○ Type: String ○ Default value: *sys_audit, sys_event, ecc_queue, ecc_event, syslog, syslog_transaction, sys_journal_field, sys_audit_relation, ecc_agent_log, ecc_mi_result, sysrule_escalate_history, sys_user_token, sys_time_dimension, sys_attachment_doc, sys_audit_delete*
Number of Global Text Search matches returned per table (*glide.ui.text_search.rowcount*)	Maximum number of matching records previewed for each table on the global text search results page. Users can select a table to see all matching results for that table. ○ Type: Integer ○ Default value: 10 **Note:** This system property only controls the global text search preview limit in Core UI. In Next Experience UI, you can set the global search preview limit by editing the **Now Experience Search Configuration** record in the Workspace Global Search Configurations [sys_aw_global_search_config] table.
Global Text Search background color for Knowledge Base results (*css.textsearch.kb.background.color*)	Background color for Knowledge article global text search results. This property only affects display of results on the legacy global search results page. ○ Type: String ○ Default value: *#f0f7f9*

Property	Description
Global Text Search background color for Catalog results (*css.textsearch.catalog.background.color*)	Background color for Catalog global text search results. This property only affects display of results on the legacy global search results page. ◦ Type: String ◦ Default value: *#ffffdd*
Form view to use for Global Text Search exact match - blank means Default view (*glide.ui.text_search.match_view*)	Name of the form view that global text search uses to display exact match results. If no value is specified, global text search uses the default form view. For information on creating and managing views, see View management. ◦ Type: String ◦ Default value: (empty)
List and form view to use for Global Text Search results. Blank means default view (*glide.ui.text_search.view*)	Name of the list view that global text search uses to format results for each searched table. If no value is specified, global text search uses the default list view. For information on creating and managing views, see View management. ◦ Type: String ◦ Default value: *text_search*
Query method for global text search. Run a query for each table, for each base class (parent table), against a single index regardless of table (requires text index regeneration), or let GlideRecord handle the query (simple). With the Zing text search engine, the preferred value is "simple" (*glide.ir.query_method*)	Query method to use for global text search. **Note:** The *base*, *single*, and *table* values aren't supported. Leave this system property set to its default *simple* value. ◦ Type: Choice list ◦ Default value: *simple* ◦ Available values: ▪ *base*: Run a search query for each base class. This value is no longer supported. ▪ *simple*: Let GlideRecord handle the search query. ▪ *single*: Run a search query against a single text search index. To use this method, you must regenerate text search indexes. This value is no longer supported. ▪ *table*: Run a search query for each searchable table. This value is no longer supported.
The maximum number of results that will be returned during a text index search. This limits	Maximum number of results that global text search will return when searching a text search index.

Property	Description
the number of records that will be extracted from the text index and displayed (*glide.lucene.base_result_max*)	**Note:** When the query method is set to *simple*, this property is ignored. ◦ Type: Integer ◦ Default value: 100
The maximum number of entries that will be retrieved from the text index. Only the table that is involved will be returned from the text index for all entries that are between the value specified for maximum results and this value. This is done so that it can be indicated that more results exists for any given table that are going to be displayed (*glide.lucene.base_hits_max*)	Maximum number of entries that global text search will retrieve from a text search index. **Note:** When the query method is set to *simple*, this property is ignored. ◦ Type: Integer ◦ Default value: 500
Number of simultaneous processes (1 to 16) used when searching though multiple groups in a global search (*glide.ts.global_search.parallelism*)	Number of concurrent threads global text search uses when searching multiple search groups. Each search group requires one thread. Set this value to yield optimal results for your search group configuration. For example, if you have five search groups and four threads, the first four groups run in parallel and the fifth group starts when one of the first four groups finishes. This setup may work well if one of the groups is much larger than another. Similarly, if you have five search groups, setting this value higher than five yields no benefits. ◦ Type: Integer ◦ Default value: 4

3. Select **Save**.
   The business rule **Text Search Property Change Rationally** validates the new values and aborts the update if they are outside its acceptable ranges.
4. To set the maximum number of matches to display for a wildcard search term, follow these steps.

   a. Open the System Properties table and search for the property **glide.ts.max_wildcard_expansion**.

Field	Description
Name	glide.ts.max_wildcard_expansion
Description	Maximum number of matches for a wildcard search term that are allowed as unambiguous.
Type	Integer

Field	Description
Default value	500

b. Edit the **Value**.
If a wildcard search matches more words than this value, the system displays a message asking the user to refine the search.

c. Select **Update**.

5. To set the Zing text search stemming language, follow these steps.

a. Install the internationalization (i18n) plugin for the language you want to stem text search terms in.

b. Add a new system property, *glide.ts.stemming_language*:

Field	Description
Name	glide.ts.stemming_language
Description	Language to use for text search stemming. For details on stemming, see Zing matches derived words with stemming.
Type	String
Values supported	▪ English ▪ French ▪ German
Default value	English

c. To make tables use the new stemming language for text search, manually regenerate their text search indexes.
For details on this procedure, see Regenerate a text index for a table.

6. **Optional:** To enable and configure fallback number search for exact matches, follow these steps.

Note:
Fallback number search is disabled by default. When it's enabled, if a global text search finds no exact match for a record number, the system searches a list of specified tables for matches. This additional search can make global text searches take longer, so enable fallback number search only if you need it.

a. To enable fallback number search, add a new system property, *glide.ui.text_search.enable_fallback_number_search*:

Field	Description
Name	glide.ui.text_search.enable_fallback_number_search

Field	Description
Description	Option to enable fallback number search for use when global text search finds no exact matches for a record number. **Note:** Fallback number search can make global text searches take longer. Enable the option only if you need it.
Type	Boolean
Default value	false

b. To configure the list of tables searched by fallback number search, add a new system property, *glide.ui.text_search.fallback_table_list*:

Field	Description
Name	glide.ui.text_search.fallback_table_list
Description	(Optional) Comma-separated list of names of tables to search when performing fallback number search. Use the special value *all* to search all tables that have prefixes defined in the Number [sys_number] table. **Note:** Increasing the number of tables to search may increase the time required for fallback number search. Test any changes to this system property before putting them into production.
Type	String
Default value	task,kb_knowledge

Configure parallel processing of search groups

To improve performance, only activate search groups and tables that are necessary to meet business needs.

Before you begin

Role required: admin

About this task

For example, if you don't need Change Task results, deactivate that table in the Tasks search group. If only one group of users needs Change Task results, set up a separate search group that includes Change Tasks. Other users search using a group that doesn't contain Change Tasks.

Global text search can render results in parallel to improve performance. To configure the number of parallel processes:

Procedure

1. Navigate to **All** > **System Properties** > **Global Text Search**.
2. Locate the property called **Number of simultaneous processes(1 to 16) used when searching though multiple groups in a global search**.
3. Enter the number of processes to run in parallel.
   Each search group uses one thread to render results. Set this value to yield optimal results for your search group configuration. For example, if you have five search groups and four threads, the first four groups run in parallel and the fifth group starts when one of the first four groups finishes. This setup may work well if one of the groups is much larger than another. Similarly, if you have five search groups, setting this value higher than five yields no benefits.
4. Select **Save**.

Revert to the legacy global search UI

Disable the dynamic global search UI to preserve customizations you made to the legacy global search UI.

Before you begin

You must have Next Experience disabled on your instance. The legacy global search UI is only available in Core UI.

Role required: admin

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Select **New**.
3. Enter these field values.

Field	Value
Name	glide.ui.use_legacy_global_search
Type	true | false
Value	true

4. Select **Submit**.

Result

The system displays the global search UI and uses the legacy display logic.

Legacy global search user preferences

The global text search results page allows users to set several user preferences for global text search.

The following search preferences are available from the **Search tips and preferences** link.

Search preferences

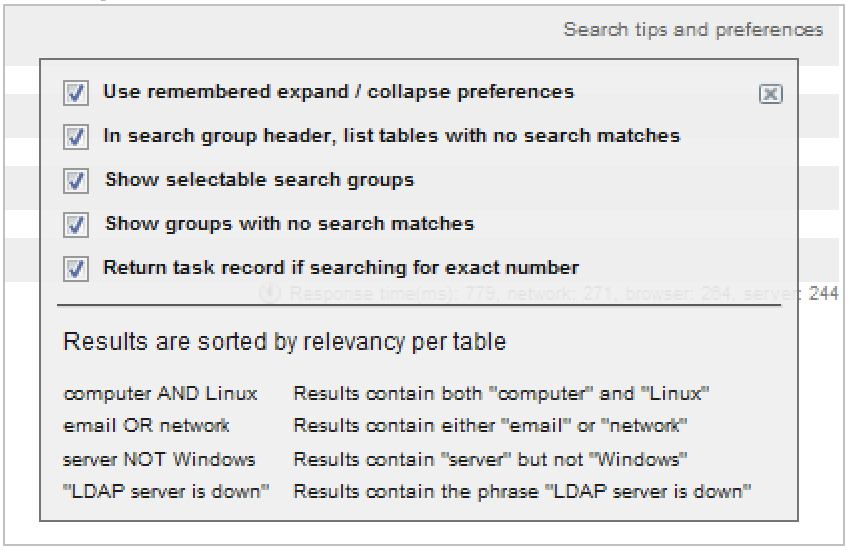

Search preferences

User preference	Description
Label Use remembered expand / collapse preferences **Preference** ts.remember.expanded	Specifies whether you want the collapse state of any search groups/tables remembered for your next search. If selected and you collapse the search results for the People & Places search group for example, that group will be collapsed on your next search. If cleared, all groups and tables are expanded for every search.
Label Show list of tables with no search matches **Preference** ts.show_negative_result_info	Specifies whether or not you want to see a summary of tables that had no search matches. For example: No matches for Requests, Catalog Tasks, Tickets.
Label Show selectable search groups **Preference** ts.show_search_groups	Shows or hides the search group check box row.
Label Show groups with no search matches **Preference** ts.show_empty_groups	Shows or hides a search group if it returns no matches. If selected, the search results display a placeholder row stating there were no matches for that group.
Label Return task record if searching for exact number **Preference**	Shows or hides the full search results page if the search term matches a task record number. When selected, an exact match to a record number causes the system to hide

Search preferences (continued)

User preference	Description
ts.match	the full search results and instead display the record with an exact match message. Exact search matches ignore most additional search conditions normally applied to the record such as the currently selected search groups or the knowledge article published state. For example, an exact search match for a knowledge article returns the article in any state as long as the current user has read access to it. **Note:** The system filters exact search results to display only knowledge records that are in the same language as the preferred language setting. **Sample exact match search query** 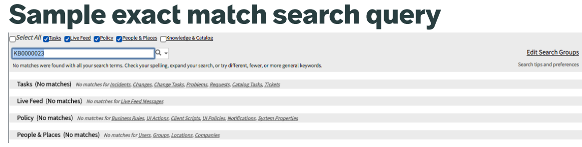 In this sample search, the **Knowledge & Catalog** search filter is cleared, which would normally prevent returning any knowledge articles. **Sample exact match search results** 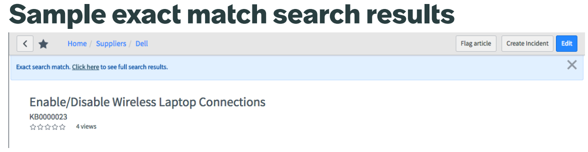 In this sample exact match, the system displays the matching article and an information message in the form header reading `Exact search match. Click here to see full search results.`

Related topics

User preferences

Set the preview limit for global or workspace search in Next Experience UI

Override the default limit of 5 preview results per search source on the Next Experience UI search results page for global or workspace search.

Before you begin

You must be using Next Experience UI.

Role required: workspace_admin

About this task

The preview limit for a search application specifies how many preview results can appear for each search source on the search results page. By default, the system displays up to 5 preview results per search source. Workspace administrators can override this default limit separately for global search and for all workspace searches.

Procedure

1. Navigate to the Workspace Global Search Configuration [sys_aw_global_search_config] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_aw_global_search_config.list`.
   c. Press Enter.
2. Open the configuration record for the search application that you want to modify the preview limit for:
   - To modify the preview limit for global search, open the **Now Experience Search Configuration** record.
   - To modify the preview limit for a workspace application, open the **Workspace search settings** record.
3. On the Workspace Global Search Configuration form, enter your desired preview limit into the **Tab overview items per section** field, then select **Update**.

Result

The search results page for global or workspace search displays preview results up to the new limit for each search source.

Hide per-table search result counts in Next Experience UI

Suppress display of per-table search result counts in the Next Experience UI search results page.

Before you begin

You must be using Next Experience UI.

Role required: ui_builder_admin or admin

About this task

By default, the search results page displays the total number of search results along with per-table result counts for each searchable table that contains matching records. You can hide the per-table result counts for search users by following this procedure.

Procedure

1. Navigate to **All** > *Next Experience* **Framework** > *UI Builder*.
   UI Builder opens in a new browser tab.
2. In the list of available experiences, open the *Unified Navigation* **App** record.
3. In the Pages and variants listing, open the **Search Results** > **search Default** record.
   UI Builder reloads and displays settings for the Search Results (search Default) page.
4. In the Content pane, select the **Body (Grid)** > **Search Result Wrapper 1** component.
   Properties, styles, and events for the selected component appear in the configuration panel.
5. If a `Different application scope` informational message appears, select **Edit in original scope**.
   The system displays an informational message indicating that you are editing in the application scope.
6. On the configuration panel's Config tab, enable the **Hide search results count** option.
7. Select **Save**.

Result

Per-table search result counts no longer display on the search results page.

Add a workspace application to the Unified Navigation search context menu

Enable users to view results for their Unified Navigation searches in workspace applications as well as in global search.

Before you begin

You must be using Next Experience UI.

The workspace application that you want to add to the search context menu must be installed. A search application configuration record must exist for the workspace, with AI Search specified as the search engine. If you need to create a new search application configuration record for the workspace, see Create a search application configuration for AI Search.

Role required: admin

About this task

You can add a workspace application, such as CSM/FSM Configurable Workspace, to the Unified Navigation search context menu. Adding the workspace application to the search context menu enables users to switch between global search results and search results from the specified workspace.

Procedure

1. Copy the sys_id of the record for the search application configuration for the chosen workspace:
   a. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
   b. Open the record for the chosen workspace's search application configuration.
      As an example, for Agent Workspace, open the **Agent Workspace Search Config** record.
   c. Select the Additional actions icon (≡), then select **Copy sys_id**.
2. Define a globalSearchDataConfigId UX page property to map the search application configuration record to the workspace's UX application:
   a. Navigate to the UX Page Property [sys_ux_page_property] table's list view by selecting **All**, entering `sys_ux_page_property.list` in the **Filter** field, and pressing Enter.
   b. Select **New**.
   c. On the UX Page Property form, fill in the fields as follows:

Field	Description
Page	Use the Lookup using list icon (🔍) to select the UX Application record for your workspace.
Name	Enter `globalSearchDataConfigId`.
Type	Select **string**.

Field	Description
Value	Enter the search application configuration record sys_id that you copied in step 1.c.

d. Select **Submit**.

3. Copy the sys_id of the record for the Workspace Global Search Configuration that you want workspace search to use:

   a. Navigate to the Workspace Global Search Configuration [sys_aw_global_search_config] table's list view by selecting **All**, entering `sys_aw_global_search_config.list` in the navigation filter, and pressing Enter.

   b. Open the record for the Workspace Global Search Configuration that you want workspace search to use.
   If you have a custom Workspace Global Search Configuration record that you want to use, open that record. Otherwise, open the Workspace search settings record.

   c. Select the Additional actions icon ( ), then select **Copy sys_id**.

4. Define a global_search_configurations UX page property to configure the workspace's search settings:

   a. Navigate to the UX Page Property [sys_ux_page_property] table's list view by selecting **All**, entering `sys_ux_page_property.list` in the **Filter** field, and pressing Enter.

   b. Select **New**.

   c. On the UX Page Property form, fill in the fields as follows:

Field	Description
Page	Use the Lookup using list icon ( ) to select the UX Application record for your workspace.
Name	Enter `global_search_configurations`.
Type	Select **json**.
Value	Enter the following JSON: `{ "globalSearchViewConfigId": "00000000000000000000000000000000", "globalSearchRoute": "search" }` Replace the `globalSearchViewConfigId` value of 00000000000000000000000000000000 with the Workspace Global Search

Field	Description
	Configuration record sys_id that you copied in step 3.c. For example, if you selected the Workspace search settings record, replace the value with `97a574ea53c0130084acddeeff7b12a6`.

d. Select **Submit**.

5. Enable Unified Navigation search for the chosen workspace:

a. Navigate to the UX Page Property [sys_ux_page_property] table's list view by selecting **All**, entering `sys_ux_page_property.list` in the **Filter** field, and pressing Enter.

b. Look for a UX Page Property record with name **chrome_header** and page matching the chosen workspace.
As an example, for CMDB Workspace, look for a record with name **chrome_header** and page *CMDB* **Workspace**.

c. If a **chrome_header** record exists for the workspace, edit it and change the **searchEnabled** JSON element to true in its **Value** field, then select **Update**.

d. If no **chrome_header** record exists for the workspace, select **New** and fill in the UX Page Property form fields as follows, then select **Submit**.

Field	Description
Page	Use the Lookup using list icon ( ) to select the UX Application record for your workspace.
Name	Enter `chrome_header`.
Type	Select **json**.
Value	Enter the following JSON: `{ "privatePage": { "searchEnabled": true } }`

6. Refresh your browser tab to reload the ServiceNow AI Platform®.

Result
The chosen workspace appears as an option in the Unified Navigation search context menu for users who have it assigned. Users can switch to the workspace's search context. For more details on using the search context menu to switch application contexts, see View and open search results using workspace applications in Next Experience Unified Navigation.

Global text search suggestions

Global text search offers two types of search suggestions.

The knowledge base and global text searches provide suggestions as you type. Type-ahead suggestions appear under the search box. Suggestions are based on similar searches that begin with the same characters.

Type-ahead search suggestions

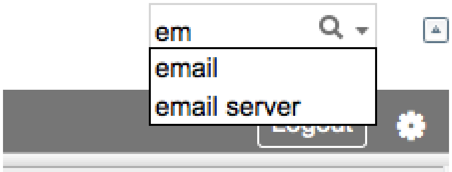

The knowledge base and global text searches also provide suggestions for alternate search terms. Suggestions appear if your original search doesn't return any results and an alternate spelling or similar recent search does. For example, if you misspell a search term (such as eail), the correct spelling (email) may appear as an alternate search suggestion. Suggestions appear beneath the search box on the results page.

Alternate search suggestions

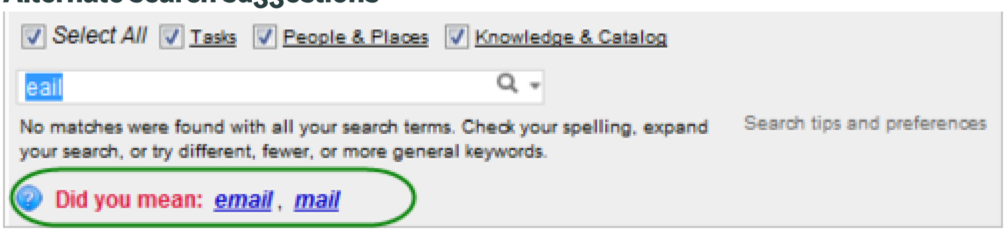

Note:
The type-ahead search suggestions feature is only available in the UI15 interface.

Update a type-ahead suggestion

The knowledge base and global text searches provide suggestions as you type. These type-ahead suggestions are compiled on a nightly basis by a scheduled job.

Before you begin

The type-ahead search suggestions feature is only available in the UI15 interface.

Role required: admin

About this task

Use the following procedure if you need to refresh the type-ahead suggestions list before the next run of the scheduled job.

Note:
Type-ahead suggestions aren't configurable. For more about how suggestions are generated and maintained, see the blog post Global Text Search Suggestions by a ServiceNow Technical Support Engineer in the ServiceNow Community.

Procedure

1. Navigate to **All** > **System Scheduler** > **Scheduled Jobs**.
2. Open **TS Search Stats**.
3. Run the scheduled job.

Configure "Did you mean?" suggestions

The Knowledge Base and global text searches can provide "Did you mean?" suggestions. Suggestions appear if a search doesn't return any results and an alternate spelling or similar recent search does. This feature is disabled by default.

Before you begin

Role required: admin

About this task

Set system properties to enable "Did you mean?" suggestions in the Knowledge Base and legacy global search tools.

Procedure

1. Navigate to **All** > **System Properties** > **Text Search**.
2. Select any of the following check boxes:
   - **Suggest alternate search spellings for knowledge or global search** (the *glide.ts.dym.enable_spell_correct* property).
   - **Suggest related searches for knowledge or global search** (the *glide.ts.dym.enable_chain_suggest* property); uses search chains, which generate suggestions by tracking occurrences of similar searches, in order, over time.

   Note:
   The global search references in these check boxes only apply to the legacy global search functionality. They don't apply to the global search functionality introduced in the New York release, because the "Did you mean?" functionality isn't available.
3. Select **Save**.

Zing generates search results in four phases

Generate search results by dividing documents among shards, searching shard documents, scoring shard documents, and merging scores into an index.

Zing generates search results in four phases:

1. Divide searchable documents among index shards.
2. Search and filter shard documents.
3. Score shard documents.
4. Merge shard document scores into an index.

Phase 1: Divide searchable documents among index shards

When a search is executed, the system equally divides all searchable documents among 10 index shards. Each shard has a unique list of documents.

To maximize search efficiency, the system creates 10 query threads to simultaneously search each index shard.

Divide documents among index shards

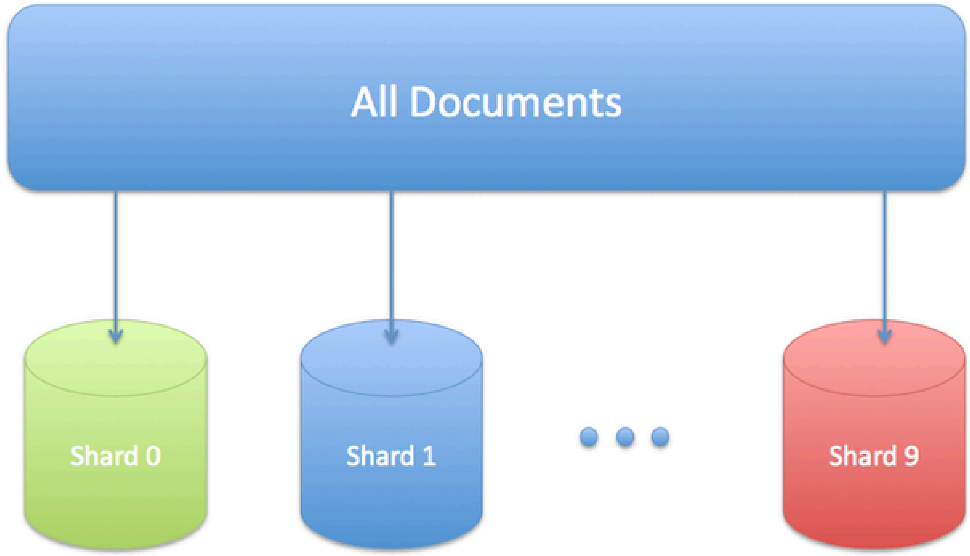

Phase 2: Search and filter shard documents

Each query thread searches through its list of shard documents to identify documents that match the search criteria. For example, if you search for "`wifi network`," the system returns all documents containing both "`wifi`" and "`network`" in no particular order.

Note:
A multiple word search is equivalent to multiple single word searches connected by an AND operator. For example, **[contains][wifi][AND][contains][network]**.

The query thread assigns each matching document a numeric value (a document ID) to uniquely identify it. All other documents are ignored.

Phase 3: Score shard documents

The query thread scores each matching shard document.

Query threads search, filter, and score shard documents

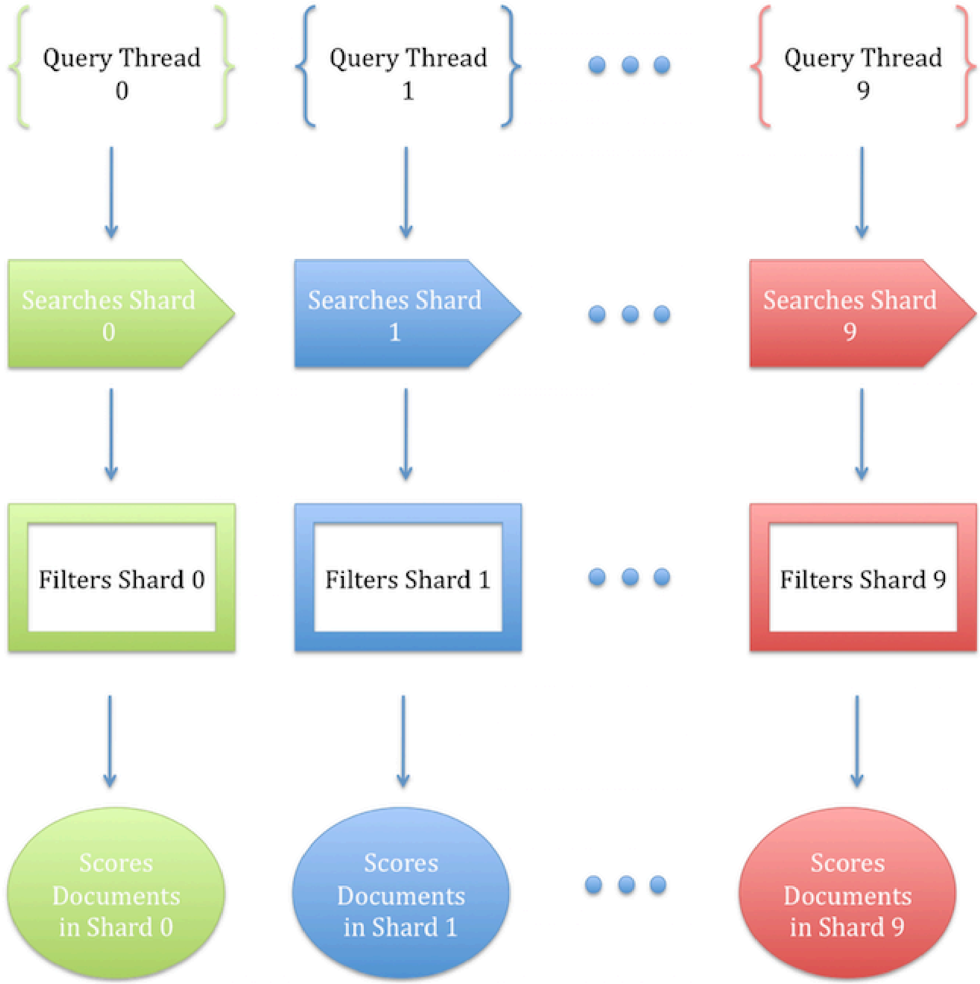

Phase 4: Merge and sort shard document scores

The system merges the document scores into a single index and sorts the documents from highest to lowest document score. The documents with the highest document score are most relevant to the search query.

Merge and sort shard document scores

Doc	Score
60	10
50	8
100	3
10	1

Doc	Score
81	11
91	6
101	5
11	4

Doc	Score
99	12
19	6

Doc	Score
99	12
81	11
60	10
50	8
91	6
19	6
101	5
11	4
100	3
10	1

Zing filters search results with access controls

Zing filters search results to only display records the user can access.

For example, suppose you index the System Property [sys_properties] table. When the ITIL user searches for a term in the System Property table, Zing returns no search results because the ITIL user doesn't meet the ACL rule requirements.

ITIL User Search

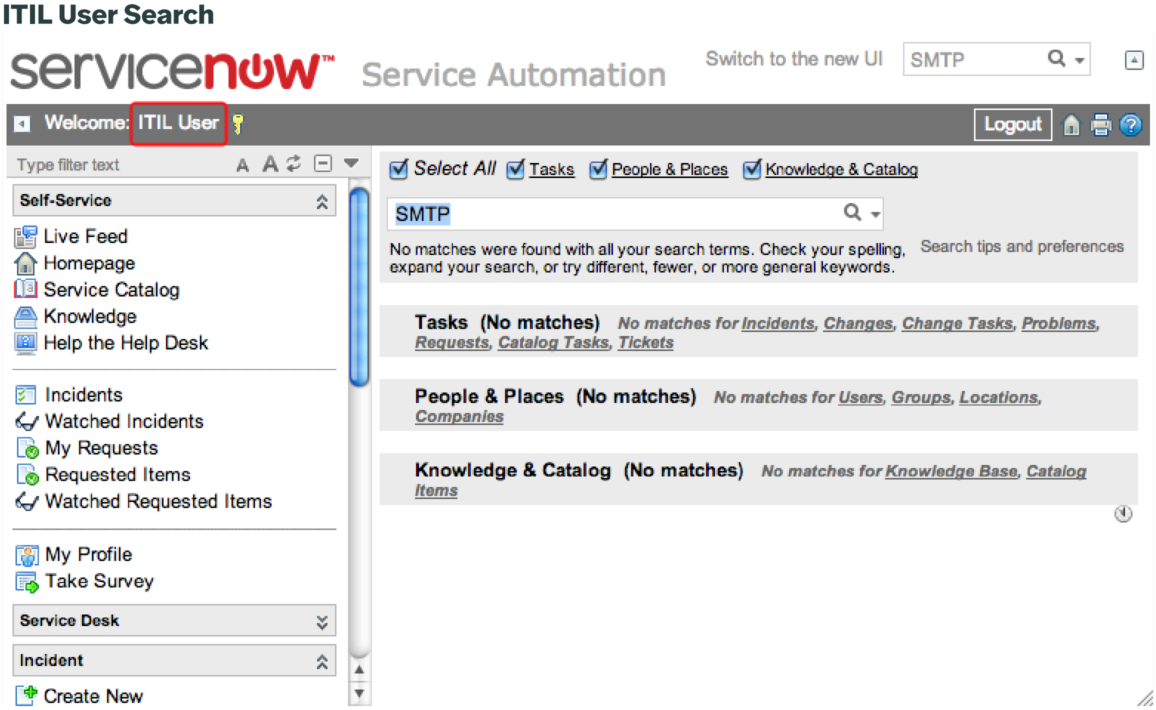

When a system administrator searches for the same property, Zing returns search results from the System Property table because the administrator meets the ACL rule requirements.

System Administrator Search

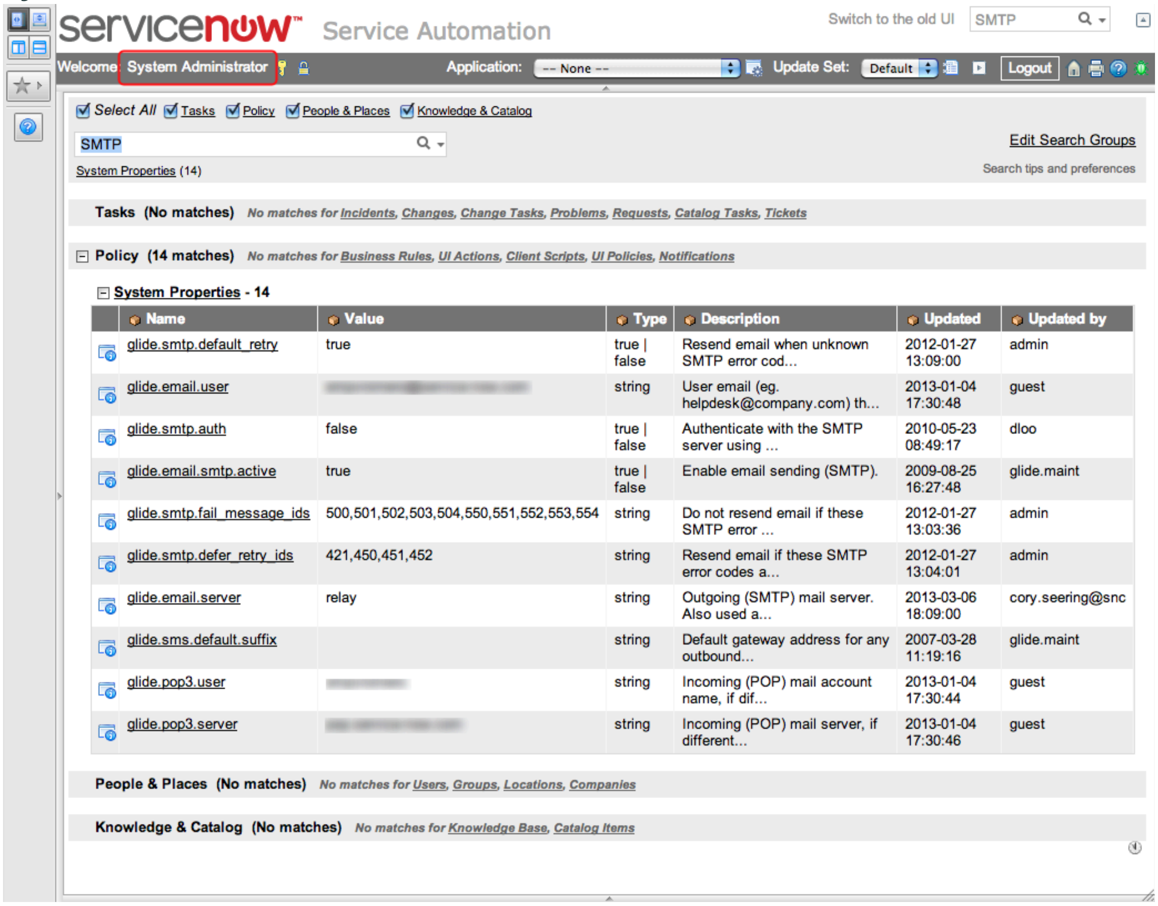

Zing computes document scores using three components

The Zing search engine computes document scores based on the frequency, sequence, and weight of search terms in the document.

Document scores

The components of a document score for a search query are:

- Frequency: how often the search terms appear in the document.
- Sequence: how often the search terms appear in the same order as the search query.
- Weight: how heavily weighted the source field is in which the search terms appear.

Sample document score computation

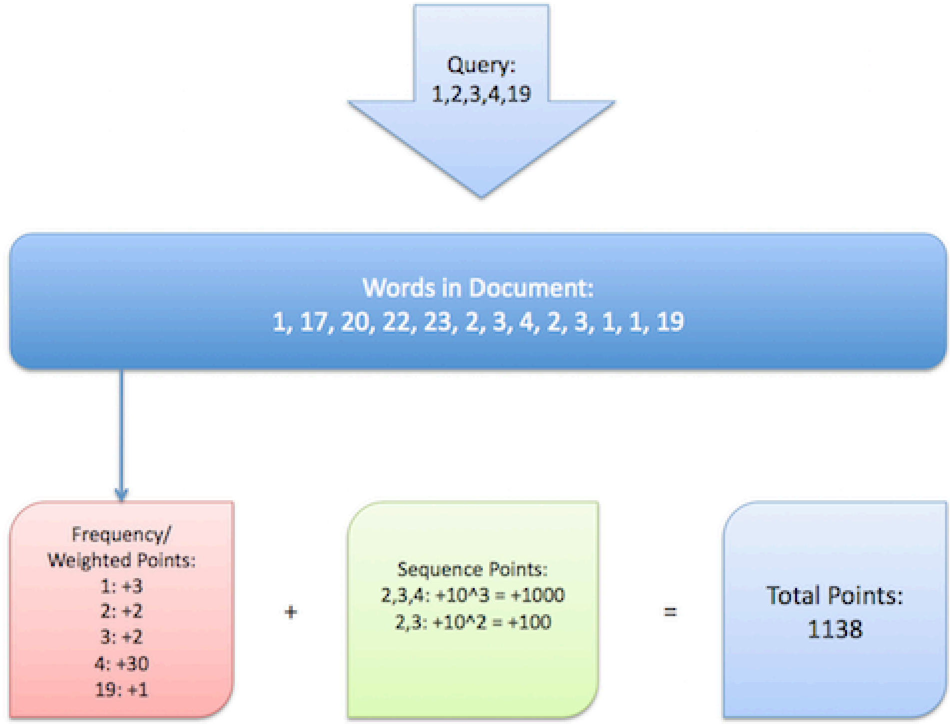

Frequency points

Zing awards one point whenever a search term appears anywhere in the document. For example, when searching for `distributed database server`, a document that contains `distributed` three times, `database` five times, and `server` 17 times would have 25 frequency points.

To increase search result scores of search terms that appear more frequently in a document, but less frequently in a document set, you can Score search terms by inverse document frequency (IDF). When TF-IDF is enabled, search term scores are calculated by multiplying the term frequency score by the inverse document frequency score. Because enabling TF-IDF increases the weight of less common search terms, search results for that table are more likely to be relevant. For example, when searching for `distributed database server`, the term `distributed` might receive a higher score than `server` if it appears frequently in one document but less frequently in the document set as a whole.

Zing applies a multiplier to frequency points based on the value of the *`ts_weight`* attribute for the field in which the search term appears. A field with a text search scoring weight of 30 (*`ts_weight=30`*) would add 30 points for each inclusion of a search term.

Sequence points

Zing awards a document more points when it contains the search terms in the same order in which they were typed. The more search terms in sequence there are, the exponentially higher the score becomes. Zing awards sequence points as 10^x, where x is the number of search terms that appear in sequence.

In the `distributed database server` search example, Zing awards a document 100 (10^2) sequence points for each time it includes the two-term string `database server`.

Likewise, Zing awards a document 1000 (10^3) sequence points each time it includes the three-term string `distributed database server`.

Zing applies a multiplier to sequence points based on the value of the *`ts_weight`* attribute for the field in which the sequence appears. The sequence points use the calculation (10^x * field *`ts_weight`* attribute).

Field scoring weights

The system elevates the default scoring weight of Knowledge record numbers, Knowledge short descriptions and metadata, task record numbers, and task short descriptions. Default *`ts_weight`* attributes for these fields are as follows:

- *`kb_knowledge.number = 50`*
- *`kb_knowledge.short_description = 10`*
- *`kb_knowledge.meta = 10`*
- *`task.number = 50`*
- *`task.short_description = 10`*

All other fields have a default *`ts_weight`* attribute of 1. The maximum possible weight value is 255.

Score search terms by inverse document frequency (IDF)

Enable term frequency–inverse document frequency (TF-IDF) to increase the search result scores of search terms that appear more frequently in a document, but less frequently in the whole collection of searchable documents.

Before you begin

To enable TF-IDF scoring, a text index must be present for the table. Make sure that text indexing is enabled for the table before changing the query mode. See Configure a table for indexing and searching.

Role required: ts_admin or admin

About this task

Administrators can enable TF-IDF scoring on a table to improve search results. The Knowledge table [kb_knowledge] has TF-IDF scoring enabled by default.

Note:
This task requires that you regenerate an index and therefore should be done with system performance in mind. Regenerating an index may take a while to complete depending on table size and other factors. You may notice performance degradation or incomplete search results while the system regenerates the index.

Procedure

1. Navigate to the system dictionary and open the entry for the desired table.
2. In the **Attributes** related list, select **New**.
3. In the **Attributes** field, select **Text search Enable IDF**.
4. Select **Submit**.
   The **Enable IDF Score** check box in the table's text index record can now be configured.
5. Navigate to **All** > **System Definition** > **Text Indexes**.
6. Open the text index created for the table.

7. Select **Enable IDF Score**.
8. Under **Related Links**, select **Regenerate Document Frequency**.
   Regenerating document frequency is a one-time task. Regenerate document frequency when TF-IDF scoring is first enabled for a table.
9. Configure the desired notification and select **OK**.

Set the relative weight of a field

To improve search results, the Zing search engine assigns to each potential match a numeric score that represents its relevancy to the query.

Before you begin

Text indexing must be enabled for the table containing the fields you want to modify the relative weight for. To learn about enabling text indexing for a table, see Configure a table for indexing and searching.

Role required: admin

About this task

Administrators can control the relative importance of a match for each field in a table with the ts_weight attribute. The majority of fiends have no ts_weight attribute. Zing treats these fields as having weight 1 when scoring search matches.

The following fields have ts_weight attribute values by default:

- kb_knowledge.number = 50
- kb_knowledge.short_description = 10
- kb_knowledge.meta = 10
- task.number = 50
- task.short_description = 10

When scoring records, Zing compares the relative weights of fields that produced search term matches. For example, if the **Title** field has `ts_weight=50` and the **Description** field has `ts_weight=10`, then a match in **Title** is 5 times as relevant as a match in **Description** and 50 times as relevant as a field without the ts_weight attribute.

Note:
To view the complete scoring system, see Document Scoring.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the system dictionary record for the field's table and column.
   For example, to modify the ts_weight for the kb_knowledge.number field, open the dictionary record with **Table** "kb_knowledge" and **Column name** "number".
3. In the Dictionary Entry form's **Attributes** field, perform one of the following steps:
   - If a ts_weight attribute appears in the list of attributes, change its integer value to the relative weight you want to set for the field.
   - If no ts_weight attribute appears in the list, add a new `ts_weight=<value>` comma-separated list entry, where <value> is the relative weight you want to set for the field.
4. Select **Update**.
5. Refresh the table's text index:

a. Open the system dictionary record for the field's table.
The table's system dictionary record is the one with the table name as its **Table** value, an empty **Column name** field, and "Collection" as its **Type** value.

b. On the Dictionary Entry form, select the **Generate Text Index** link in the Related Links section.

Result
Once Zing finishes regenerating the text index for the table, search result scoring uses your new relative weights for fields.

Zing indexes words

Index documents by dividing them into words. Depending on the languages your instance supports, a word may be a single character such as a Chinese or Japanese pictogram or a sequence of characters separated by spaces such as with Latin, Arabic, and Pinyin languages.

Zing uses spaces to define the basic unit of word separation. Zing uses the following rules to index words.

Zing word indexing rules

Rule	Behavior
Punctuation	Zing indexes some punctuation marks as part of some words to improve search results for common search terms. Zing converts any non-indexed punctuation characters to spaces. See Zing indexes punctuation as part of some words.
Pictograms	Zing supports the full range of Unicode pictogram characters. Each Chinese or Japanese pictogram is indexed as a separate word as if it were a single Latin-1 character surrounded by spaces.
Letter	Zing treats Latin-1, Arabic, and Pinyin characters as individual letters of space-separated-words. Sequences of letters define indexable words.

Zing indexes punctuation as part of some words

Zing indexes some punctuation marks as part of some words to improve search results for common search terms.

Punctuation indexed as part of a word

Punctuation	Indexed to find	Search behavior
Ampersands [&]	Company names	Zing locates terms containing punctuation characters that are common in company names: ampersands (&), plus signs (+), or hyphens (-). For example, a search for `H&R Block` locates exact matches to the search term.
Apostrophes [']	Proper names	Zing ignores trailing possessives and retains most others to find company names. For example, a search for `O'Reilly's` locates matches to O'Reilly and O'Reilly's.

Punctuation indexed as part of a word (continued)

Punctuation	Indexed to find	Search behavior
Hyphens [-]	Company names	Zing locates terms containing punctuation characters that are common in company names: ampersands (&), plus signs (+), or hyphens (-). For example, a search for `Coca-Cola` locates exact matches to the search term.
	Product numbers	Zing locates product numbers that follow a typical pattern. To meet this pattern, search terms must contain frequent numbers and only include underscores (_), hyphens (-), or periods for punctuation. For example, a search for `PROD-10-987` locates exact matches to the search term.
Numbers [123]	Product numbers	Zing locates product numbers that follow a typical pattern. To meet this pattern, search terms must contain frequent numbers and only include underscores (_), hyphens (-), or periods for punctuation. For example, a search for `PROD10987` locates exact matches to the search term.
	Record numbers	Zing locates record numbers that follow a typical pattern. To meet this pattern, search terms must start with a recognized record number prefix. For example, a search for `INT1234567` locates exact matches to the Incident record.
Periods [.]	Acronyms	Zing locates acronyms whether they are separated by periods or not. For example, a search for `u.s.a.` locates matches to usa or u.s.a.. Note that wildcard searches may affect acronym handling. In another example, a search for `u.s.*` may yield better results than a search for `u.s*`.
	Host names	Zing locates sequences of letters and numbers separated only by periods. For example, a search for `en.myhost123.com` locates exact matches to the search term.
	IP addresses	Zing locates numbers that follow a typical IPv4 address pattern.

Punctuation indexed as part of a word (continued)

Punctuation	Indexed to find	Search behavior
		For example, a search for `10.0.0.1` locates exact matches to the search term.
	Product numbers	Zing locates product numbers that follow a typical pattern. To meet this pattern, search terms must contain frequent numbers and only include underscores (_), hyphens (-), or periods for punctuation. For example, a search for `PROD.10.987` locates exact matches to the search term.
Plus signs [+]	Company names	Zing locates terms containing punctuation characters that are common in company names: ampersands (&), plus signs (+), or hyphens (-). For example, a search for `Google+` locates exact matches to the search term.
Underscores [_]	Product numbers	Zing locates product numbers that follow a typical pattern. To meet this pattern, search terms must contain frequent numbers and only include underscores (_), hyphens (-), or periods for punctuation. For example, a search for `PROD_10_987` locates exact matches to the search term.

Note:
The punctuation handling described here doesn't affect wildcard and boolean operators. These operators provide separate Zing functions.

Zing indexes some HTML elements

Zing indexes the contents of certain HTML elements to improve document search results. This indexing is addition to the normal indexing by word.

Indexed HTML elements

HTML element	Description
Title attribute	Zing indexes the contents of the *`title`* attribute of any HTML element. This is separate from the *`title`* element, which Zing already indexes when identifying the words of the document.
Anchor link target	Zing indexes the target URL specified in the *`href`* attribute of any anchor element.

Indexed HTML elements (continued)

HTML element	Description
Alt text for an image	Zing indexes the alternative text of any image element.

Configure a table for indexing and searching

Enable Zing text indexing for a table to allow users to search its records for string values. Configure the table's text indexing attributes, stop words, and synonym dictionaries to control search behavior.

Before you begin

Role required: ts_admin or admin

About this task

Enabling text indexing for a table allows users to search for string values in records on that table. You can configure attributes, stop words, and synonym dictionaries to control indexing and search behavior for an indexed table.

By default, the system creates text indexes for the tables with a Text index record (**All** > **System Definition** > **Text Indexes**). For example:

- Knowledge-related tables
- Core-data-related tables
- Connect-related tables
- Task table
- User table

Enabling text indexing for a table also enables it for all child tables. If you want to disable text indexing for a child table, see Remove the text index for a child table.

> **Note:**
> Whenever you create an index table, the system also creates Access Control Roles (ACLs) for new tables.

Text indexing can be a resource-intensive task that may take a while to complete. You may notice performance degradation or incomplete search results during index generation. To estimate text indexing duration, you can view historical text indexing statistics.

> **Note:**
>
> New tables created by update set transfers have text indexing disabled, even if it's enabled for the tables included in the update set.
>
> To enable users to search tables created by update set transfers, you must manually enable text indexing for those tables.

Procedure

1. Navigate to **All** > **System Definition** > **Text Index Configurations** and select **New**.
2. In **Applies to**, select **Table**.
3. In **Table**, select the table that you want to configure text indexing for.

4. **Optional:** If you intend to use the table in multi-table searches or if you want to use the BM25 scoring algorithm on the table, upgrade its indexing format to V4.

   a. If the value in **Format** isn't grayed out, select **v4**, then select **Submit** and skip to step 5.

   b. If the value in **Format** is grayed out, select **Submit**.
   The grayed-out value means that the table has already been indexed and needs its indexing upgraded to V4.

   c. Navigate to **All** > **System Definition** > **Text Index Configurations** and select the table you submitted.

   d. Under **Related Links**, select **Show index name record**.

   e. Under **Related Links**, select **Upgrade to V4** and select **Submit**.
   The specified table's indexing is upgraded from V3 to V4. Now, this table can be part of multiple-table indexing and BM25 is used to score the table's relevance score. Optionally, select **Show Configuration** under **Related Links**to show the table indexing format on the **Text Index Configuration** page.

5. Configure indexing and search attributes on the selected table.

   a. Navigate to **All** > **System Definition** > **Text Index Configurations** and select the table you configured.

   b. In the Text Index Table Attribute Maps related list, select **New**.
   Attributes pertain to the way Zing indexes documents and ranks documents returned in search results.

   c. On the Text Index Table Attribute Maps form, fill in the fields.

Field	Description
Table	Table to define attributes for.
Attribute name	Search engine attribute. For a list of the attributes and valid values, see Zing index and search attributes.
Value	Value for attribute.

   d. Repeat this procedure to define additional attributes for the table.

   e. In the Related Links section, select **Generate Text Index**.

6. Configure column-level attributes.

   a. Navigate to **All** > **System Definition** > **Text Index Configurations** and select a table.

   b. In the Text Index Column Attribute Maps related list, select **New**.

   c. On the Text Index Column Attribute Map form, fill in the fields.

Field	Description
Table	Table that contains the column you want to define attributes for. Only tables in the index group and their child tables display.
Attribute name	Attribute. For a list of the attributes and valid values, see Zing index and search attributes.
Column	The name of the column in the selected table you want to apply the search attribute to.
Value	Value for attribute.

d. Select **Submit**.

e. Repeat this procedure to define additional search engine parameters for table columns in the index group.

f. In the Related Links section, select **Generate Text Index**.

7. **Optional:** Add stop words.
Stop words are words that you don't want to index, such as `the`, `a`, or `an`. There is a global list of common stop words, so you only have to add stop words specific to your situation. To see the global list of stop words, navigate in your browser to `https://<instance name>.service-now.com/ts_stop_list.do`.
   a. On the **Index Stop Words** tab, select **New**.
   b. In **Index**, select the index to add the stop word to.
   c. In **Word**, enter a single word you don't want to index and select **Submit**.
8. **Optional:** Add a synonym dictionary.
Specifying a synonym dictionary searches on the search term as well as synonyms of the search term, such as car, and its synonym, auto. Zing provides a default synonym dictionary.
   a. On the **Index synonym Dictionaries** tab, select **New**.
   b. In **Index**, select the index to add the stop words to.
   c. In **Synonym dictionary**, select the synonym dictionary to use on the specified table.
   To add a dictionary, select **New** and follow the procedure in Create synonym dictionaries.

Result

The list search *for text* option appears on the list views for the table and its child tables.

What to do next

You can search records on the indexed table using list search in form view. To learn more about using list search, see List search finds records from the current table.

You can also query the indexed table for string matches using the `GlideRecord.addQuery(name, value)` or `GlideRecord.addEncodedQuery(name, value)` methods with the `123TEXTQUERY321` reserved name. For more information about querying for string matches, see Querying tables in script.

Fields excluded from text indexing

Text indexing excludes fields that meet certain name, type, attribute, or encryption conditions. User searches don't match or return values from excluded fields.

Text indexing excludes a field if it meets any of the conditions in the following table.

Category	Description
Field name	Text indexing includes the **sys_class_name** and **sys_tags** fields but excludes all other fields whose names begin with **sys_**.
Field type	Text indexing excludes fields of these types: • Date • Date/Time • Duration • True/False • Workflow For more information on these types, see Field types.
Field attributes	Text indexing excludes fields that have the **no_text_index** attribute set to true. For more information on this attribute, see Remove an index for a specific field.
Field encryption	Text indexing excludes fields configured for Edge Encryption. For more information on configuring fields with Edge Encryption, see Encrypt fields using encryption configurations.

Configure a text index group to search across multiple tables

Update tables to the V4 indexing format and include them in a text index group to facilitate searching across multiple tables.

Before you begin

Role required: ts_admin or admin

About this task

Text index groups enable you to index and search across multiple tables.

Performing searches across multiple tables requires normalizing the tables first so the relevance scores mean the same thing for all tables. V4 indexing provides that normalization. Zing uses the Okapi BM25 ranking function to calculate the relevancy score of search results.

The V4 table indexing format is optimized for documentation, such as knowledge base articles. You don't need to reindex tables that use this format when you change the value of an indexing

attribute, such as the weight of a field. If you want to use the BM25 ranking function on a single table, you can convert it to the V4 format.

To enable search across multiple tables, you:

1. Configure all of your chosen tables to use the V4 index format if they don't already use it.

   Note:
   In new instances, the Catalog Item [sc_cat_item], Knowledge [kb_knowledge], Social Q&A Questions [kb_social_qa_question], and User [sys_user] tables use the V4 index format. If you upgrade your instance from a previous release, the indexing format of these tables remains the same as it was in the older release.

2. Add the tables to a text index group so they can be indexed and searched together.

   Note:
   You can't add the Task [task] table or any tables that extend it to a text index group.

Procedure

1. Configure the tables that you want to include in multiple table searches to use the V4 index format.

   Note:
   Upgrading from V3 to V4 is easy. Downgrading from V4 to V3 is more difficult and discouraged.

   a. Navigate to **All** > **System Definition** > **Text Index Configurations**.

   b. Open the record for the table that you want to upgrade to the V4 index format.

   c. If **Format** isn't disabled, select **V4**, select **Submit**, and then skip to step 2.
   You can choose index formatting in this step because no index exists for the table.

   d. If **Format** is disabled, select the **Show index name record** related link.
   The disabled field means that an index already exists for the table and you must explicitly update the index format.
   The Text Index page displays.

   e. Select the **Upgrade to V4** related link, then select **OK**.
   The system schedules a task to upgrade the text index from V3 to V4, enabling this table to participate in multiple-table indexing.

2. Create a text index group that specifies the tables you want to index together so you can search them together.

   a. Navigate to **All** > **System Definition** > **Text Index Groups**.

   b. Select **New**.

   c. Enter the following field values:

Field	Value
Label	Name for the group of tables that you want to index and search together.
Tables	List of tables that you want to include in the text index group.

You can add tables formatted with V3 and V4 to text index groups. If you add V3 tables, the text index group remains inactive.

d. Select **Submit**.

3. Configure the indexing and search attributes on tables in a text index group.

a. Navigate to **All** > **System Definition** > **Text Index Configurations**.

b. Select **New**.
The Text Index configuration form displays. If you already have a text index group in the list, skip to step 3.d.

c. In **Applies to**, select **Index Group**.

d. In **Index group**, select the magnifying glass and select the text index group you created, then select **Submit**.
The only text index groups that display are the ones you have already created.
The Text Index Configuration table displays.

e. Open the record for your text index group.

f. On the **Text Index Table Attribute Maps** tab, select **New**.
The Text Index Table Attribute Map form displays. Attributes pertain to the way Zing indexes or searches individual tables in the text index group.

g. In **Table**, select a table to define attributes for.
Only tables in the specified text index group and their child tables display.

h. In **Attribute name**, select the magnifying glass and select an attribute.
To see the list of attributes, enter `ts_configuration_attribute.list` in the application navigator. For attribute definitions, see Zing index and search attributes.

i. In **Value** enter a value and select **Submit**.
Valid values depend on the attribute. For information on valid values, see Zing index and search attributes.

j. Repeat this procedure to define additional search engine parameters for other tables in the text index group.

k. Under **Related Links**, select **Generate Text Index**.

4. Configure search parameters on individual table columns in the text index group.

a. Navigate to **All** > **System Definition** > **Text Index Configurations**.

b. Open the record for your text index group.

c. On the **Text Index Column Attribute Maps** tab, select **New**.
The Text Index Column Attribute Map form displays. Attributes pertain to the way Zing indexes or searches individual table columns in the text index group.

d. In **Table**, select a table to define attributes for.
Only tables in the text index group and their child tables display.

e. In **Attribute name**, select a search engine parameter, enter a value for that parameter in **Value**, and select **Submit**.
To see the list of attributes, enter `ts_configuration_attribute.list` in the application navigator. For attribute definitions, see Zing index and search attributes.

f. In **Column**, specify the name of the column in the selected table you want to apply the search attribute to.

g. In **Value**, specify the attribute's value that governs the search.

h. Select **Submit**.

i. Repeat this procedure to define additional search engine parameters for table columns in the text index group.

j. Under **Related Links**, select **Generate Text Index**.

5. **Optional:** Add a synonym dictionary.
Specifying a dictionary enables searches on the search term and synonyms of the search term, such as `car` and its synonym `auto`. Zing provides a default synonym dictionary.

a. Navigate to **All** > **System Definition** > **Text Index Configurations**.

b. Open the record for your text index group.

c. On the **Index synonym Dictionaries** tab, select **New**.

d. In **Index**, select the table index to add the stop words to.

e. In **Synonym dictionary**, select the dictionary to use.
To add a dictionary, select **New** and follow the procedure in Create synonym dictionaries.

f. Select **Submit**.

g. Under **Related Links**, select **Generate Text Index**.

6. Select **Update**.

What to do next

You can query the tables in the text index group for string matches using the `GlideRecord.addQuery(name, value)` or `GlideRecord.addEncodedQuery(name, value)` methods with the `123TEXTINDEXGROUP321` reserved name. For more information about querying text index groups for string matches, see Querying tables in script.

Zing index and search dictionary attributes

The following dictionary attributes either affect the way Zing indexes tables or ranks search results.

Index and search dictionary attributes

Zing uses dictionary attributes set at the table or field level to configure the following settings:

- indexing behavior
- ranking of search results
- the BM25 relevancy ranking algorithm

Zing index and search dictionary attributes

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
attachment_index	boolean	Table	Both	If true, attachments on the table are indexed for search purposes.	Attachment Index	TRUE	FALSE
no_text_index	boolean	Both	Both	This attribute on a field prevents the field from being included in a text index.	No Text Index	TRUE	FALSE
text_index_attachment_body_weight	float	Table	V4	Specifies how much consideration (weight, a float value) to give a search term found in the body of an attachment. For more	Text Index Attachment Body Weight	FALSE	TRUE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
				information on how weight value affect search result relevancy, see Zing computes document scores using three components.			
text_index_attachment_displayed	boolean	Table	Both	Displays attachments in text search results.	Text Index Attachment Displayed	FALSE	FALSE
text_index_attachment_title_weight	float	Table	V4	Specifies how much consideration (weight, a float value) to give a search term found in the title of an attachment. For more information on how weight value affect search result relevancy, see Zing computes document scores using three components.	Text Index Attachment Title Weight	FALSE	TRUE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
text_index_auto_stop	boolean	Table	Both	Automatically identifies stop words and doesn't index them. Based on the threshold specified in text_index_auto_stop_threshold.	Text Index Auto Stop	TRUE	TRUE
text_index_auto_stop_threshold	int	Table	Both	Threshold (integer, number of occurrences) for a word to become a stop word automatically. See text_index_auto_stop.	Text Index Auto Stop Threshold	TRUE	TRUE
text_index_default_partial_matching_rule	String	Table	Both	String for partial matching rule configuration. Input should be based on Change the query mode of an indexed table.	Text Index Default Partial Matching Rule	FALSE	TRUE
text_index_default_query_mode	String	Table	Both	String for query mode configuration. Input should be based on Change the query mode of an	Text Index Default Query Mode	FALSE	TRUE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
				indexed table			
text_index_disable_synonym	boolean	Table	Both	TRUE disables searching on synonyms as well as the search term.	Text Index Disable Synonym	FALSE	TRUE
text_index_enable_idf	boolean	Table	V3	TRUE enables TF-IDF scoring.	Text Index Enable IDF	FALSE	TRUE
text_index_filter_junk	boolean	Table	Both	TRUE by default, FALSE disables the junk filter for a table. By default, Zing doesn't index or search for 2-digit numbers and single character words (unless they are Chinese or Japanese characters). Regenerate the index after disabling the junk filter. This attribute results in a larger table index. For optimal	Text Index Filter Junk	TRUE	TRUE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
				performance, don't apply it unless it's required.			
text_index_filter_query	An encodedQuery on the table	Table	Both	Use this attribute when you want to index the rows matching specific filter conditions. The attribute value must be an encodedQuery on the table.	Index Filter	TRUE	FALSE
text_index_filterable	boolean	Column	Both	If a field with this attribute set to true, is changed, the row is evaluated for re-indexing (even if the field itself isn't indexable).	Index Filterable	TRUE	FALSE
text_index_include_stop_word	boolean	Table	Both	if you need to index stop words in this table, set this attribute to true.	Text Index Include Stop Word	TRUE	TRUE
text_index_scorer_bm25_b	float	Table	V4	This attribute is the b parameter	bm25_b	FALSE	TRUE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
				used in the BM25 algorithm. See text_index_scorer_bm25_b.			
text_index_scorer_bm25_k1	Float	Table	V4	This attribute is the K1 parameter used in the BM25 algorithm. See text_index_scorer_bm25_k1.	bm25_k1	FALSE	TRUE
text_index_tags	String	Table	Both	Should be set to specific String: empty/ everyone_only/ all_shared, depending on type of tags that wants to be indexed.	Text Index Tags	TRUE	FALSE
text_index_tokenizer_language	String	Table	Both	Put "ja" (without double quote) when there is need to support tokenization in Japanese.	Text Index Tokenizer Language	TRUE	TRUE
text_index_translations	Boolean	Table	Both	If present or true, forces indexes to be recalculated when translated strings are added. Requires	Text Index Translations	TRUE	FALSE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
				sys-admin role to modify. Automatically set for indexed fields that are translated, and to fields that have a translation and are being indexed. The glide.i18n.force_index system property, which defaults to true, overrides this attribute.			
text_search_and_or_limit	int	Table	Both	Add table level custom attribute text_search_and_or_limit that overrides the global property.	Text Search And Or Limit	FALSE	FALSE
ts_language	String	Column	Both	Added to allow for text indexing of only specific languages.	Text Search Language	TRUE	FALSE
ts_multi_value_field_splitter	String	Column	Both	If one field is a multi-value field, add the splitter in the field in this	Text Search Multi-value Field Splitter	TRUE	TRUE

Zing index and search dictionary attributes (continued)

Name	Data Type	Applies To	Applies to V3 or V4	Description	Label	Needs re-indexing?	Root Table Only
				attribute to spread out index positions of the values. For example, to separate multi:value, put : for this attribute.			
ts_weight	int	Column	Both	Controls the relative importance of a match in the field for text search. For more information on how weight value affect search result relevancy, see Zing computes document scores using three components. To learn how to set text search weight for a field, see Set the relative weight of a field.	Text Search Weight	FALSE	FALSE

no_text_index

For Customer Service and Support guidelines on setting the **no_text_index** option for different field types, see KB0859922 on the Now Support site.

text_index_scorer_bm25_b

This dictionary attribute is part of the BM25 algorithm. The attribute impacts the importance of the length of a document in search results. If this attribute is 0, the length has no bearing. In this case, a document of two words would be valued equally as a document with 1000 words. Very short documents are often weighted lower than longer documents. The default is 0.75.

text_index_scorer_bm25_k1

This dictionary attribute is part of the BM25 algorithm. The attribute helps determine term frequency saturation characteristics, that is, having the search term appear in a document increases the relevancy of the document, but the increase is less impactful as the numbers climb. For example, a search term occurring 10 times in a document makes it more relevant than a document that has only one occurrence. However, a document that has 100 occurrences of the search term may not be ten times more impactful than a document that has 10 occurrences. The relevance approaches a maximum effect asymptote. The higher the attribute value, the lower the maximum effect of multiple occurrences. Typical values are between 1.25 - 2.0.

Regenerate a text index for a table

You can regenerate a table text index when you change table stop words or display values.

Before you begin

The table that you want to regenerate the text index for must already be configured for indexing and searching. For details on this configuration process, see Configure a table for indexing and searching.

Role required: ts_admin or admin

About this task

By default, the system maintains text indexes on a daily schedule. Typically, you only need to manually regenerate a text index when you change these values.

- You change the list of table-specific stop words.
- You change the display value of a record such as changing a user or group name. Until you regenerate the index, text searches for old display values will still produce results and searches for the new display value won't show results.

You can also regenerate a text index if you observe incorrect search results for the indexed table. This is rare and usually only occurs if text indexing was interrupted.

Text indexing can be a resource-intensive task that may take a while to complete. You may notice performance degradation or incomplete search results during index generation. To estimate text indexing duration, you can view historical statistics.

Note:
This index regeneration process purges the existing text search index for the table before it regenerates it. While it's processing, no search results are returned if you perform a text search before the regeneration is complete. An alternate method is available that doesn't impact use of text searches for the table while the regeneration is in process. See Reindex a table without impacting text search results.

Procedure

1. Navigate to **All** > **System Definition** > **Text Indexes**.
2. Open the text index for the table.

   Example
   For example, select **task**.
   The system displays the Text Index record for the table.
3. Select the **Regenerate Text Index** related link, then select **OK**.

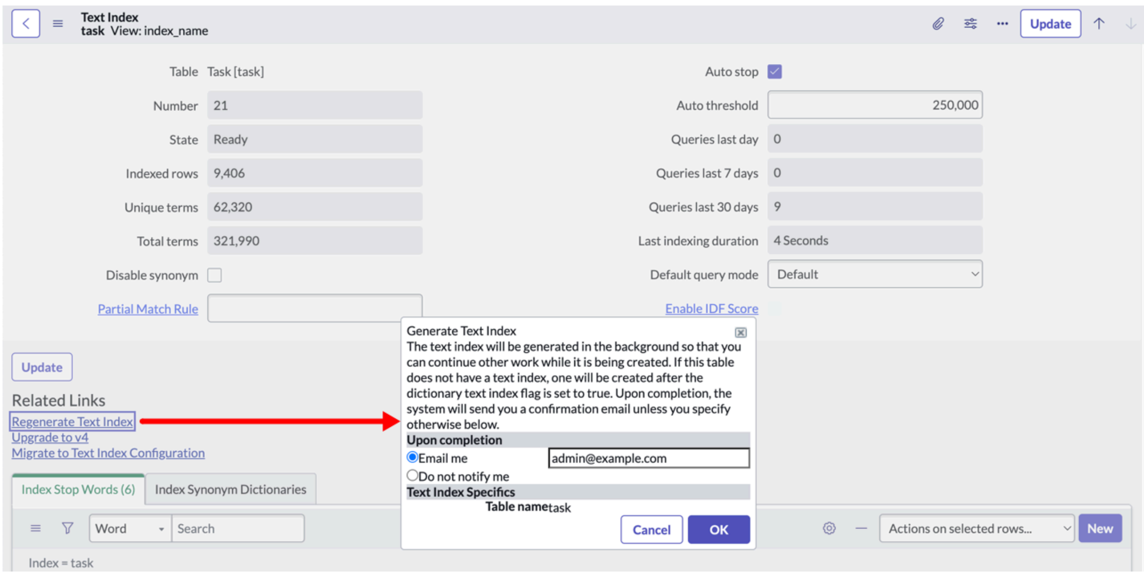

Result

The system schedules the selected table for text indexing.

Reindex a table without impacting text search results

Rebuild text search indexes without adversely impacting search results. You can continue to perform text searches on a table while the index regeneration takes place.

Before you begin

The table that you want to rebuild the index for must already be configured for indexing and searching. For details on this configuration process, see Configure a table for indexing and searching.

Role required: admin

Note:
If the `glide.script_processor.admin` system property's value has been changed from the default value, you need the role specified by that property value. For details on this property value and its effects, see Restrict access to background script [Updated in Security Center 1.3].

About this task

This reindexing method doesn't purge the entire text search index before the rebuild takes place. It rebuilds the text search table index one record at a time, and reindexes the target table hierarchy only. For example, if you select the sc_cat_item table, only the sc_cat_item table is reindexed.

Procedure

1. Navigate to **All** > **System Definition** > **Scripts - Background**.

   Note:
   If the *glide.script_processor.admin* system property's value is set to security_admin, you must elevate to that privileged role to access this module. For more information on elevating privileges, see Elevated privilege roles.

2. In the **Run script (JavaScript executed on server)** text field, enter the following background script, being sure to enclose both arguments in quotes as shown:

   ```
   new GlideTextIndexEvent ( ).indexUpdate('TABLE_NAME',
    'EMAIL_ADDRESS');
   ```

   Where:

   - `TABLE_NAME` is the name of the table you want to reindex.

     Note:
     Only the specified table and its child tables are reindexed.

   - `EMAIL_ADDRESS` is the address to which an email notification should be sent when the table reindex completes. Type `null` for no email notification.

   Example
   For example, to regenerate the kb_knowledge table text search index, and specify that no email notification is required after the rebuild completes, create the following background script:

   ```
   new GlideTextIndexEvent ( ).indexUpdate('kb_knowledge',
    'null');
   ```

   The script creates an event that the index handler processes, so the script completes without blocking on completion of the reindexing job.

3. Select **Run script**.

Result

The system reindexes the text search for the specified table. For large tables, such as sys_metadata, reindexing can take a couple of hours, but you can still perform text searches on the table while the script is running.

Regenerate the text index for a single record

Update the text search index for a single record. Use this approach to quickly verify whether text indexing is the cause of a search issue without rebuilding the full text index for an entire table.

Before you begin

Identify the following:

- A search term that doesn't return the expected results
- A record containing the identified search term that isn't returned when searching for the term
- The table that includes the identified record

The identified table must already be configured for indexing and searching. For details on this configuration process, see Configure a table for indexing and searching.

Role required: admin

Note:
If the *glide.script_processor.admin* system property's value has been changed from the default value, you need the role specified by that property value. For details on this property value and its effects, see Restrict access to background script [Updated in Security Center 1.3].

About this task

This reindexing method rebuilds the text search index for a single record. It takes less time than rebuilding the text index for all records in a table.

You can use this method to quickly determine whether text indexing is the cause of a search issue. Doing so avoids the delay associated with rebuilding the entire table's text search index.

Procedure

1. Copy the sys_id for the identified record by selecting **Copy sys_id** from its context menu.
2. Navigate to **All** > **System Definition** > **Scripts - Background**.

   Note:
   If the *glide.script_processor.admin* system property's value is set to security_admin, you must elevate to that privileged role to access this module. For more information on elevating privileges, see Elevated privilege roles.

3. In the **Run script (JavaScript executed on server)** text field, enter the following script, replacing `TABLE_NAME` with the name of the table containing the identified record and `SYS_ID` with the record's sys_id that you copied in step 1:

   ```
   var gr = new GlideRecord('TABLE_NAME');
   gr.get('SYS_ID');
   gs.eventQueue("text_index", gr,
    '[TABLE_NAME]','update','text_index');
   ```

4. Select **Run script**.
   The script creates a text_index update event for the specified table in the text_index queue.
5. Navigate to **All** > **System Policy** > **Events** > **Event Log** and wait for the new text_index update event's state to change to **processed**.
   To find the new event, you can apply a **[Name][contains][text_index]** filter and sort the Event list view by **Created** date/time in descending (z to a) order.

Result

The system updates the text index for the specified record in the specified table.

What to do next

Repeat your sample search. If this search returns the identified record, you may want to update the text index for the entire table that contains that record. For details, see Regenerate a text index for a table.

Remove an index

You can remove the index for a table if you no longer want the search engine to return results for that table. This procedure also removes the index for all tables that extend the specified table.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the dictionary record for the table.

   The dictionary record for a table is the record with **Table** matching the table's name, an empty column name, and a **Type** value of **Collection**, as shown in the following example image.

   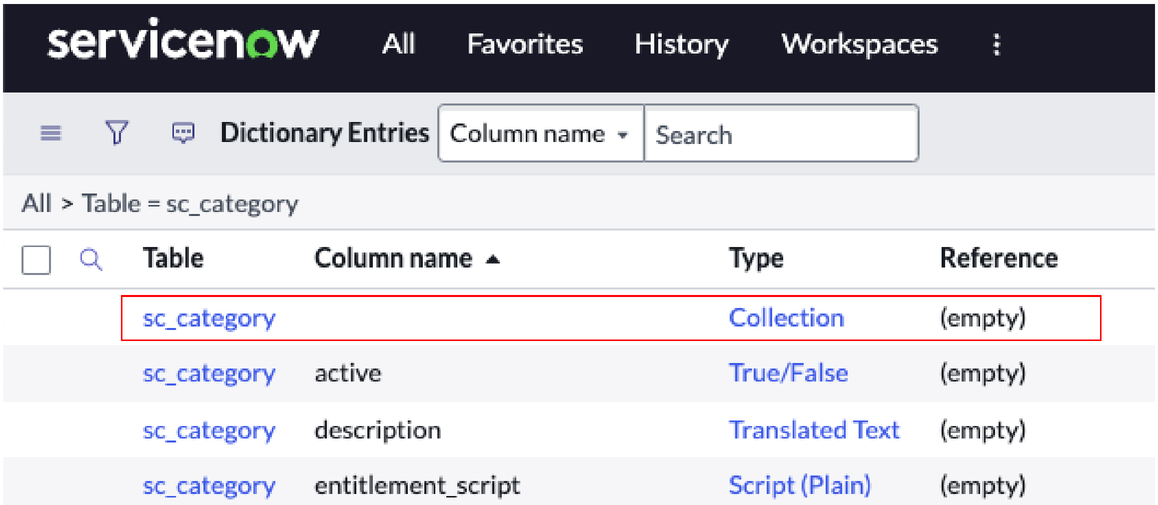

3. **Optional:** If the **Text index** option field is hidden, configure the form layout to show it.
   For details on showing and hiding fields on a form, see Configuring the form layout.
4. Clear the record's **Text index** option, then select **Update**.

Result

The system no longer indexes text from the specified table or queries it for text search results. This change also disables text indexing and search for all tables that extend the specified table.

Remove an index for a specific field

You can remove the index for a specific field in a table if you no longer want the search engine to return results for that field.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the dictionary record for the table field.

   The dictionary record for a table field is the record with **Table** matching the table's name and **Column name** matching the field's name, as shown in the following example image.

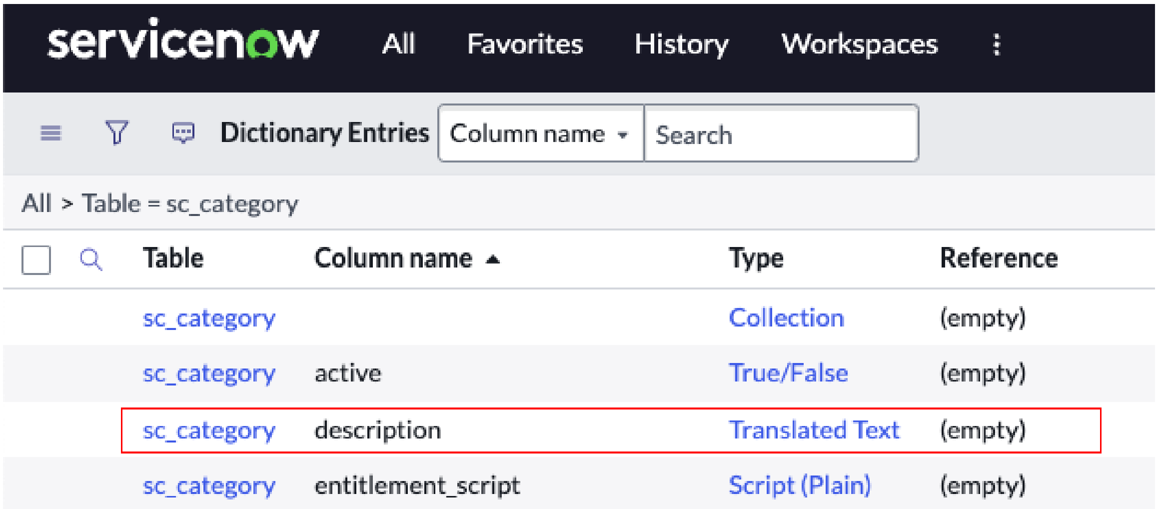

3. In the Attributes related list, select **New**.
4. On the Dictionary Attribute form, fill in the fields.

Field	Value
Attribute	No text index
Value	true

5. Select **Submit**.
6. Select **Update**.

Result

The system no longer indexes text from the specified table field or queries it for text search results. This change also disables text indexing and search for the specified field in all tables that extend the specified table.

Remove the text index for a child table

You can remove the text index for a child table if you no longer want the search engine to return results for that child table.

Before you begin

Role required: admin

About this task

The **Text index** option for a table is inherited by all of that table's child tables. Clearing this option for an individual child table has no effect.

For example, **Text index** is selected for the Task table in the base system. This setting enables text indexing for the Task table. All tables that extend Task, such as Catalog Task, Incident, and Request, inherit this option. Clearing the **Text index** option in the Incident child table doesn't disable text indexing for Incident records.

To disable text indexing for a specific child table, set that table's *`No text index`* attribute. This attribute overrides the table's inherited **Text index** option without affecting indexing behavior for other tables in its hierarchy.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Open the dictionary record for the child table.

   The dictionary record for a table is the record with an empty column name and a **Type** value of **Collection**, as shown in the following example image.

   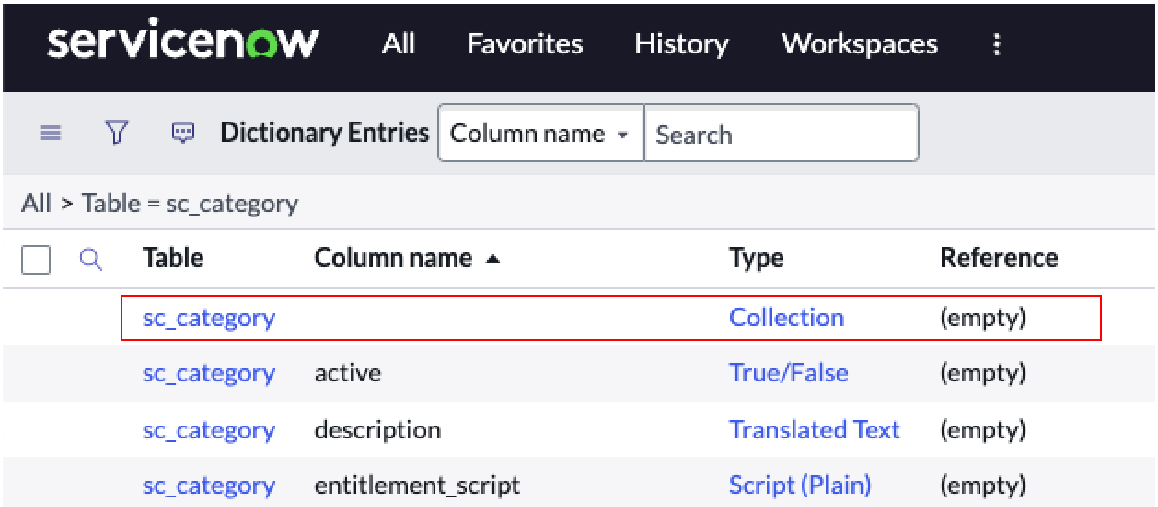

3. In the Attributes related list, select **New**.
4. On the Dictionary Attribute form, fill in the fields:

Field	Value
Attribute	No text index
Value	true

5. Select **Submit**.
6. Select **Update**.

Result

The system no longer indexes or queries the child table for text search results. Indexing continues for other tables in the table's hierarchy.

Change the query mode of an indexed table

Set the query mode to specify whether searches on an indexed table must match all or only some of the search terms.

Before you begin

To change the query mode of an indexed table, a text index must be present for the table. Make sure that text indexing is enabled for the table before changing the query mode. See Configure a table for indexing and searching.

Role required: ts_admin or admin

About this task

Once a table is indexed, the data within the table can be searched through the global search bar. Change the query mode for the indexed table to allow global search to match all or only some of

the searched terms. Enable a partial match rule to return results that partially match the searched terms.

Procedure

1. Navigate to **All** > **System Definition** > **Text Indexes**.
2. Select the table you want to configure from the list of indexed tables.
3. Use the combo box to select a **Default query mode**.

Option	Description
Default	The default query mode for a table is an AND query.
AND	Matches all search terms.
OR	Matches one or more search terms.
AND_OR	Matches all search terms. If no results are found, matches one or more search terms.

4. To enable partial matching of search terms, configure the partial match rule in the **Partial Match Rule** field.
Partial match rules return results that partially match the search terms. Use a partial match rule if queries often include many keywords (if, for example, users cut and paste phrases into the global search bar).

Partial match rules

Input type	Example	Description
Integer	3	Indicates the minimum number of search terms to be matched.
Negative integer	-2	Indicates that the number of search terms, minus this number, should be matched.
Percentage	75%	Indicates the minimum percentage of search terms to be matched. The minimum number of search terms is the number computed from the percentage rounded down.
Negative percentage	-25%	Indicates that the number of search terms, minus this percentage, should be matched. The minimum number of search terms is the number computed from the percentage rounded down and subtracted from the total number of search terms.
Combination	3<70%	A positive integer, followed by a less than symbol, followed by any of the previously mentioned input types. Indicates that if the number of search terms is equal to or less than the integer, all search terms are required. If the number of search terms is greater than the integer, the specification applies. In this example: if there are 1–3 search terms, all are required. If there are 4 or more search terms, only 70% are required.

Input type	Example	Description
Multiple combinations	3<-25%, 9<-5	Multiple conditional specifications separated by commas. In this example: if there are 1–3 search terms, all are required. If there are 4–9 search terms, all but 25% are required. If there are more than 9 search terms, all but 5 are required.

Enable indexing of text in multi-row variable sets

Make text from multi-row Service Catalog variable sets searchable in global search.

Before you begin

Role required: admin

About this task

By default, Zing doesn't index text from multi-row Service Catalog variable sets. Global search doesn't return results for text in these variable sets.

Administrators can override this default behavior by adding the *glide.ts.index.variableset* system property and setting its value to **true**. With this value set, Zing indexes text in multi-row variable sets, and global search returns results for that text.

To learn more about variable sets, see Service catalog variable sets.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Search the table for a system property with name *glide.ts.index.variableset*.
   - If the property already exists, change its value to **true** and skip the following steps.
   - If the property doesn't exist, go on to the following steps.
3. Select **New**.
4. On the System Property form, fill in the fields.

Field	Value
Name	glide.ts.index.variableset
Type	true | false
Value	true

5. Select **Submit**.

Result

Zing indexes text from multi-row Service Catalog variable sets so that you can find it using global search.

What to do next

To see text from multi-row variable sets immediately in global searches, regenerate the text indexes for the affected tables. For details on this procedure, see Regenerate a text index for a table.

If you don't regenerate the text indexes for the affected tables, records on those tables will be reindexed as they are updated or added. Global searches will only return text from multi-row variable sets on reindexed records.

Text indexing statistics and status

To view text indexing statistics and status, navigate to **All** > **System Definition** > **Text Indexes**.

- If text indexing is in progress, view the status of each index in the **State** field. Refresh the list to view updates. Tables are indexed one at a time.
- To see how long it takes to index a specific table, view the **Last indexing duration** field for the table entry. Although each process varies based on activity and data, historical data can provide a good estimate.
- View statistics (such as the number of indexed rows, terms, and queries in the past week) in the list or on the form for a specific table.

Text Index Status

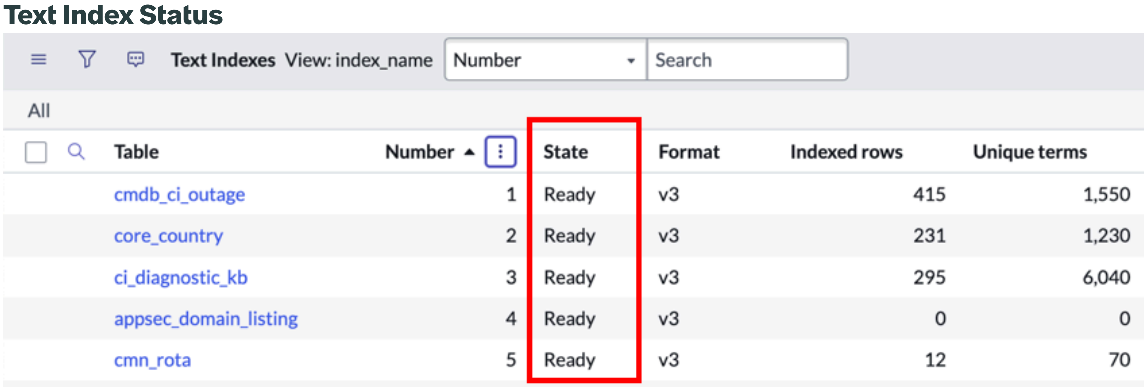

Text Index Statistics

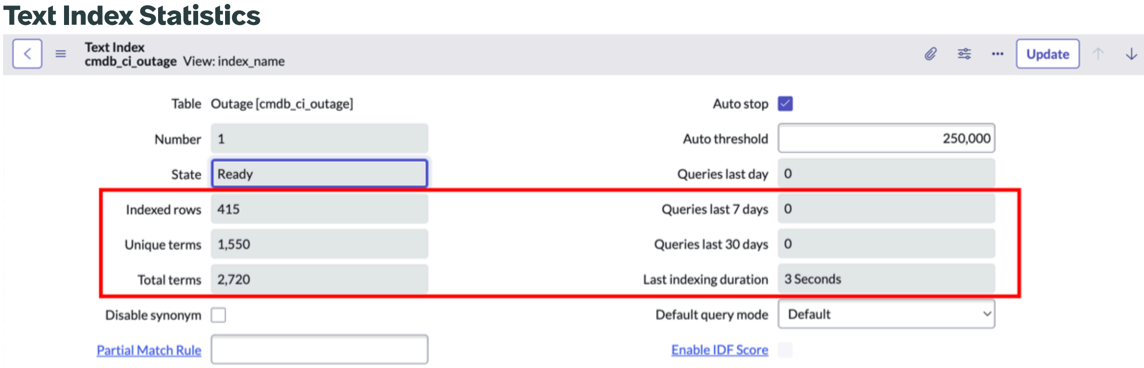

Configure tables to use the Japanese tokenizer

Improve search results in Japanese language searches by configuring individual tables to use the Japanese tokenizer.

Before you begin

Role required: admin

About this task

Japanese language searches use the Kuromoji Japanese tokenizer in Zing. The Japanese tokenizer attribute is active by default, but you must configure individual tables to use it and then regenerate the text index.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. Search for the table name.
3. In the search results, open the table, which is the list item that has an empty **Column name** and a **Type** of **Collection**.
4. Scroll down the page and select the **Attributes** tab, and then select **New**.
5. Fill in the fields on the form.

   Attributes fields

Field	Description
Attribute	Enter **Text index tokenizer language**.
Dictionary entry	Defaults to the name of the table.
Value	Enter **ja** for the Japanese tokenizer.

6. Select **Submit**.
7. Regenerate a text index for a table.

   Note:
   Text indexing can be a resource-intensive task that may take a while to complete. You may notice performance degradation or incomplete search results during index generation. To estimate text indexing duration, you can view historical statistics.

Result

When the text index is generated and you go to the table you configured, the table now uses the Japanese tokenizer.

Zing can include attachments in search results

Search content from attachments on indexed tables. Display attachments for search results from the Knowledge [kb_knowledge] table.

Sample Knowledge search showing attachments with matching content

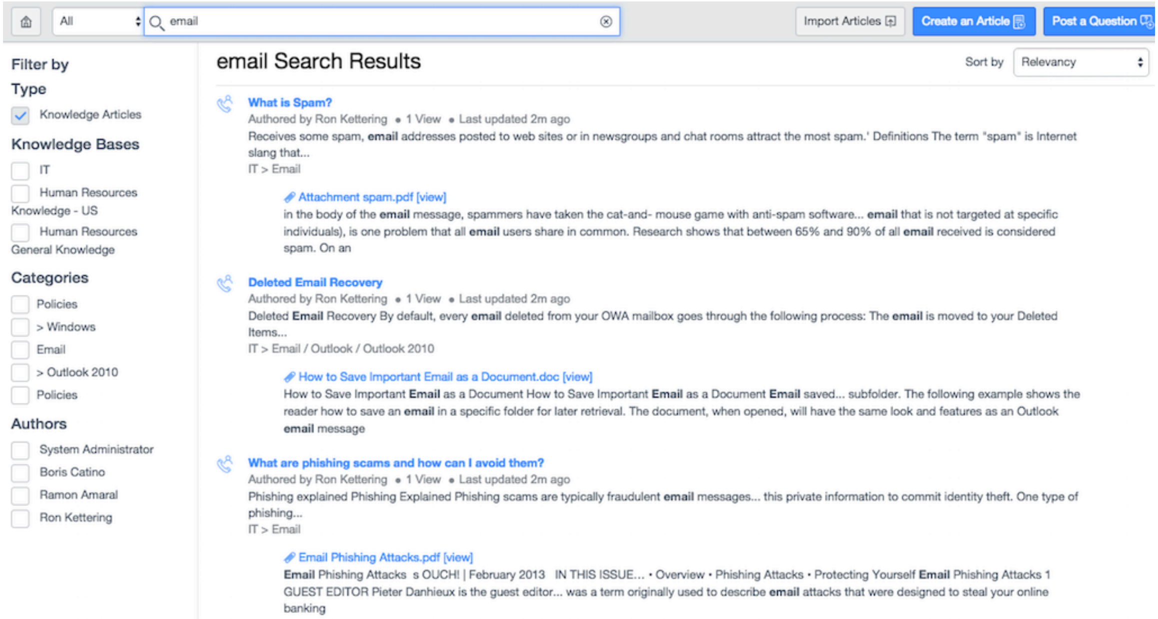

By default, search only matches content from attachments on Knowledge [kb_knowledge] records. Administrators can enable search for attachments on other tables, but doing so causes the system to re-index the selected table, its parent table, and any children of the parent table.

> **Warning:**
> For large tables, such as the Task table, re-indexing can take several hours and slows down the system until complete. Re-indexing is best performed during non-peak times.

The search results page only displays attachments for search results from the Knowledge [kb_knowledge] table, even if other tables have attachment search enabled.

Global search displays attachments for Knowledge search results but not for search results from other tables

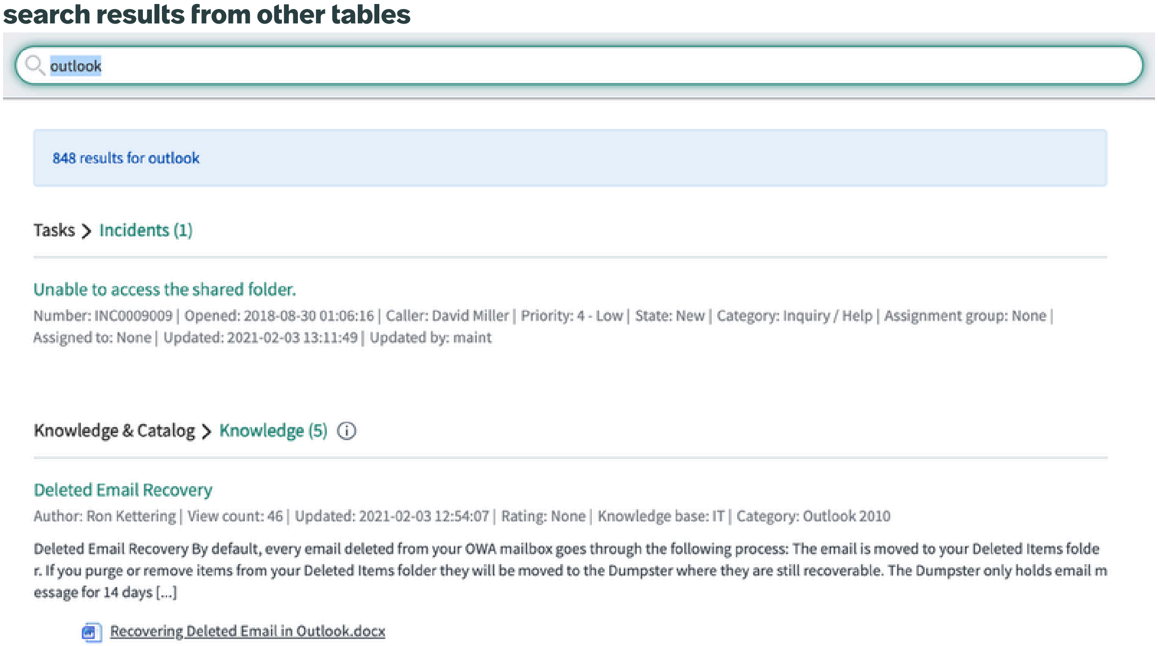

Zing supports indexing and searching these attachment file types.

- .doc
- .docx
- .dot
- .dotx
- .htm
- .html
- .ini
- .pdf
- .pot
- .potx
- .ppt
- .pptx
- .reg
- .txt
- .xls
- .xlsx
- .xlt
- .xltx

Related topics

Index attachments on a table

Index attachments on a table

You can enable attachment indexing for a table so text searches can return matches from the record and its file attachments.

Before you begin

Role required: admin

About this task

By default, attachment indexing is enabled for the Knowledge Base. You can enable attachment indexing for other tables. Enabling attachment indexing causes the platform to reindex the selected table, its parent table, and any children of the parent table. For large tables, such as the Task table, reindexing can take several hours and slows down the system until complete. Reindexing is best performed during non-peak times.

Procedure

1. Navigate to **All** > **System Definition** > **Text Index Configurations**.
2. Edit the record for the table that you want to enable attachment indexing for.
3. In the Text Index Table Attribute Maps related list, select **New**.
4. On the Text Index Table Attribute Map form, fill in the fields.

Field	Value
Table	<table name>

Field	Value
Attribute name	Attachment index
Value	true

Note:
The attachment index attribute only applies to the tables on which you explicitly add it. It does not cascade to child tables. For example, enabling indexing of attachments on the Task table does not enable indexing of attachments on the Incident table. Not all file types are supported for attachment indexing. For a list of supported file types, see Zing can include attachments in search results.

5. Select **Submit**.
   A message appears indicating that the attachment index attribute change will not take effect until you reindex the selected table.
6. Dismiss the reindexing message by selecting **OK**.
7. On the Text Index Configuration form, select the **Generate Text Index** related link.
8. In the confirmation window, perform one of these two steps:
   - To receive an email notification from the system when text indexing for the table is complete, enter your email address and select **OK**.
   - If you do not want to receive an email notification, select **Do not notify me**, then select **OK**.

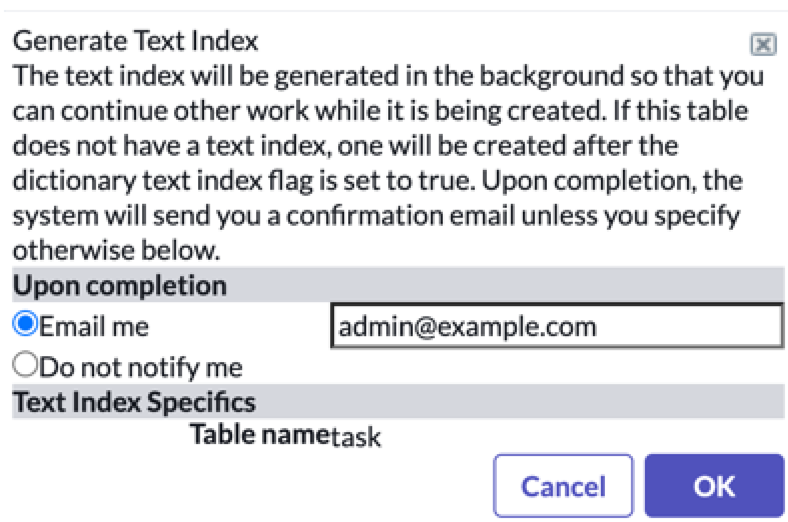

9. Dismiss the confirmation message by selecting **OK**.

Result

The system begins text indexing for the table. When it is complete, attachments can be searched on that table.

Zing removes stop words from queries

Remove common words from search queries that don't produce meaningful results.

Stop words are common words that aren't indexed because they aren't meaningful in search results. Articles, conjunctions, personal pronouns, and prepositions are examples of stop words that aren't used in keyword searches. Administrators can configure stop words for all indexed tables and for specific tables.

By default, the system maintains two types of stop words.

Types of stop words

Stop word type	Description
System-wide text index stop words	The system always ignores system-wide text index stop words when generating text indexes. Any search for a system-wide stop word returns no search results.
Table-specific stop words	The system uses the table-specific Text Index record to determine whether to index the stop word or to just remove it from keyword search queries against the table.

By default, the system has stops words for common English words. Search administrators typically create stop words from search terms that produce too many search results such as articles, conjunctions, personal pronouns, and prepositions.

Configure a global stop word

Configure stop words that shouldn't be indexed by the search.

Before you begin

Role required: ts_admin or admin

Procedure

1. Navigate to **All** > **System Definition** > **Text Index Stop Words**.
2. Add or remove stop words from the list.
3. If a message appears at the top of the list, contact Technical Support to regenerate all indexes.
   You must regenerate indexes whenever words may be missing from an index. For example, if you delete, inactivate, or change an active global stop word, the word may be missing from the index. An after business rule checks these conditions and generates the notification message when index regeneration is necessary.

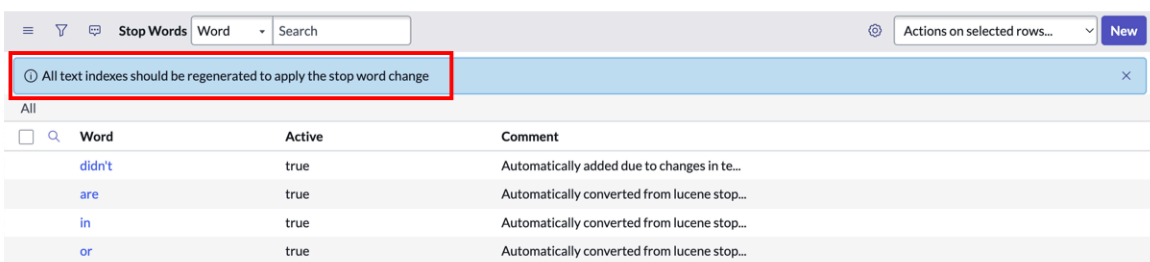

Configure a table-specific stop word

You can configure stop words for a specific table.

Before you begin

Role required: ts_admin or admin

Procedure

1. Navigate to **All** > **System Definition** > **Text Indexes**.
2. Open the text index entry for the table.
3. Add or remove stop words from the **Index Stop Words** related list.
   For each word, select the desired **Stop mode** (Not a Stop Word (inactivates the stop word), Neither Index nor Query, or Index but do not Query).
4. If a message appears at the top of the form, select **Regenerate Text Index**, then select **OK**.

You must regenerate indexes whenever words may be missing from an index. For example, if you delete or change a stop word for which the mode was Neither Index nor Query, the word may be missing from the index. An after Business rules checks these conditions and generates the notification message when index regeneration is necessary.

Note:
Text search uses the global list plus the table-specific list of stop words when indexing the table.

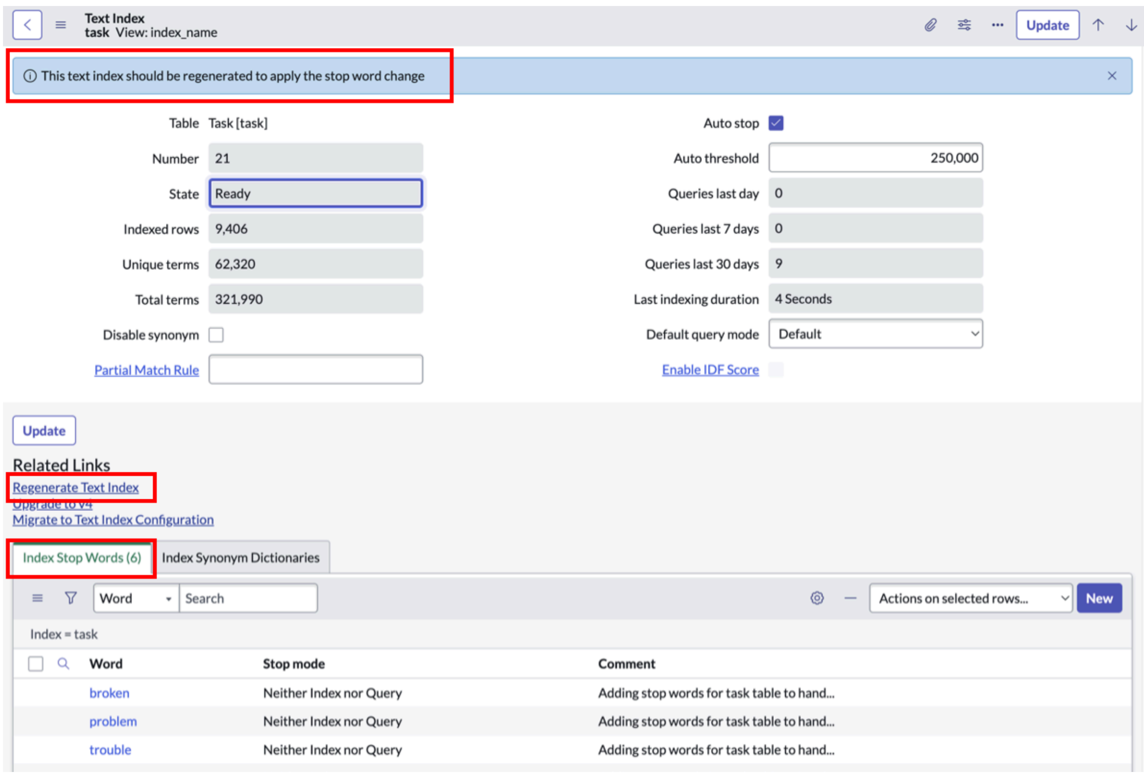

Enable automatic stop words for a table

The system can identify and generate stop words when a search term exceeds an occurrence threshold.

Before you begin

Enable text indexing for the table.

Role required: ts_admin or admin

About this task

By default, the TS Index Stats scheduled job identifies and creates stop words for tables with text indexes on a nightly basis. Automatically adding stop words improves your search results by removing search terms that return too many search results.

Procedure

1. Navigate to **All** > **System Definition** > **Text Indexes**.
2. Open the text index entry for the table.
3. Select the **Auto stop** check box.
4. In **Auto threshold**, enter the maximum number of occurrences a search term can have in search results.

When a search term generates more search results than the threshold, the system automatically creates a stop word for the search term.

Example
For example, to create stop words for the task table, see the blog post Configuring auto stop words and regenerating text indexes from a ServiceNow employee on the ServiceNow Community.

5. Select **Update**.

Result

The job identifies and creates stop words for the table with a **Stop mode** value of `Index but do not Query` and inserts a **Comment** to indicate that the stop word was generated automatically.

Disable a stop word in Zing

Configure Zing to return results for text search terms that match a stop word's stemmed form.

Before you begin

Role required: admin

About this task

Zing's stemming feature can cause search to return no results for an acronym or short word if its stemmed form is defined as a stop word.

Indications that a search term may be stemmed to a stop word include:

- Text search returns no results for the affected term.
- The system displays an informational message: `Your text query contained only common words or ambiguous wildcards, please refine your search and try again.`

As an example, imagine that a user searches for the acronym `AMS`. Zing stems this search term to `am`. Because `am` is defined as a stop word, Zing ignores the search term. The user's `AMS` search returns no results, and the system displays the informational message.

To enable correct search results for an affected term, administrators can disable the relevant stop word in the global and table-specific stop word lists.

Note:
Disabling a stop word can impact search performance and relevancy. Text search will return matching results for the disabled stop word as well as for its stemmed forms.

For more details on stemming in Zing, see Zing matches derived words with stemming.

Procedure

1. To disable a stop word in the global stop word list, perform the following steps:
    a. Navigate to **All** > **System Definition** > **Text Index Stop Words**.
    b. Open the record for the stop word you want to disable.
    c. On the Stop Word form, clear the **Active** option, then select **Update**.
2. To disable a stop word in table-specific stop word lists, perform the following steps for each table that includes the stop word:

a. Navigate to **All** > **System Definition** > **Text Indexes**.

b. Open the table's text index record.

c. In the Index Stop Words related list, open the record for the stop word.

Note:
If the Index Stop Words related list isn't visible, open the form context menu and select **View** > **Default view**.

d. On the Index Stop Word form, set the stop word's stop mode to **Not a Stop Word** and select **Update**.

e. On the Text Index form, select the **Regenerate Text Index** related link.

Result
Zing no longer ignores the disabled stop word and its stemmed forms in text searches.

Zing matches derived words with stemming

Convert any multiple-character search keyword to its stem form to find derived versions of the word.

Stemming removes a variety of common word inflections, such as plurals and past tense forms. Zing treats all words with a common stem as synonyms of the original search term. Stemming doesn't apply to searches involving single-character words such as Chinese and Japanese pictograms.

Zing uses the Porter Stemming Algorithm, which is most effective for English text, but supports stemming in these languages.

- English
- French
- German

The stemming language determines how the system generates indexes of record data. The system only supports one stemming language at a time regardless of how many languages the instance supports. When you change the stemming language, the system creates any new index entries based on the new language stemming rules, but doesn't regenerate any existing index records. You must manually regenerate indexes for tables you want to use the new stemming language.

Zing can expand search results with synonyms

Expand search results to include additional search terms.

By default, Zing includes a sample synonym dictionary consisting of a list of synonym sets (synsets). Zing supports two types of synonym sets.

Synonym set types

Synonym set type	Description	Example
Synonym list	A synonym list is a list of words separated by commas. When a search uses any word	Cyber,Internet,Cybercafe

Synonym set types (continued)

Synonym set type	Description	Example
	from the list, Zing expands the search to include all words in the list.	
Synonym map	A synonym map is two word lists separated by the equal and greater than characters => . When a search uses any word from the list on the left, Zing expands the search to include all words from the list on the right. When a search uses a word from the list on the right, Zing doesn't expand the search to include terms from the list on the left.	IOT,Internet of Things=>Chip,Nest thermostat,RFID,IOT

Search results, synonym expansion and weighting

In search results, documents containing the original search keywords have a higher score than documents containing the search synonyms. For example, documents with the search keyword `IOT` have a higher score than documents with the synonym `Internet of Things`.

- One-way synonym expansion to/from multiple terms `(blork, cork => malork, elastic)` works as expected. `Blork` and `cork` both expand to `malork` and `elastic`.
- Multi-level one-way expansion `(blork => cork => malork)` doesn't work transitively. Neither `blork` nor `cork` expands to `malork`.
- Multi-level one-way expansion split between multiple synonym set records `(blork => cork)` `(cork => malork)` doesn't work transitively. Only `cork` expands to `malork`, `blork` only expands to `cork` but not `malork`.

By default, synonyms don't use the same weight as regular search terms. Synonyms are only weighted at 10% of the default search weight. Additionally, multiple instances of a synonym in a single doc/field don't increase the search weight. The weight of synonyms is controlled by the property *`glide.ts.synonym.expanded.boost`*. When this property value = 1, synonyms are weighted equally to normal search terms. 0.5 = 50% weight etc.

Wildcard characters and synonym matching

Zing doesn't expand the `%` and `*` wildcard characters when looking for synonyms that match your search term. These characters are treated as literal characters when finding matching synonyms. As an example, if you search for `issu*`, Zing doesn't find synonyms for `issue`, but it does find synonyms for `issu*`.

Enable search synonyms

Enabling text index synonyms allows Zing to expand text searches with additional search keywords.

Before you begin
Role required: ts_admin or admin

About this task
By default, text searches don't use text index synonyms.

Procedure

1. Navigate to **All** > **System Properties** > **Text Search**.
2. Scroll down to **Enable Synonym**, and select **Yes**.

Result

Text searches use active synonym dictionaries to expand keyword searches.

Create synonym dictionaries

You can create synonym dictionaries to expand Zing text searches with additional search keywords.

Before you begin

Role required: ts_admin or admin

About this task

By default, the system provides a sample Zing synonym dictionary called **Example synonym dictionary**. You can add synonyms to this dictionary or create your own synonym dictionaries.

Procedure

1. Navigate to **All** > **System Definition** > **Text Index Synonym Dictionaries**.
2. Select **New**.
3. On the Synonym Dictionary form, fill in the fields.

   Synonym Dictionary form

Field	Description
Name	Unique name for the synonym dictionary.
Active	Option to make the synonym dictionary active. Zing only uses active synonym dictionaries when expanding searches.
Global	Option to use the synonym dictionary for searches against all tables that have not been individually configured.
State	Publication state for the synonym dictionary. This field is read-only.
Order	Order of evaluation for the synonym dictionary. Zing evaluates synonym dictionaries with lower `Order` values before dictionaries with higher values.
Description	Description of the synonym dictionary.

4. For each set of synonyms you want to add to the synonym dictionary, insert a new Synonym Set row and fill in the fields.

   Synonym set fields

Field	Description
Synset	Comma-separated set of synonym expansions. Zing expands searches

Field	Description
	containing any one of these terms to include all of these terms. As an example, if you define the synonym set `medicalise,médicalisé` and search for `medicalise`, Zing expands your search to match records containing either `medicalise` or `médicalisé`.
Active	Option to make this synonym set active. Zing only uses active synonym sets when expanding searches.
Description	Description of the synonym set.

5. Select **Submit**.
6. Select the **Publish All Dictionaries** related link.
   A message indicates that publication of your synonym dictionaries has been scheduled. Refresh the page to see whether publication is complete. The dictionary is published when its state changes from **Draft** to **Published**.

Select synonym dictionaries for a table

Search administrators can specify one or more synonym dictionaries available to a particular table.

Before you begin

Role required: ts_admin or admin

Procedure

1. Navigate to **All** > **System Definition** > **Text Indexes**.
2. Select the table for which you want to change search behavior.
   The system displays the Text index record for the table.
3. Select the **Index Synonym Dictionaries** tab.
4. Select **New**.
5. Select a synonym dictionary in the **Synonym Dictionary** field on the **Index Synonym Dictionary** form.
6. Select **Submit**.

Result

Text searches on this table use the selected synonym dictionaries.

Disable synonyms for a table

By default, Zing uses all active synonym dictionaries when a user performs a text search. Search administrators can disable synonym dictionary usage on a per-table basis.

Before you begin

Role required: ts_admin or admin

Procedure

1. Navigate to **All** > **System Definition** > **Text Indexes**.
2. Select the table for which you want to change search behavior.

The system displays the Text index record for the table.

3. Select **Disable synonym**, so the check appears.
4. Select **Update**.

Debug synonym searches

You can view debug messages to determine if synonym searches are working as expected.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug Text Search**.
2. Perform a search using a term in a synonym dictionary.
3. View the session debug output.

   Note:
   You can only view the debug output on a table view (_list.do). You can't view it on a UI page itself.

Zing displays search suggestions as users enter search terms

Display possible search query completions as users enter search terms.

To use the Search Suggestions application, activate the com.glide.search.suggestions plugin if it's not already installed. For information, see Request a plugin and Activate a plugin.

Once you enable Search Suggestions, suggestions appear automatically when users enter text into Zing search fields in the Service Portal or Now Mobile.

Automatic search suggestions

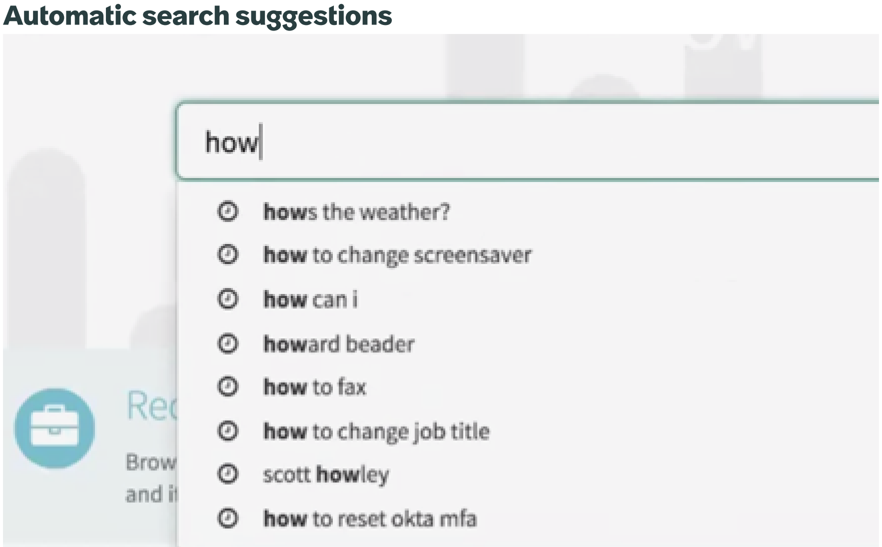

If there are any matching suggestions from the user's own previous searches, those suggestions appear first, with a clock icon. The remaining suggestions are those suggestions coming from the searches of other users.

In the Service Portal, you can use Search Suggestions or type-ahead functionality, but not both at the same time. In Now Mobile, only Search Suggestions is available. By default, new ServiceNow instances use Search Suggestions in the Service Portal. Upgraded instances use

whatever was previously enabled. You can easily enable Search Suggestions on upgraded instances. To learn more, see Enable and disable Search Suggestions in Zing.

For full details on configuration and use of the Search Suggestions application, see Search Suggestions.

Enable and disable Search Suggestions in Zing

Enable the Search Suggestions application to improve the Zing search user experience.

Before you begin

ServiceNow personnel installed and activated the com.glide.search.suggestions plugin.

Role required: admin

About this task

The Search Suggestions application is available on new instances by default. On upgraded instances, Search Suggestions is enabled for Now Mobile and disabled for Service Portal by default.

The following system properties enable and disable Search Suggestions:

- *`glide.search.suggestions.enabled`*: Set to **true** (default) to enable Search Suggestions in Zing, or set to **false** to disable Search Suggestions in Zing.
- *`glide.service_portal.search_as_you_type_behavior`*: Select **Suggestions** or **Typeahead** functionality in the Service Portal.

On new instances, the first parameter is set to **true** and the second is set to **Suggestions**. In upgraded instances, the parameters are set to whatever was previously enabled.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. To switch between Search Suggestions and type-ahead functionality in Service Portal, follow these steps.
   a. In the **Name** search column, enter *`glide.service_portal.search_as_you_type_behavior`*.
   b. On the System Property page, in the **Value** field, enter `Suggestions` to enable Search Suggestions, or `Typeahead` to enable type-ahead functionality.
      You can only enter one of these options.
3. To turn Search Suggestions on or off globally, follow these steps.
   a. In the **Name** search column, enter *`glide.search.suggestions.enabled`*.
   b. In the **Value** field, double-click the value and select **true** to enable suggestions, or **false** to disable Search Suggestions globally.

Set the maximum number of search suggestions Zing displays

Set the maximum number of search suggestions Zing displays when users enter search strings in search applications.

Before you begin

You must be using a search application that has Zing specified as its Search Engine in its search application configuration record. For details on search application configuration record fields, see Search Application Configuration form.

Role required: admin

About this task

Zing displays suggestions generated by the Search Suggestions application when users enter text in a search field in a supported search application. You can control the number of suggested search strings displayed by editing the search application configuration records for your search applications that use Zing.

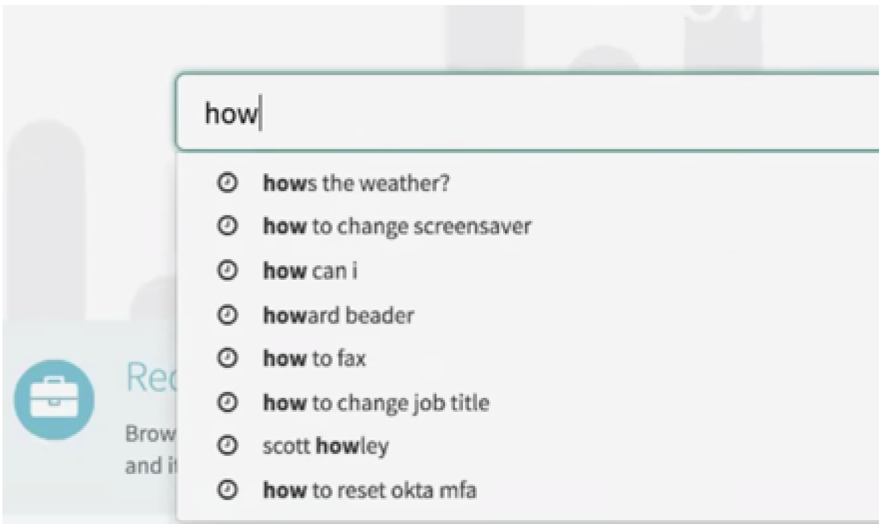

Note:
You can also configure a search suggestion limit for Knowledge and System Status searches in Service Portal by editing the instance options for the Typeahead Search widget. For details on these instance options, including the Limit option, see Typeahead Search widget. To learn how to configure instances of the Typeahead Search widget for the Knowledge and System Status portal pages, see Configure widget instances.

Procedure

Set the maximum number of search strings displayed in a search application that uses Zing.

a. Navigate to the Search Application Configurations [sys_search_context_config] table's list view by selecting **All**, entering `sys_search_context_config.list` in the navigation filter, and pressing Enter.
b. Open the Search Application Configuration record for your Zing search application.
c. If a warning message about the current application scope appears, select the link in the message to switch to the correct application scope for the record.
d. In the **Suggestions Limit** field, enter the new maximum number of search suggestions you want the application to display.
e. Select **Update**.

Installed with Zing

Several types of components are installed with Zing.

Tables installed with Zing

Table	Description
Text Index [ts_index_name]	Stores the tables the system indexes.
Index Stop Word [ts_index_stop]	Stores the stop words for a specific table.
Stop Word [ts_stop]	Stores the global stop words.
Text Search Groups [ts_group]	Stores search groups for global text search.
• ts_attachment • ts_chain_summary • ts_chain • ts_deleted_doc • ts_document • ts_index_stats • ts_phrase • ts_search_stats • ts_search_summary • ts_word_roots • ts_word	System tables that support Zing. Extending or modifying these tables isn't recommended.

Business rules installed with Zing

Business Rule	Description
Text Search Property Change Rationally	Ensures that valid values are entered for Zing text search properties.
Text Index Stop Reminder	Warns the user of stop word changes that require the index to be rebuilt (table-specific). The warning is issued if record is deleted that had a stop mode "neither index nor query", if record's stop mode is updated to something else and was "neither index nor query", and if record's word is updated to something else and stop mode is "neither index nor query".
Stop Word Reminder	Warns the user of stop word changes that require the index to be rebuilt (global). The warning is issued if record is deleted and it was active, if record is inactivated, and if record's word is changed and was active.

Scheduled jobs installed with Zing

Scheduled job	Description
TS Search Stats	Compiles type-ahead suggestions each night. See Update A Type-Ahead Suggestion.
TS Index Stats	Collects statistics and performs maintenance for text search and indexing. Runs nightly.
text index events process	Collects statistics and performs maintenance for text search and indexing. Runs every 30 seconds.
TS Chain Summary	Compiles search chain statistics each hour.

UI action installed with Zing

UI action	Description
Regenerate Text Index	Displays the **Regenerate Text Index** link on Text Index forms.

Zing roles

Zing is installed with these roles.

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

Search application administrator [search_application_admin]

Create, read, update, and delete search application configurations for Zing and AI Search.

Contains Roles

List of roles contained within the role.

- ais_admin
- personalize_dictionary

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Grant this role to users who need to configure AI Search applications.

Text search administrator [ts_admin]

Configure text indexes, text index groups, text index stop words, and text index synonyms for Zing.

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.

AI Search

The ServiceNow® AI Search application provides a consumer-grade search engine for ServiceNow® Service Portal, ServiceNow® Now Mobile, and ServiceNow® Virtual Agent . Intelligent query features help you quickly find the answers you need.

AI Search overview video, approximately three minutes and six seconds long.

Watch this short video to see how users can find answers using AI Search.

Get started

Explore Learn about AI Search features.	**Configure** Configure searchable content, search experiences, and ServiceNow AI Platform® search applications.	**Integrate** Integrate AI Search functionality into other ServiceNow AI Platform applications.
Use Find your answers using AI Search.	**Administer** Gain deeper insights into AI Search behavior.	**Reference** Get details on AI Search system properties and features.

Note:
AI Search isn't available in personal developer instances (PDIs).

Helpful resources

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

ServiceNow Community

ServiceNow Community

Customer Success Center

Impact

Developer

developer.servicenow.com

Impact

http://impact.servicenow.com

ServiceNow University

ServiceNow University

NowCreate

Best Practices

Partner

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

ServiceNow

https://www.servicenow.com

ServiceNow Store

https://store.servicenow.com/

Support

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

Extend AI Search with apps from the ServiceNow® Store

Analyze trends for search query traffic, review indexed record counts and configuration settings, and preview search query results with Advanced AI Search Management Tools.

Exploring AI Search

Learn about AI Search features that enable you to provide consumer-grade search experiences for your users.

AI Search architecture

The AI Search index stores data from ServiceNow AI Platform® records or external sources and makes that data available for users to search in multiple applications. Search query features use ServiceNow AI Platform technologies to improve the search user experience.

AI Search architecture diagram

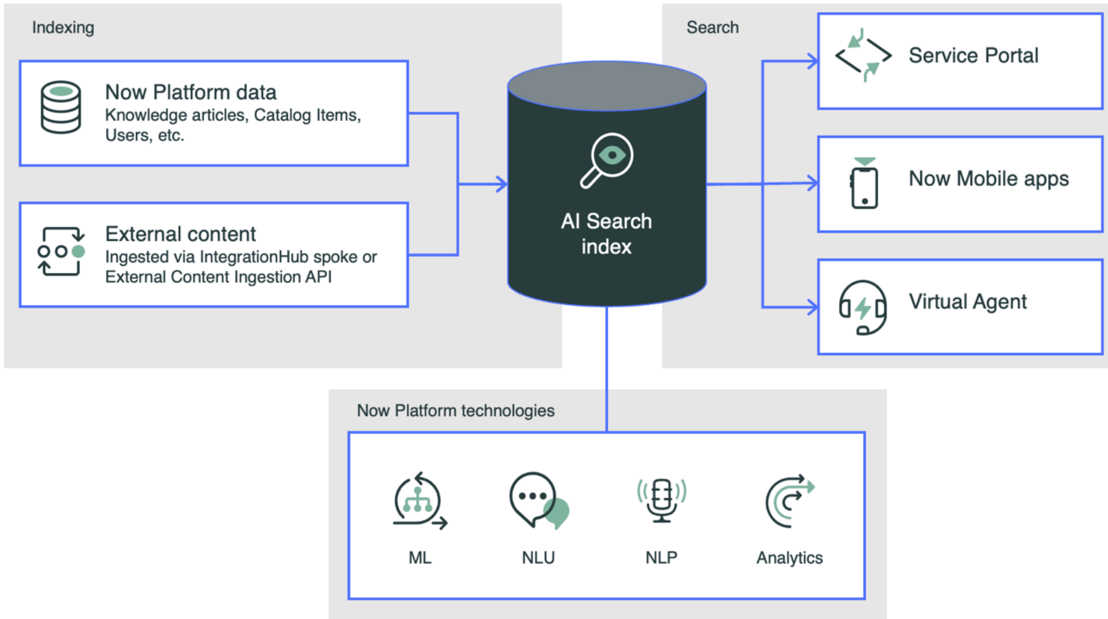

AI Search configuration

Search administrators configure AI Search indexing and search settings in indexed sources, search sources, search profiles, and search application configurations.

AI Search configuration diagram

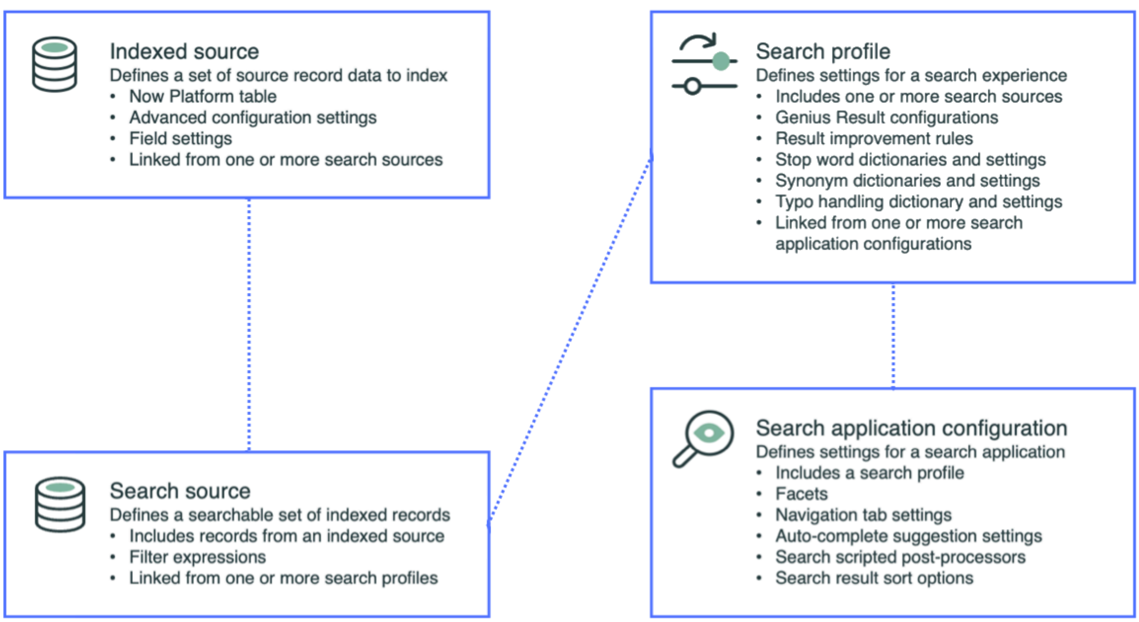

Index searchable content from ServiceNow AI Platform® tables and external sources

Make record content searchable by configuring an indexed source to add it to the AI Search index.

Define an indexed source for each ServiceNow AI Platform table that you want to index records from. Search administrators can configure settings for each indexed source:

- Enable indexing for child tables of the source table.
- Configure indexing of source record tags, attachments, translated fields, and referenced fields.
- Map source table fields to AI Search index fields for ease of search.

The base system includes preconfigured indexed sources for Knowledge Base articles, Catalog Items, and user records. You can create your own indexed sources to index records from other ServiceNow AI Platform tables.

To learn more about configuring and using indexed sources for ServiceNow AI Platform table records, see Indexed sources.

By activating the External Content for AI Search plugin (com.glide.ais.external_content), you can enable indexing and search for record content from external sources. Supported sources include SharePoint collections, Confluence sites, and remote tables. For more details, see Indexing and searching external content in AI Search.

Control access to searchable content using search sources

A search source defines a subset of indexed content that users can search using a search profile. Create search sources by applying filter conditions to content from indexed sources.

Search administrators link search sources to search profiles, enabling precise control of the content that users can search. For details on creating search sources, see Search sources.

Define search experience settings in search profiles

Search profiles define search sources and settings for linguistic search query features. Each search profile represents an individual search experience that you want to offer users in a supported ServiceNow AI Platform® application.

Each search profile can include the following settings:

- One or more search sources defining content that users can search using the profile.
- Dictionaries of synonyms, stop words, and typo handling auto-correction terms applied to searches using the profile.
- Genius Result configurations applied to searches using the profile.
- Result improvement rules applied to searches using the profile.

Search administrators create search profiles, link search sources to them, configure their settings, and publish them. For information on creating, configuring, and publishing search profiles, see Search profiles.

Enable AI Search in supported ServiceNow AI Platform® applications

Provide multiple search entry points across your organization by enabling AI Search as the search engine for search applications. The base system enables AI Search as the search engine for global and workspace search, Customer Service Management, Now Mobile, Service Portal, and Virtual Agent .

Search administrators define search application configurations for ServiceNow AI Platform applications. Each configuration links to a search profile defining the search experience for the application. Administrators can configure the application's behavior in the following ways:

- Enable or disable typo handling auto-correction for search queries.
- Configure navigation tab and facet filters to enable users to refine search results.

- Limit the maximum number of search suggestions and Genius Result cards the application can display for a search.
- Limit the number of search results the application displays on each result page.

To learn more about configuring and enabling AI Search in supported ServiceNow AI Platform applications, see Search application configurations and Enabling and configuring AI Search in ServiceNow AI Platform applications.

Provide users with powerful and flexible search

AI Search includes search features that help users find the answers they need.

Powerful query language

Query for indexed terms and phrases. Control query logic with Boolean operators. Match a range of indexed terms using wildcard operators.

For a complete guide to supported query language syntax and operators, see AI Search query language.

Auto-complete search queries

Display suggestions in the input field as users compose their searches. Suggestions include recent and popular search queries and results as well as entries from the user's personal search history.

Auto-complete suggestions are linguistic features that search administrators configure in search application configurations. For more information on auto-complete suggestion type settings, see Configure an auto-complete suggestion type in an AI Search application configuration.

Language-sensitive lemma and Unicode normalization of search query terms

Support natural language search by expanding search query terms to match alternate inflections and forms. AI Search supports language-specific lemma normalization for Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese. It supports Unicode normalization for content in all ServiceNow AI Platform® languages.

Lemma and Unicode normalization are linguistic features that don't require configuration. For more details on normalization behavior, see Lemma and Unicode normalization.

Expand search query terms using configurable synonyms

Improve search recall by configuring language-specific dictionaries of terms with equivalent meanings or usage. When a search query includes any of these terms, AI Search expands it to include all equivalent terms, providing users with a more natural search experience.

Synonyms are linguistic features that search administrators configure in search profiles. For details on configuring synonyms and synonym dictionaries, see Synonyms. AI Search supports synonym dictionaries for Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese.

Remove excessively frequent search terms using configurable stop words

Increase search precision by configuring language-specific dictionaries of frequently appearing terms that dilute search relevancy. AI Search removes these stop words from user search queries to exclude irrelevant results.

Stop words are linguistic features that search administrators configure in search profiles. For details on configuring stop words and stop word dictionaries, see Stop words. AI Search supports stop word dictionaries for Brazilian Portuguese, Dutch, English, French - Canada, French, German, Italian, Portuguese, Spanish, and Swedish.

Auto-correct typos in search query terms

Automatically replace misspelled search query terms with spellings found in indexed content. Typo corrections are displayed above search results. Users can always choose to repeat the search with their original query terms.

Typo handling auto-correction is a linguistic feature that search administrators configure in search profiles. For details on controlling the set of available auto-correction terms, see Typo handling. AI Search supports derivation of auto-correction terms for Brazilian Portuguese, Dutch, English, French - Canada, French, German, Italian, Portuguese, Spanish, and Swedish.

Enable search result refinement filters

AI Search enables users to refine searches by filtering their results. Use navigation tabs to filter results by source, or select facets to apply filters dynamically generated from the result set.

Filter results by source using navigation tabs

Display pre-defined categories that users can select to immediately exclude unwanted search results.

Search administrators configure navigation tabs in search application configurations. For details on available settings, see Configure navigation tabs in an AI Search application configuration.

Filter results dynamically using facets

Display field value-based filter selections generated dynamically from the search query's result set. Users can select multiple facet values to narrow their area of focus. Facets only display when they produce results, so users never need to worry about dead ends.

Search administrators define facets in search application configurations. For details on creating and configuring facets, see Create a facet in an AI Search application configuration.

Display relevant and actionable search results

AI Search provides users with clear answers for their search queries.

Hit highlighting

AI Search highlights search query terms that appear in search results. This highlighting enables users to see which query terms contributed to record matches.

Display the most relevant results first with machine learning relevancy

AI Search orders search results in decreasing order of relevancy. Machine learning automatically tunes and improves search result relevancy scoring for each search profile based on aggregated user interaction data and A/B testing evaluations of live search traffic.

Machine learning relevancy is automatically enabled and not configurable, though search administrators can exclude specific search profiles from automatic relevancy-model updates. For more details on automatic relevancy score tuning, see Machine learning relevancy in AI Search.

Display the best answers as actionable Genius Result cards

Configure Genius Results to analyze search query intent and put the best answers first. Search users can read a Knowledge article, view a user's profile or organization chart, or request a Catalog Item directly from a Genius Result answer card.

Search administrators can link Genius Results to search profiles, modify their settings, and create new configurations. To learn more about Genius Results configurations, see Genius Results.

Boost, block, or promote search results

Define result improvement rules with configurable trigger conditions. A rule can boost search result relevancy based on the search query or user context, or can block or promote specific results for a particular search query.

Search administrators define result improvement rules in search profiles. For details on creating rules and configuring their boost, block, or promote actions, see Result improvement rules. AI Search supports result improvement rules for Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese search queries.

Content security in AI Search

When indexing records, AI Search preserves their ServiceNow AI Platform® security features. At query time, AI Search filters search query results and displays only records that the user can access.

For more information on how AI Search handles security features for ServiceNow AI Platform table records, see Content security in AI Search.

Internationalization support

AI Search supports indexing and search for all languages offered by the ServiceNow AI Platform®. Search administrators can configure AI Search search features for languages activated in your instance.

For full details on language support in AI Search search features, see Internationalization support for AI Search.

Search features

The following table displays the set of AI Search search features available for specific ServiceNow AI Platform languages.

Search features supported by language

Language	Search features supported
English	• Genius Results • Language identification • Lemma and Unicode normalization • Result improvement rules

Search features supported by language (continued)

Language	Search features supported
	• Stop words • Synonyms • Typo handling
Brazilian Portuguese, Dutch, French, French - Canada, Italian, and Spanish	• Language identification • Lemma and Unicode normalization • Result improvement rules • Stop words • Synonyms • Typo handling
German and Swedish	• Language identification • Lemma and Unicode normalization (including decompounding) • Result improvement rules • Stop words • Synonyms • Typo handling
Japanese, Simplified Chinese, and Traditional Chinese	• Language identification (including text-region and search-term identification) • Lemma and Unicode normalization • Result improvement rules • Stop words • Synonyms
Korean	• Language identification • Lemma and Unicode normalization (including decompounding) • Result improvement rules • Stop words • Synonyms
All other ServiceNow AI Platform languages activated in your instance	• Unicode normalization • Result improvement rules

Search features supported by language (continued)

Language	Search features supported
	• Stop words • Synonyms

Translated content search

By default, AI Search only returns results for records with translated content, such as Knowledge articles, that are in the user's ServiceNow AI Platform session language. You can modify this default behavior in several ways:

Assign Knowledge article search languages by country

Improve multilingual search recall by defining a set of Knowledge article search languages for each user country. User searches can match Knowledge articles in any of the search languages specified for their country as well as in the language associated with their ServiceNow AI Platform session.

Configure globally searchable Knowledge articles

Define an encoded query to match Knowledge articles that you want to make searchable regardless of the user's session language.

Configure a fallback language

Set a fallback language for the user's session language. AI Search returns translated content results from the selected fallback language as well as the user's session language.

Enable a global fallback locale for translated content

Configure a global fallback locale to use for all translated content searches. AI Search returns translated content results from the global fallback locale's language as well as the user's session language.

Note:
The global fallback locale works best if you want to make all English-language records globally searchable.

Use AI Search for global search and configurable workspace search in Next Experience

AI Search is the default search engine for Unified Navigation searches in Next Experience, including global search and configurable workspace search.

Search administrators can configure global search and configurable workspace search settings using the AI Search for Next Experience application.

Note:
The AI Search for Next Experience application is automatically installed and enabled in new Washington DC instances. If you upgrade to Washington DC from a previous release, your administrators can manually install and enable the application.

For details on AI Search for Next Experience, including configuration steps and a usage guide, see AI Search for Next Experience.

Deflect incidents in Service Portal with the AI Search Assist widget

Improve incident deflection for Service Portal users by displaying the most relevant search results within a record producer.

AI Search Assist can help an agent or user find the most relevant items in a search on your portal. For example:

- An agent, working on an Incident record, can use AI Search Assist to find and attach Knowledge articles, order a catalog item on behalf of the caller, or search for an article and link it to a related incident or problem.
- A user, creating an Incident via the Service Portal, can see related Catalog Items that the user can order. The user can also view related Knowledge articles to help them resolve the issue on their own without involving the service desk operators.

AI Search Assist helps with incident deflection for the following areas:

- Incident record producers: Deflect incidents by helping end users resolve issues before they raise an incident.
- Incident forms: Help service desk staff resolve incidents quickly by providing relevant knowledge.
- Incident email notifications: Help end users resolve their incidents themselves without requiring manual intervention from service desk staff.

To learn more about configuring and using the AI Search Assist widget in Service Portal, see AI Search Assist record producer integration for Service Portal.

Configuring AI Search

Plan and configure your AI Search implementation.

AI Search configuration overview

Configuration stage	Description
Assign roles to AI Search administrators and users	Assign the ais_admin role to users who need to create, read, update, and delete content indexing and search settings for AI Search. Assign the ais_high_security_admin elevated privilege role to users who need to access High Security settings for AI Search. Assign the search_application_admin to users who need to create, read, update, and delete search application configurations.
Define indexed sources	Create and configure an indexed source for each ServiceNow AI Platform® table that has content you want to make searchable. Index source tables to add their searchable content to the AI Search index. The base system includes pre-configured indexed sources for the Catalog Item, Knowledge, and User tables.
Define search sources	Apply filter conditions to indexed sources to define search sources that users can access in your search experiences.

Configuration stage	Description
	The base system includes pre-configured search sources that reference the default indexed sources.
Define search profiles	Create a search profile for each distinct search experience that you want to offer in search applications. Specify the data that users can search by linking search sources to the search profile. Control search behavior by configuring synonyms, stop words, typo handling settings, Genius Results, and result improvement rules. The base system includes pre-configured search profiles for use with Service Portal, Now Mobile, and the AI Search Assist Service Portal widget.
Define search application configurations	Create a configuration for each ServiceNow AI Platform application that supports AI Search. Select a search profile for the application to use. Configure application-specific search settings. The base system includes pre-configured search application configurations for Service Portal, Virtual Agent, and the AI Search Assist Service Portal widget.
Enable and configure AI Search in ServiceNow AI Platform applications	Specify AI Search as the search engine in supported ServiceNow AI Platform applications. Select a search application configuration for each application to use and configure available search UX settings.
Create a custom search matcher for global search	Optionally configure exact-match global searches that only match results from a specified column (field) on a particular ServiceNow AI Platform table.

Using guided setup to implement AI Search

Guided setup provides a sequence of tasks that help you configure AI Search on your ServiceNow instance. To open AI Search guided setup, navigate to **All** > `AI Search` > **Guided Setup**.

For more information about using the guided setup interface, see Using guided setup.

Assign roles to AI Search administrators and users

Grant users the ais_admin, ais_external_content, ais_high_security_admin, and search_application_admin roles so they can configure and manage settings and content for AI Search and search applications.

Before you begin

Role required: admin

About this task

Users with the following roles can access and configure settings and content for the AI Search application.

Role	Description
AI Search administrator [ais_admin]	Manages configuration settings for the AI Search application, including the following: • Creates, reads, updates, and deletes indexed sources • Creates, reads, updates, and deletes search sources • Creates and deletes mappings between search sources and search profiles • Creates, reads, updates, and deletes search profiles • Creates, reads, updates, and deletes synonym and stop word dictionaries • Creates and deletes mappings between dictionaries and search profiles • Reads and updates typo handling auto-correction dictionaries linked to search profiles • Creates, reads, updates, and deletes Genius Result configurations • Creates and deletes mappings between Genius Result configurations and search profiles • Creates, reads, updates, and deletes result improvement rules • Creates and deletes mappings between result improvement rules and search profiles
AI Search high security administrator [ais_high_security_admin]	Accesses High Security settings for AI Search, including the following: • Bypasses all search filters from search sources and content security in the Search Preview UI, viewing all search query results available in the AI Search index • Reads external content user mapping import history records This is an elevated privilege role. Elevated privilege roles aren't assigned to users or groups, and must be used by elevation. For more information on elevated privilege roles, see Elevated privilege roles. To learn about elevation, see Elevate to a privileged role.

Role	Description
	The Instance Security Center **Users with High Privilege Roles** user metric displays the count of users assigned this role. For more details, see User metrics.
Search application administrator [search_application_admin]	Creates, reads, updates, and deletes search application configurations for Zing and AI Search. This role includes the ais_admin and personalize_dictionary roles.
AI Search external content API user [ais_external_content]	Accesses endpoints for AI Search external content APIs, including the following: • Feeds external documents for AI Search to index as search results, or deletes records for external documents from the index, using the External Content Ingestion API. • Imports user mappings for external content security using the AI Search External User Mapping API. Assign this role to non-interactive users and integration users who require access to external content API endpoints. For more information on non-interactive users, see Non-interactive sessions. For details on integration users, see Mark service accounts as internal integration users. Users with the admin role don't need this role to access external content API endpoints.

Assign these roles to users and groups to enable them as search administrators or external content API users for AI Search.

Procedure

- Assign a role to a user
- Assign a role to a group

Indexed sources

Indexed sources designate ServiceNow AI Platform® tables and external document sets with content that you want to make searchable. AI Search ingests records or documents from these sources and stores their searchable content in its search index.

For instructions on creating an indexed source, see Create an indexed source.

Indexed source types

AI Search supports the following indexed source types.

Internal indexed source

An internal indexed source retrieves content and metadata from ServiceNow AI Platform records. It includes a unique name and a reference to a ServiceNow AI Platform table with records that you want to make searchable. AI Search extracts and indexes searchable content and metadata from records in this table and in any of its child tables that you configure for indexing.

AI Search excludes some ServiceNow AI Platform tables from indexing. You can't define indexed sources for these excluded tables or their derived tables. For a list of excluded tables, see ServiceNow AI Platform tables excluded from AI Search indexing.

You can't index remote tables with internal indexed sources. To index content from a remote table, create an external indexed source.

External indexed source

An external indexed source retrieves document content and metadata from an external repository or a remote table. It includes a unique name and a reference to an external content schema table instead of a ServiceNow AI Platform table. For more details on configuring indexed sources for external content, see Indexing and searching external content in AI Search.

Indexed source retention policies and filter conditions

Indexing large source tables, such as the Task [task] table and tables that extend it, can add large numbers of records to the AI Search index. To limit the set of records indexed from a source table, you can configure retention policies and filter conditions for your indexed sources. AI Search also uses these settings to automatically purge stale records from the index. For more information on retention policies and filter conditions, see Indexed source retention policies and filter conditions.

Note:
Retention policies are required for indexed sources that index records from the Task [task] table or tables that extend it. They are optional for other indexed sources.

Indexed source attributes and field settings

You can configure attributes and field settings for an indexed source to control indexing behavior for source records. Attributes control the indexed source's behavior at the record level, while field settings define its behavior for individual fields on indexed records. For more information, including lists of available attributes and field settings, see Indexed source attributes and Field settings.

Indexing content from an indexed source

Once you define an indexed source, AI Search begins automatically indexing to reflect changes to records in the selected source table and its specified child tables. The results of all record create, update, and delete operations in these tables are reflected in the search index. AI Search doesn't index content from unmodified records in these tables until you perform a full table index. For more information on indexing behavior, including steps for full table indexing, see Indexing content from indexed sources.

Note:
The pre-configured indexed sources included with AI Search only index content from string fields on source records. When searching records from these indexed sources, you can use numeric fields to facet or filter your search results, but you can't find records using their numeric field values. To search on a record's numeric field values, copy them into a string field so they can be indexed.

Indexing content from knowledge articles

Starting in Washington DC Patch 9, when indexing content from records in the Knowledge [kb_knowledge] table, AI Search defaults to including content defined in knowledge blocks. Administrators can override this default behavior and configure AI Search to exclude content from knowledge blocks when indexing knowledge articles. For details on making this change, see Exclude knowledge block content from the AI Search index.

Multiple indexed sources for the same ServiceNow AI Platform table

You can't create more than one indexed source for a single ServiceNow AI Platform table. However, plugins and applications may define duplicate indexed sources for a ServiceNow AI Platform table. For example, the base system includes an indexed source defined for the User [sys_user] table, but a plugin or application might define a second indexed source for this table under a different name.

> **Note:**
> Only one indexed source can be active at a time for a given ServiceNow AI Platform table. The system makes duplicate indexed sources for a table inactive by default. Before you can make one of these duplicate sources active, you must edit the currently active source and make it inactive. AI Search only indexes content and metadata from active indexed sources.

Create an indexed source

Define an indexed source to make content and metadata from ServiceNow AI Platform® table records searchable using AI Search.

Before you begin

Role required: ais_admin

About this task

When you create an indexed source, you specify a ServiceNow AI Platform table. AI Search indexes content and metadata from the specified source, making it available for inclusion in search sources.

By default, AI Search indexes all records from the specified source table. You can limit which records are indexed by configuring a retention policy and filter conditions for the indexed source. For more information about retention policies and filter conditions, see Indexed source retention policies and filter conditions.

> **Note:**
> Retention policies are required for indexed sources that index records from the Task [task] table or tables that extend it. They are optional for other indexed sources.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Select **New**.
3. On the Indexed Source form, fill in the fields.

   For a description of the field values, see Indexed Source form.

   > **Note:**
   > You can't create more than one indexed source for a single ServiceNow AI Platform table.

4. Select **Submit**.

The new indexed source appears in the AI Search Indexed Sources list with the **State** set to **New**. If the new indexed source requires manual indexing, an informational message and an alert triangle icon (⚠) appear.

5. **Optional:** To configure indexing for child tables of the selected source table, open the new indexed source and create a related list record for each child table:

   a. In the Child Tables related list, select **New**.

   b. On the Child Table form, fill in the fields.
      For a description of the field values, see Child Table form.

   c. Select **Submit**.
      The new child table appears in the Child Tables related list.

What to do next

To control how your new indexed source processes source records and fields during indexing, configure its indexed source attributes and field settings. For more information, see Indexed source attributes and Field settings.

To make content from source records searchable, perform a full table index for your new indexed source. For details on this procedure, see Perform a full table index or reindex for a single indexed source.

Indexed source retention policies and filter conditions

AI Search uses settings to automatically purge stale records from the index and optimize search performance. To limit the set of records indexed from source tables, you can configure retention policies and filter conditions for your indexed sources.

Indexed source retention policies

Indexing large source tables, such as the Task [task] table and tables that extend it, can add significant numbers of records to the AI Search index. To limit the set of source table records indexed based on the time since they were last updated, select a retention policy for your indexed source. AI Search only indexes source records updated within the time period defined for the retention policy. For example, if you select a two-year retention policy, AI Search excludes source records that were last updated more than two years ago.

When the time since a source record's last update exceeds the limit from the indexed source's retention policy, AI Search marks the corresponding indexed record as stale.

> **Note:**
> Retention policies are required for indexed sources that index records from the Task [task] table or tables that extend it. They are optional for other indexed sources.

Indexed source filter conditions

To limit the set of records indexed from a source table, define filter conditions for your indexed source. AI Search only indexes records that match all defined filter conditions.

When a source record no longer satisfies the indexed source's filter conditions, AI Search marks the corresponding indexed record as stale.

Purging stale records

AI Search automatically purges stale records from the index daily. Users with the admin role can manually purge stale records. For details on this procedure, see Purge stale records from the AI Search index.

Purge stale records from the AI Search index

Execute a scheduled job to delete stale records from the AI Search index.

Before you begin
Role required: admin

> **Note:**
> If the *glide.script_processor.admin* system property's value has been changed from the default value, you need the role specified by that property value. For details on this property value and its effects, see Restrict access to background script [Updated in Security Center 1.3].

About this task
If either of these conditions is satisfied for a record in the AI Search index, AI Search marks the record as stale.

- The time since the record's last update exceeds the age limit from its indexed source's retention policy
- The record no longer satisfies the filter conditions defined for its indexed source

AI Search automatically purges stale records from the index daily. Administrators can purge records manually by running a background script.

For more information on running background scripts, see Scripts - Background module.

Procedure

1. Navigate to **All** > **System Definition** > **Scripts - Background**.

   > **Note:**
   > If the *glide.script_processor.admin* system property's value is set to security_admin, you must elevate to that privileged role to access this module. For more information on elevating privileges, see Elevated privilege roles.

2. In the **Run script (JavaScript executed on server)** text field, enter one of the following background scripts.
   - To purge stale records indexed from a specific table, enter this script, replacing `TABLE_NAME` with the name of your chosen table:

     ```
     new sn_ais.IndexEvent().purgeFilteredDocuments('TABLE_NAME');
     ```

     For example, to purge all stale records indexed from the Knowledge Table indexed source, replace `TABLE_NAME` with `kb_knowledge` in the script.
   - To purge all stale records in the index, regardless of which table they were indexed from, enter this script:

     ```
     new sn_ais.IndexEvent().purgeFilteredDocuments();
     ```

3. Select **Run script**.
   The script creates an event that the index handler processes, so the script completes without blocking on completion of the purge job.

Result

The background script executes, deleting stale records from the AI Search index.

Exclude knowledge block content from the AI Search index

Prevent AI Search from indexing content found in your knowledge blocks.

Before you begin

Role required: admin

About this task

Starting in Washington DC Patch 9, AI Search includes content specified in your published knowledge blocks when it indexes content and metadata from your knowledge articles.

Administrators can override this default behavior and exclude your knowledge block content from the search index by setting the value of the *`glide.ais.disable_kbb`* system property. If this system property's value is set to true, AI Search ignores knowledge block content when indexing content and metadata from your knowledge articles.

To learn more about creating and publishing knowledge blocks for reference in your knowledge articles, see Using knowledge blocks.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Search for a system property record with name *`glide.ais.disable_kbb`*.
   - If a system property record with this name exists, open the record by selecting it.
   - If no system property record with this name exists, select **New**, then enter `glide.ais.disable_kbb` in the System Property form's **Name**field.
3. In the System Property form's **Value** field, enter *`true`*.
4. Save the modified System Property record by selecting **Submit** or **Update**.

What to do next

To make the new setting take effect, reindex the Knowledge Table indexed source. For details on reindexing, see Perform a full table index or reindex for a single indexed source.

Indexed source attributes

An indexed source attribute defines indexing behavior for all records from an indexed source.

Examples of how you might use indexed source attributes include the following.

- Define a filter to limit the set of records indexed from the source table
- Modify default indexing behavior for attachments or tags found on source table records
- Specify whether to index translated fields from source table records

An indexed source's attributes appear in its Advanced Configuration related list.

Create an indexed source attribute

Define an indexed source attribute to configure the AI Search content indexing behavior for all records on a source table.

Before you begin

Role required: ais_admin

About this task

For details on available indexed source attributes and values, see List of AI Search indexed source attributes.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Open the indexed source that you want to define an attribute for.
3. In the Advanced Configuration related list, select **New**.
4. On the Indexed Source Attribute form, fill in the fields.
   For a description of the field values, see Indexed Source Attribute form.
5. Select **Submit**.

Result

The new indexed source attribute appears in the Advanced Configuration related list.

What to do next

To make the new indexed source attribute take effect, perform a full table reindex for the indexed source. For details on this procedure, see Perform a full table index or reindex for a single indexed source.

Enable indexing of tags for an indexed source

Configure an indexed source attribute to enable indexing of content from tags found on records from an indexed source.

Before you begin

Role required: ais_admin

About this task

By default, AI Search indexing ignores tags from source records. You can enable indexing of tags for all records from an indexed source. Choose whether you want to index all shared and globally visible tags or only globally visible tags. Private tags are never indexed.

For more general instructions on defining indexed source attributes, see Create an indexed source attribute.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Open the indexed source that you want to enable tag indexing for.
3. In the Advanced Configuration related list, select **New**.
4. On the Indexed Source Attribute form, fill in the fields.
   a. In the **Attribute** field, enter `index_tags`.
   b. In the **Value** field, enter one of the following options:

Option	Description
everyone_only	Only tags shared with everyone are indexed.

Option	Description
all_shared	All shared tags (**Everyone**, **Groups and Users**) are indexed.

For a description of the field values, see Indexed Source Attribute form.

5. Select **Submit**.

Result

The new *index_tags* indexed source attribute appears in the Advanced Configuration related list.

What to do next

To make the new *index_tags* attribute take effect, perform a full table reindex for the indexed source. For details on this procedure, see Perform a full table index or reindex for a single indexed source.

Enable indexing of attachments for an indexed source

Configure an indexed source attribute to enable indexing of content from attachments found on records from an indexed source.

Before you begin

Role required: ais_admin

About this task

By default, AI Search indexes searchable content and metadata (such as file size and date) for attachments found on source table records. Indexing supports the following attachment file formats:

- Active Server Page Extended (`.aspx`)
- Hypertext Markup Language (`.html`, `.htm`)
- Microsoft Excel (`.xls`, `.xlsx`)
- Microsoft PowerPoint (`.pot`, `.potx`, `.ppt`, `.pptm`, `.pptx`)
- Microsoft Word (`.doc`, `.docx`, `.dot`, `.dotx`)
- Plain Text (`.txt`)
- Portable Document Format (`.pdf`) with searchable text

Note:
AI Search only indexes content from the first 1 MB of an attachment's data. It ignores attachments that are larger than 25 MB in size.

AI Search provides limited support for encrypted record attachments:

- When indexing an encrypted attachment, AI Search extracts metadata such as the attachment file's size and date, but it doesn't extract searchable content.
- The encrypted attachment's MIME type is detected as `application/octet-stream`.
- No feedback for the encrypted attachment appears in ingestion log messages.

Attachment indexing is controlled by an indexed source's *index_attachments* attribute:

- true (default value): Attachments are indexed for records from the indexed source.
- false: Attachments aren't indexed for records from the indexed source.

The following procedure explains how to make sure this attribute is set to true for an indexed source. For more general instructions on defining indexed source attributes, see Create an indexed source attribute.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Open the indexed source that you want to verify or change attachment indexing for.
3. In the Advanced Configuration related list, check for an `index_attachments` attribute with value `false`.

   Note:
   If no `index_attachments` attribute exists in the related list, or the attribute exists with value `true`, attachment indexing is already enabled for the indexed source. No further steps are required.

4. Open the `index_attachments` attribute and change its value from `false` to `true`.

Result

Attachment indexing is enabled for the indexed source.

What to do next

If you changed the `index_attachments` attribute's value from `false` to `true`, reindex content for the indexed source. For details on reindexing, see Perform a full table index or reindex for a single indexed source.

List of AI Search indexed source attributes

You can adjust indexing behavior for an AI Search indexed source by configuring indexed source attributes and values.

For instructions on defining AI Search indexed source attributes, see Create an indexed source attribute.

Attribute	Description
filter	Specify a filter condition that applies on the indexed source table and all of its child tables configured for indexing. Indexing ignores source records that don't pass this filter condition. If you change the value of this attribute for an indexed source, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
index_attachments	Control indexing behavior for attachments from indexed records. Supported values:

Attribute	Description
	• `true`: Enable indexing of attachments for this table. • `false`: Disable indexing of attachments for this table. Default value: `true` If you change the value of this attribute for an indexed source, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
index_tags	Control indexing behavior for tags from indexed records. Supported values: • `all_shared`: Index shared tags and globally visible tags. • `everyone_only`: Only index globally visible tags. • `none`: Disable indexing of tags. Default value: `none` **Note:** Search results display indexed tags based on the visibility of the result record instead of the visibility of the tag. If you change the value of this attribute for an indexed source, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
index_translated_fields	Control indexing behavior for translated fields from indexed records. Supported values: • `true`: Enable indexing of translated fields for this table. • `false`: Disable indexing of translated fields for this table. Default value: `false`

Attribute	Description
	Note: For tables with manually mapped records, such as kb_knowledge, indexing ignores this attribute and honors the manual mapping. If you change the value of this attribute for an indexed source, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
user_read_takes_precedence_over_group_deny	Control precedence of external user read access permissions and external group deny access permissions for external document search results from the indexed source. Supported values: • `true`: External user read access permissions take precedence over external group deny access permissions for external document search results from the indexed source. • `false`: External group deny access permissions take precedence over external user read access permissions for external document search results from the indexed source. Default value: `true` on indexed sources for external content schema tables. Not set on indexed sources for internal tables. **Note:** This attribute only applies to indexed sources for external content schema tables. If you apply this attribute to an indexed source for an internal table, it has no effect.

Note:
Indexed source attributes are records on the AI Search Configuration Attribute [ais_configuration_attribute] table that have Table as their **Applies To** field value. Records on this table that have Column as their **Applies To** field value are Field settings.

Field settings

A field setting controls indexing behavior for a specified field (column) on all records from an indexed source.

Examples of how you might use field settings include the following.

- Disable text indexing or searchability for a field found on source table records
- Modify default mappings between source table fields and fields in the AI Search index
- Add a referenced table field to the index by dot-walking from a reference field on source table records

An indexed source's field settings appear in its Field Settings & Mapping related list.

Create a field setting for an indexed source

Define an indexed source field setting to configure AI Search content indexing behavior for a specific field on source table records.

Before you begin

Role required: ais_admin

About this task

For details on available field settings and values, see List of AI Search indexed source field settings.

Procedure

1. Navigate to **All** > *AI Search* > *AI Search* **Index** > **Indexed Sources**.
2. Open the indexed source that you want to define a field setting for.
3. In the Field Settings & Mapping related list, select **New**.
4. On the Field Setting form, fill in the fields.
5. Select **Submit**.

Result

The new field setting appears in the Field Settings & Mapping related list.

What to do next

To make the new field setting take effect, perform a full table reindex for the indexed source. For details on this procedure, see Perform a full table index or reindex for a single indexed source.

Enable indexing of referenced table fields for an indexed source

Configure an indexed source to index field values from tables referenced by fields in the source table. You can index these field values for use in filters and EVAM search result configurations, index them for search, or both.

Before you begin

You must have an indexed source configured for the source table. For details on creating an indexed source for a table, see Create an indexed source.

Role required: ais_admin

About this task

When an indexed source table includes a reference field, AI Search defaults to indexing values for the reference field but not for other fields on the table that it references. To index values from these referenced table fields, you can create `dot_walk_fields` and `searchable_dot_walk_fields` field settings on your indexed source table.

- To use field values from referenced tables in search source filters, facet filters, and EVAM search result configurations, create a `dot_walk_fields` field setting.
- To make field values from referenced tables searchable, create a `searchable_dot_walk_fields` field setting.

As the value for either of these field settings, list the names of the referenced table fields you want to index, separated by commas. For example, to index the **name** and **city** fields from the table referenced by the **company** field on the indexed source table, select the **company** field and enter value `name,city`.

To dot-walk across multiple tables, you can enter field name values with dot-separated reference field prefixes. For example, enter value `company.contact.name` to index the **name** field in the table referenced by the **contact** field on the table referenced by the indexed source table's **company** reference field.

Note:
Each dot-walk reference level imposes a performance impact on indexing. Avoid using multi-level references unless necessary.

You can create both *dot_walk_fields* and *searchable_dot_walk_fields* field settings for the same reference field.

Procedure

1. Navigate to **All** > *AI Search* > *AI Search* **Index** > **Indexed Sources**.
2. In the Field Settings & Mapping related list, select **New**.
3. On the Field Setting form, enter the field values shown for your use-case.
   - To index field values from referenced tables for use in search source filters, facet filters, and EVAM search result configurations, enter the following field values.

Field	Value
Attribute	dot_walk_fields
Field	<name of indexed source table reference field that you want to use to dot-walk to another table>
Value	<comma-separated list of names for fields that you want to index from the referenced table>

   - To index field values from referenced tables as searchable text, enter the following field values.

Field	Value
Attribute	searchable_dot_walk_fields
Field	<name of indexed source table reference field that you want to use to dot-walk to another table>
Value	<comma-separated list of names for fields that you want to index from the referenced table>

Note:
If the *`dot_walk_fields`* and *`searchable_dot_walk_fields`* attributes don't appear in the **Attribute** selection list, ensure that your **Field** selection is a reference field.

4. Select **Submit**.
   The attribute and value appear in the Field Settings & Mapping related list.

List of AI Search indexed source field settings

You can adjust indexing behavior for source record fields in an AI Search indexed source by configuring field setting attributes and values.

For more information on creating field settings, see Create a field setting for an indexed source.

Attribute	Description
dot_walk_fields	Index reference and display values from fields on the selected reference field's source table for use in search source filters, facet filters, and EVAM search result configurations. AI Search automatically updates the indexed field values to reflect changes made to the referenced table's field values. For more information about search source filters, see Search sources. For details on facet filters, see Create a facet in an AI Search application configuration. • Field: Name of the reference field on your indexed source table. As an example, if your indexed source table contains a **company** reference field that references the Company [core_company] table, and you want to index field values from that table's **name** and **city** fields, you'd set `company` as the field. • Type: string • Value: Comma-separated list of referenced table fields to index values from. As an example, if your indexed source table contains a **company** reference field that references the Company [core_company] table, and you want to index field values from that table's **name** and **city** fields, you'd set `name,city` as the value.

Attribute	Description
	Note: You can't search field values indexed with this field setting. To index field values from referenced tables for search, create a *searchable_dot_walk_fields* field setting. You can create both *dot_walk_fields* and *searchable_dot_walk_fields* field settings for the same reference field. If you change the value of this setting for an indexed source's field, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
index_calculated_field	Option to enable indexing of searchable content from calculated field values. If your indexed source contains a calculated field, set this field setting to *true* for that field to correctly index its values. With a field setting value of *false* or no value, AI Search ignores values from the calculated field when indexing content. • Field: Name of a calculated field from the indexed source table. (If you specify a non-calculated field, the system rejects the field setting submission and displays a warning message.) • Type: true | false • Supported values: ○ *true*: Apply special handling to correctly index calculated field values from the selected field. ○ *false*: Don't apply special handling when indexing field values from the selected field. Calculated field values aren't indexed. If you change the value of this setting for an indexed source's field, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
map_to	Map the selected field from the indexed source table to an AI Search index field. When

Attribute	Description
	indexing records from the source table, AI Search populates the specified index field with the value of the selected source field. For example, the base system's Knowledge Table indexed source maps the **kb_knowledge.short_description** source field to the **title** index field. When AI Search indexes a record from the Knowledge [kb_knowledge] table, it populates the indexed record's searchable title with the value from the source record's **short_description** field. • Field: Name of a field from the indexed source table. • Type: string • Value: Name of the AI Search field that you want to map the selected field's display values to. For details on AI Search index fields, see AI Search index fields. **Note:** You can only define one **map_to** field setting for an indexed source field. If you try to define multiple **map_to** settings for the same field, the system displays an error message. If you change the value of this setting for an indexed source's field, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
no_text_index	Option to disable indexing of searchable content from the selected field on records from the indexed source. • Field: Name of a field from the indexed source table. • Type: true | false • Supported values: ◦ `true`: Disable searchable content indexing for the selected field. Search and filters can't match the field's value. AI Search doesn't generate an index event when the field is updated. ◦ `false`: Enable searchable content indexing for the selected field. Search and filters can match the field's value. AI

Attribute	Description
	Search generates an index event when the field is updated, adding the affected record to its indexing queue. • Default value: `false` If you change the value of this setting for an indexed source's field, the new value immediately affects index event generation for updates to that field, but it doesn't affect searches or filters for previously indexed records until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source. **Warning:** Don't set this option to `true` for the sys_updated_on field on the Task [task] table, tables that extend Task, or any other tables that you've configured retention policies for. The retention policies for these tables rely on indexing of sys_updated_on field values. For more information on retention policies, see Indexed source retention policies and filter conditions. For Customer Service and Support guidelines on setting the **no_text_index** option for different field types, see KB0859922 on the Now Support site.
not_searchable	Option to disable search matching for content indexed from the selected field. • Field: Name of a field from the indexed source table. • Type: true | false • Supported values: ◦ `true`: Disable search matching for content indexed from the selected field. Search can't match the field's value. Filters (such as for content security, navigation tabs, and facets) can still match the field's value. ◦ `false`: Enable search matching for content indexed from the selected field. Search and filters can match the field's value. • Default value: `false` If you change the value of this setting for an indexed source's field, the change doesn't

Attribute	Description
	take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.
searchable_dot_walk_fields	Index reference and display values from fields on the selected reference field's source table as searchable text. AI Search doesn't automatically update the indexed field values to reflect changes made to the referenced table's field values. To update the indexed field values, you must reindex the indexed source. For details on this procedure, see Perform a full table index or reindex for a single indexed source. • Field: Name of the reference field on your indexed source table. As an example, if your indexed source table contains a **company** reference field that references the Company [core_company] table, and you want to index searchable text from that table's **name** and **city** fields, you'd set `company` as the field. • Type: string • Value: Comma-separated list of referenced table fields to index as searchable content. As an example, if your indexed source table contains a **company** reference field that references the Company [core_company] table, and you want to index searchable text from that table's **name** and **city** fields, you'd set `name,city` as the value. **Note:** You can't use field values indexed with this field setting in search source filters. To index field values from referenced tables for use in search source filters, create a *`dot_walk_fields`* field setting. You can create both *`dot_walk_fields`* and *`searchable_dot_walk_fields`* field settings for the same reference field. If you change the value of this setting for an indexed source's field, the change doesn't take effect until you reindex content from the indexed source. For reindexing steps, see Perform a full table index or reindex for a single indexed source.

Note:
Field settings are records on the AI Search Configuration Attribute [ais_configuration_attribute] table that have Column as their **Applies To** field value. Records on this table that have Table as their **Applies To** field value are Indexed source attributes.

This example shows how the base system's Catalog Item Table indexed source maps **short_description** field values from the Catalog Item [sc_cat_item] table to the AI Search index's **text** field for indexed records.

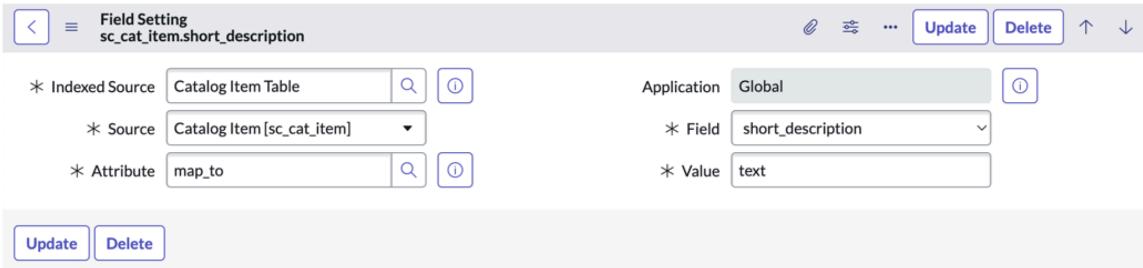

Security for searchable referenced table field values

When a user searches referenced table field values that you have indexed for search with the `searchable_dot_walk_fields` field setting, only field values that the user can view appear in the search results. The system uses a field value's role-based access control list rules (ACLs) to determine whether the search user can view that field value.

Searchable referenced table field values that have condition-based or script-based ACLs or user criteria don't appear in search results.

For searchable field values indexed through multiple dot-walk reference levels, the system only considers role-based ACLs on the final field value. As an example, if you index `company.contact.name` referenced table field values for search, user access to **name** field values isn't affected by role-based ACLs on **contact** field values.

To bypass all ACLs and allow users to search for all searchable referenced table field values, you can set the `glide.ais.query.allow_indexlookup_for_dotwalk` system property to true. This bypasses ACL evaluation for field values indexed via the `dot_walk_fields` and `searchable_dot_walk_fields` field settings.

For more details on ACL types and configuration, see Access control list rules.

Indexing content from indexed sources

AI Search indexes records on indexed source tables to make their content searchable.

Once you define an indexed source, AI Search begins automatically indexing to reflect changes to records in the selected source table and its specified child tables. The results of all record create, update, and delete operations in these tables are reflected in the search index. AI Search doesn't index content from unmodified records in these tables until you perform a full table index.

Indexing content from referenced tables

When an indexed source table includes a reference field that stores a reference to another table, AI Search defaults to indexing values for the reference field but not for other fields on the referenced table. For example, the User [sys_user] table includes a **company** reference field that stores a reference to the Company [core_company] table. When indexing records from the User table, AI Search indexes values for the **company** reference field, but doesn't index values for other Company table fields such as **city** or **website**.

To index additional fields from referenced tables, you can configure *dot_walk_fields* and *searchable_dot_walk_fields* field settings for reference fields on your indexed source table.

- To use field values from referenced tables in search source filters, facet filters, and EVAM search result configurations, create a *dot_walk_fields* field setting.
- To make field values from referenced tables searchable, create a *searchable_dot_walk_fields* field setting.

You can create both *dot_walk_fields* and *searchable_dot_walk_fields* field settings for the same reference field.

For instructions on configuring *dot_walk_fields* and *searchable_dot_walk_fields* field settings, see Enable indexing of referenced table fields for an indexed source. For more information on field settings, see Field settings. To learn about search source filters, see Search sources. For details on facet filters, see Create a facet in an AI Search application configuration. To understand EVAM configurations, see Entity View Action Mapping.

Indexing content from source record attachments

By default, AI Search indexes searchable content and metadata (such as file size and date) for attachments found on source table records. Indexing supports the following attachment file formats:

- Active Server Page Extended (`.aspx`)
- Hypertext Markup Language (`.html`, `.htm`)
- Microsoft Excel (`.xls`, `.xlsx`)
- Microsoft PowerPoint (`.pot`, `.potx`, `.ppt`, `.pptm`, `.pptx`)
- Microsoft Word (`.doc`, `.docx`, `.dot`, `.dotx`)
- Plain Text (`.txt`)
- Portable Document Format (`.pdf`) with searchable text

Note:
AI Search only indexes content from the first 1 MB of an attachment's data. It ignores attachments that are larger than 25 MB in size.

AI Search provides limited support for encrypted record attachments:

- When indexing an encrypted attachment, AI Search extracts metadata such as file size and date, but doesn't extract searchable content.
- The encrypted attachment's MIME type is detected as `application/octet-stream`.
- No feedback for the encrypted attachment appears in ingestion log messages.

Attachment indexing is controlled by an indexed source's *index_attachments* attribute. To verify that attachment indexing is enabled for an indexed source, see Enable indexing of attachments for an indexed source.

Indexing content from source record tags

AI Search can index tags found on source table records. Search results display indexed tags based on the visibility of the result record instead of the visibility of the tag.

By default, indexing ignores tags from source records. You can enable indexing of tags for each indexed source. Choose whether you want to index all shared and globally visible tags or just globally visible tags.

For details on enabling tag indexing for an indexed source, see Enable indexing of tags for an indexed source.

Indexing content from the Task table and its child tables

When defining an indexed source that indexes records from the Task table or any of its child tables, you must specify a retention policy. For details on retention policies, see Indexed source retention policies and filter conditions.

Reindexing content

You must perform a full table reindex for an indexed source under the following circumstances:

- After you add, modify, or delete an indexed source attribute or a field setting for the indexed source.
- If you want to update indexed values for fields on tables referenced by source table records. As an example, the **company** field on the User [sys_user] table stores a reference to the Company [core_company] table. If you configure indexing for values for fields on the referenced Company table, such as **city** or **website**, you must reindex the User table to update those indexed Company table field values. For details on configuring indexing of referenced table fields, see Enable indexing of referenced table fields for an indexed source.

Perform a full table index or reindex for a single indexed source

Make content from an internal indexed source searchable by performing a full table index. This procedure indexes existing records from the source table and any child tables configured for indexing. You can manually reindex content from an internal indexed source by repeating this procedure.

Before you begin

Your indexed source must be active and internal (referring to a ServiceNow AI Platform® table). To index content from an indexed source for external documents, see Indexing and searching external content in AI Search.

Role required: ais_admin

About this task

Once you define an indexed source, AI Search begins automatically indexing to reflect changes to records in the selected source table and its specified child tables. The results of all record create, update, and delete operations in these tables are reflected in the search index. AI Search doesn't index content from unmodified records in these tables until you perform a full table index.

Use this procedure to index searchable content from existing records in an indexed source, or to manually reindex existing content for an indexed source that has already been indexed.

If an indexed source contains unindexed records, the Indexed Sources list displays an informational message and marks the new source with an alert triangle icon (⚠). A similar informational message appears on the Indexed Source form. Make sure to perform a full table index for each indexed source that displays the alert triangle icon.

Procedure

1. Navigate to **All** > *AI Search* > *AI Search* **Index** > **Indexed Sources**.
2. Open the indexed source with content that you want to make searchable.
3. Select **Index All Tables**.
   The Generate Text Index dialog box appears.
4. **Optional:** If you want AI Search to allow indexing of updates to individual records in the source table while the table indexing event is running, select the **Non blocking index** option.

   Note:
   With this option deselected, AI Search defers indexing of updates to individual records in the source table until after the table indexing event finishes. The index may not reflect the latest contents for updated records until the deferred record indexing events finish.

5. Select **OK**.
   The Indexed Source History page appears. A message reports that AI Search has queued an indexing task for the indexed source.
6. **Optional:** To monitor the progress of the indexing task, refresh the Indexed Source History form page.
   When the task completes, the **Ingestion State** field shows *indexed*.
7. To return to the list of indexed sources, select the back icon (<).

Result

AI Search indexes content from existing records in the indexed source. Indexing continues on a scheduled basis. You don't need to repeat this task for the indexed source unless you want to reindex its content before the next scheduled indexing run.

What to do next

Define search sources to make searchable content from the indexed source available in user search experiences. For details on creating search sources, see Create a search source.

Perform a full table index or reindex for multiple indexed sources

Make content from multiple internal indexed sources searchable by performing a full table index. This procedure indexes existing records from the source tables and any child tables configured for indexing. You can manually reindex content from internal indexed sources by repeating this procedure.

Before you begin

Your indexed sources must be active and internal (referring to ServiceNow AI Platform® tables). To index content from indexed sources for external documents, see Indexing and searching external content in AI Search.

Role required: ais_admin

About this task

Once you define an indexed source, AI Search begins automatically indexing to reflect changes to records in the selected source table and its specified child tables. The results of all record create, update, and delete operations in these tables are reflected in the search index. AI Search doesn't index content from unmodified records in these tables until you perform a full table index.

Use this procedure to index searchable content from existing records in indexed sources, or to manually reindex content for indexed sources that have already been indexed.

If an indexed source contains unindexed records, the Indexed Sources list displays an informational message and marks the new source with an alert triangle icon (⚠). A similar informational message appears on the Indexed Source form. Make sure to perform a full table index for each indexed source that displays the alert triangle icon.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Select the indexed sources with content that you want to make searchable.
   To select all indexed sources on the page, follow the instructions for your list version in Perform actions on selected items in a list.
3. In the **Actions on selected rows...** menu, select **Index All Tables**.
   The Indexed Source History page for the last selected index appears. A message reports that AI Search has queued an indexing task for the indexed source.
4. **Optional:** Refresh the Indexed Source History form page to monitor the progress of the indexing task.
   When the task completes, the **Ingestion State** field shows `indexed`.
5. To return to the list of indexed sources, select the back icon (<).

Result

AI Search indexes content from existing records in the indexed sources. Indexing continues on a scheduled basis. You don't need to repeat this task for the indexed sources unless you want to reindex their content before the next scheduled indexing run.

What to do next

Define search sources to make searchable content from indexed sources available in user search experiences. For details on creating search sources, see Create a search source.

Search sources

Define search sources to make your searchable content available in search experiences. Choose an indexed source and configure optional filter conditions to limit the content that users can search.

Search source structure

A search source includes two elements:

- A reference to an indexed source. By default, search queries using the search source can return all records from the referenced indexed source. For more information on indexed sources, see Indexed sources.
- An optional set of filter conditions applied to the indexed source's records. Indexed source records must pass all of these filter conditions to be returned as search results for queries using the search source.

You can define multiple search sources that reference the same indexed source, specifying distinct filter conditions for each.

Filter conditions

To limit the set of indexed source table records that AI Search returns as search results, you can apply filter conditions to your search source.

Note:
Search source filter conditions can only operate on source record fields that are indexed. Fields on referenced records aren't indexed by default. To use a referenced record field in a search source filter condition, you must first configure indexing for the field in the indexed source. For details on this procedure, see Enable indexing of referenced table fields for an indexed source.

You can use static and dynamic filter conditions to filter search source records. As an example, you could apply a dynamic **[Assigned to] [is (dynamic)] [Me]** user filter to make your search source only return results corresponding to indexed source table records assigned to the search user.

Search sources don't support any of the following dynamic filter options:

- Dynamic filters with scripts that use the `current` keyword or global variable
- Dynamic filters with scripts that reference any global business rules
- Dynamic filters with scripts that use script includes

For more information on filter conditions, see Filters. To learn about dynamic filter options and their associated scripts, see Create a dynamic filter option.

Linking search sources to search profiles

To make your search source's content available to users, link it to one or more search profiles. You can link multiple search sources to a single search profile.

For more information, see Link a search source to a search profile.

Viewing search sources

You can view or edit search sources from any of the following application modules:

- To view all defined search sources, navigate to **All** > `AI Search` > **Search Experience** > **Search Sources**.
- To view search sources that reference a particular indexed source, navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**. Open the indexed source and select its Search Source Mappings related list.
- To view search sources linked to a particular search profile, navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**. Open the search profile and select its Search Sources related list.

Create a search source

Make searchable content from an indexed source available in user search experiences by defining a search source.

Before you begin

You must have an indexed source defined. For more details on setting up indexed sources, see Create an indexed source.

Role required: ais_admin

About this task

Define at least one search source for each of your indexed sources. You can define multiple search sources that reference the same indexed source, specifying distinct filter conditions for each.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. In the Search Source Mappings related list, select **New**.
3. On the Search Source form, fill in the fields.
   For a description of the field values, see Search Source form.
4. Select **Submit**.
   The new search source appears in the Search Sources related list.

Preview matching records for a search source

View the ServiceNow AI Platform® table records that match a search source's indexed source table and filter conditions.

Before you begin

Role required: ais_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Sources**.
2. Open the search source that you want to preview matching records for.
3. On the Search Source form, select the **preview** link in the informational message that reports the number of matching records.
   For a description of the field values, see Search Source form.

Result

The ServiceNow AI Platform table specified in the search source's indexed source appears, with filter conditions applied to match conditions from the search source. This filtered view displays all records that are searchable through the search source.

Note:
Preview uses the saved conditions from the search source. If you change a search source's conditions, preview doesn't reflect your changes until you save them.

Search profiles

AI Search stores search experience settings in search profiles. Define a search profile for each unique search experience that you want to offer users in a ServiceNow AI Platform® application. Create and populate synonym and stop word dictionaries. Configure typo handling dictionary settings. Define Genius Results and result improvement rules.

A search profile contains settings that determine how AI Search generates search results for a given search. The search application's configuration specifies which search profile to use, and also determines how the generated results are displayed to users. To learn about search application configuration display settings, see Search application configurations.

The following table describes the search experience settings that you can configure in a search profile.

Search profile settings

Setting	Description
Search sources	Link one or more search sources to a search profile.

Search profile settings (continued)

Setting	Description
	Search sources give administrators precise control over the searchable content available to users.
Synonyms	Define language-specific dictionaries of terms that AI Search treats as equivalent in search queries. Synonyms improve search recall and provide a more natural search experience for users.
Stop words	Define language-specific dictionaries of common terms for AI Search to remove from search queries. Stop words improve search precision and result relevancy by discarding irrelevant matches.
Typo handling	Configure settings for derivation of auto-correction terms from search source indexed content. AI Search replaces misspelled search query terms with auto-correction terms.
Genius Results	Analyze search query intent and display the best answers as Genius Result answer cards. Genius Results enable users to act directly on their search results.
Result improvement rules	Define rules with configurable triggers to modify search results for specific search queries. Result improvement rules enable administrators to boost, block, or promote records in a search query's result set.

Publishing search profile settings

To make a search profile's settings active in search queries, you must publish it. Publishing propagates the following settings to search application configurations that reference the search profile.

- The set of synonym and stop word dictionaries linked to the search profile
- Settings, blocked terms, and derived auto-correction terms for typo handling dictionaries linked to the search profile
- Result improvement rules and actions defined in the search profile

A search profile can have one of these states:

- **New**: The search profile has never been published.
- **Draft**: The search profile is saved with unpublished changes.
- **Published**: The search profile is saved with all changes published.

In the **All** > `AI Search` > **Search Experience** > **Search Profiles** list, search profiles with the *New* state display an alert triangle icon (⚠). An informational message notifies you to publish these search profiles.

Create a search profile

Create a new search profile to store search experience settings for a search application.

Before you begin

Role required: ais_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Select **New**.
3. On the Search Profile form, fill in the fields.
   For a description of the field values, see Search Profile form.
4. Select **Submit**.

Result

The new search profile appears in the AI Search Search Profiles list with **State** set to *New*. An informational message and an alert triangle icon (⚠) indicate that the new search profile has never been published.

What to do next

To make content searchable using your new search profile, link one or more search sources to it. For instructions on linking search sources, see Link a search source to a search profile.

Your new search profile includes sample synonym and stop word dictionaries and a default typo handling dictionary. For more details on configuring and populating these dictionaries, see Synonyms, Stop words, and Typo handling.

Link a search source to a search profile

Link search sources to a search profile to specify the content users can search through that profile.

Before you begin

You must have already created at least one search source. For instructions on creating search sources, see Create a search source.

Role required: ais_admin

About this task

Linking a search source to a search profile makes its filtered content available for searches using that search profile. AI Search only uses content from the search profile's linked search sources when finding and generating standard search results and Genius Result answers.

When linking a search source to a search profile, you can specify options that control how the search source's content is accessed for searches using that search profile. You can exclude the search source's content from being considered when finding standard search results, or prevent it from being considered during Genius Result answer generation. For a search source that includes internal content (from a ServiceNow AI Platform® table), you can also exclude indexed attachment content from being considered during Genius Result answer generation.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the search source to.
3. In the Search Sources related list, select **Link Existing**.
4. On the Search Profile - Search Source Mapping form, fill in the fields.
   For a description of the field values, see Search Profile - Search Source Mapping form.
5. Select **Submit**.

Result

The new search source appears in the Search Sources related list. Search query results immediately reflect changes to the set of records searchable through the search profile.

What to do next

Derived auto-correction term lists for the search profile's typo handling dictionary don't automatically update to reflect changes to the set of searchable records. To update these auto-correction term lists, publish the search profile you edited. For details on publishing a search profile, see Publish a search profile.

Publish a search profile

Publish a search profile to make pending changes to its settings take effect in linked search applications.

Before you begin

Role required: ais_admin

About this task

To make a search profile's settings active in search queries, you must publish it. Publishing propagates the following settings to search application configurations that reference the search profile.

- The set of synonym and stop word dictionaries linked to the search profile
- Settings, blocked terms, and derived auto-correction terms for typo handling dictionaries linked to the search profile
- Result improvement rules and actions defined in the search profile

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.

   Note:

   Search profiles with the *New* state display an alert triangle icon (⚠). An informational message notifies you to publish these search profiles.

2. Open the search profile that you want to publish.
3. Select **Publish**.

Result

The search profile appears in the AI Search Search Profiles list with **State** set to `Published`.

Clone a search profile

Create a copy of an existing search profile. The cloned search profile retains search sources, stop word and synonym dictionaries, Genius Result configurations, and result improvement rules from the original search profile.

Before you begin

Role required: ais_admin

About this task

Multiple search applications can link to the same search profile. For example, in the base system, Service Portal and Virtual Agent both link to the **Service Portal search profile** record.

As a search administrator, you may wish to configure separate search profiles for two applications that share a search profile. Cloning the existing search profile enables you to create a new record with the same search sources, stop word and synonym dictionaries, Genius Result configurations, and result improvement rules. This approach saves time if you want the two applications to use similar or identical search experience settings.

Note:
Cloning a search profile does not preserve its existing typo handling dictionary settings or blocked terms.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Select the record for the search profile that you want to clone.
   If the system prompts you to change the application scope so that you can edit the record, select the provided link.
3. Select **Clone**.
4. In the Set Profile Label dialog box, enter a label for your new search profile record, then select **OK**.
5. **Optional:** Make changes to the cloned search profile's settings.
   For details on search profile settings, see Create a search profile, Synonyms, Stop words, Typo handling, Genius Results, and Result improvement rules.
6. Publish your cloned search profile by selecting **Publish**.
   For more information on the effects of publishing a search profile, see Publish a search profile.

Result

Your cloned search profile record appears in the Search Profile table.

What to do next

Edit the search application configuration record of your choice and link it to your cloned search profile. For details on search application configurations, see Search application configurations.

Synonyms

Synonyms expand search queries to include additional terms with equivalent meaning or usage. Improve search recall by configuring synonym dictionaries and defining synonyms.

To define synonyms, modify the default synonym dictionary for a search profile or create a new synonym dictionary linked to one or more search profiles. Create and publish synonyms for each dictionary. Publish the dictionary to make your synonyms take effect for searches that use linked search profiles.

You can delete synonym dictionaries that aren't linked to any search profiles.

Synonym usage in search

When a search query includes a term defined in a synonym, AI Search expands the search to include all equivalent terms defined in the synonym.

For example, a synonym might expand the search query term `vacation` to `vacation OR holiday OR pto`. The expanded query returns records that contain any of the three equivalent terms.

Conditions for synonym eligibility

To be eligible for expanding a search query's terms, a synonym must satisfy the following conditions:

- It must be active.
- The synonym dictionary that includes it must be published.
- Its length doesn't exceed the dictionary term-length limit.

The system locates eligible synonyms in synonym dictionaries as follows:

Condition	Result
The search profile specified for the search query links to a published synonym dictionary for the search query's language.	The system only considers synonyms from the linked synonym dictionary. As an example, suppose a search application's search profile links to a French synonym dictionary. When a user searches in French in this application, only synonyms from the linked French synonym dictionary are eligible to expand search query terms.
The search profile specified for the search query doesn't link to a published synonym dictionary for the search query's language.	The system only considers synonyms from the default (English) synonym dictionary. As an example, suppose a search application's search profile has no linked Spanish synonym dictionary. When a user searches in Spanish in this application, only synonyms from the default (English) synonym dictionary are eligible to expand search query terms.

For more details on language support in synonym dictionaries, see Language dependence.

Synonym terms and payloads

Each synonym includes two or more equivalent terms, defined in two fields. The **Term** field contains one term designated as the synonym's label. The **Payload** field contains the remaining equivalent terms, formatted as a JSON array.

Synonym term expansion is bidirectional. AI Search treats all terms defined in the **Term** and **Payload** fields as equivalent, as illustrated in the following table.

Use case	Synonym expansion behavior
Search query term matches the term from a synonym's **Term** field	AI Search expands the search query term to match all equivalent terms from the **Payload** field.
Search query term matches a term from the **Payload** field	AI Search expands the search query term to match the term from the **Term** field and all remaining equivalent terms from the **Payload** field.

Dictionary term-length limit

By default, AI Search ignores synonym terms that exceed 256 characters in length. You can override this default limit by adding the *`query.maxDictionaryTermLength`* system property. The integer value of this system property indicates the maximum allowable character length for stop word and synonym dictionary terms.

For information on adding a system property, see Add a system property.

Warning:
Exercise caution when increasing the dictionary term-length limit beyond its default value. Processing long dictionary terms increases AI Search's memory usage, and in some cases search queries may fail with out-of-memory errors.

Multi-word phrases

A synonym term can include either a single word or a multi-word phrase. AI Search expands multi-word phrases when the search query terms exactly match the order and proximity of the phrased words in the synonym definition. If a user searches with these words in a different arrangement, no expansion occurs.

For example, the following table describes AI Search expansion behavior for a synonym that includes the single word `401k` and the multi-word phrase `retirement account`.

Search query terms	Synonym expansion behavior
401k	AI Search expands the search query to match either `401k` or `retirement account`.
retirement account	AI Search expands the search query to match either `401k` or `retirement account`.

Search query terms	Synonym expansion behavior
retirement account bank	AI Search expands the search query to match either `401k` or `retirement account` in addition to `bank`.
account retirement	AI Search doesn't expand the search query terms.
retirement bank account	AI Search doesn't expand the search query terms.

When search query terms match a multi-term synonym, AI Search treats those terms as a quoted phrase. For example, if you define `401k` and `retirement account` as synonyms and a user searches for the unquoted terms `retirement account`, AI Search returns results that match the quoted phrase `"retirement account"`.

Case sensitivity

By default, synonym expansion ignores letter case when comparing synonyms with search query terms. You can enable case sensitivity for a synonym dictionary. This option affects all synonyms defined in the dictionary.

When case sensitivity is enabled, synonyms in the dictionary only match search query terms or phrases that include the same combination of uppercase and lowercase letters.

For example, if you define a synonym including the word CASE and the phrase `Computer-Aided Software Engineering`, AI Search expands searches for CASE but not searches for `case`.

Punctuation sensitivity

By default, synonym expansion ignores punctuation when comparing synonyms with search query terms. You can enable punctuation sensitivity for a synonym dictionary. This option affects all synonyms defined in the dictionary.

When punctuation sensitivity is enabled, synonyms in the dictionary only match search query terms or phrases that include the same punctuation.

For example, suppose you define a synonym equating `U-S-A` with `University of South Australia` and another synonym equating `U.S.A.` with `United States of America`. Without punctuation sensitivity enabled, a search for `U.S.A.` expands to match both `University of South Australia` and `United States of America`. When punctuation sensitivity is enabled, the same search only expands to match `United States of America`.

Language dependence

Synonym expansion is language-dependent. You specify a language for each synonym dictionary, and this language applies to all synonyms defined in the dictionary. Synonyms from dictionaries with the same language as the search query are eligible for expansion.

Note:
If the search profile specified for a search query has no linked synonym dictionary for the search query language, AI Search uses the default (English) synonym dictionary as a fallback.

AI Search supports synonym dictionaries for all languages activated in your instance. For the list of languages you can activate, see Activate a language.

Interaction with other search features

The following table describes interactions between synonyms and other search features.

Feature	Interaction with synonyms
Genius Results	Synonym expansion terms can't trigger Genius Result configurations with Term trigger conditions. When computing Q&A Genius Results for a search query, AI Search expands synonyms in the query.
Lemma and Unicode normalization	If a search query term is defined as a synonym, AI Search doesn't normalize it.
Result improvement rules	By default, a synonym expansion term can activate a result improvement rule if it matches the rule's Query trigger. You can prevent synonym expansion terms from activating a rule with a Query trigger by clearing the **Expand Synonyms For Query Conditions** option that appears when you define the Query trigger.
Stop words	If a synonym definition includes a term defined as a stop word, AI Search retains that term when checking whether the synonym expands a search query. For example, suppose you define `as` and `a` as stop words and SAAS↔#`software as a service` as a synonym. • If you search for `saas`, AI Search retains the `as` and `a` terms in the synonym expansion. As a result, the search query matches the following records: ◦ records that contain the original search query term `saas` ◦ records that contain the synonym expansion phrase `software as a service`

Feature	Interaction with synonyms
	• If you search for `software as a service`, AI Search retains the `as` and `a` query terms when checking for synonym expansions. As a result, the search query matches the following records: ◦ records that contain the phrase `software as a service` ◦ records that contain the synonym expansion term `saas`
Typo handling	If a misspelled search query term is defined as a synonym, AI Search expands that synonym without auto-correcting it.

Publishing synonym dictionaries

Publishing a synonym dictionary makes its settings and active synonyms take effect in all linked search profiles.

Publishing a search profile updates the profile's synonym dictionary links but doesn't update the dictionary's settings or active synonyms in AI Search.

Create a synonym dictionary linked to a search profile

Define a new synonym dictionary and link it to a search profile.

Before you begin

Role required: ais_admin

About this task

To make a synonym dictionary affect search behavior, you must link it to a search profile, either at creation time or afterward. To learn about linking a saved dictionary to a search profile, see Link a synonym dictionary to a search profile.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the new synonym dictionary to.
3. In the Synonyms related list, select **Create and Link**.
4. On the Dictionary form, fill in the fields.
   For a description of the field values, see Dictionary and AI Search Dictionary forms.
5. Select **Submit**.

Result

The new synonym dictionary appears in the Synonyms related list.

What to do next

To make new or updated synonym dictionary settings take effect, publish the search profile that you edited. For details on publishing a search profile, see Publish a search profile.

Create synonyms

Define terms to treat as equivalent in search queries.

Before you begin
Role required: ais_admin

About this task
You must have already created a synonym dictionary. For information on creating synonym dictionaries, see Create a synonym dictionary linked to a search profile.

Procedure
1. Navigate to **All** > `AI Search` > `AI Search` **Engine** > **Dictionaries**.
2. Open the synonym dictionary that you want to add the new synonym to.
3. In the AI Search Dictionary Terms related list, select **New**.
4. On the Dictionary Term form, fill in the fields.
   For a description of the field values, see Dictionary Term form.
5. Select **Submit**.
   The new synonym appears in the AI Search Dictionary Terms related list.
6. **Optional:** Repeat the preceding steps for each additional synonym that you want to create.
7. Choose one of the following actions.
   - To save the modified dictionary and refresh its active synonyms in all linked search profiles, select **Publish**.
   - To save the modified dictionary without refreshing its active synonyms in all linked search profiles, select **Update**.

Example: Define a synonym with equivalent words and phrases

This example shows a synonym defining `vacation`, `holiday`, `PTO`, and the phrase `time off` as equivalent.

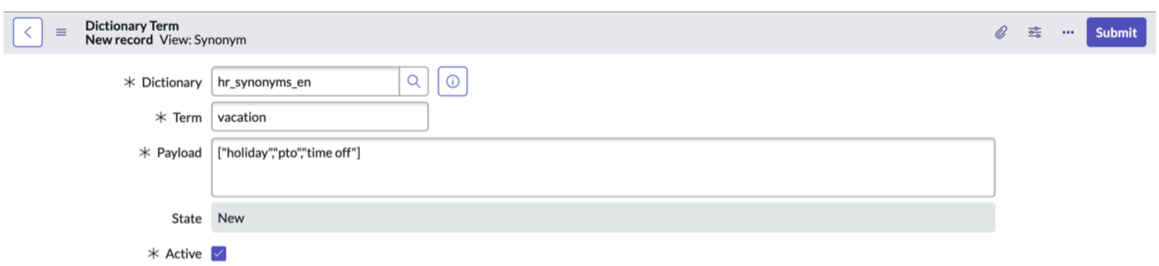

Clone a synonym dictionary

Create a copy of an existing synonym dictionary.

Before you begin
Role required: ais_admin

About this task
Cloning a synonym dictionary copies the source dictionary's settings and all of its all defined synonym terms. This process can save you time when you need multiple synonym dictionaries that have many synonym terms in common.

Procedure
1. Navigate to **All** > `AI Search` > **Search Query Settings** > **Dictionaries**.
2. Expand the Synonym record group if necessary and open the synonym dictionary that you want to clone.

3. Select **Clone**.
4. In the Set Dictionary Label dialog box, enter a label for the new synonym dictionary in the **Cloned dictionary label** field and select **OK**.

Result

The new synonym dictionary form opens. All synonym terms defined in the source dictionary appear in the AI Search Dictionary Terms related list.

Link a synonym dictionary to a search profile

Link an existing synonym dictionary to one or more search profiles.

Before you begin

Role required: ais_admin

About this task

You can link a single synonym dictionary to multiple search profiles.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the synonym dictionary to.
3. In the Synonyms related list, select **Link Existing**.
4. On the Search Profile - Dictionary Mapping form, select the synonym dictionary and search profile to link.
   a. Select the reference lookup icon ( ) for the **Dictionary** field.
   b. Select the synonym dictionary that you want to link to the search profile.
   c. Select **Submit**.
   For a description of the field values, see Search Profile - Dictionary Mapping form.

Result

The selected synonym dictionary appears in the search profile's Synonyms related list.

What to do next

To make the new synonym dictionary linking take effect, publish the search profile. For more information on publishing a search profile, see Publish a search profile.

Unlink a synonym dictionary from a search profile

Remove a synonym dictionary from a search profile.

Before you begin

Role required: ais_admin

About this task

Unlinking a synonym dictionary from a search profile doesn't delete the dictionary. To delete a synonym dictionary, you must first unlink it from all search profiles. For steps to delete a synonym dictionary, see Delete a synonym dictionary.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to unlink the synonym dictionary from.

3. In the Synonyms related list, select the option for the synonym dictionary that you want to unlink from the search profile.
4. Select **Unlink Selected**.

Result

The selected synonym dictionary disappears from the search profile's Synonyms related list.

What to do next

To make the change in synonym dictionary linking take effect, publish the search profile. For more information on publishing a search profile, see Publish a search profile.

Delete a synonym dictionary

Remove a synonym dictionary that isn't linked from any search profile.

Before you begin

Role required: ais_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Query Settings** > **Dictionaries**.
   The AI Search Dictionaries list appears with a **[Type] [!=] [Spell Check]** filter condition applied and records grouped by dictionary type.
2. Expand the Synonym record group if necessary and open the synonym dictionary that you want to delete.
3. Select **Delete**.

   Note:
   If the synonym dictionary is still linked to any search profiles, deletion fails and an error message lists the linked profiles. Unlink the dictionary from these search profiles, then repeat the preceding steps. For details on unlinking a synonym dictionary from a search profile, see Unlink a synonym dictionary from a search profile.

Result

The selected synonym dictionary disappears from the AI Search Dictionaries list.

Stop words

Stop words remove search query terms that don't produce meaningful results. Improve search precision and relevancy by configuring stop word dictionaries and defining stop words.

Stop words are common terms that appear on a large proportion of indexed records. Examples of such terms can include articles, conjunctions, personal pronouns, and prepositions. By defining these terms as stop words, you can improve search precision and result relevancy for users.

You can define stop words in the default stop word dictionary for a search profile, or create your own stop word dictionaries linked to search profiles. Create and publish stop words for each dictionary. Publish the dictionary to make your stop words take effect for searches that use linked search profiles.

You can delete stop word dictionaries that aren't linked to any search profiles.

Stop word usage in search

When a search query includes a term defined as a stop word, AI Search removes the term from the search query.

For example, a stop word might remove the term `my` from a search query for `upgrade my laptop`. This modified search query only matches records containing the more relevant terms `upgrade` and `laptop`.

Conditions for stop word eligibility

To be eligible for removal from a search query's terms, a stop word must satisfy the following conditions:

- It must be active.
- The stop word dictionary that includes it must be published.
- Its length doesn't exceed the dictionary term-length limit.

The system locates eligible stop words in stop word dictionaries as follows:

Condition	Result
The search profile specified for the search query links to a published stop word dictionary for the search query's language.	The system only considers stop words from the linked stop word dictionary. As an example, suppose a search application's search profile links to a French stop word dictionary. When a user searches in French in this application, only stop words from the linked French stop word dictionary are eligible for removal from the search query terms.
The search profile specified for the search query doesn't link to a published stop word dictionary for the search query's language.	The system only considers stop words from the default (English) stop word dictionary. As an example, suppose a search application's search profile has no linked Spanish stop word dictionary. When a user searches in Spanish in this application, only stop words from the default (English) stop word dictionary are eligible for removal from the search query terms.

For more details on language support in stop word dictionaries, see Language dependence.

Dictionary term-length limit

By default, AI Search ignores stop word terms that exceed 256 characters in length. You can override this default limit by adding the *`query.maxDictionaryTermLength`* system property. The integer value of this system property indicates the maximum allowable character length for stop word and synonym dictionary terms.

For information on adding a system property, see Add a system property.

Warning:
Exercise caution when increasing the dictionary term-length limit beyond its default value. Processing long dictionary terms increases AI Search's memory usage, and in some cases search queries may fail with out-of-memory errors.

Multi-word phrases

A stop word can include either a single word or a multi-word phrase. AI Search removes multi-word phrases from the search query when the query terms exactly match the order and proximity of the phrased words in the stop word definition. If a user searches with these words in a different arrangement, no term removal occurs.

For example, the following table describes AI Search stop word removal behavior for a stop word that includes the multi-word phrase `retirement account`.

Search query terms	Stop word removal behavior
my retirement account	AI Search removes the terms `retirement account` from the search query.
my bank retirement account	AI Search removes the terms `retirement account` from the search query.
my account retirement	AI Search doesn't remove any terms from the search query.
my retirement bank account	AI Search doesn't remove any terms from the search query.

Case sensitivity

By default, AI Search ignores letter case when comparing stop words with search query terms. You can enable case sensitivity for a stop word dictionary. This option affects all stop words defined in the dictionary.

When case sensitivity is enabled, stop words in the dictionary only match search query terms or phrases that include the same combination of uppercase and lowercase letters.

For example, if you define a `RAID` stop word in a dictionary with case sensitivity enabled, AI Search ignores the search query term `RAID` but treats the search query term `raid` normally.

Punctuation sensitivity

By default, AI Search ignores punctuation when comparing stop words with search query terms. You can enable punctuation sensitivity for a stop word dictionary. This option affects all stop words defined in the dictionary.

When punctuation sensitivity is enabled, stop words in the dictionary only match search query terms or phrases that include the same punctuation.

For example, if you define a `one-half` stop word in a dictionary with punctuation sensitivity enabled, AI Search removes the search query term `one-half` but doesn't remove the search query term `one half`.

Language dependence

Stop word removal is language-dependent. You specify a language for each stop word dictionary, and this language applies to all stop words defined in the dictionary. Stop words from dictionaries with the same language as the search query are eligible for removal.

Note:
If the search profile specified for a search query has no linked stop word dictionary for the search query language, AI Search uses the default (English) stop word dictionary as a fallback.

AI Search supports stop word dictionaries for all languages activated in your instance. For the list of languages you can activate, see Activate a language.

You can only link one stop word dictionary per language to a given search profile.

Interaction with other search features

The following table describes interactions between stop words and other search features.

Feature	Interaction with stop words
Genius Results	If a search query term that matches a Genius Result configuration's Term trigger is defined as a stop word, AI Search evaluates the trigger before removing that stop word from the search query. When computing Q&A Genius Results for a search query, AI Search ignores stop words in the query.
Lemma and Unicode normalization	If a search query term is defined as a stop word, AI Search removes that term without normalizing it.
Result improvement rules	If a search query term that matches a result improvement rule's Query trigger is defined as a stop word, AI Search evaluates the rule trigger before removing the stop word from the search query.
Synonyms	If a synonym definition includes a term defined as a stop word, AI Search retains that term when checking whether the synonym expands a search query. For example, suppose you define `as` and `a` as stop words and SAAS↔`#software as a service` as a synonym.

Feature	Interaction with stop words
	• If you search for `saas`, AI Search retains the `as` and `a` terms in the synonym expansion. As a result, the search query matches the following records: ◦ records that contain the original search query term `saas` ◦ records that contain the synonym expansion phrase `software as a service` • If you search for `software as a service`, AI Search retains the `as` and `a` query terms when checking for synonym expansions. As a result, the search query matches the following records: ◦ records that contain the phrase `software as a service` ◦ records that contain the synonym expansion term `saas`
Typo handling	If a misspelled search query term is defined as a stop word, AI Search auto-corrects that term normally instead of removing it. AI Search ignores terms defined as stop words when determining an auto-correction for a misspelled search query term. For example, if you define `email` as a stop word, AI Search doesn't offer `email` as an auto-correction when you search for `emial`.
Wildcard operators	When expanding search terms that contain `%` or `*` wildcard operators, AI Search ignores terms defined as stop words. For example, suppose you define `the` and `their` as stop words. A search for `the*` won't expand to match `the` or `their`, but will still match non-stop word terms such as `there` and `these`.

Publishing stop word dictionaries

Publishing a stop word dictionary makes its settings and active stop words take effect in all linked search profiles.

Publishing a search profile updates the profile's stop word dictionary links but doesn't update the dictionary's settings or active stop words in AI Search.

Create a stop word dictionary linked to a search profile

Define a new stop word dictionary and link it to a search profile.

Before you begin

Role required: ais_admin

About this task

To make a stop word dictionary affect search behavior, you must link it to a search profile, either at creation time or afterward. You can only link one stop word dictionary per language to a given search profile. To learn about linking a saved stop word dictionary to a search profile, see Link a stop word dictionary to a search profile.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the new stop word dictionary to.
3. In the Stop Words related list, select **Create and Link**.
4. On the Dictionary form, fill in the fields.
   For a description of the field values, see Dictionary and AI Search Dictionary forms.
5. To save the new stop word dictionary and link it to the specified search profile, select **Submit.**

Result

The new stop word dictionary appears in the Stop Words related list.

What to do next

To make new or updated stop word dictionary settings take effect, publish the search profile that you edited. For details on publishing a search profile, see Publish a search profile.

Create stop words

Define terms to remove from search queries.

Before you begin

Role required: ais_admin

About this task

You must have already created a stop word dictionary. For information on creating stop word dictionaries, see Create a stop word dictionary linked to a search profile.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Engine** > **Dictionaries**.
2. Open the stop word dictionary that you want to add the new stop word to.
3. In the AI Search Dictionary Terms related list, select **New**.
4. On the Dictionary Term form, fill in the fields.
   For a description of the field values, see Dictionary Term form.
5. Select **Submit**.
   The new stop word appears in the AI Search Dictionary Terms related list.
6. **Optional:** Repeat the preceding steps for each additional stop word that you want to create.
7. Choose one of the following actions.
   - To save the modified dictionary and refresh its active stop words in all linked search profiles, select **Publish**.
   - To save the modified dictionary without refreshing its active stop words in all linked search profiles, select **Update**.

Clone a stop word dictionary

Create a copy of an existing stop word dictionary.

Before you begin

Role required: ais_admin

About this task

Cloning a stop word dictionary copies the source dictionary's settings and all of its all defined stop word terms. This process can save you time when you need multiple stop word dictionaries that have many stop word terms in common.

Procedure

1. Navigate to **All** > `AI Search` > **Search Query Settings** > **Dictionaries**.
2. Expand the Stop Word record group if necessary and open the stop word dictionary that you want to clone.
3. Select **Clone**.
4. In the Set Dictionary Label dialog box, enter a label for the new stop word dictionary in the **Cloned dictionary label** field and select **OK**.

Result

The new stop word dictionary form opens. All stop word terms defined in the source dictionary appear in the AI Search Dictionary Terms related list.

Link a stop word dictionary to a search profile

Link an existing stop word dictionary to one or more search profiles.

Before you begin

Role required: ais_admin

About this task

You can link a single stop word dictionary to multiple search profiles. You can only link one stop word dictionary per language to a given search profile.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the stop word dictionary to.
3. In the Stop Words related list, select **Link Existing**.
4. On the Search Profile - Dictionary Mapping form, select the stop word dictionary and search profile to link.

    a. Select the reference lookup icon ( ) for the **Dictionary** field.

    b. Select the stop word dictionary that you want to link to the search profile.

    c. Select **Submit**.

    For a description of the field values, see Search Profile - Dictionary Mapping form.

Result

The selected stop word dictionary appears in the search profile's Stop Words related list.

What to do next

To make the new stop word dictionary linking take effect, publish the search profile. For more information on publishing a search profile, see Publish a search profile.

Unlink a stop word dictionary from a search profile

Remove a stop word dictionary that isn't linked from any search profile.

Before you begin

Role required: ais_admin

About this task

Unlinking a stop word dictionary from a search profile doesn't delete the dictionary. To delete a stop word dictionary, you must first unlink it from all search profiles. For steps to delete a stop word dictionary, see Delete a stop word dictionary.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to unlink the stop word dictionary from.
3. In the Stop Words related list, select the option for the stop word dictionary that you want to unlink from the search profile.
4. Select **Unlink Selected**.

Result

The selected stop word dictionary disappears from the search profile's Stop Words related list.

What to do next

To make the change in stop word dictionary linking take effect, publish the search profile. For more information on publishing a search profile, see Publish a search profile.

Delete a stop word dictionary

Remove a stop word dictionary that isn't linked from any search profile.

Before you begin

Role required: ais_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Query Settings** > **Dictionaries**.
   The AI Search Dictionaries list appears with a **[Type] [!=] [Spell Check]** filter condition applied and records grouped by dictionary type.
2. Expand the Stop Word record group if necessary and open the stop word dictionary that you want to delete.
3. Select **Delete**.

   Note:
   If the stop word dictionary is still linked to any search profiles, deletion fails and an error message lists the linked profiles. Unlink the dictionary from these search profiles, then repeat the preceding steps. For details on unlinking a stop word dictionary from a search profile, see Unlink a synonym dictionary from a search profile.

Result

The selected stop word dictionary disappears from the AI Search Dictionaries list.

Typo handling

AI Search replaces misspelled search query terms with auto-correction terms derived from indexed content. Configure typo handling dictionary settings to control derivation of auto-correction terms.

To define auto-correction terms for a search profile, modify settings in the profile's linked typo handling dictionary. The typo handling dictionary derives a list of auto-correction terms from indexed content found in search sources linked to its search profile. Update the typo handling

dictionary to save its settings. Publish the linked search profile to make auto-correction terms available for searches that use the profile.

AI Search automatically creates a typo handling dictionary for a search profile when you link search sources to that profile. You can't delete a typo handling dictionary, unlink it from its search profile, or link it to any other search profile.

To disable typo handling auto-correction for a search application, deselect the **Enable Typo Handling** option in the search application configuration. For more details on search application configuration settings, see Create a search application configuration for AI Search.

Conditions for auto-correction term eligibility

When deriving its list of auto-correction terms, the typo handling dictionary only considers indexed terms that satisfy the following conditions:

- The term includes at least 3 characters and no more than 255 characters.
- The term appears at least once in indexed records from the search profile's linked search sources.

Auto-correction of search query terms

When a search query includes a term not present in the typo handling dictionary, AI Search looks for terms in the dictionary with similar spelling. If it finds a good match, it automatically replaces the original search query term with the dictionary term.

For example, a search for `emial`, which isn't present on any indexed record, might auto-correct to the similarly spelled term `email` which exists on multiple indexed records and thus appears in the typo handling dictionary.

A notification message displays the auto-corrected term. Users can select the link in this message to repeat their search with the original query term.

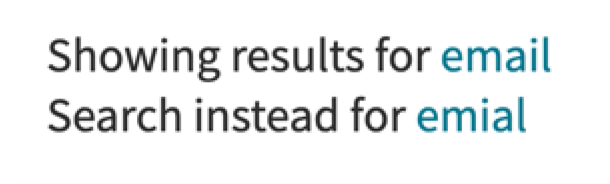

Terms excluded from auto-correction

Auto-correction ignores search query terms that satisfy any of the following conditions:

Condition	Description
Search query term includes only 1 or 2 characters	Auto-correction ignores search query terms below 3 characters in length. For example, AI Search doesn't auto-correct the search query term `gi`.
Search query term includes only uppercase letters	To enable search for abbreviations, auto-correction ignores search query terms that contain only uppercase letters.

Condition	Description
	For example, AI Search doesn't auto-correct the search query term `PINN`, even if it auto-corrects the lowercase form `pinn` to `pin`.
Search query term includes one or more numeric digits	To enable search for numbers and alphanumeric identifiers, auto-correction ignores search query terms that include numeric digits. For example, AI Search doesn't auto-correct the search query term `402K`.
Search query term excluded from typo handling dictionary	Auto-correction ignores search query terms excluded from the typo handling dictionary. For instructions on excluding terms from a typo handling dictionary, see Exclude a term from a typo handling dictionary.

Automatic refresh of auto-correction term lists

Each typo handling dictionary automatically refreshes its list of auto-correction terms after a set interval. The refresh operation has the same effect as republishing the dictionary's linked search profile.

By default, typo handling dictionaries refresh every 24 hours. If you want the term list refreshed less frequently, you can increase this interval in the dictionary settings.

Language dependence

Typo handling is language-dependent. AI Search derives a separate list of auto-correction terms for each supported language found in the search sources' indexed content.

AI Search supports derivation of auto-correction terms for Brazilian Portuguese, Dutch, English, French - Canada, French, German, Italian, Portuguese, Spanish, and Swedish. Typo handling auto-correction term derivation isn't supported for Japanese, Korean, Simplified Chinese, or Traditional Chinese.

Interaction with other search features

The following table describes interactions between typo handling and other search features.

Feature	Interaction with typo handling
Genius Results	Auto-corrected search query terms can't trigger Genius Result configurations with Term trigger conditions. When computing Q&A Genius Results for a search query, AI Search corrects terms in the query.

Feature	Interaction with typo handling
Lemma and Unicode normalization	AI Search performs lemma and Unicode normalization on auto-corrected search query terms.
Result improvement rules	An auto-corrected search query term can trigger a result improvement rule if it matches the rule's Query trigger.
Stop words	If a misspelled search query term is defined as a stop word, AI Search auto-corrects that term normally instead of removing it. AI Search ignores terms defined as stop words when determining an auto-correction for a misspelled search query term. For example, if you define `email` as a stop word, AI Search doesn't offer `email` as an auto-correction when you search for `emial`.
Synonyms	If a misspelled search query term is defined as a synonym, AI Search expands that synonym without auto-correcting it.

Modify settings for a typo handling dictionary

Modify typo handling dictionary settings for a search profile.

Before you begin

Role required: ais_admin

About this task

Each typo handling dictionary is linked to a search profile. For more information about search profiles, see Search profiles.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that contains the typo handling dictionary that you want to modify.
3. In the Typo Handling related list, open the record for the typo handling dictionary.
4. On the AI Search Dictionary form, update the fields.
   For a description of the field values, see Dictionary and AI Search Dictionary forms.
5. Select **Update**.

What to do next

To make updated settings for a typo handling dictionary take effect, publish the search profile linked to the dictionary. For details on publishing a search profile, see Publish a search profile.

Exclude a term from a typo handling dictionary

Prevent AI Search from auto-correcting a term in user search queries by excluding the term from a typo handling dictionary.

Before you begin

You must have already defined a search profile with one or more linked search sources.

Role required: ais_admin

Procedure

1. Navigate to **All** > *AI Search* > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to exclude a typo handling term for.
3. In the Typo Handling related list, open the typo handling dictionary.
4. In the AI Search Dictionary Terms related list, select **New**.
5. On the Dictionary Term form, fill in the fields.
   For a description of the field values, see Dictionary Term form.
6. Select **Submit**.
   The new term appears in the AI Search Dictionary Terms related list with the **State** set to *New*.
7. **Optional:** Repeat the preceding steps for each additional term that you want to exclude.
8. Select **Publish**.
   The typo handling dictionary appears in the Typo Handling related list with the **State** set to *Published*.

What to do next

To make updated typo handling settings take effect, publish the search profile that you edited. For details on publishing a search profile, see Publish a search profile.

Genius Results

Optionally display the best answers for a search query as actionable Genius Result cards included with search results.

A Genius Result configuration includes a trigger condition and logic for creating answers. The trigger condition determines whether the Genius Result configuration activates for a search query. When activated, the configuration checks for search results with extremely high relevancy. If it finds any, its answer logic populates one or more Genius Result cards with information and actions derived from these results. The Genius Result cards appear alongside normal search results, giving the user immediate access to the most relevant answers and actions.

To use a Genius Result configuration in searches, link it to one or more of your search profiles. For details on this procedure, see Link a Genius Result configuration to a search profile.

Genius Result evaluation order and result limit

AI Search evaluates active Genius Result configurations for a search profile in an order determined by their numeric **Order** field values, as defined in the search profile's Genius Results related list. Genius Result configurations with lower **Order** field values are evaluated before those with higher field values. As an example, a Genius Result configuration with **Order** 100 is evaluated before one with **Order** 200. To learn how to change the evaluation order, see Set the evaluation order for Genius Result configurations in a search profile.

Note:
Order field values for linked Genius Result configurations are set separately for each search profile. A Genius Result configuration that's linked to multiple search profiles can have a different **Order** field value in each profile. For instance, you might set an **Order** value of 100 for the Q&A Genius Result configuration in one search profile but set a value of 750 for the same Genius Result configuration in a second search profile that it's linked to.

A search application configuration's **Genius Results Limit** field value specifies the maximum number of Genius Result cards the application can display. For more details on search application configuration settings, see Search Application Configuration form.

Genius Result trigger conditions

You can select any one of the following trigger conditions for a Genius Result configuration:

- *Always*: The Genius Result configuration activates for all search queries.
- *NLU*: The Genius Result configuration activates when the search query's detected intent matches a specified Natural Language Understanding (NLU) model intent. For more information on NLU models and intents, see Natural Language Understanding.
- *Term*: The Genius Result configuration activates when the search query consists of a single specified term or phrase.

Note:
A Genius Result configuration can trigger and still not return answers. Genius Result answers only display when AI Search has extremely high confidence that there's a correct response for the search query. You should typically expect to see Genius Result answers for a small percentage of search queries.

Genius Result answer types

A Genius Result configuration can include either of the following answer types:

- *Search*: When triggered, displays Genius Result cards for results from a search query. The search query must be defined in a server-side script included in the Genius Result configuration. Each search result displays as a separate Genius Result card, to a maximum of three cards. If the search query returns no results, no Genius Result card is displayed.
- *Script*: When triggered, displays a Genius Result card containing key-value pairs. The object containing these key-value pairs must be defined in a server-side script included in the Genius Result configuration. The script can use methods from any ServiceNow AI Platform® public API to construct this object.

Language support

In the base system, Genius Results only support English search queries.[store-future: BEGIN review] To add support for French, German, Japanese, and Spanish search queries, install and activate the Genius Results for AI Search ServiceNow® Store application. For details on the NLU models and Genius Result configurations included with this application, see Genius Results for AI Search.[End]

Table record limit for Genius Result configurations with NLU trigger conditions

If an NLU model includes a table vocabulary source that references a table containing more than 100,000 records, sync fails for the vocabulary source. When sync fails, the model can't be successfully trained, and Genius Result configurations that use the model in their NLU trigger conditions don't operate correctly.

> **Warning:**
> Don't activate the People Genius Result configuration if your instance has more than 100,000 records in any of the following tables:
>
> - Catalog Item [sc_cat_item]
> - Department [cmn_department]
> - Location [cmn_location]
> - User [sys_user]
>
> These tables are configured as vocabulary sources in the Default AI Search Genius Results Model specified in the People configuration's NLU trigger condition. If your instance has more than 100,000 records in any of the listed tables, training will fail for the NLU model and the People Genius Result configuration won't operate correctly.

For more information on NLU vocabulary sources, see NLU vocabulary.

Genius Result review feedback

You can review individual Genius Result answers and provide feedback about their accuracy. To provide review feedback on a Genius Result answer, select one of the thumb icons shown on the answer card:

- If the answer answers your question accurately, select the thumbs up icon ().
- If the answer answers your question inaccurately, select the thumbs down icon ().

Your feedback helps ServiceNow improve future versions of Genius Result configurations.

Interaction with other search features

The following table describes interactions between Genius Results and other search features.

Feature	Interaction with Genius Results
Lemma and Unicode normalization	Search query terms added by lemma or Unicode normalization can't trigger Genius Result configurations with Term trigger conditions.
Stop words	If a search query term that matches a Genius Result configuration's Term trigger is defined as a stop word, AI Search evaluates the trigger before removing that stop word from the search query.
Synonyms	Synonym expansion terms can't trigger Genius Result configurations with Term trigger conditions.
Typo handling	Auto-corrected search query terms can't trigger Genius Result configurations with Term trigger conditions.

Genius Result configurations in the base system

Preconfigured AI Search Genius Result configurations display concise, actionable answers derived from high-relevancy search results. These configurations return Genius Result answers for Catalog Items, tables, users, and Knowledge articles that match your search.

The base system includes multiple Genius Result configurations.

Note:
You can link the base system's Genius Result configurations to your own search profiles. You can't modify these Genius Result configurations, but you can clone them and modify the cloned configurations.

Catalog Item Genius Results

Catalog Item Genius Results display top search results from the Catalog Item [sc_cat_item] table. Each answer card shows a single Catalog Item record. You can create a request for the Catalog Item directly from the answer card.

Sample Catalog Item Genius Result answer card

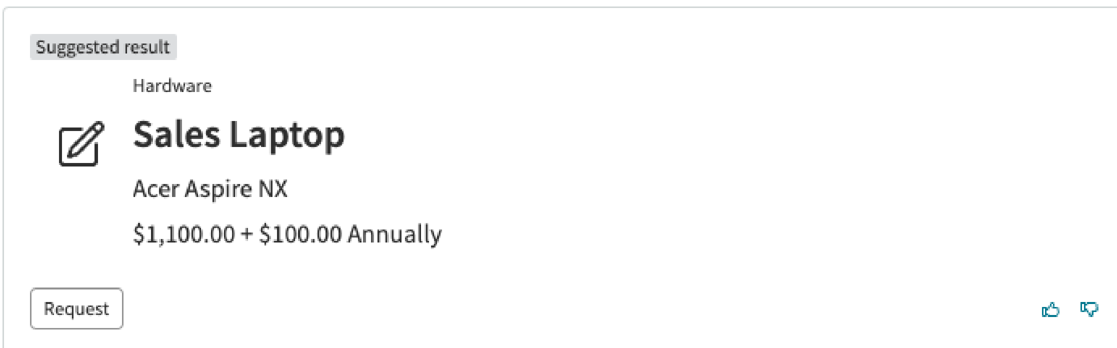

Select **Request** to create a new request for the specified Catalog Item.

Catalog Item Genius Results only support English searches.

NLQ Genius Results

NLQ (Natural Language Query) Genius Results use NLQ processing to surface relevant results from tables that match your search query. Each NLQ Genius Result answer card displays a preview of records from matching tables. You can navigate to a matching table's list view or the CMDB workspace directly from the Genius Result answer card.

NLQ Genius Results don't use indexed sources when finding tables that match your query. You don't need to define a table as an indexed source to make it eligible to appear in NLQ Genius Result answers.

When computing NLQ Genius Result answers, AI Search identifies tables that match your search using the following NLQ query types:

- Natural language queries similar to the ones made using Analytics Center. To learn more about Analytics Center table query capabilities, see Analytics Center.
- Natural language queries similar to the ones made using the Configuration Management Database (CMDB). These CMDB queries only match tables that aren't defined as indexed sources. For more information about CMDB table query capabilities, see Configuration Management.

AI Search combines the results from these query types when creating NLQ Genius Result answer cards. As a result, a single NLQ Genius Result answer card may contain tables returned by Analytics Center queries, CMDB queries, or both.

For more information about NLQ processing, see Natural Language Query.

Answer card format

The NLQ Genius Result answer card shows different results depending on how many tables your search matched.

Search matches a single table

When your search only matches one table, the answer card's preview list displays up to five records from that table. To see more records from the table, select **View full table**.

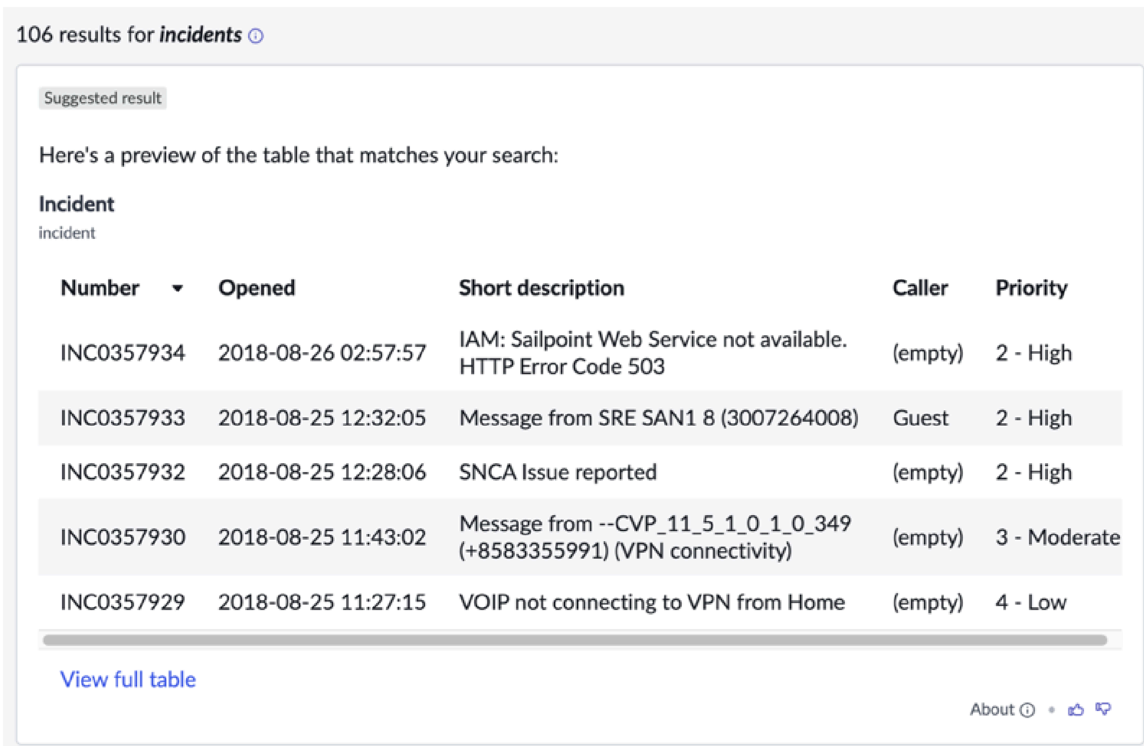

Number	Opened	Short description	Caller	Priority
INC0357934	2018-08-26 02:57:57	IAM: Sailpoint Web Service not available. HTTP Error Code 503	(empty)	2 - High
INC0357933	2018-08-25 12:32:05	Message from SRE SAN1 8 (3007264008)	Guest	2 - High
INC0357932	2018-08-25 12:28:06	SNCA Issue reported	(empty)	2 - High
INC0357930	2018-08-25 11:43:02	Message from --CVP_11_5_1_0_1_0_349 (+8583355991) (VPN connectivity)	(empty)	3 - Moderate
INC0357929	2018-08-25 11:27:15	VOIP not connecting to VPN from Home	(empty)	4 - Low

Search matches multiple tables

When your search matches multiple tables, the answer card displays the matching tables as pills or as entries in a drop-down list. You can select any of the matching tables to preview.

The preview list displays up to five records from the selected table. To see more records from the table, select **View full table**.

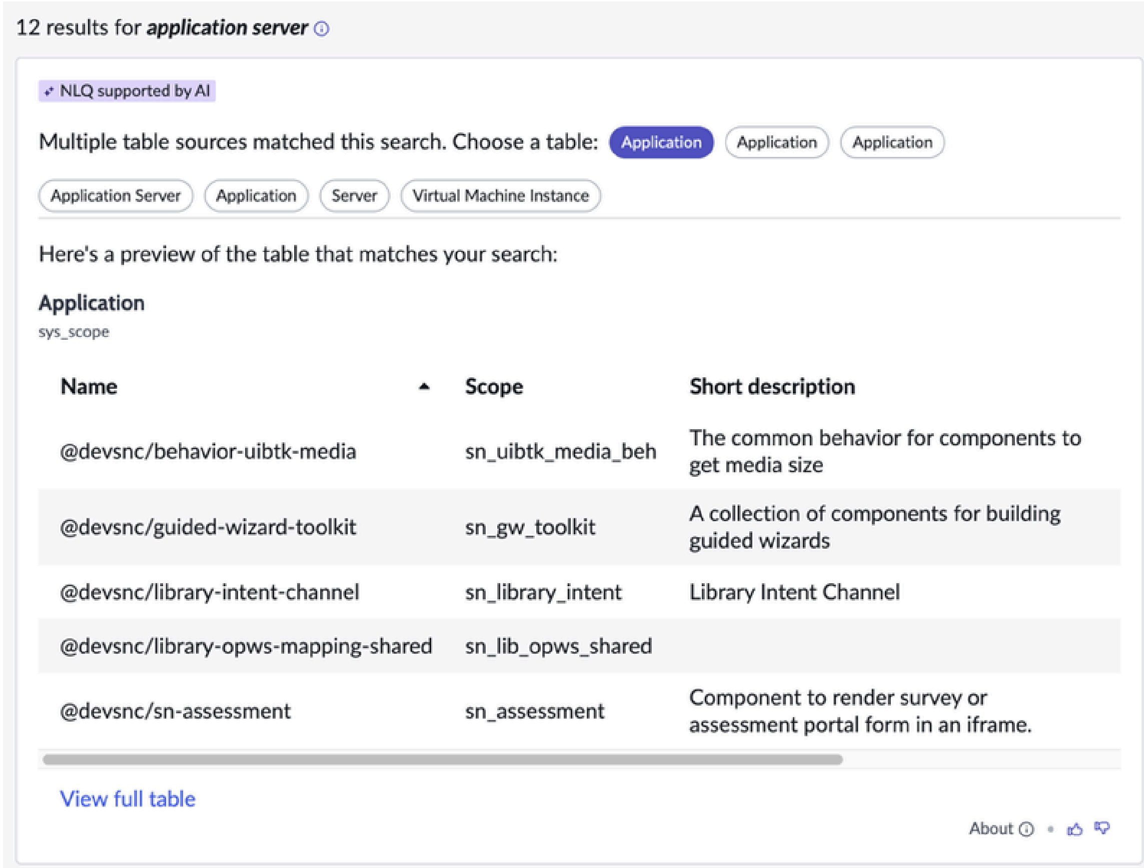

Viewing full tables

When you select **View full table** in an NLQ Genius Result answer card, the action performed depends on how the matching table was found. The default actions are as follows.

- If the table was found by an Analytics Center table query, the system displays it in list view.
- If the table was found by a CMDB table query, the system displays it in the CMDB Workspace.

Search administrators can customize these actions. For more information, see the Configure navigation actions for NLQ Genius Result answer cards section.

Review feedback

Some NLQ Genius Result answers show `NLQ supported by AI` as their answer type, indicating that they used artificial intelligence to help find tables that matched your search query. Because these answers are automatically generated, you should review them for accuracy.

To provide review feedback on an NLQ Genius Result answer, select one of the thumb icons shown on the answer card. If the answer's information is accurate, select the thumbs up icon ( ). If the answer's information isn't accurate, select the thumbs down icon ( ).

Your feedback helps ServiceNow improve future versions of this Genius Result configuration.

Limitations

Like the other base system Genius Result configurations, NLQ Genius Results only support searches in English.

NLQ Genius Results rely on Natural Language Query processing, which can increase the time required to return results on the first search for a given set of terms. Consider this performance impact before enabling NLQ Genius Results in any of your AI Search applications.

NLQ Genius Results don't support the following table query options:

- Multi-table CMDB queries
- Grouping of table query results
- Rendering of table query results as pie charts or bar charts
- Trend and single score visualizations
- Keywords that return results in non-list formats, such as `total`, `count`, or `average`

Display NLQ Genius Result answers in Service Portal search applications

Configure the system to display search result preview lists for NLQ Genius Result answers in Service Portal search applications.

Before you begin

You only need to perform this task if you want to use NLQ Genius Results in your Service Portal search applications. If you only use NLQ Genius Results in Next Experience search applications, you don't need to perform this task.

Role required: admin

About this task

The *glide.service_portal.ais_nlq_enabled* system property determines whether the system displays search result preview lists for NLQ Genius Result answers in Service Portal search applications.

Administrators should set this system property's value to *true* if the NLQ Genius Result configuration is active in any Service Portal application.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
    a. Select **All**.
    b. In the **Filter** field, enter `sys_properties.list`.
    c. Press Enter.
2. Select **New**.
3. On the System Property form, enter these values.

Field	Value
Name	glide.service_portal.ais_nlq_enabled
Type	true | false
Value	true

    For details on this system property's values, see the System properties for NLQ Genius Results section.

4. Press **Submit**.
    The new *glide.service_portal.ais_nlq_enabled* record appears in the System Property [sys_properties] table list view.

Configure navigation actions for NLQ Genius Result answer cards

Configure the navigation actions that trigger when a user selects **View full table** on an NLQ Genius Result action card. Each navigation action includes a redirection payload and can optionally include filters for a table's name, the NLQ source that returned the table, or the search profile used by the search application.

Before you begin

Role required: ais_admin

About this task

When you select **View full table** on an NLQ Genius Result answer card, the action performed depends on how the matching table was found. The default actions are as follows.

- For a table found by an Analytics Center table query, the table displays in list view.
- For a table found by a CMDB table query, the CMDB workspace opens.

As a search administrator, you can customize the actions available for NLQ Genius Result answer cards.

Procedure

1. Navigate to the AI Search NLQ Navigation [ais_nlq_navigation] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `ais_nlq_navigation.list`.
   c. Press Enter.
2. Edit an existing navigation record or create a new one by selecting **New**.
3. On the AI Search NLQ Navigation form, fill in the fields for your navigation action.

Field	Description
Source	Optional NLQ source filter for the navigation action. The action only operates for tables that satisfy this source filter. When this field is left empty, the navigation action operates for tables from any NLQ source that isn't specified in another navigation action's **Source** field. As an example, suppose you have two navigation actions: one with **Source** set to `CMDB_WS` and a second one with no **Source** field value. ◦ When an NLQ Genius Result answer card previews a table returned by the CMDB NLQ source, selecting **View full table** triggers the first navigation action because the table name satisfies the action's NLQ source filter. ◦ When an NLQ Genius Result answer card previews a table from the Analytics Center NLQ source, selecting **View full table** instead triggers the second navigation

Field	Description
	action because the table's NLQ source doesn't satisfy the first action's NLQ source filter. ◦ Type: string ◦ Default value: none. ◦ Supported values: ▪ *AC*: The navigation action only operates for tables returned by the Analytics Center NLQ source. ▪ *CMDB_WS*: The navigation action only operates for tables returned by the CMDB NLQ source.
Action Payload	JSON payload for the navigation action. This field's value determines where AI Search redirects when you select **View full table** on an NLQ Genius Result answer card. When submitting the JSON payload for redirection, AI Search replaces the following references with relevant values: ◦ *${logID}*: Replaced with the NLQ transaction's ID as shown in system logs. ◦ *${query}*: Replaced with the query filter from the original search query that triggered the NLQ Genius Result. ◦ *${table}*: Replaced with the name of the table previewed on the NLQ Genius Result answer card. ◦ *${utterance}*: Replaced with the original search query that triggered the NLQ Genius Result. As an example, the default navigation action that opens a table in the CMDB workspace specifies this action payload that makes use of all four references: `{ "external": { "url": "now/cmdb/home/sub/smart-search-results/${table}/${utterance}/params/encoded-query/${query}/nlq-query-log/${logID}", "target": "_self" } }` ◦ Type: JSON ◦ Default value: none

Field	Description
Profile	Optional reference to a search profile record. The navigation action only operates in search applications that use the referenced search profile. When this field is left empty, the navigation action operates for any search application that doesn't use a search profile referenced in another navigation action's **Profile** field. As an example, suppose you have two navigation actions: one with **Profile** referencing the `Service Portal Default Search Profile` record and a second one with no **Profile** reference field value. ◦ When you search in an application that uses the Service Portal Default Service Profile, selecting **View full table** on an NLQ Genius Result answer card triggers the first navigation action because the search profile in use is the one referenced in the action. ◦ When you search in any other search application, selecting **View full table** instead triggers the second navigation action because the search profile in use isn't referenced in the first navigation action. ◦ Type: reference ◦ Default value: none
Action	Name for the navigation action. ◦ Type: string ◦ Default value: none
Table	Optional table filter for the navigation action. The action only operates for the table whose name matches this filter. When this field is left empty, the navigation action operates for any table that doesn't have its name specified in another navigation action's **Table** field. As an example, suppose you have two navigation actions: one with **Table** set to `my_custom_table` and a second one with no **Table** field value.

Field	Description
	○ When an NLQ Genius Result answer card previews your custom table, selecting **View full table** triggers the first navigation action because the table name satisfies the action's table filter. ○ When an NLQ Genius Result answer card previews any other table, selecting **View full table** instead triggers the second navigation action because the table name doesn't satisfy the first action's table filter. ○ Type: table_name ○ Default value: none

4. Select **Update** if you're editing an existing navigation action record, or **Submit** if you're creating a new one.

System properties for NLQ Genius Results

A system property determines whether the system displays search result preview lists for NLQ Genius Result answers in Service Portal search applications.

Property	Description
glide.service_portal.ais_nlq_enabled	Specifies whether the system displays search result preview lists for NLQ Genius Result answers in Service Portal search applications. If you activate the NLQ Genius Result configuration in any Service Portal search application, set this system property's value to `true`. • Type: true | false • Default value: false • Location: System Property [sys_properties] table • Supported values: ○ `true`: The system displays search result preview lists for NLQ Genius Result answers in Service Portal search applications. ○ `false`: The system doesn't display search result preview lists for NLQ Genius

Property	Description
	Result answers in Service Portal search applications. **Warning:** This value prevents NLQ Genius Result answers from displaying in Service Portal search applications.

People Genius Results

People Genius Results display top search results from the User [sys_user] table. Each People Genius Result answer card shows a single user record. You can view the user's full profile directly from the answer card.

Sample People Genius Result answer card

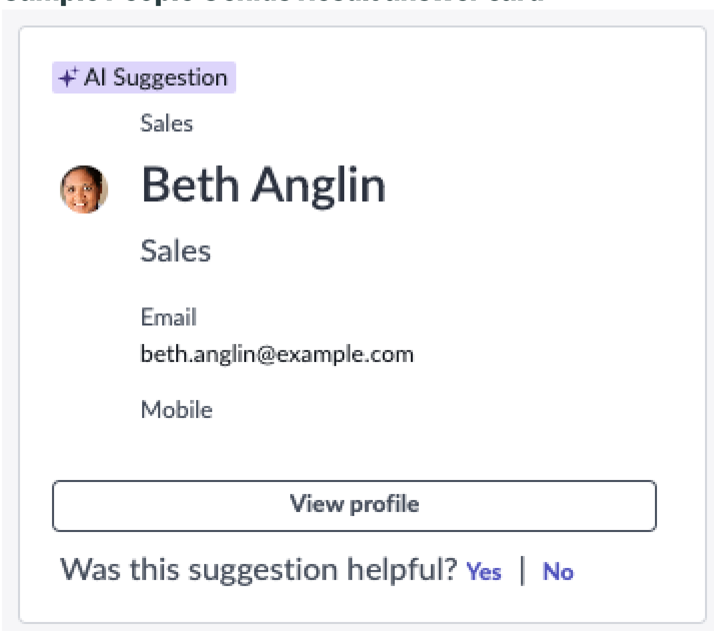

Select **View profile** to see the user's profile.

People Genius Results only support English searches.

The People Genius Result configuration is inactive in the base system. To learn how to activate it, see Activate People Genius Results.

Warning:
Don't activate the People Genius Result configuration if your instance has more than 100,000 records in any of the following tables:

- Catalog Item [sc_cat_item]
- Department [cmn_department]
- Location [cmn_location]
- User [sys_user]

These tables are configured as vocabulary sources in the Default AI Search Genius Results Model specified in the People configuration's NLU trigger condition. If your instance has more than 100,000 records in any of the listed tables, training will fail for the NLU model and the People Genius Result configuration won't operate correctly.

Activate People Genius Results

Enable AI Search's People default Genius Result configuration for use in your search profiles.

Before you begin

Role required: ais_admin

Warning:
Don't activate the People Genius Result configuration if your instance has more than 100,000 records in any of the following tables:

- Catalog Item [sc_cat_item]
- Department [cmn_department]
- Location [cmn_location]
- User [sys_user]

These tables are configured as vocabulary sources in the Default AI Search Genius Results Model specified in the People configuration's NLU trigger condition. If your instance has more than 100,000 records in any of the listed tables, training will fail for the NLU model and the People Genius Result configuration won't operate correctly.

About this task

The People default Genius Result configuration is inactive in the base system. To use this Genius Result configuration in your search profiles, you must activate it.

Note:
If you upgraded from the Quebec release, this Genius Result configuration is already activated. You don't need to perform the steps from this task.

As part of activation, you need to run a scheduled job that trains and publishes the Default AI Search Genius Results Model. This *NLU* model contains intents used by the People configuration's *NLU* trigger.

For more information on the People default Genius Result configuration, see People Genius Results.

Procedure

1. Train and publish the Default AI Search Genius Results Model.

   a. Navigate to **All** > **System Definition** > **Scheduled Jobs**.

   b. Open the `AI Search` **Create Train Model** job.

   c. Select **Execute Now**.

   d. Check the publication status of the NLU model by navigating to **All** > `AI Search` > **Search Query Settings** > **Search Intents**.
   When a **Last Published** field value is displayed, the NLU model is trained and published. You can now activate the People default Genius Result configuration.

2. Activate the People default Genius Result configuration.

a. Navigate to **All** > `AI Search` > **Search Query Settings** > **Genius Results**.

b. Open the People Genius Result configuration.

c. On the AI Search Genius Result Configuration form, select **Change Active State**.
For a description of the field values, see AI Search Genius Result Configuration form.
The Genius Result configuration's **Active** field value changes to `true`.

Result
The People Genius Result configuration is available for use in search profiles.

Q&A Genius Results

Q&A Genius Results display top search results extracted from HTML fields of records on the Knowledge [kb_knowledge] table and tables that extend it. Each Q&A Genius Result answer card shows a topic snippet and an answer snippet extracted from a single Knowledge article. You can view the full article directly from the answer card.

A snippet is a short, typically paragraph-length, section of text that contains useful and relevant information. A Genius Result answer's topic snippet indicates the subject for the answer, while the answer snippet provides details on that subject.

Sample Q&A Genius Result answer card

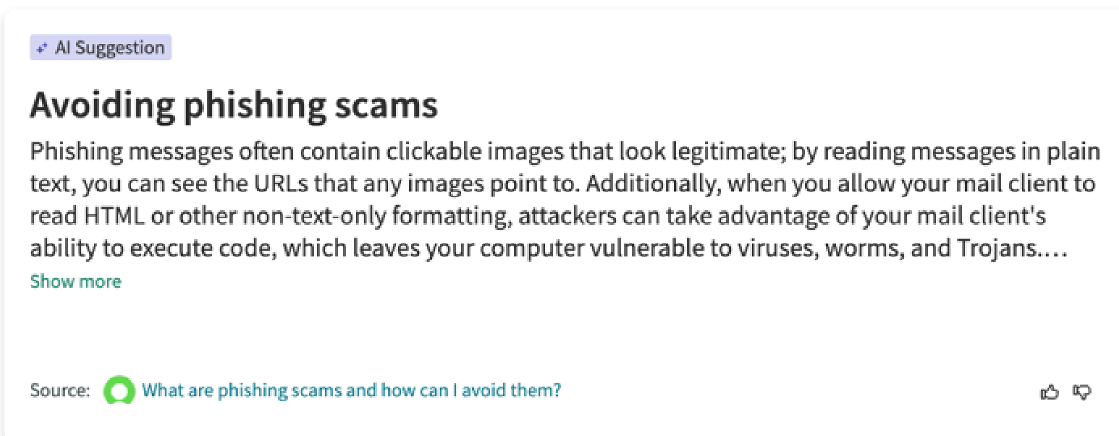

Select **View Article** to view the full Knowledge article.

Note:
Knowledge search property settings don't affect Q&A Genius Results. For more information on these settings, see Knowledge search properties.

Q&A Genius Results only support English searches.

Because the Q&A implementation is based on deep neural networks that are continually improving, Q&A results for specific queries may vary from release to release. Articles boosted or promoted by result improvement rules are more likely to appear as Q&A Genius Results, but aren't guaranteed to appear.

Topic and answer snippet extraction for Q&A Genius Results

By default, AI Search automatically determines which elements of a Knowledge article's HTML field values to extract as the Q&A topic and answer snippets. You can't configure the criteria for this extraction behavior.

Topic and answer snippet extraction behavior varies depending on the source table for the Knowledge article:

Knowledge article source	Q&A topic and answer snippet extraction
FAQ [kb_template_faq] table **Note:** This table is enabled when you activate the Knowledge Management Advanced plugin (com.snc.knowledge_advanced). For more information, see Activate the Knowledge Management Advanced plugin	AI Search extracts the Q&A topic snippet from the Knowledge article record's **Question** HTML field. It extracts the Q&A answer snippet from the record's **Answer** HTML field.
Knowledge [kb_knowledge] table and all other tables that extend it	AI Search extracts the Q&A topic and answer snippets from the Knowledge article record's HTML fields.

Because the default extraction behavior favors search precision, it may not extract all desired Q&A topic and answer snippets. You can override the extraction behavior for a Knowledge article by manually designating topic and answer snippets in that article's HTML fields. Manually designating snippets increases the likelihood that they will appear as Q&A results, but doesn't guarantee that they will.

To manually designate snippets for an article, set the following HTML classes on elements in the article's HTML fields.

HTML Class	Description
`class="sn_snippet_topic"`	Set this class on the HTML field element that contains text that you want to designate as the Q&A Genius Result topic snippet for the Knowledge article.
`class="sn_snippet_answer"`	Set this class on the HTML element that contains text that you want to designate as the Q&A Genius Result answer snippet for the Knowledge article. **Note:** Answer snippets should be from 250 to 300 characters in length. Longer answer snippets may be cut off when displayed on the Q&A Genius Results card.

Machine reading comprehension model

AI Search uses a machine reading comprehension (MRC) model to validate extracted answer snippets. The MRC model analyzes each extracted snippet to determine whether it answers your

question. An answer snippet must pass this validation filter step before it can be displayed as a Q&A Genius Result answer.

When you enable factoid extraction for Q&A Genius Results, the MRC model also finds the exact span of text within the extracted snippet that represents the answer to your question.

You can't configure settings for MRC models or choose which model to use for Q&A Genius Results. AI Search performs A/B testing with different search query parameter values to automatically select the best MRC model for you. For more details on this A/B testing process, see Search query parameter evaluation framework.

Knowledge Management and Q&A Genius Results

When you activate the Knowledge Management Advanced plugin (com.snc.knowledge_advanced), Q&A Genius Results can return answers from records on the FAQ [kb_template_faq] table as well as records on the Knowledge [kb_knowledge] table and its child tables.

Additionally, if a search query satisfies one of the following conditions, AI Search bypasses the usual Q&A routines and returns the FAQ record's **Answer** field value as a Q&A Genius Result answer:

- The user's search query exactly matches the **Question** field value from the FAQ record.
- A result improvement rule promotes the FAQ record to be the first result for the user's search

As an example, suppose you create an FAQ record with **Question** `How do I reset my password?`. When a user searches for `how do I reset my password`, AI Search returns a Q&A Genius Result card with an answer snippet extracted from the FAQ record's **Answer** field value.

Interaction with other search features for Q&A Genius Results

The following table describes interactions between Q&A Genius Results and other search features.

Feature	Interaction with Q&A Genius Results
Result improvement rules with block actions	When computing Q&A Genius Results for a search query, AI Search ignores blocked search results.
Stop words	When computing Q&A Genius Results for a search query, AI Search ignores stop words in the query.
Synonyms	When computing Q&A Genius Results for a search query, AI Search expands synonyms in the query.
Typo handling	When computing Q&A Genius Results for a search query, AI Search corrects terms in the query.

Factoid extraction for Q&A Genius Results

Factoid extraction uses the machine reading comprehension (MRC) model to find the exact span of text within a longer extracted snippet that represents the answer to your question.

Enabling factoid extraction

To enable factoid extraction for Q&A Genius Results, set the *`glide.ais.genius_result.qna_mode`* system property to **sentence** or **snippet**.

For details on setting this system property, see the Set the factoid extraction mode for Q&A Genius Results section. To learn about the effects of this system property's values, see the System properties for factoid extraction in Q&A Genius Results section.

Examples of factoid extraction

The following images illustrate how the *`glide.ais.genius_result.qna_mode`* system property setting affects the display of Q&A Genius Result answer cards for a sample `What is the Windows key?` search in Service Portal.

- With the system property set to *`none`*, the Q&A Genius Result answer card displays the full extracted text snippet with no highlighting:

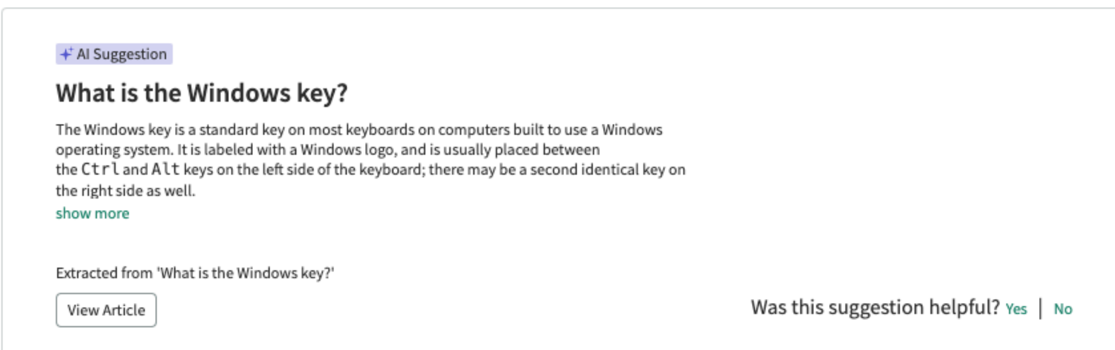

- With the system property set to *`sentence`*, the Q&A Genius Result answer card displays only the sentence containing the factoid, and highlights the factoid:

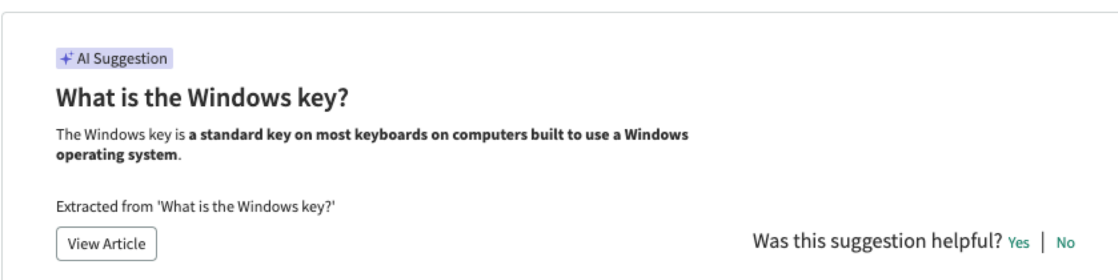

- With the system property set to *snippet*, the Q&A Genius Result answer card displays the full extracted text snippet and highlights the factoid:

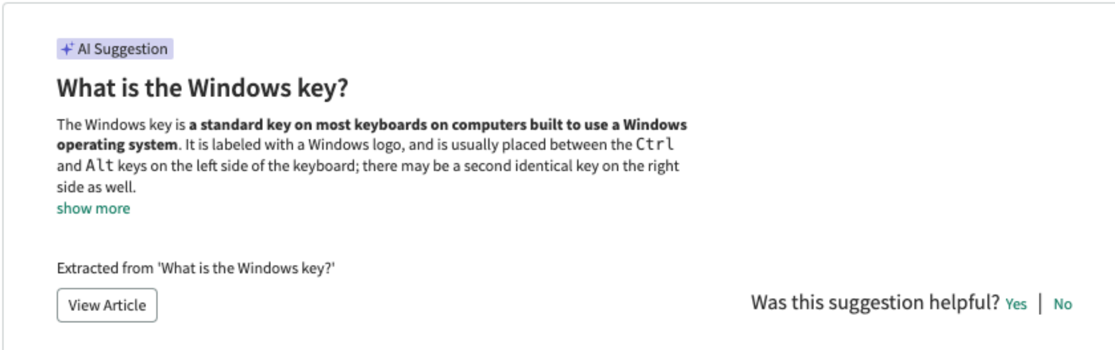

Set the factoid extraction mode for Q&A Genius Results

Specify how you want factoid extraction to display and highlight answers for Q&A Genius Results.

Before you begin

Role required: ais_admin

About this task

As a search administrator, you can specify the factoid extraction mode for Q&A Genius Results by setting the value of the *glide.ais.genius_result.qna_mode* system property. This system property does not exist in the base system, so you may have to add it.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Search for a system property record with name *glide.ais.genius_result.qna_mode*.
   - If a system property record with this name exists, open the record by selecting it.
   - If no system property record with this name exists, select **New**, then enter `glide.ais.genius_result.qna_mode` in the System Property form's **Name**field.
3. In the System Property form's **Value** field, enter one of the following values:
   - *none* if you want to disable factoid extraction and highlighting for Q&A Genius Results.
   - *sentence* if you want Q&A Genius Result answer cards to display only the sentence containing the factoid (and to highlight the factoid).
   - *snippet* if you want Q&A Genius Result answer cards to display the full extracted text snippet containing the factoid (and to highlight the factoid).

   For details on these system property values, see the System properties for factoid extraction in Q&A Genius Results section.
4. Save the modified System Property record by selecting **Submit** or **Update**.

System properties for factoid extraction in Q&A Genius Results

A system property determines how the system extracts factoids from Knowledge articles when computing Q&A Genius Results.

Property	Description
glide.ais.genius_result.qna_mode	Specifies the factoid extraction mode that you want Q&A Genius Results to use. For more information on factoid extraction, see the Factoid extraction for Q&A Genius Results section. • Type: string • Default value: *none* • Location: System Property [sys_properties] table • Supported values: ◦ *none*: Q&A Genius Result answer cards display the full extracted snippet from the source HTML article without highlighting. ◦ *sentence*: Q&A Genius Result answer cards display the sentence from the source HTML article that contains the factoid, with the factoid highlighted if possible. ◦ *snippet*: Q&A Genius Result answer cards display the full extracted snippet from the source HTML article, with the factoid highlighted if possible.

Link a Genius Result configuration to a search profile

Link a Genius Result configuration to one or more search profiles.

Before you begin

Role required: ais_admin

About this task

To use a Genius Result configuration, link it to a search profile. AI Search evaluates the Genius Result configuration for all searches made with the search profile.

You can link a single Genius Result configuration to multiple search profiles. A search profile can include multiple Genius Result configuration links.

AI Search evaluates active Genius Result configurations for a search profile in an order determined by their numeric **Order** field values, as defined in the search profile's Genius Results related list. Genius Result configurations with lower **Order** field values are evaluated before those with higher field values. As an example, a Genius Result configuration with **Order** 100 is evaluated before one with **Order** 200.

Procedure

1. Navigate to **All** > *AI Search* > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the Genius Result configuration to.
3. In the Genius Results related list, select **Link Existing**.
4. On the Search Profile - Genius Result Mapping form, fill in the fields.
   For a description of the field values, see Search Profile - Genius Result Mapping form.
5. Select **Submit**.

Result

The new Genius Result configuration link appears in the search profile's Genius Results related list.

Set the evaluation order for Genius Result configurations in a search profile

Change the order in which AI Search evaluates Genius Result configurations linked to a search profile.

Before you begin

Role required: ais_admin

About this task

AI Search evaluates active Genius Result configurations for a search profile in an order determined by their numeric **Order** field values, as defined in the search profile's Genius Results related list. Genius Result configurations with lower **Order** field values are evaluated before those with higher field values. As an example, a Genius Result configuration with **Order** 100 is evaluated before one with **Order** 200.

> **Note:**
> **Order** field values for linked Genius Result configurations are set separately for each search profile. A Genius Result configuration that's linked to multiple search profiles can have a different **Order** field value in each profile. For instance, you might set an **Order** value of 100 for the Q&A Genius Result configuration in one search profile but set a value of 750 for the same Genius Result configuration in a second search profile that it's linked to.

A search application configuration's **Genius Results Limit** field value specifies the maximum number of Genius Result cards the application can display. For more details on search application configuration settings, see Search Application Configuration form.

Procedure

1. Navigate to **All** > *AI Search* > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to set the Genius Result configuration evaluation order for.
3. In the Genius Results related list, edit the **Order** field values for Genius Result configurations linked to the search profile.
   For details on editing field values in the related list, see Use the list editor.

What to do next

To make changes in the **Order** field values take effect, publish the search profile. For more information on publishing a search profile, see Publish a search profile.

Create a Genius Result configuration

Create a Genius Result configuration, set its trigger condition, and configure the results displayed on its answer card.

Before you begin

Role required: ais_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to link the new Genius Result configuration to.
3. In the Genius Results related list, select **Create and Link**.
4. On the AI Search Genius Result Configuration form, fill in the fields.
   For a description of the field values, see AI Search Genius Result Configuration form.
5. Select **Submit**.
   The new Genius Result configuration appears in the Genius Results related list.
6. **Optional:** Open the **Order** for the new Genius Result configuration to control its display order for the search profile.

What to do next

Changes to Genius Result configurations take effect in search as soon as you save them. You don't need to publish the linked search profile to make the new configuration take effect.

If you selected an *NLU* trigger condition, you must link an NLU model and its intents to your Genius Result configuration. For instructions, see Link an NLU model and intents to a Genius Result configuration.

Link an NLU model and intents to a Genius Result configuration

Specify the NLU model and intents to use for a Genius Result configuration with an NLU trigger condition.

Before you begin

Role required: ais_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Query Settings** > **Genius Results**.
2. Open the Genius Result configuration that you want to link an NLU model to.
3. In the NLU Model related list, select **Link Existing**.
4. On the AI Search Genius Result Configuration NLU Model Mapping form, fill in the fields.
5. Select **Submit**.

Result

The NLU model and its selected intents appear in the NLU Models related list.

What to do next

For an NLU trigger condition to function, the NLU model must be published. You can verify NLU model status by navigating to **All** > `AI Search` > **Search Query Settings** > **Search Intents**.

Result improvement rules

Define rules with configurable trigger conditions to boost, block, or promote search results for specific searches. You can also boost search results for documents matching elements of the user context, such as a user's country, OS, or device type.

Result improvement rules with boost actions override the default search result order produced by the AI Search machine learning relevancy feature. Rules with block or promote actions override the default result set for a search query.

Rule structure

A result improvement rule consists of these parts:

- A set of conditions that determine when the rule activates.
- A set of boost, block, or promote actions that determine how the rule modifies the search query's results when activated. The rule can only include actions of a single type, but can include multiple actions of that type.

Rule conditions include the following:

- A date range bounded by start and end dates. The rule can only activate during the specified date range.
- A language. The rule can only activate if this language matches the search query's language.
- Optionally, a set of one or more trigger conditions. Triggers can operate on individual search query keywords, the full search query, or elements of the search user's context such as their department, language, or browser version. The rule can only activate if all of its trigger conditions are met.

Rule actions include the following types:

- `Boost`: Increase the relevancy scores for search result records that include specific field values. You can use either of the following options:
  - Boost search result records that include field values matching a dynamically populated value from the user context.
  - Boost search result records that include field values matching a static value that you specify.
- `Block`: Exclude specific records from the search query's results.
- `Promote`: Display specific records at the start of the search query's results.

Each result improvement rule can only perform actions of a single type when activated. You can define multiple actions of the same type for a rule.

Block actions take precedence over promote actions. If two rules block and promote the same record, AI Search excludes the record from the search query's results.

Language dependence

Result improvement rule activation is language-dependent. You specify a language for each result improvement rule.

AI Search supports result improvement rules for all languages activated in your instance. For the list of languages you can activate, see Activate a language.

Interaction with other search features

The following table describes interactions between result improvement rules with Query triggers and other search features.

Feature	Interaction with result improvement rules
Lemma and Unicode normalization	A search query term added by lemma or Unicode normalization can trigger a result

Feature	Interaction with result improvement rules
	improvement rule if it matches the rule's Query trigger.
Stop words	If a search query term that matches a result improvement rule's Query trigger is defined as a stop word, AI Search evaluates the rule trigger before removing the stop word from the search query.
Synonyms	By default, a synonym expansion term can activate a result improvement rule if it matches the rule's Query trigger. You can prevent synonym expansion terms from activating a rule with a Query trigger by clearing the **Expand Synonyms For Query Conditions** option that appears when you define the Query trigger.
Typo handling	An auto-corrected search query term can trigger a result improvement rule if it matches the rule's Query trigger.

Create a result improvement rule

Create a rule to boost, block, or promote search results or to block Genius Results. Specify the conditions that activate your rule and choose the action it performs when activated.

Before you begin

Role required: ais_admin

About this task

You can configure activation conditions and triggers for a result improvement rule. Additional topics cover steps for configuring boost, block, or promote actions for your rule.

Note:
Each result improvement rule can only perform actions of a single type when activated. You can define multiple actions of the same type for a rule.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the search profile that you want to add the result improvement rule to.
3. In the Result Improvement Rules related list, select **Create**.
4. On the AI Search Results Improvement Rule form, fill in the fields.
   For a description of the field values, see AI Search Results Improvement Rule form.
5. Select **Submit**.
   The Actions related list appears.
6. Choose the action that you want the result improvement rule to perform when activated, and see the appropriate topic for configuration steps:

- To increase the relevancy scores of search result records that include specific field values, see Boost search results using a result improvement rule.
- To prevent specific records from appearing in the search results, see Block search results and Genius Results using a result improvement rule.
- To make specific records appear first in the search results, see Promote search results using a result improvement rule.

What to do next

To make new or updated result improvement rule settings take effect, publish the search profile that you edited. For details on publishing a search profile, see Publish a search profile.

Boost search results using a result improvement rule

Increase the relevancy scores for search results that satisfy a custom set of conditions.

Before you begin

Create your result improvement rule first to define your trigger conditions. For steps to create a result improvement rule, see Create a result improvement rule.

Role required: ais_admin

About this task

Boost actions define relevancy changes applied by result improvement rules. When a result improvement rule with a boost action activates, AI Search modifies the relevancy scores for search results from an indexed source specified in the action. You can choose one of these boost types:

- Boost all search results from the specified indexed source.
- Boost only search results from the specified indexed source with a field value that matches a dynamically populated field value from the searcher's user context.

  [store-future: BEGIN review]**Note:**
  For the list of available user context fields, see User context fields for boost result improvement rules.
  [End]

- Boost only search results from the specified indexed source with a field value that matches a static string value specified in the action.

You can specify a boost weight for each boost action. This is an integer value that AI Search uses to alter the relevancy scores for search results affected by your boost action. Boost weights can be positive or negative.

Positive boost weights increase the relevancy scores for boosted search results. Each boost weight of 1 increases a boosted search result's relevancy score by 0.1% of its original value. For example, if you set a boost weight of 100, you increase each boosted search result's relevancy by +10%, so a search result with an original relevancy score of 60 gets a final relevancy score of 66. Setting a boost weight of 1000 adds +100%, doubling the relevancy scores for boosted search results.

Negative boost weights reduce the relevancy scores for boosted search results. Each boost weight of –1 reduces a boosted result's relevancy score by 0.1% of its original value. As an example, if you set a boost weight of –150, you reduce a boosted search result's relevancy by 15%, so a search result with an original relevancy score of 100 gets a final relevancy score of 85. Any boost weight of –1000 or lower reduces a boosted search result's relevancy score to 0, causing it to appear at the end of the search result list.

When a single search result is affected by multiple boost actions, AI Search adds the individual boost weights together to determine the total relevancy score modifier. For example, a search result boosted by three rules with boost weights 100, 150, and −25 has a total boost weight of 225, increasing its final relevancy score by 22.5% of its original value.

Procedure

1. In the Actions related list, select **Create Boost Action**.
2. On the Rule - Action Mapping form, set the boost weight and fill in the other fields.
3. Select **Submit**.
4. **Optional:** Edit the **Order** values for the result improvement rule's boost actions.

What to do next

To make new or updated result improvement rule settings take effect, publish the search profile you edited. For details on publishing a search profile, see Publish a search profile.

List of result improvement rule boost types

This list describes the types of search result relevancy boost that you can apply using a result improvement rule.

Boost action types

Type	Description
Boost All Documents In Indexed Source	Applies relevancy boost to every search result from the selected indexed source. As an example, you might boost all results from the Knowledge Table indexed source.
Boost By User Context (dynamic)	Applies relevancy boost to search results from the selected indexed source if they include a value for the selected field that matches a field value specified in the dynamically populated user context. As an example, you might boost all results from the Knowledge Table indexed source with Short description field values that contain matches for the city specified in the current search user's location, as set in their user context. This configuration could end up boosting different records when you search from Paris than when you search from New York City. [store-future: BEGIN review]For the list of available user context fields, see User context fields for boost result improvement rules.[End]
Boost By Field Match (static)	Applies relevancy boost to search results from the selected indexed source if they include

Boost action types (continued)

Type	Description
	a value for the selected field that matches a specified static value. As an example, you might boost all results from the Knowledge Table indexed source that have Short description field values matching the static value `Dealing with Spyware and Viruses`.

Block search results and Genius Results using a result improvement rule

Configure a result improvement rule to prevent specific documents from appearing as Genius Results or search results.

Before you begin

Create your result improvement rule first to define your trigger conditions. For steps to create a result improvement rule, see Create a result improvement rule.

Role required: ais_admin

About this task

Block actions hide selected search results from the result set. When a result improvement rule with a block action activates, its effect depends on the result improvement rule's **Apply only to Genius Results** option setting:

- If this option is selected, the block action prevents selected documents from appearing as Genius Result answers. Selected documents can still appear as search results.
- If this option is cleared, the block action prevents the selected documents from appearing as Genius Result answers or search results.

Procedure

1. In the Actions related list, select **Block Documents**.
2. In the Select Documents To Block dialog box, search for results that you want to block.
3. To add a search result to the selected results pane, select the plus sign next to it.
4. After selecting all the results that you want to block, select **Add Documents and Close**.
5. **Optional:** Edit the **Order** values for the result improvement rule's block actions.

What to do next

To make new or updated result improvement rule settings take effect, publish the search profile that you edited. For details on publishing a search profile, see Publish a search profile.

Promote search results using a result improvement rule

Configure a result improvement rule to make specific records appear first in results for a search.

Before you begin

Create your result improvement rule first to define your trigger conditions. For steps to create a result improvement rule, see Create a result improvement rule.

Role required: ais_admin

About this task

Promote actions define search results to spotlight in the result set. When a result improvement rule with a promote action activates, it displays search results specified in the action ahead of other results for the search query.

Procedure

1. In the Actions related list, select **Promote Documents**.
2. In the Select Documents To Promote dialog box, search for the results that you want to promote.
3. To add a search result to the selected results pane, select the plus sign next to it.
4. After selecting all the results that you want to promote, select **Add Documents and Close**.
5. **Optional:** Edit the **Order** values for the result improvement rule's promote actions.
   The **Order** field only accepts integer values. AI Search evaluates promote actions with lower **Order** values before promote actions with higher values. As an example, a promote action with **Order** 150 is evaluated before one with **Order** 200.

What to do next

To make new or updated result improvement rule settings take effect, publish the search profile that you edited. For details on publishing a search profile, see Publish a search profile.

Share result improvement rules between search profiles

Make your search experiences more consistent by configuring a search profile to mirror result improvement rules from other search profiles.

Before you begin

You must have Now Assist in AI Search installed.

Role required: ais_admin

About this task

Result improvement rules are specific to the search profile they're created in, and can't be copied between search profiles. As a search administrator, however, you can configure a search profile to use the set of result improvement rules defined in another search profile. This approach lets you define your result improvement rules in one place while making them available in multiple different search experiences.

This procedure is called mirroring, since your affected search profile reflects result improvement rules defined in another search profile. The search profile that includes your result improvement rules is called the mirror source, and the search profile that uses those rules is called the mirror target.

> **Note:**
> You can link multiple mirror source search profiles to a single mirror target search profile. Similarly, you can link a single mirror source search profile to multiple mirror target search profiles.

When you modify result improvement rules in a mirror source search profile and then publish that search profile, AI Search checks all mirror target search profiles that link to that mirror source.

- If the mirror target search profile is published, AI Search republishes it so that it reflects your result improvement rule changes.
- If the mirror target search profile isn't published, an informational message reminds you to publish it so that it can reflect your result improvement rule changes.

This check keeps result improvement rule changes in sync across your search profiles.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the record for the mirror target search profile.
   As an example, if you've defined result improvement rules in the Service Portal Default Search Profile and now want to mirror those rules in the ESC Portal Default Search Profile, open the ESC Portal Default Search Profile record.
3. To add another search profile as a mirror source for result improvement rules, create a link to it.
   a. In the Mirrored Result Improvement Rules related list, select **Link Search Profile**.
   b. On the Mirrored Result Improvement Rule form, in the **Mirror Source** field, enter the name of the search profile that you want to mirror result improvement rules from, or use the Lookup using list ( ) icon to populate the field.
   c. Select **Submit**.
      The new mirror source search profile appears in the Mirrored Result Improvement Rules related list.
   d. **Optional:** Repeat the linking procedure for each additional search profile you want to add as a mirror source.
4. To remove one or more search profiles as mirror sources for result improvement rules, unlink them:
   a. In the Mirrored Result Improvement Rules related list, select the row checkbox for each mirror source that you want to remove.
   b. Select **Unlink Selected**.
      The selected mirror source search profiles are removed from the Mirrored Result Improvement Rules related list.
5. To make your mirror source changes take effect, publish your modified mirror target search profile.
   For details on publishing a search profile, see Publish a search profile.

Search application configurations

A search application configuration specifies the search engine and settings to use for search in a ServiceNow AI Platform® application. When you choose AI Search as an application's search engine, you can select a search profile for the application and configure its search result display and refinement settings.

A search application configuration contains a reference to a search profile, along with settings that affect how search results are displayed to users of the application. The search profile specifies settings that determine how AI Search generates search results for a given search. To learn more about search profile configuration settings, see Search profiles.

Some ServiceNow AI Platform applications install their own AI Search application configurations when you activate them. For example, activating the Customer Service Management (CSM) application installs AI Search application configurations for the Customer Service Portal and Consumer Service Portal.

AI Search includes search application configurations for use in the following ServiceNow AI Platform applications:

- Employee Center
- Virtual Agent
- Service Portal

You can modify existing search application configurations, or create and configure your own search application configurations with AI Search specified as the search engine.

For details on enabling AI Search and specifying search application configurations in supported applications, see Enabling and configuring AI Search in ServiceNow AI Platform applications.

Note:
When you view or edit search application configuration records, the system may display a warning message telling you to reindex one or more indexed sources. This reindexing process (a one-time process for each indexed source listed) is required to correctly populate search-based auto-complete suggestions. For more details on auto-complete suggestions, see Auto-complete suggestions in AI Search applications.

Create a search application configuration for AI Search

Create a search application configuration to specify AI Search as the search engine for a ServiceNow AI Platform® application. Select a search profile for the application to use. Configure the application's settings for refinement and display of search results.

Before you begin

Role required: search_application_admin

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Select **New**.
3. On the Search Application Configuration form, fill in the fields.
   For a description of the field values, see Search Application Configuration form.
4. Select **Submit**.

Result

The new search application configuration record appears in the Search Application Configurations list.

What to do next

To use your new search application configuration, select it in the configuration for a ServiceNow AI Platform application that supports AI Search. For more information on these ServiceNow AI Platform applications, see Enabling and configuring AI Search in ServiceNow AI Platform applications.

Display search result counts on the results page for a search application

Display search result counts on the results page for AI Search applications that use Seismic components. Counts show the total number of matching search results and the number of matching results for each navigation tab.

Before you begin

Role required: search_application_admin

About this task

Search administrators can enable display of search result counts in AI Search applications that use Seismic components. With this option enabled, the results page for the search application displays matching search result counts in two locations:

- The total number of results that matched the user's search
- The number of matching search results for each search source displayed as a navigation tab

Note:
Navigation tabs with 1,000 or more matching search results display **1000+** as their result count.

Sample results page with results counts

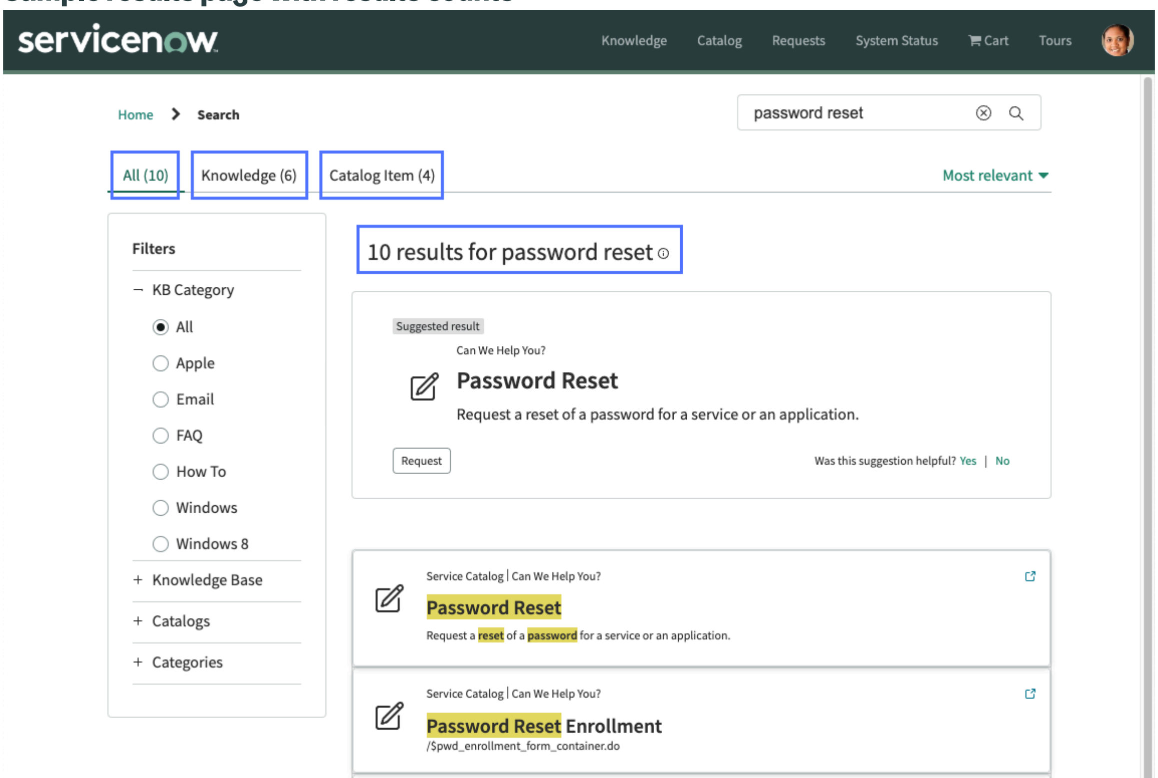

Note:
Result counts indicate the number of results that matched the user's search. Late binding security can remove results after they're counted, causing the actual number of displayed results to be lower than the count. For more information on late binding security, see Content security in AI Search.

Result counts are supported in search applications that use Seismic components. You can display search result counts in the following search applications:

- Service Portal
- Employee Service Center (ESC)
- AI Search for Next Experience global and workspace search
- Custom UI Builder applications that use Seismic UI components

Other search applications, such as Virtual Agent and Now Mobile, don't use Seismic components and so don't support display of search result counts. For best results, don't enable search result counts for these search applications.

Search result counts are enabled by default for new search application configuration records. You only need to perform this procedure to enable search result counts for existing search application configuration records after upgrading from a previous release.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Select the record for the AI Search application that you want to display result counts for.
3. Select the **Show Search Results Count** option.
4. Select **Update**.

Result

AI Search displays the total result count and navigation tab result counts on the results page for the selected search application.

Group attachment search results with their parent results

Configure your AI Search applications to display attachment search results grouped with the results for their parent records.

Before you begin

The index_attachments attribute can't be set to false for any indexed source used in the AI Search application. For details on creating and editing attributes for indexed sources, see Indexed source attributes.

> **Note:**
> If you remove a false `index_attachments` attribute value from an indexed source, perform a full reindex for the indexed source before continuing with this procedure. To learn how to reindex content for an indexed source, see Perform a full table index or reindex for a single indexed source.

Role required: search_application_admin

About this task

Search administrators can enable grouping of attachment search results with the results for their parent records. With this option enabled, search results for attachments that match a user's search display inline with the search results for their parent record.

If the parent record itself or any of its attachments matches the user's search, the parent record appears in the search result list. Grouped attachments that match the search appear in descending relevancy order within the search result for their parent record.

When attachment grouping is in use, AI Search computes relevancy scores for the parent record and each of its grouped attachments. The final relevancy score of the grouped search result is the highest of these individual relevancy scores.

Sample results page showing attachment results grouped with parent result

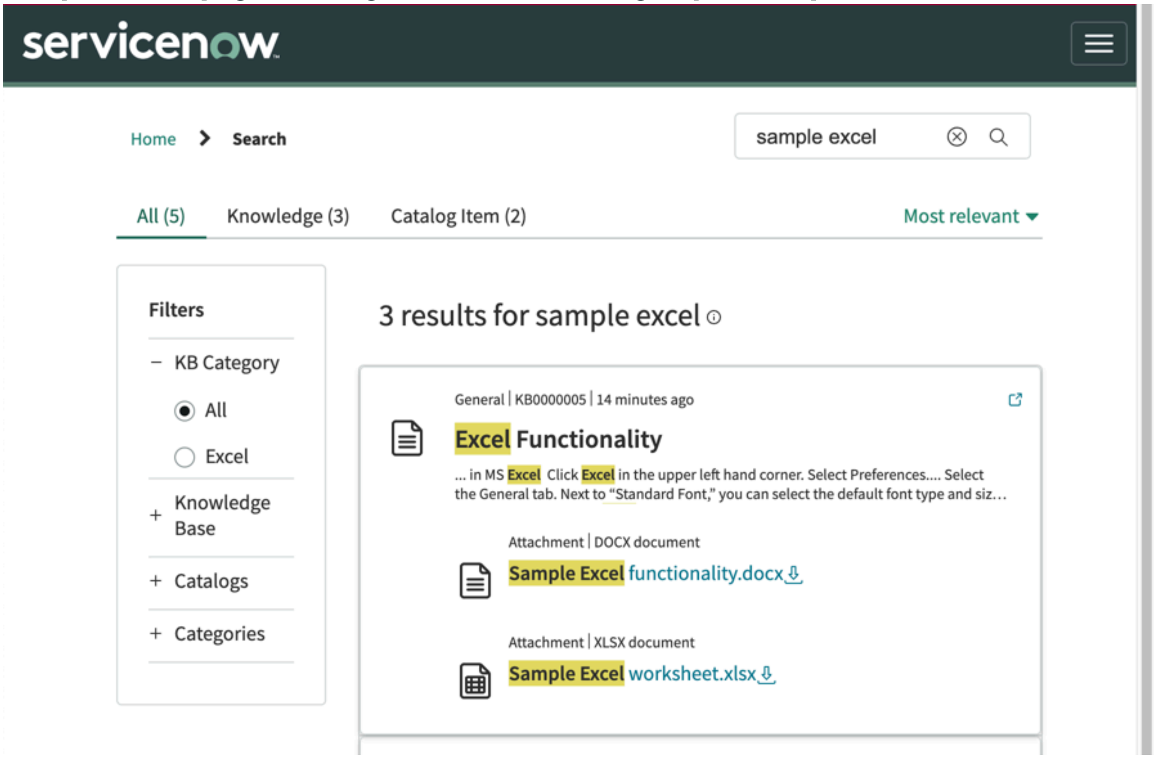

By default, AI Search displays separate search results for the parent record and each of its attachments that matches the user's search. With attachment grouping enabled, you can view all matching attachments for a record in a single search result, instead of needing to collate information from multiple search results.

Note:
When you select a grouped attachment search result, AI Search logs click rank signal data for the parent record search result. All other search signal data is specific to the selected attachment record. For more details on search signal logging, see Search signal tables.

Attachment grouping is enabled by default for new search application configuration records. You only need to perform this procedure to enable attachment grouping for existing search application configuration records after upgrading from a previous release.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Select the record for the AI Search application that you want to enable attachment grouping for.
3. Select the **Show parent for results that are attachments** option.
4. Select **Update**.
   If any indexed source used in the search application doesn't have the index_attachments attribute set to true, the system displays a warning message. For each affected indexed source, set the index_attachments attribute to true, then perform a full reindex.

Result

AI Search groups matching attachment search results with their parent record results on the search application results page.

Create a facet in an AI Search application configuration

Define facets to refine your searches using categories dynamically generated from the current search results.

Before you begin

Role required: search_application_admin

About this task

Create a facet to populate a refinement filter for search results based on indexed source field values. The facet initially displays filter options for all field values found in the current search result set. Selecting a field value refines the search result set, displaying only those results that satisfy the filter logic for the facet.

> **Note:**
> If AI Search applies late binding security to any results from an indexed source used by a facet, the list of available filter selections for the facet changes. When this happens, the facet only displays field values found in the current page of search results. For more information on late binding security, see Content security in AI Search.

Facets can have any one of the following types. Each facet type has its own filter logic.

Single Select

In a Single Select facet, you can only select a single field value as a facet refinement filter.

When you select a field value from a Single Select, AI Search filters the search results list, displaying only results that contain your chosen field value.

A Single Select facet only allows you to select field values that are found in the current result set.

Multi Select And

In a Multi Select And facet, you can select multiple field values as facet refinement filters.

When you select a field value from a Multi Select And facet, AI Search filters the search results list, displaying only results that contain all of the field values you've selected from the facet.

A Multi Select And facet only allows you to select field values that are found in the current result set.

Multi Select Or

In a Multi Select Or facet, you can select multiple field values as facet refinement filters.

When you select a field value from a Multi Select Or facet, AI Search filters the search results list, displaying only results that contain any of the field values you've selected from the facet.

A Multi Select Or facet initially only allows you to select field values that are found in the current result set. All field values that are available when you make your first selection remain available for selection.

As an example, suppose you have a Multi Select Or facet for a custom **Color** field with blue, green, and red field values. If your search returns only green and red results, the **Color** facet only displays **green** and **red** field values. When you select

the **red** field value, AI Search keeps the **green** field value available for selection because it was available before you made any **Color** selection.

> **Warning:**
> Applying multiple filters from a Multi Select And or Multi Select Or facet makes your search query URL longer. Test your multi-select facet configurations to ensure that search query URLs with facet filters don't exceed the maximum URL length limit for your web browser. (Different web browsers have different URL length limits.)

Facets work best when created from fields of the following types with relatively few unique values:

- Choice
- Reference
- String
- String (Full UTF-8)
- True/False

When sorting facet refinement filters by name, AI Search uses collation-based sort following the rules for the ServiceNow AI Platform® session language.

For tables other than Knowledge [kb_knowledge], facets display appropriate field translations for the current user session. Facets don't display translated versions of kb_knowledge translated fields.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Open the search application configuration that you want to configure a facet for.
3. In the Facets related list, select **New**.
4. On the Facet form, fill in the fields.
   For a description of the field values, see Facet form.
5. Select **Submit**.

Result

The new facet appears in the Facets related list.

Clear facet filters when selecting a new navigation tab

Specify whether AI Search applications should preserve existing facet filters when you select a new navigation tab.

Before you begin

Role required: admin

About this task

By default, AI Search preserves your facet filters when you select a new navigation tab. As an example, suppose that you narrow your search results by selecting the Email field value from the **KB Category** facet displayed on the **All** navigation tab. If you then select the **Knowledge** navigation tab, the **KB Category** facet retains your Email selection and still filters your search results accordingly.

Administrators can override this default behavior by setting the value of the *`glide.ais.facets.cross_tab_persistence`* system property. When this system property value is set to false, selecting a new navigation tab removes all facet filters. In the

preceding example, selecting the **Knowledge** navigation tab would remove your **KB Category** facet filter.

This option is global and applies to all AI Search applications on your instance.

Note:
If an existing facet filter cannot be applied to search results from a selected navigation tab, AI Search ignores the facet filter. This behavior only occurs in search applications where you've disabled search result counts. For more details on search result counts, see Display search result counts on the results page for a search application.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Open the record for the *`glide.ais.facets.cross_tab_persistence`* system property.
3. Enter the system property value according to the behavior that you want:
   - To preserve existing facet filters when a new navigation tab is selected, set the system property value to true.
   - To remove all facet filters when a new navigation tab is selected, set the system property value to false.

   The default property value is true.
4. Select **Update**.

Result

All AI Search applications preserve or discard existing facet filters when a new navigation tab is selected.

Configure navigation tabs in an AI Search application configuration

Define settings for navigation tabs in a search application configuration. Users of the application can select navigation tabs to refine their search query results by search source.

Before you begin

Role required: search_application_admin

About this task

AI Search automatically defines navigation tabs for search application configurations. Each tab is tied to a search source linked to the search application configuration's search profile. You can't create or delete these navigation tabs.

Use this procedure to configure display settings for navigation tabs in your search application configuration.

Procedure

1. Navigate to **All** > *`AI Search`* > **Search Experience** > **Search Applications**.
2. Open the search application configuration that you want to configure navigation tabs for.
3. In the Navigation Tabs related list, open the navigation tab that you want to configure.
4. On the Navigation Tab form, fill in the fields.

For a description of the field values, see Navigation Tab form.

5. Select **Update**.

Result

The navigation tab appears with updated settings in the Navigation Tabs related list.

Auto-complete suggestions in AI Search applications

Automatically display intelligent suggestions in your search application's input field, helping users formulate their searches or jump directly to relevant results. Suggestions can include relevant search results, popular queries, recent queries from the user's personal history, and popular search results.

When you create a search application configuration, AI Search automatically populates it with auto-complete suggestion types. The settings for these suggestion types control the appearance of auto-complete suggestions when users interact with its input field.

Each auto-complete suggestion type includes default settings that you can configure to suit your application's needs. It also includes one or more search suggestion readers, which generate search query and search result suggestions.

You can remove an unwanted suggestion type from your search application configuration. If you want to restore a removed suggestion type, you can relink it to your search application configuration. You can only link one auto-complete suggestion source of each type to any given search application configuration.

Auto-complete suggestions are generated by the Search Suggestions application. For full details on configuration and use of the Search Suggestions application, see Search Suggestions.

AI Search uses a dedicated relevancy model to rank records for display as auto-complete suggestions in the search field. This relevancy model scores records based on their freshness and on search query term matches in their **title** fields. The system doesn't train this auto-complete suggestion relevancy model. For more details on relevancy models, see Machine learning relevancy in AI Search.

List of auto-complete suggestion types for AI Search applications

AI Search provides multiple types of auto-complete suggestions that you can configure for display to users in the search field.

The following auto-complete suggestion types are available for linking to search application configurations that specify AI Search as their search engine.

Note:
Not all auto-complete suggestion types are available for the search application configurations included in the base system.

Auto-complete suggestion type	Description
Popular Search Reader Group for <application name>	Displays search query suggestions based on searches submitted by all users of the search application. When you create a search application configuration, this suggestion reader group isn't linked to it by default.

Auto-complete suggestion type	Description
	Default section header: `Popular searches` Readers: Popular Query Reader
Recent Search Reader Group for <application name>	Displays search query suggestions based on the user's previous searches. Default section header: `Recent searches` Included readers: User History Reader
Search-based Suggested Results Reader Group for <application name>	Displays titles and tables for search result suggestions based on a search for the terms entered in the search field. This search applies search source filters, stop word and synonym dictionaries, and result improvement rules from the search application's linked search profile. AI Search treats the last term in the search field specially, as follows: • If the last term in the search field is followed by a space, AI Search treats it as a literal search term. As an example, `act` followed by a space matches results containing the term `act`. • If the last term in the search field isn't followed by a space, AI Search treats it as a wildcard search prefix. For example, `act` without following space matches results containing the terms `act`, `action`, `activated`, `actuated`, and so on. Selecting a suggested search result opens it as follows: • Catalog Item results open in Service Catalog. • Knowledge article results open in Knowledge. • Other results open in form view. **Note:** If the suggested search result is an attachment, selecting it opens its parent record in form view. Default section header: `Suggested` Included readers: Search Based Suggestion Reader

Auto-complete suggestion type	Description
Suggested Query Reader Group for <application name>	Displays a mix of search query suggestions (based on the user's previous searches) and search result suggestions (based on results the user selected for previous searches). Default section header: none Included readers: Popular Query Reader, User History Reader, User Search Term Reader
Suggested Result Reader Group for <application name>	Displays search result suggestions based on results the user selected for previous searches. When you create a search application configuration, this suggestion reader group isn't linked to it by default. Default section header: `Suggested results` Included readers: Personal Click Reader
User History Reader Group for <application name>	Displays the user's current search query, plus search query and result suggestions from the user's previous search activity that match the current input. When you create a search application configuration, this suggestion reader group isn't linked to it by default. Default section header: `Recently viewed` Included readers: User History and Click Reader, User Search Term Reader

Note:
A search application configuration can't link to multiple search suggestion types that include the same reader and have the same click-in mode. If you try to link a second suggestion type with the same reader and click-in mode to a search application configuration, the system displays an error message.

Configure an auto-complete suggestion type in an AI Search application configuration

Modify default settings for display of auto-complete suggestions of a selected type in your search application's input field.

Before you begin

Role required: search_application_admin

About this task

When you define a new search application configuration, AI Search automatically populates it with auto-complete suggestion types. You can alter the default settings for these auto-complete suggestion types.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Open the search application configuration that contains the auto-complete suggestion type that you want to modify.
3. In the Autocomplete Suggestions related list, open the auto-complete suggestion type that you want to modify settings for.
4. On the Suggestion Reader Group form, fill in the fields.
   For a description of the field values, see Suggestion Reader Group form.
5. Select **Update**.

Result

The updated settings for the auto-complete suggestion type appear in the search application configuration's Autocomplete Suggestions related list.

Remove an auto-complete suggestion type from an AI Search application configuration

Disable the display of auto-complete suggestions of a selected type in your search application's input field.

Before you begin

Role required: search_application_admin

About this task

When you define a search application configuration, AI Search automatically populates it with auto-complete suggestion types. You can remove unwanted auto-complete suggestion types from the search application configuration.

> **Note:**
> Removing an auto-complete suggestion type from your search application configuration doesn't delete it. To relink the removed suggestion type to your search application configuration, see Relink an auto-complete suggestion type to an AI Search application configuration.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Open the search application configuration that contains the auto-complete suggestion type you want to remove.
3. In the Autocomplete Suggestions related list, open the auto-complete suggestion type that you want to remove.
4. On the Suggestion Reader Group form, select **Remove from Search Application**.
   For a description of the field values, see Suggestion Reader Group form.

Result

The selected auto-complete suggestion type no longer appears in the search application configuration's Autocomplete Suggestions related list.

Relink an auto-complete suggestion type to an AI Search application configuration

Re-enable the display of auto-complete suggestions of a selected type in your search application's input field.

Before you begin

Role required: search_application_admin

About this task

When you define a search application configuration, AI Search automatically populates it with auto-complete suggestion types. If you removed one of these auto-complete suggestion types from the search application configuration, you can relink it.

Note:
A search application configuration can't link to multiple search suggestion types that include the same reader and have the same click-in mode. If you try to link a second suggestion type with the same reader and click-in mode to a search application configuration, the system displays an error message.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Open the search application configuration that you want to relink the auto-complete suggestion type to.
3. In the Autocomplete Suggestions related list, select **Add Existing**.
4. In the Suggestion Reader Groups list, open the auto-complete suggestion type that you want to relink to your search application configuration.
5. On the Suggestion Reader Group form, enter or look up the search application configuration by name in the **Search Application** field.
   For a description of the field values, see Suggestion Reader Group form.
6. Select **Update**.

Result

The selected auto-complete suggestion type appears in the search application configuration's Autocomplete Suggestions related list.

Using search scripted post-processors in AI Search application configurations

Use scripted post-processors to modify search result records for display in a search application.

For example, a search scripted post-processor can modify search result date or currency field values to display in the user's regional formats.

Create a search scripted post-processor

Define a script to post-process AI Search search results before displaying them to the user.

Before you begin

Role required: search_application_admin

Procedure

1. **Optional:** Copy code from an existing post-processor script to use as a starting point for your own script.
   a. Select **All**.
   b. In the Filter field, enter `sys_search_scripted_processor.list`.

c. Press Enter.

d. Open a post-processor script of your choice.

e. Copy the code from the **Post processing script** field.

2. Navigate to **All** > *AI Search* > **Search Experience** > **Search Applications**.
3. Open the search application configuration that you want to configure a scripted post-processor for.
4. In the Search Scripted Post-processors related list, select **New**.
5. On the Search Scripted Post-processors form, fill in the fields.
   For a description of the field values, see Search Scripted Post-processors form. If you copied code from an existing post-processor script in step 1.e, paste it into the new script's **Post processing script** field and modify it as needed.
6. Select **Submit**.

Result

The new search scripted post-processor appears in the search application configuration's Search Scripted Post-processors related list.

What to do next

To make your new search scripted post-processor take effect, link it to the search application configuration for the application whose results you want to post-process.

Link a search scripted post-processor to an AI Search application configuration

Link a search scripted post-processor to an AI Search search application configuration. The search application executes the linked script for each search result before displaying it to the user.

Before you begin

Role required: search_application_admin

You must have already created the search script post-processor that you want to link to the search application configuration.

Procedure

1. Navigate to **All** > *AI Search* > **Search Experience** > **Search Applications**.
2. Open the search application configuration that you want to link a search scripted post-processor to.
3. In the Search Scripted Post-processors related list, select **New**.
4. On the Search Scripted Post-processors form, select the reference lookup icon ( ) for the **Search scripted processor** field and select the search scripted post-processor that you want to link to the search application configuration.
   For a description of the field values, see Search Scripted Post-processors form.
5. Select **Submit**.

Result

The new search scripted post-processor appears in the Search Scripted Post-processors related list.

Search result sort options in AI Search application configurations

A search result sort option enables users to display search results in an order determined by their field values. Applying a search result sort option overrides the default AI Search

relevancy-based result order. You can define custom search result sort options for your AI Search applications.

When sorting search results on a string field, AI Search uses collation-based sort following the rules for the ServiceNow AI Platform® session language.

For details on how users apply search result sort options, see Change the sort order for your search results.

Create a search result sort option in an AI Search application configuration

Define a sort option to enable users of an AI Search application to order their search results using field values.

Before you begin

Role required: search_application_admin

About this task

A search result sort option enables users to sort their search results on a field, overriding the default AI Search relevancy-based result order. Each search result sort option specifies a field to sort results on, and a sort order (ascending or descending) order for the sort.

The sort menu in an AI Search application displays all active search result sort option defined for that application, plus the default **Most relevant** option.

> **Note:**
> If no search result sort options are active for an AI Search application, the sort menu doesn't display for users of that application and results are automatically sorted by relevance.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Applications**.
2. Open the search application configuration that you want to create a search result sort option for.
3. In the Sort Options related list, select **New**.
4. On the Sort Option form, fill in the fields.
   For a description of the field values, see Sort Option form.
5. Select **Submit**.

Result

The new search result sort option appears in the search application's Sort Options related list, and its title displays in the sort menu for users of the search application. Users can select the new option from the sort menu to apply the defined sort to their search results.

Remove a search result sort option from an AI Search application configuration

Remove a search result sort option from the sort menu in an AI Search application. The removal of the search result sort option prevents users of the application from selecting that option to reorder their search results with.

Before you begin

Role required: search_application_admin

servicenow

Procedure

1. Navigate to **All** > *AI Search* > **Search Experience** > **Search Applications**.
2. Open the search application configuration that you want to disable a search result sort option for.
3. In the Sort Options related list, open the search result sort option that you want to disable.
4. On the Sort Option form, clear the **Active** option and select **Update**.
   For a description of the field values, see Sort Option form.

Result

The selected search result sort option no longer appears in the search application's sort menu.

> **Note:**
> If no search result sort options are active for an AI Search application, the sort menu no longer displays for users of that application and results are automatically sorted by relevance.

Configure EVAM display settings for search results in AI Search applications

Customize display settings for Search Result EVAM (Entity View Action Mapping) cards in portal and non-portal AI Search applications.

Before you begin

Role required: evam_admin

About this task

EVAM administrators can modify settings on Search Result EVAM cards to customize the appearance and layout of search results in AI Search applications.

String fields on Search Result EVAM cards can be populated with literal text or with values from AI Search indexed fields.

Procedure

1. Open the EVAM Definition record associated with your chosen AI Search application.
   - For portal applications, navigate to **All** > **Service Portal** > **Portals**, select the search application's portal record from the list, then use the Preview this record icon (ⓘ) to preview and open the portal record's linked Search Results Configuration record.
   - For non-portal applications, navigate to **All** > *AI Search* > **Search Experience** > **Search Applications**, select the search application's configuration record from the list, then use the Preview this record icon (ⓘ) to preview and open the search application configuration record's linked EVAM Definition record.
2. In the EVAM View Config Bundle M2Ms related list, select the record for the EVAM definition's config bundle.
3. In the EVAM View Configs list, select the record for the specific Search Result category that you want to customize.
   The record name should include `Search Result`. Don't select a record with `Genius Result` in its name, as Genius Results have different display settings than standard search results do.
4. Use the Preview this record icon (ⓘ) to preview and open the EVAM View Config record's linked View Template record.
5. In the Template field, edit the JSON for the Search Result card properties you want to customize.

For details on available properties and their supported values, see List of Search Result EVAM card properties.

6. Select **Update**.

Result

Changes made to Search Result EVAM card settings take effect immediately in AI Search applications.

List of Search Result EVAM card properties

You can customize the appearance of standard search results in your AI Search application by changing these properties in the application's Search Result EVAM (Entity View Action Mapping) cards.

> **Note:**
> The following properties are only supported for EVAM cards with `component` property value `sn-search-result-evam-card`. Don't set these properties for EVAM cards with any other `component` value.

Property	Description
avatarURL	Defines the URL for the image you want to show on the Search Result card when **imageType** is set to **avatar**. • Type: string • Default value: (no value) **Note:** To specify literal text for this property, include the property in the `staticValues` object and specify the text as its `key` value. To populate this property with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields.
detailLabelOne, detailLabelTwo, detailLabelThree, detailLabelFour, detailLabelFive, detailLabelSix, detailLabelSeven, detailLabelEight, detailLabelNine, and detailLabelTen	Defines label content for the Search Result card's detail label-value pairs. You can show up to 10 label-value pairs on the Search Result card. • Type: string • Default value: (no value)

Property	Description
	Note: To specify literal text for one of these properties, include the property in the `staticValues` object and specify the text as its `key` value. To populate one of these properties with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields. To define value content for these label-value pairs, see **detailValueOne** through **detailValueTen**. **detailValueOne** sets the value for **detailLabelOne**, **detailValueTwo** sets the value for **detailLabelTwo**, and so on up through **detailValueTen** and **detailLabelTen** For additional configuration of how a search result's details are shown on the Search Results card, see **detailLabelType**, **detailLabelSeparator**, and **detailPosition**.
detailLabelSeparator	Defines the character used to separate inline label-value pairs. Only applies when **detailLabelType** is set to **inline**. • Type: string • Default value: (no value) **Note:** To specify literal text for this property, include the property in the `staticValues` object and specify the text as its `key` value. To populate this property with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields.
detailLabelType	Defines how to render label-value pairs on the Search Result card. • Type: choice list • Default value: inline • Supported values:

Property	Description
	◦ *`inline`*: Render label-value pairs using the Label-value component, with each label appearing on the same line as its value string. ◦ *`stacked`*: Render label-value pairs using the Label value stacked component, with each label appearing on a separate line from its value string.
detailPosition	Defines where search result detail label-value pairs are shown relative to the Search Result card's summary field. • Type: choice list • Default value: below • Supported values: ◦ *`above`*: Search result details are shown above the summary field. ◦ *`below`*: Search result details are shown below the summary field.
detailValueOne, detailValueTwo, detailValueThree, detailValueFour, detailValueFive, detailValueSix, detailValueSeven, detailValueEight, detailValueNine, and detailValueTen	Defines value content for the Search Result card's detail label-value pairs. You can show up to 10 label-value pairs on the Search Result card. • Type: string • Default value: (no value) **Note:** To specify literal text for one of these properties, include the property in the `staticValues` object and specify the text as its `key` value. To populate one of these properties with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields. To define label content for these label-value pairs, see **detailLabelOne** through **detailLabelTen**. **detailLabelOne** sets the label for **detailValueOne**, **detailLabelTwo** sets the label for **detailValueTwo**, and so on up through **detailLabelTen** and **detailValueTen** For additional configuration of how a search result's details are shown on the

Property	Description
	Search Results card, see **detailLabelType**, **detailLabelSeparator**, and **detailPosition**.
hideIcon	Defines whether you want to show or hide the icon on the Search Result card. • Type: true | false • Default value: false • Supported values: ◦ `false`: Show the configured icon on the Search Result card. ◦ `true`: Do not show an icon on the Search Result card.
icon	Defines the icon you want to show on the Search Result card when **imageType** is set to **icon**. This icon is also shown as a fallback when **imageType** is set to **image** but no **imageURL** value is provided. • Type: string • Default value: (no value) **Note:** To specify literal text for this property, include the property in the `staticValues` object and specify the text as its `key` value. To populate this property with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields. For the set of available icon names, see the Icon Gallery in the Next Experience reference section on the ServiceNow® Developer Site.
iconSize	Defines the size of the icon or avatar image shown on the Search Result card. • Type: choice list • Default value: xl (if **imageType** is set to **icon**) or md (if **imageType** is set to **avatar**) • Supported values:

Property	Description
	◦ *sm*: Show a small icon or avatar image on the Search Result card. ◦ *md*: Show a medium-sized icon or avatar image on the Search Result card. ◦ `lg`: Show a large icon or avatar image on the Search Result card. ◦ `xl`: Show an extra-large icon or avatar image on the Search Result card.
imageType	Defines the type of marker you want to show on the Search Result card. Supported values: • avatar • icon • image
imageURL	Defines the URL for the image you want to show on the Search Result card when **imageType** is set to **image**. • Type: string • Default value: (no value) **Note:** To specify literal text for this property, include the property in the `staticValues` object and specify the text as its `key` value. To populate this property with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields.
summary	Defines the content of the Search Result card's summary field. • Type: string • Default value: (no value)

Property	Description
	Note: To specify literal text for this property, include the property in the `staticValues` object and specify the text as its `key` value. To populate this property with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields.
summaryFontSize	Defines the font size for the Search Result card's summary field. • Type: choice list • Default value: sm • Supported values: ◦ *sm*: Show the Search Result card's summary field in small font size. ◦ *md*: Show the Search Result card's summary field in medium font size. ◦ *lg*: Show the Search Result card's summary field in large font size. ◦ *xl*: Show the Search Result card's title summary in extra-large font size.
textHeaderLabelOne, textHeaderLabelTwo, and textHeaderLabelThree	Defines the content of the Search Result card's text header labels. If you specify multiple text header labels, they appear separated by a pipe **|** character. • Type: string • Default value: (no value) **Note:** To specify literal text for one of these properties, include the property in the `staticValues` object and specify the text as its `key` value. To populate one of these properties with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields.
titleFontSize	Defines the font size for the Search Result card's title field.

Property	Description
	• Type: choice list • Default value: lg • Supported values: ◦ *sm*: Show the Search Result card's title field in small font size. ◦ *md*: Show the Search Result card's title field in medium font size. ◦ *lg*: Show the Search Result card's title field in large font size. ◦ *xl*: Show the Search Result card's title field in extra-large font size.
titleLabel	Defines the content of the Search Result card's title field. • Type: string • Default value: (no value) **Note:** To specify literal text for this property, include the property in the `staticValues` object and specify the text as its `key` value. To populate this property with the current search result's value from an AI Search indexed field, include the property in the `mappings` object with the field name as its value. For a list of AI Search indexed field names, see AI Search index fields.

Enabling and configuring AI Search in ServiceNow AI Platform® applications

Configure and use the AI Search engine in the Customer Service Management, Now Mobile, Service Portal, and Virtual Agent applications.

For each supported ServiceNow AI Platform application that you want to use AI Search in, define a search application configuration and specify AI Search as the search engine. To learn more about search application configurations, see Search application configurations.

Some applications support additional configuration for AI Search, as described in the following sections.

Customer Service Management (CSM)

Customer Service Management enables you to provide service and support for external customers through communication channels such as web, email, chat, telephone, and social media. The CSM application includes Customer and Consumer Service Portals that support AI Search.

For details on configuring the Customer and Consumer Service Portals to use AI Search, see Configure AI Search for the Customer and Consumer Service Portals. To activate AI Search

widgets so that guest users on these portals can use AI Search, see Enabling AI Search in the Customer and Consumer Service Portals for unauthenticated users.

Now Mobile

Enable your users to submit incidents and requests, manage tasks, and access company resources from anywhere using the Now Mobile app.

For details on configuring and using AI Search as the Now Mobile search engine, see AI Search in mobile.

Service Portal

Service Portal is a portal framework that enables administrators to build a mobile-friendly self-service experience for users. It interacts with parts of the ServiceNow AI Platform, so users can access specific platform features using Service Portal.

For details on enabling and using AI Search as the Service Portal search engine, see Enable and configure AI Search in Service Portal.

Virtual Agent

Virtual Agent is a platform for providing user assistance through conversations within a messaging interface. Use Virtual Agent to design and build automated conversations that help your users quickly obtain information, make decisions, and perform common work tasks.

For information on how Virtual Agent uses AI Search, see Improving the user experience with AI Search.

AI Search for Next Experience

The AI Search for Next Experience ServiceNow® Store application enables administrators to configure AI Search as the search engine for Unified Navigation searches in Next Experience, including global search and configurable workspace search.

For more information on how you can configure and use AI Search for Next Experience in global search and configurable workspace search, see AI Search for Next Experience.

Create a custom search matcher for global search

Customize the global search experience to return exact match search results from a specified table and column (field). You can configure custom search matchers for search terms with prefixes followed by numbers, such as `INT470271`, or for numeric search terms with delimiter characters, such as `525-326-3827` or `2016/01/08`.

Before you begin
You must have Next Experience enabled.

AI Search must be configured as the search engine for global search. For details on configuring AI Search as the search engine for global search, see AI Search for Next Experience.

Note:
If your instance was originally created running Utah or a later family release, AI Search is automatically configured as the search engine for global search.

Role required: ais_admin

About this task

To customize exact-matching behavior in global search, you can define a custom search matcher. The matcher specifies a Java regular expression pattern that AI Search compares search queries to, plus a table and a column (field) from that table. When a search query exactly matches the regular expression pattern, the custom matcher triggers and global search only returns exact-match results from the specified table and column.

Custom search matchers take precedence over the default AI Search exact match functionality for tables with prefixes defined in the Number [sys_number] table. If a global search query doesn't match any custom search matcher's regular expression pattern, AI Search falls back to evaluating the query for exact matches in tables with defined prefixes.

> **Note:**
> Users with the admin role can disable this fallback behavior by setting the value of the *`glide.ais.custom_matcher.legacy_fallback_enabled`* system property to *`false`*. With this value set, if a search query doesn't match any custom search matcher's regular expression patterns, AI Search treats it as a normal full-text search, ignoring exact matches in tables with defined prefixes.

To learn how AI Search evaluates global search queries for exact matches, see Using AI Search for Next Experience.

Procedure

1. Navigate to *`AI Search`* > **Search Experience** > **Global Search Custom Matchers**.
2. Select **New**.
3. On the Custom Matcher form, fill in the **Name**, **Table**, and **Column** fields.
   For a description of the field values, see Custom Matcher form.
4. Populate the form's **Regex** field.
   a. Select **Generate Regex**.
   b. In the pop-up window, select the option that more closely matches the values you expect to see in the specified **Column**.
   c. When prompted, enter the value prefix or select the value numeric delimiter.
   d. Close the pop-up window by selecting **Submit**.
5. Select **Submit.**
6. Navigate to the Custom matchers for Search Applications table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `m2m_search_context_config_custom_matcher.list`.
   c. Press Enter.
      The Custom matchers for Search Application [m2m_search_context_config_custom_matcher] table opens in list view.
7. For each search application that you want to use your custom matcher with, including global search, link the matcher to the application.
   a. Select **New**.
   b. On the Custom matchers for Search Application form, fill in the fields.

Field	Description
Custom matcher	The name of your custom search matcher.
Order	Evaluation order for your custom search matcher in the context of the specified AI Search application. Matchers with lower **Order** values are evaluated first.
Search context config	The name of the search application configuration record for the AI Search application that you want to link your custom search matcher to. For global search, select the **[AIS] Next Experience Search Configuration** record.

c. Select **Submit**.
A linkage record for your custom search matcher and the selected AI Search application appears in the Custom matchers for Search Application [m2m_search_context_config_custom_matcher] table.

Domain separation and AI Search

Domain separation is unsupported for the AI Search application tables, but user search queries respect domain restrictions from indexed ServiceNow AI Platform® table records. 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.

Note:
Refer to the Overview and following sections for details on how AI Search search queries respect domain settings from indexed records.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

AI Search domain separation overview

AI Search enables users to index and search content from records in ServiceNow AI Platform tables. The AI Search application tables and properties don't support domain separation, but user search queries respect domain settings from indexed content as described in the following sections.

Domain separation in indexed content and search queries

When indexing searchable content from a ServiceNow AI Platform table, AI Search stores the sys_domain field values for records on the source table and referenced tables.

If domain separation is enabled, AI Search applies a filter for the current session domain to every search query. This filter excludes records that aren't visible in the session domain. Only records visible in the session domain appear as search query results.

For child tables that have the *domain_master* attribute set, both indexing and search use the appropriate reference field to derive the child record domain from a referenced parent record.

Restrictions apply to referenced records in search queries, as summarized in the following table.

Condition	Behavior
Referenced record isn't visible in current session domain	• Search queries don't match content from the referenced record • Filtering on the reference field has no effect on search results • Faceting on the reference field has no effect
Referenced record domain differs from parent record domain	Search result records don't display the reference field

Indexing changes to record domains

AI Search automatically updates indexed domains for records in tables configured as indexed sources.

AI Search doesn't automatically update indexed domains for referenced tables. You can update the indexed domain for a referenced table in either of the following ways.

Use case	Action
Perform a one-time update of domain field values for a referenced table	Perform a full table reindex for the indexed source that includes the affected table. For details on this procedure, see Perform a full table index or reindex for a single indexed source.
Enable automatic updating of domain field values for a referenced table	Create a *dot_walk_fields* field setting for the reference field, specifying the `sys_domain,sys_domain_path` domain fields as the attribute value. For full steps for this task, see Enable automatic domain updates for a referenced table.

Related topics

Domain separation for service providers

Enable automatic domain updates for a referenced table

Configure a field setting to automatically index domain field value updates for records from a referenced table.

Before you begin

Role required: ais_admin

About this task

This topic provides steps for configuring a `dot_walk_fields` field setting to enable automatic indexing of domain field values for records in referenced tables. For more general instructions on defining field settings, see Create a field setting for an indexed source.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Open the indexed source that you want to define a field setting for.
3. In the Field Settings & Mapping related list, select **New**.
4. On the Field Setting form, enter the following field values.

Field	Value
Attribute	dot_walk_fields
Field	<reference field>
Value	sys_domain,sys_domain_path

   For a description of the field values, see Field Setting form.

5. Select **Submit**.

Result

The new `dot_walk_fields` field setting appears in the Field Settings & Mapping related list.

What to do next

To make the new `dot_walk_fields` attribute take effect, perform a full table reindex for the indexed source. For details on this procedure, see Perform a full table index or reindex for a single indexed source.

Indexing and searching external content in AI Search

With a subscription for the External Content for AI Search plugin, you can enable AI Search users to search content and metadata from documents in external repositories and remote tables.

The External Content for AI Search plugin enables you to index documents located outside of your ServiceNow AI Platform® instance, such as those stored in SharePoint collections, Confluence sites, and remote tables. Indexing these documents makes their content and metadata searchable in AI Search.

AI Search displays search results for indexed external documents just as it does for indexed ServiceNow AI Platform records. Users don't need to perform any special searches to match content and metadata from external documents.

Indexing doesn't duplicate external documents as records in your ServiceNow AI Platform instance. When a user selects an external document search result, they are directed to the source document in its original location.

External content indexing tools

The following diagram illustrates how AI Search provides indexing and search for external document content and metadata:

AI Search external content architecture diagram

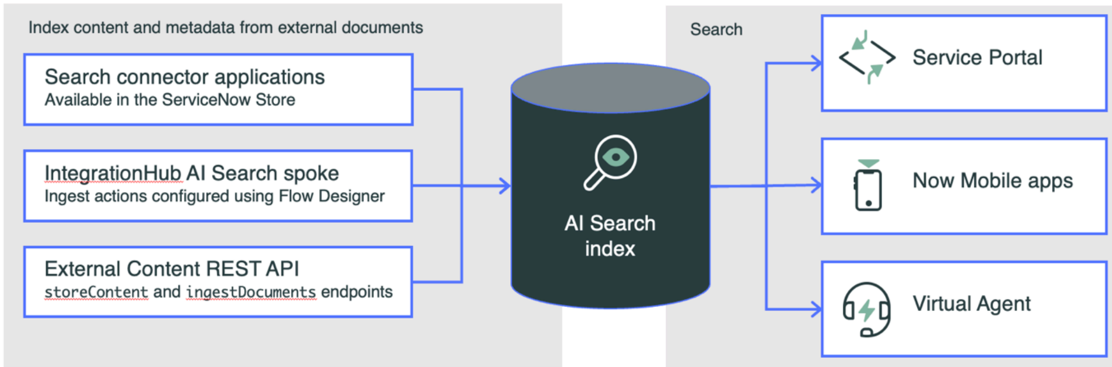

As shown in the architecture diagram, you can index content and metadata from external documents using any or all of the following tools:

External content indexing tool	Description
Search connector applications	Search connector applications available in the ServiceNow Store provide simple tools for indexing content and metadata from documents in supported external sources. They handle some of the configuration work, such as creation of external content schema tables, for you. Each search connector application supports a specific external document source (or set of sources). For example, the SharePoint Online Search Connector only supports documents in SharePoint collections, while the Raytion Connector Hub for AI Search supports multiple document sources. If no search connector application is available for your external document source, you must use the Integration Hub AI Search spoke or the External Content REST API to index content and metadata for documents from that source. Search connector applications provide the most straightforward way to index document content and metadata from supported external sources.
Integration Hub AI Search spoke	The Integration Hub AI Search spoke provides Flow Designer actions that you can use to ingest external documents for indexing:

External content indexing tool	Description
	• *`Ingest Document`*: Define property name-value pairs for an external document and submit the document to AI Search for indexing. • *`Ingest Document Using Download URL`*: Specify an external document's URL and have AI Search download it for indexing. You can specify security principal access permissions for documents ingested using these actions. To learn more about access permissions and security for external document content, see External content security for AI Search. For the full list of Flow Designer actions provided by the AI Search spoke, see AI Search spoke. The Integration Hub AI Search spoke combines flexible configuration and handling with the simple drag-and-drop Flow Designer interface.
External Content REST API	The External Content REST API provides endpoints that you can use to index content and metadata from external documents: • Send external document text content and metadata to AI Search for indexing using the *`ingestDocuments`* API endpoint. • Upload external document binary content to AI Search using the *`storeContent`* API endpoint. You can associate the stored binary content with a document fed using the *`ingestDocuments`* endpoint. **Note:** To access the External Content REST API endpoints, you need either the ais_external_content role or the admin role. You can specify security principal access permissions for documents ingested using the *`ingestDocuments`* endpoint. To learn more about access permissions and security for external document content, see External content security for AI Search. For full descriptions of the External Content REST API endpoints, see External Content Ingestion REST API.

External content indexing tool	Description
	The External Content REST API is the most complex tool for indexing content and metadata from your external documents, but it offers the most flexibility.

External content search prerequisites

Before you can index external document content and metadata for search in AI Search, you must satisfy the following requirements. Except as noted in the table, these requirements apply regardless of which external content indexing tools you use.

Requirement	Description
Have an active subscription for the External Content for AI Search plugin (com.glide.ais.external_content)	The External Content for AI Search plugin provides functionality for indexing content from external data sources. For instructions on purchasing a subscription for the plugin and activating it, see Request the External Content for AI Search plugin.
Define an external content schema table for the external document source	An external content schema table doesn't store external document content in the database. Instead, its columns serve as a map of AI Search index fields to populate when you index content from external data sources. For instructions on creating a new external content schema table, see Create an external content schema table. When documents from multiple external data sources share the same set of fields and have unique document identifiers, you can use a single external content schema table for all of them. To avoid data loss, the external data sources must satisfy these conditions: • All documents in the external data sources must include the same fields, matching the columns defined in the content schema table. • Every document in the external data sources must have a unique identifier so that documents don't overwrite one another in the AI Search index. If the external data sources can't satisfy these conditions, create a separate external content schema table for each data source.

Requirement	Description
	Note: Search connector applications from the ServiceNow Store automatically create external content schema tables as needed. You only need to create your own external content schema tables when indexing external documents using the Integration Hub AI Search spoke or the External Content REST API.
Create an indexed source with your external content schema table specified as the table to index	The indexed source sends external documents' content and metadata to the AI Search index. For instructions on creating an indexed source, see Create an indexed source.
Create one or more search sources for the indexed source	Each search source specifies a subset of external documents from the indexed source whose content you want to make searchable in AI Search. For instructions on creating search sources, see Create a search source.

Request the External Content for AI Search plugin

Request activation of the External Content for AI Search plugin (com.glide.ais.external_content) to enable indexing of searchable content and metadata from records in external data sources.

Before you begin

You must purchase a paid subscription before requesting activation of the plugin. For details on purchasing a paid subscription for a plugin, see ServiceNow plugins.

Role required: admin

About this task

You can't activate the External Content for AI Search plugin from your instance or the ServiceNow Store. Use this procedure to submit a Now Support request for activation of the plugin for your instance.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. On the All Applications page, select **Request Plugin** to open the **Activate Plugin** form on Now Support.

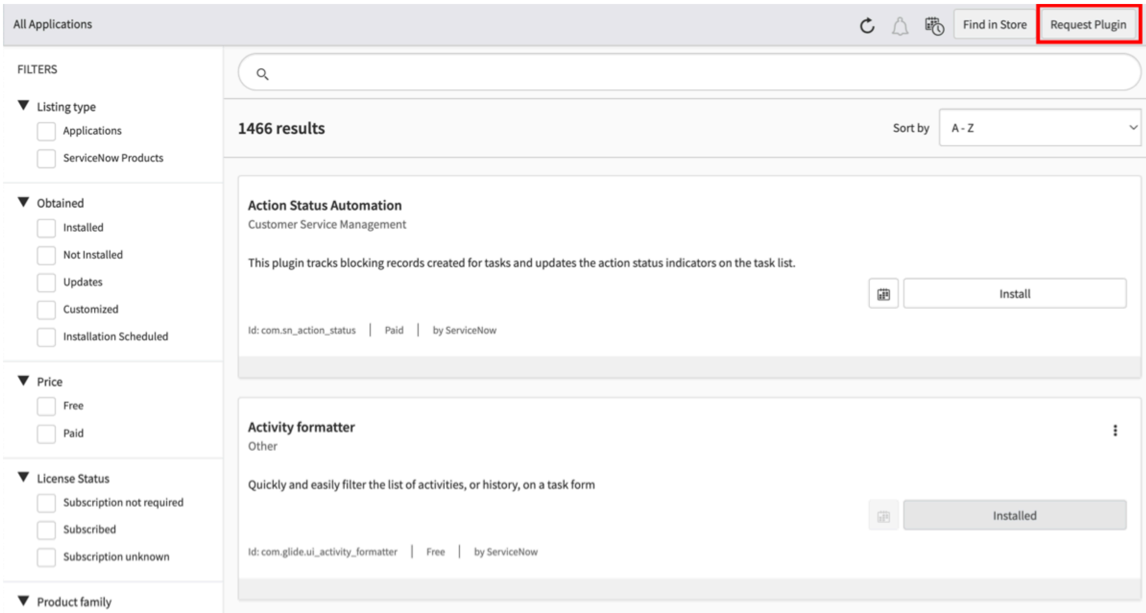

3. On Now Support, select the link to access the Now Support Service Portal Service Catalog.

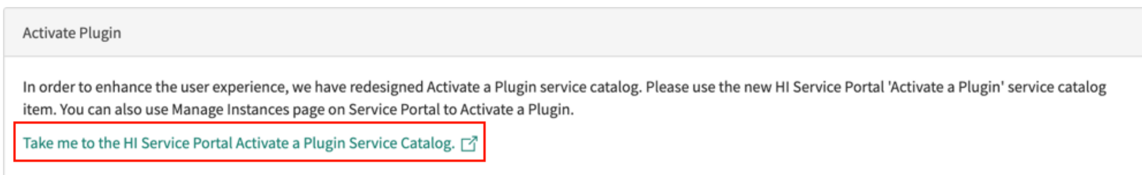

4. Select your instance.
5. Select **Actions > Activate Plugin**.
6. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Instance on which to activate the plugin.
Which plugin would you like to activate	Name of the plugin to activate. **Note:** If the system does not list the plugin you want or if you are activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	The date and time to activate the plugin.

Field	Description
	Note: Plugins are activated in two batches, once in the morning and once in the evening, on every business day in the US Pacific time zone. If the plugin must be activated at a specific time, enter the request in the **Reason/Comments** field.

Example

For example, see the following form to activate the CSM Workspace plugin on an instance named My Instance.

Activate Plugin form

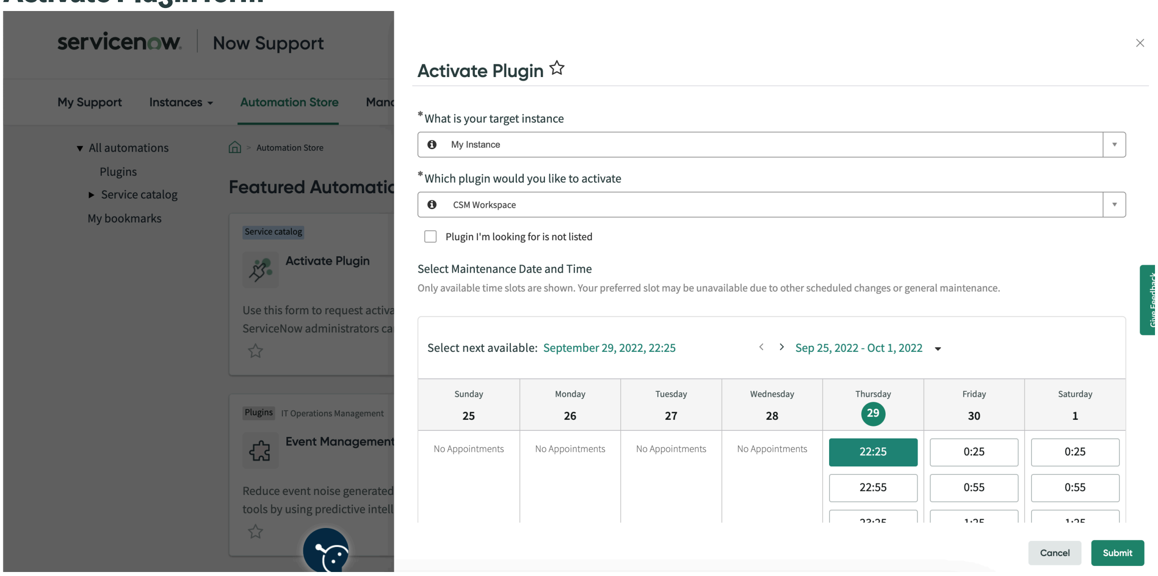

7. Select **Submit**.
   For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

Create an external content schema table

Define a schema table with columns corresponding to fields on records from an external data source. AI Search uses the schema when indexing content from the external data source.

Before you begin

Role required: ais_admin

About this task

An external content schema table doesn't store data in the database. Instead, its columns serve as a map of AI Search index fields to populate when you index content from external data sources.

When you index or delete records associated with this content schema using endpoints from the External Content Ingestion REST API, you must specify the schema table name as a path parameter.

Procedure

1. Navigate to **All** > `AI Search` > **External Content** > **Create Schema**.
2. Enter a name for your new external source schema table.
3. Select **Submit**.
   The new external content schema table opens in list view.
4. **Optional:** If you need to define additional field mappings for records from the external data source, add a new column for each field.
   a. In the Columns related list, select **New**.
   b. On the Dictionary Entry form, fill in the fields.
      For descriptions of these form fields, see Dictionary Entry form.

      Note:
      When filling in the **Type** field, select **String**, **Date**, **Date/Time**, or **Integer**. AI Search doesn't support other column types on external content schema tables.

   c. Select **Submit**.
5. If you modified the default column set, select **Update** to save your changes.

What to do next

Create a new indexed source with your external content schema table defined as its source. For instructions on creating an indexed source, see Create an indexed source

External content security for AI Search

AI Search preserves user and group security access permissions specified for documents indexed from external sources. You can control access to external content search results by mapping these externally defined users and groups to ServiceNow AI Platform® users.

For an overview of AI Search content security, see Content security in AI Search.

Requirements

To use external content security, include the following two steps in your ingestion and indexing process for external documents:

Step	Description
Specify access permissions on each external document fed for ingestion	Access permissions for an external document can allow or deny access to the document globally (for all users), or can include lists of specific externally defined users and groups who are allowed or denied access to the document. AI Search preserves the external document's security access permissions during indexing. Additional information:

Step	Description
	• For an overview of specifying access permissions on external documents, see Defining access permissions for external documents. • To learn about the set of access permissions supported for external content security, see External content access permissions. • For details on special access permissions granted by certain user roles, see Special external content access permissions by role. • To learn how to reverse the default precedence of `users.read` and `groups.deny` access permissions for an external content indexed source, see Change the precedence of user read and group deny permissions for an external content indexed source.
Define user mappings in tables linked to external content indexed sources	A user mapping specifies externally defined user and group aliases for a ServiceNow AI Platform user. Link these user mappings to indexed sources for external content. AI Search uses an indexed source's user mappings in conjunction with indexed records' access permissions to determine ServiceNow AI Platform user access for search results from the indexed source. Additional information: • For an overview of user mappings, see Mapping external users and groups to ServiceNow AI Platform users. • To learn about the set of externally defined security principal types supported in user mappings, see External user mapping security principal types. • For instructions on creating a table to store your user mappings, see Create a user mapping table. • For instructions on linking your user mapping tables to external content indexed sources, see Link a user mapping table to an external content indexed source. • To learn about importing user mappings via the AI Search External User Mapping API, see Importing user mappings. • For instructions on viewing user mapping import history records to confirm that user

Step	Description
	mappings imported correctly, see View history records for user mapping import operations.

Security implementation for search queries

When a user's search query matches an indexed record created from an external document, AI Search performs these steps:

1. Examines the user mappings linked to the record's indexed source and retrieves the set of all externally defined users and groups aliased to the current ServiceNow AI Platform user's account.
2. Compares the mapped set of externally defined user and group aliases with the access permissions on the indexed search result record to see whether the ServiceNow AI Platform user should be allowed to view the search result.

Defining access permissions for external documents

You can specify access permissions for externally defined users and groups when ingesting external documents. AI Search preserves these permissions during indexing so that user content security filters can operate on them at search time.

For details on including access permissions for externally defined users and groups in ingested external documents, see the `POST /ais/external_content/ingestDocument/{schema_table_name}` endpoint of the External Content Ingestion API.

External content access permissions

External content security includes permissions that describe user and group access for an external document. When indexing an external document, AI Search stores these permissions so that content security filters can limit user access to the indexed search result.

Access permissions

AI Search supports the following access permissions on ingested external documents.

Security principal	Description
everyone	Boolean option that indicates whether access to the external document is allowed for all users. AI Search applies this global access permission to the indexed record created from the document. Set this permission for an ingested document via the `[array].principals.everyone` request body parameter in a request to the `POST /ais/external_content/ingestDocument/{schema_table_name}` endpoint of the External Content Ingestion API. When set to **true**, this permission overrides all **groups** and **users** permissions.

Security principal	Description
	This permission is mutually exclusive with **none**. Only one of these two permissions can be set to **true** for any external document.
groups.deny	List of externally defined groups that are denied access to the external document. ServiceNow AI Platform® users mapped to any of these external groups can't view the indexed search result record created from the document. Set this permission for an ingested document via the *[array].principals.groups.deny* request body parameter in a request to the *POST /ais/external_content/ingestDocument/{schema_table_name}* endpoint of the External Content Ingestion API. This permission takes precedence over **groups.read**. If the same group has both read and deny access permissions for a document, AI Search denies members of the group access to the indexed record. By default, **users.read** takes precedence over this permission. To reverse this precedence order for an indexed source, see Change the precedence of user read and group deny permissions for an external content indexed source.
groups.read	List of externally defined groups that are allowed access to the external document. ServiceNow AI Platform users mapped to any of these external groups can view the indexed search result record created from the document. Set this permission for an ingested document via the *[array].principals.groups.read* request body parameter in a request to the *POST /ais/external_content/ingestDocument/{schema_table_name}* endpoint of the External Content Ingestion API. **groups.deny** permission takes precedence over this permission. If the same group has both read and deny access permissions for a document, AI Search denies members of the group access to the indexed record.

Security principal	Description
none	Boolean option indicating whether access to the external document is denied for all users. AI Search applies this global denial permission to the indexed record created from the document. Set this permission for an ingested document via the *`[array].principals.none`* request body parameter in a request to the *`POST /ais/external_content/ingestDocument/{schema_table_name}`* endpoint of the External Content Ingestion API. When set to **true**, this permission overrides all **groups** and **users** permissions. This permission is mutually exclusive with **none**. Only one of these two permissions can be set to **true** for any external document.
users.deny	List of externally defined users that are denied access to the external document. ServiceNow AI Platform users mapped to any of these external users can't view the indexed search result record created from the document. Set this permission for an ingested document via the *`[array].principals.users.deny`* request body parameter in a request to the *`POST /ais/external_content/ingestDocument/{schema_table_name}`* endpoint of the External Content Ingestion API. This permission takes precedence over **users.read**. If the same user has both read and deny access permissions for a document, AI Search denies that user access to the indexed record.
users.read	List of externally defined users that are allowed access to the external document. ServiceNow AI Platform users mapped to any of these external users can view the indexed search result record created from the document. Set this permission for an ingested document via the *`[array].principals.users.read`* request body parameter in a request to the *`POST /ais/external_content/`*

Security principal	Description
	ingestDocument/{schema_table_name} endpoint of the External Content Ingestion API. **users.deny** takes precedence over this permission. If the same user has both read and deny access permissions for a document, AI Search denies that user access to the indexed record. By default, this permission takes precedence over **groups.deny**. To reverse this precedence order for an indexed source, see Change the precedence of user read and group deny permissions for an external content indexed source.

Precedence order for principal permissions

The precedence order for *[array].principals* permissions on an ingested external document depends on the value of the *user_read_takes_precedence_over_group_deny* attribute for the document's indexed source.

Attribute value	Precedence order for principal permissions
true	From highest precedence to lowest: 1. *everyone* and *none* 2. *users.deny* 3. *users.read* 4. *groups.deny* 5. *groups.read* **Note:** This is the default attribute value for external content indexed sources.
false	From highest precedence to lowest: 1. *everyone* and *none* 2. *users.deny* and *groups.deny* 3. *users.read* and *groups.read* **Note:** For instructions on setting this attribute value, see Change the precedence of user read and group deny permissions for an external content indexed source.

For details on how content security permissions from certain user roles interact with these external content security principals, see Special external content access permissions by role.

Special external content access permissions by role

Certain user roles provide special access permissions for external content indexed records.

Role	Permissions
AI Search administrator [ais_admin]	An AI Search administrator can access all external content indexed records in a search application. **Note:** To bypass all search source and content security filtering in the Search Preview UI, you also need the impersonator and AI Search high security administrator [ais_high_security_admin] roles. For details on this procedure, see Diagnose search result access issues using the Search Preview UI.
Guest user [public]	Non-authenticated guest users can only access external content indexed records that have the *everyone* permission set to *true*.
Self-registered external user [snc_external]	Self-registered external users that belong to groups can access external content indexed records based on their group memberships. External users that don't belong to any group can only access external content indexed records that have the *everyone* permission set to *true*. For more details on self-registered external users, see Self-register to ServiceNow instance.

Change the precedence of user read and group deny permissions for an external content indexed source

Make external group deny access permissions take precedence over external user read access permissions for all external documents ingested through an indexed source.

Before you begin

The External Content for AI Search plugin (com.glide.ais.external_content) must be activated in your instance.

The source table for the indexed source must be an external content schema table.

Role required: ais_admin

About this task

By default, external user read access permissions (*users.read*) on an external document take precedence over external group deny access permissions (*groups.deny*) on the same document.

For example, suppose you ingest external content through an indexed source with a user mapping that maps ServiceNow AI Platform® user `beth.anglin@example.com` to external user `ad|beth-anglin` and external group `report-users`. If an external document grants read access to `ad|beth-anglin` and denies access to `report-users`, AI Search allows `beth.anglin@example.com` to view the indexed search result record for the external document.

To reverse this default behavior for an indexed source, making external group deny permissions take precedence over external user read permissions for all of its indexed records, change the value of the indexed source's `user_read_takes_precedence_over_group_deny` attribute. In the preceding example, making this change would prevent `beth.anglin@example.com` from viewing the indexed search result record for the external document.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. If the Advanced Configuration related list doesn't appear on the form, follow the steps in Add a related list to a form, selecting the **AI Search Indexed Source Attribute->Indexed Source** list in the slushbucket.
3. In the Advanced Configuration related list, select **New**.
4. On the Indexed Source Attribute form, enter the following field values.

Field	Value
Attribute	user_read_takes_precedence_over_group_deny
Value	false

   For a description of the field values, see Indexed Source Attribute form.
5. Select **Submit**.
   The attribute and value appear in the Advanced Configuration related list.

Result

The change in permission preference takes effect for search results from the external content indexed source.

Mapping external users and groups to ServiceNow AI Platform® users

AI Search external user mappings enable you to configure externally defined user and group aliases for ServiceNow AI Platform users. AI Search uses these aliased security principals to apply content security filters that control user access to search results indexed from external documents.

You can use mapped external user and group names when defining access permissions for external documents that you ingest for indexing via the External Content Ingestion API.

External content security user mapping overview

To configure user mappings for external content security, you must perform the following steps:

1. Create one or more user mapping tables. For instructions on this procedure, see Create a user mapping table.
2. Link a user mapping table to each of your external content indexed sources. For details on this task, see Link a user mapping table to an external content indexed source.

Note:
If no mapping is found for the current ServiceNow AI Platform user, AI Search only allows the user to view search results for external documents that have the *everyone* permission set to true. This can occur if no user mapping table is linked to a search result's indexed source, or if no mapping for the current ServiceNow AI Platform user exists in the indexed source's linked user mapping table.

3. Import user mappings into your user mapping tables with the AI Search External User Mapping API.

Users with the ais_high_security_admin role can review history records for import operations to verify that user mappings imported correctly. For instructions on this procedure, see View history records for user mapping import operations.

External user mapping security principal types

AI Search user mapping aliases can include externally defined group and user security principals.

Mapped security principal type	Description
Externally defined group	The ServiceNow AI Platform® user specified in the user mapping inherits group-level access permissions for all aliased external groups. For example, if you define external group *report-users* as an alias for ServiceNow AI Platform user *beth.anglin@example.com*, AI Search allows *beth.anglin@example.com* to view indexed search results for external documents that include read access permission for *report-users*.
Externally defined user	The ServiceNow AI Platform user specified in the user mapping inherits user-level access permissions for all aliased external users. For example, if you define external user *ad|abel-tuter* as an alias for ServiceNow AI Platform user *abel.tuter@example.com*, AI Search allows *abel.tuter@example.com* to view indexed search results for external documents that include read access permission for *ad|abel-tuter*.

Create a user mapping table

Define a table to store externally defined user and group aliases for ServiceNow AI Platform® users.

Before you begin

The External Content for AI Search plugin (com.glide.ais.external_content) must be activated in your instance.

Role required: ais_admin

About this task

User mapping tables store user mappings for use with external documents that have access permissions. Each mapping specifies externally defined user and group aliases for an individual ServiceNow AI Platform user. AI Search uses the mappings linked to an external document's indexed source to determine whether a ServiceNow AI Platform user can view the indexed record for the document.

All user mapping tables extend the AI Search External Search User Mapping Base [ais_external_search_user_mapping] table.

Procedure

1. Navigate to **All** > *AI Search* > **External Content** > **Create User Mapping**.
2. Enter a name for your user mapping table.
3. Select **Submit**.
   The dictionary form for your new user mapping table appears.

What to do next

Import user mappings into your new user mapping table. For an overview of this process, see Importing user mappings. For details on the relevant API endpoint, see AI Search External User Mapping API.

Link a user mapping table to an external content indexed source

Specify the user mapping table you want AI Search to use for external documents ingested through an indexed source.

Before you begin

The External Content for AI Search plugin (com.glide.ais.external_content) must be activated in your instance.

The source table for the indexed source must be an external content schema table.

Role required: ais_admin

About this task

A user mapping table linked to an indexed source defines external user and group aliases for ServiceNow AI Platform® users. AI Search uses these aliases to determine a user's effective access permissions when their search matches a record from the indexed source.

A user mapping table can be linked to any number of indexed sources.

> **Note:**
> If no mapping is found for the current ServiceNow AI Platform user, AI Search only allows the user to view search results for external documents that have the `everyone` permission set to true. This can occur if no user mapping table is linked to a search result's indexed source, or if no mapping for the current ServiceNow AI Platform user exists in the indexed source's linked user mapping table.

Procedure

1. Navigate to **All** > *AI Search* > *AI Search* **Index** > **Indexed Sources**.
2. Open the indexed source record that you want to link a user mapping to.

3. In the AI Search External Search User Mapping Table Data Source Relationships related list, create a new relationship record for the user mapping table that you want to link to the indexed source:

   a. Select **New**.

   b. On the AI Search External Search User Mapping Table Data Source Relationship form, fill in the fields.

   c. Select **Submit**.
   The user mapping table's name appears in the related list.

Importing user mappings

You can use the AI Search External User Mapping API to import mappings between ServiceNow AI Platform® users and externally defined users and groups.

To define a user mapping for import, you need the following information:

- The email address that uniquely identifies a ServiceNow AI Platform user in the User [sys_user] table.
- The list of externally defined groups that you want to alias to the ServiceNow AI Platform user. This list can be empty.
- The list of externally defined users that you want to alias to the ServiceNow AI Platform user. This list can be empty.

For more information on importing user mappings, see AI Search External User Mapping API.

View history records for user mapping import operations

Review the results for your user mapping import operations, including data, log messages, and any errors. Use this procedure to verify that your user mappings are imported correctly.

Before you begin

Role required: ais_admin, ais_high_security_admin, and import_transformer.

About this task

AI Search creates an import set for each batch of user mappings that you import via the AI Search External User Mapping API. The transform histories for these import sets provide insight into the results of your user mapping import operations.

Procedure

1. Elevate to the ais_high_security_admin role:

   a. Perform the appropriate action for your version of the UI:

UI version	Action
Next Experience UI	In the banner frame, select the icon for your account to open the user menu, then select **Elevate role**.

UI version	Action
	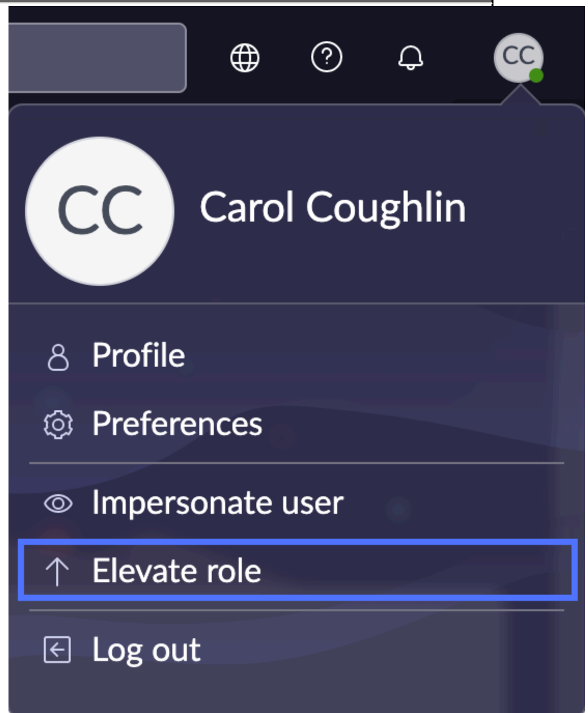
Core UI	In the banner frame, select your name to open the user menu, then select **Elevate Roles**. 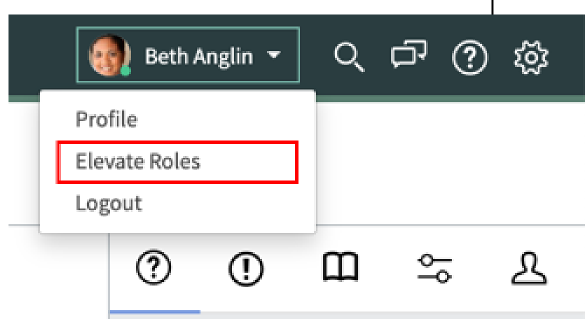

A dialog box appears, displaying a checklist of your available privileged roles.

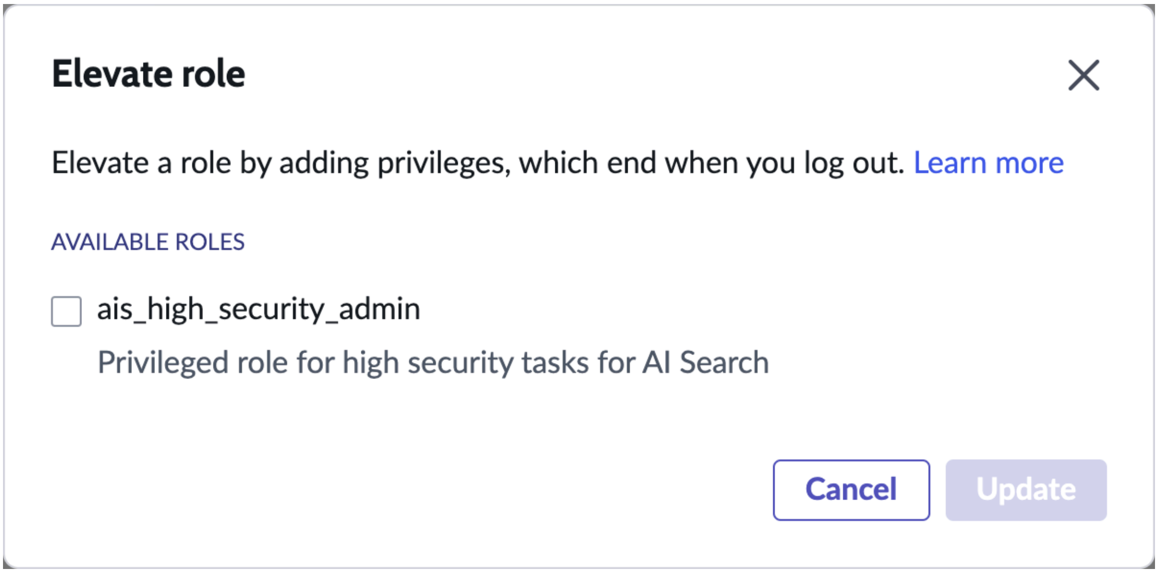

b. In the dialog box, select the **ais_high_security_admin** option, then select **Update** (in Next Experience UI) or **OK**.
The page reloads and an elevated role indicator appears next to your user name in the user menu. In Next Experience UI, this indicator displays the names of the active privileged roles. In Core UI, the indicator displays the elevated role icon ().

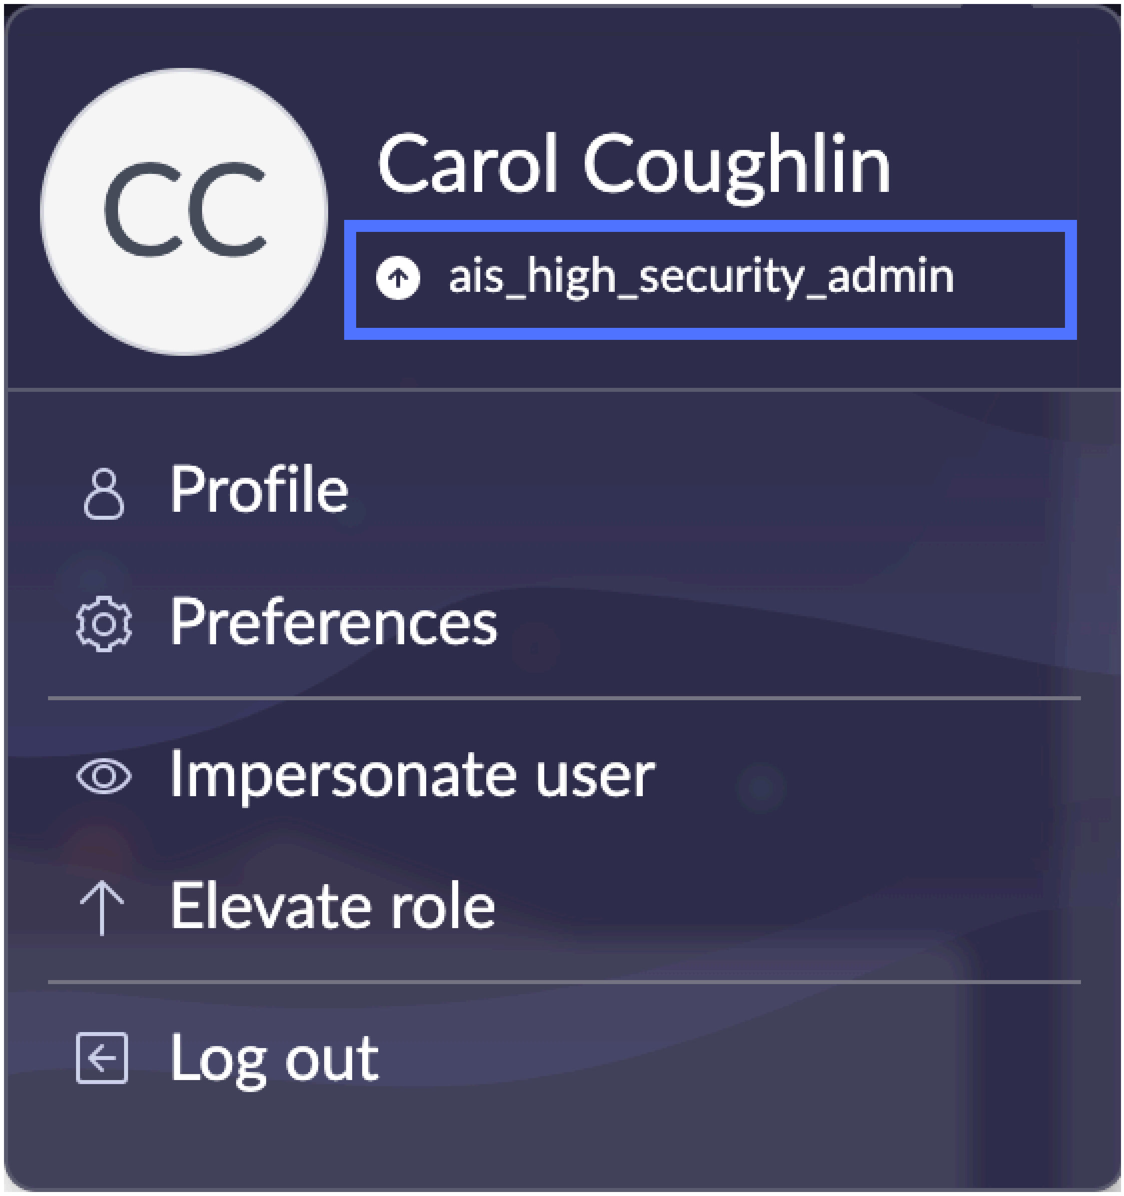

Note:
When the page reloads, any unsaved edits are lost.

2. Navigate to **All** > `AI Search` > **External Content** > **User mapping import history**.
3. Review each user mapping import history record that you're interested in.
    a. Open the user mapping import history record.
    b. On the AI Search External Content User Mapping History Record form, select the **Show Transform History** related link.
    The Transform History [sys_import_set_run] table opens, with a filter applied for the import set from the selected user mapping import history record.
    c. Open the transform history record.
    d. Review import set results on the form, errors in the Import Set Row Errors related list, and log messages in the Import Log related list.

Components installed with External Content for AI Search

When you activate the External Content for AI Search plugin (com.glide.ais.external_content), the system creates a new virtual table.

Tables installed

Table	Description
AI Search External Search Content [v_ais_external_search_content]	Stores the base set of field mapping columns for inclusion in external content schema tables. **Note:** All user-created external content schema tables extend this table.
AI Search External Search User Mapping Table Data Source Relationship [ais_external_search_user_mapping_data_source_m2m]	Stores links between user mapping tables and indexed sources.

Integrating AI Search into other ServiceNow AI Platform® applications

AI Search integrates with other ServiceNow AI Platform applications. Populate your UI Builder pages with search components and customize display of search results and Genius Results with EVAM definitions. Improve incident deflection in Service Portal by configuring AI Search Assist for record producers.

UI Builder search components

AI Search provides components for search input, search results, search facets, and navigation tab filters. You can add these components to your pages in UI Builder.

The AI Search components are described in detail on the ServiceNow® *Developer Site*:

- Search facets
- Search input
- Search results
- Tab filter

For details on adding components to your pages in UI Builder, see Add and configure components.

Entity View Action Mapping (EVAM) definitions

Entity View Action Mapping (EVAM) is a ServiceNow AI Platform application that standardizes how different data sources display records in cards and lists.

AI Search provides EVAM configurations that you can use to customize display of search results and Genius Results in Service Portal and your UI Builder pages.

You can customize EVAM settings to display AI Search results as a card grid view or as a list view. Users can page through large data sets of search results and see different views based on filtering.

For details on EVAM definitions, configurations, templates, and configuration bundles, see Entity View Action Mapping (EVAM).

Service Portal record producer integration

AI Search Assist is a Service Portal widget that uses AI Search to improve incident deflection by displaying the most relevant set of related search results within a record producer. You can use AI Search Assist instead of Contextual Search.

For more details on the AI Search Assist widget, see AI Search Assist record producer integration for Service Portal.

AI Search Assist record producer integration for Service Portal

AI Search Assist is a Service Portal widget that uses AI Search to improve incident deflection by displaying the most relevant set of related Catalog Item and knowledge article search results within a record producer. You can use AI Search Assist instead of Contextual Search.

AI Search Assist overview

AI Search Assist can help a user find the most relevant Catalog Items and knowledge articles in a search on your portal. For example, a user, creating an Incident via the Service Portal, can see related Catalog Items that the user can order. The user can also view related knowledge articles to assist them to resolve the issue on their own without involving the service desk.

AI Search Assist helps with incident deflection for incident record producers, deflecting incidents by helping end users resolve issues before they raise an incident.

AI Search Assist supports ServiceNow mobile.

> **Note:**
> AI Search Assist doesn't support external users. If an external user tries to use an Assist UI action, the system returns an error.

Activation information

The AI Search Assist plugin (com.snc.ai_search_assist) is a ServiceNow AI Platform® application that is activated by default when AI Search is active.

Configure AI Search Assist for a record producer

Configure AI Search Assist settings for a record producer to help users view related Catalog Item and knowledge article search results.

Before you begin

Role required: aisa_admin or admin

About this task

AI Search Assist is a Service Portal widget that displays relevant Catalog Item and knowledge article search results within a record producer configuration.

> **Note:**
> If both AI Search Assist and Contextual Search are configured for the same record producer, see `Results` at the end of this task for how to control which one is displayed.

Procedure

1. Navigate to **All** > *AI Search* > *AI Search* **Assist** > **Record Producer Configuration**.
2. Select **New** and fill in the form to create a configuration for Catalog Item and knowledge article search results.
3. On the **Search application information** tab, select a Search application configuration.
   You can choose from a list of pre-defined configurations for specific features. For example, *[AI Search Assist] - Incident Deflection*.

   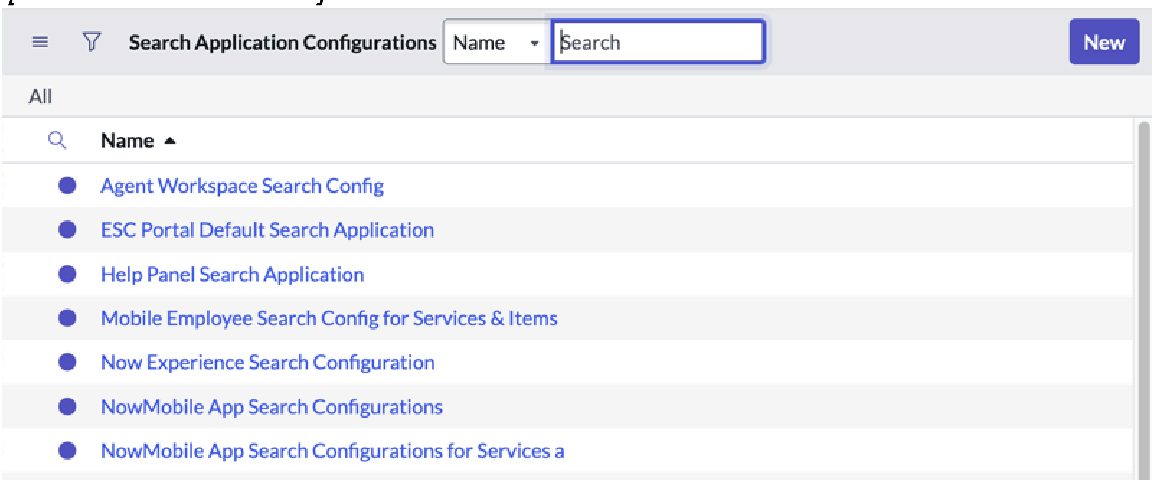

4. On the **Record producer information** tab, select a Record producer and Search variable.
   For example, *Service Category Request* as a record producer can have a search variable of *Category name*.

   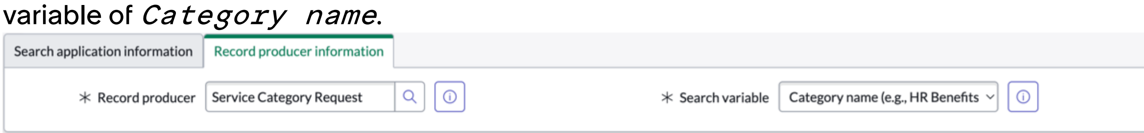

5. **Optional:** To adjust the visibility of your card or detail view, select **Show UI Actions** under the new Record Producer Configuration record's related links.
   AI Search Assist *UI actions* can hide from view using a simple condition or advanced condition via a script.
6. Select **Submit**.

Result

AI Search Assist is active.

Note:
AI Search Assist doesn't support external users. If an external user tries to use an Assist UI action, the system returns an error.

If Contextual Search has been configured for the same record producer, and is showing results instead of AI Search Assist, check the following conditions:

- The *Enable AI Search* flag on the Portal (sp_portal) is active
- The AI Search Assist record configuration is active
- The AI Search Profile is active and published
- The record producer variable for AI Search Assist is created for the associated record producer and the variable is active

If these conditions aren't all met, then Contextual Search is the fallback for the same record producer to make sure you see related search results.

AI Search Assist roles

AI Search Assist includes the following roles.

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

AI Search Assist administrator [aisa_admin]

Configure and manage AI Search Assist settings.

- Manages AI Search Assist on record producers, including the associated UI actions
- Adds variables to a record producer
- Administers AI Search Assist capability

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

AI Search Assist properties

The AI Search Assist properties determine when searching begins based on characters entered and the time that has elapsed. These properties reduce the impact of searches for systems with many Catalog Items or Knowledge articles.

AI Search Assist properties

Property	Description
com.snc.ai_search_assist.min_length	The minimum number of characters entered before a search initiates. The default is 3.
com.snc.ai_search_assist.wait_time	The time elapsed after the user finishes entering characters in the search field before the search begins.

Property	Description
	Note: If you set the time to '-1', the search begins when the user leaves the search field. This behavior is known as 'Search on tab'.

Domain separation and AI Search Assist

Domain separation is supported for AI Search Assist. 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.

AI Search Assist overview

The goal of AI Search Assist is to give a Service Portal record producer AI Search capabilities, which improve search results over time. Domain separation is enabled in AI Search Assist.

How domain separation works in AI Search Assist

Domain separation for AI Search Assist is set in the record producer configuration. A related link for each record producer configuration enables you to expand or collapse the domain scope. If you expand the domain scope, any child domain record producer configurations can override a global or parent domain record configuration.

Related topics

Domain separation for service providers

Using AI Search

See how AI Search displays search results, refinement filter options, and Genius Result answers. Learn how to search using the AI Search query language. Understand when and how AI Search automatically resubmits queries.

> **Note:**
> For more details on using AI Search features in global and workspace search, see Using AI Search for Next Experience.

AI Search query language

Learn how to construct search queries using terms, phrases, and AI Search query operators.

Search query terms and phrases

Specify terms and quoted phrases in the search query to find records that contain the same terms or phrases.

Search query terms	Description
`conference`	Single-term search. Finds records that contain the term `conference`.

Search query terms	Description
	Note: Search terms don't match substrings of indexed terms. As an example, searching for `exam` finds records that contain the term `exam`, but not records that contain the term `example`. To search for substring matches, you can use the `%` or `*` wildcard matching operators.
`conference room`	Conjunctive multi-term search. Finds records that contain both of the terms `conference` and `room` (an intersection of sets). The search terms can appear in any order or proximity in the record and may appear in different fields. **Note:** If you set the **Boolean search operator to use when a search query includes multiple terms** (*`glide.ais.query.search_operator`*) system property to *`OR query`*, this search finds records that contain either of the terms `conference` or `room` (a union of sets).
`"conference room"`	Quoted phrase search. Finds records that contain the term `conference` followed immediately by the term `room`. The search terms must appear in this order and in the same field. **Note:** Quoting a phrase doesn't disable linguistic features (such as normalization, synonym expansion, and stop word removal) or wildcard operators. For example, a search for the quoted phrase `"email acc*"` expands the wildcard operator normally to find records that contain `email account` and `email access`.

AI Search ignores letter case for search query terms and phrases. For example, searching for `PIN` finds records containing `PIN` and `pin`.

Note:
When processing search query terms, AI Search automatically strips out HTML and XML content by removing strings that start with `<` and end with `>`. As an example, if you enter `user <beth.anglin@example.com>` as your search terms, AI Search only searches for `user`. If you search for `<abel.tuter@example.com>`, AI Search treats the search as empty and returns no results. This behavior isn't configurable.

Boolean search operators

Separate query terms and phrases with Boolean operators to override the default matching logic for the search query.

Operator	Description
`OR`	Boolean disjunction operator. Finds records that contain any of the terms or phrases separated by `OR`. For example, searching for `"conference room" OR suite` finds records that contain the phrase `conference room` and records that contain the term `suite`. **Note:** If you set the **Boolean search operator to use when a search query includes multiple terms** (*`glide.ais.query.search_operator`*) system property to *`OR query`*, this operator has no effect because search already treats all query terms and phrases as though separated by the `OR` operator.
hyphen (`-`)	Boolean negation operator. Excludes records containing the term or phrase that immediately follows the hyphen. For example, searching for `email -"email signature"` finds records that contain the term `email` but don't contain the phrase `email signature`.

Wildcard matching operators

Use wildcard operators to find records that contain indexed terms matching a wildcard pattern.

Operator	Description
`%`	Single-character wildcard operator. Include `%` in a search query term to match any one character in an indexed term. As an example, searching for `the%e` finds results containing wildcard matches such as `theme`, `there`, or `these`. AI Search also evaluates the search with the wildcard operator ignored, to look for literal matches. The non-wildcard operator parts

Operator	Description
	of the search term are treated as a phrase, meaning that they must appear in the same order in the result as they do in the search. As an example, a search for `5000%000` returns results that contain `5000` followed by `000` as well as results that contain wildcard matches like `50007000`. **Note:** The `%` operator must follow a string of three or more non-wildcard characters. For search query terms that don't satisfy this condition, AI Search only returns literal matches. As an example, if you search for `99%`, with only two non-wildcard characters preceding the `%` operator, AI Search only returns results with literal matches for `99`.
`*`	String wildcard operator. Include `*` in a search query term to match any string of zero or more characters in an indexed term. As an example, searching for `acc*` finds results containing wildcard matches such as `access`, `account`, `accrue`, `accumulate`, and `accuracy`. AI Search also evaluates the search with the wildcard operator ignored, to look for literal matches. The non-wildcard operator parts of the search term are treated as a phrase, meaning that they must appear in the same order in the result as they do in the search. As an example, a search for `753*268` returns results that contain `753` followed by `268` as well as results that contain wildcard matches like `7539268`. **Note:** The `*` operator must follow a string of three or more non-wildcard characters. For search query terms that don't satisfy this condition, AI Search only returns literal matches. As an example, if you search for `ad*`, with only two non-wildcard characters preceding the `*` wildcard operator, AI Search only returns results with literal matches for `ad`.
`***`	Universal wildcard operator. Specify `***` as a search query to find all indexed terms and thus all records.

Operator	Description
	Note: AI Search doesn't apply relevancy ranking to `***` queries. Results from `***` queries appear in an unspecified order.

Note:
When expanding search terms that contain `%` or `*` wildcard operators, AI Search ignores terms defined as stop words. For example, suppose you define `the` and `their` as stop words. A search for `the*` won't expand to match `the` or `their`, but will still match non-stop word terms such as `there` and `these`.

Viewing search results

Learn about the components that AI Search uses to display results for search queries.

Search result page elements

The layout and appearance of AI Search UI elements varies from application to application. This example shows the elements as they might appear on a Service Portal search results page.

Sample search result page

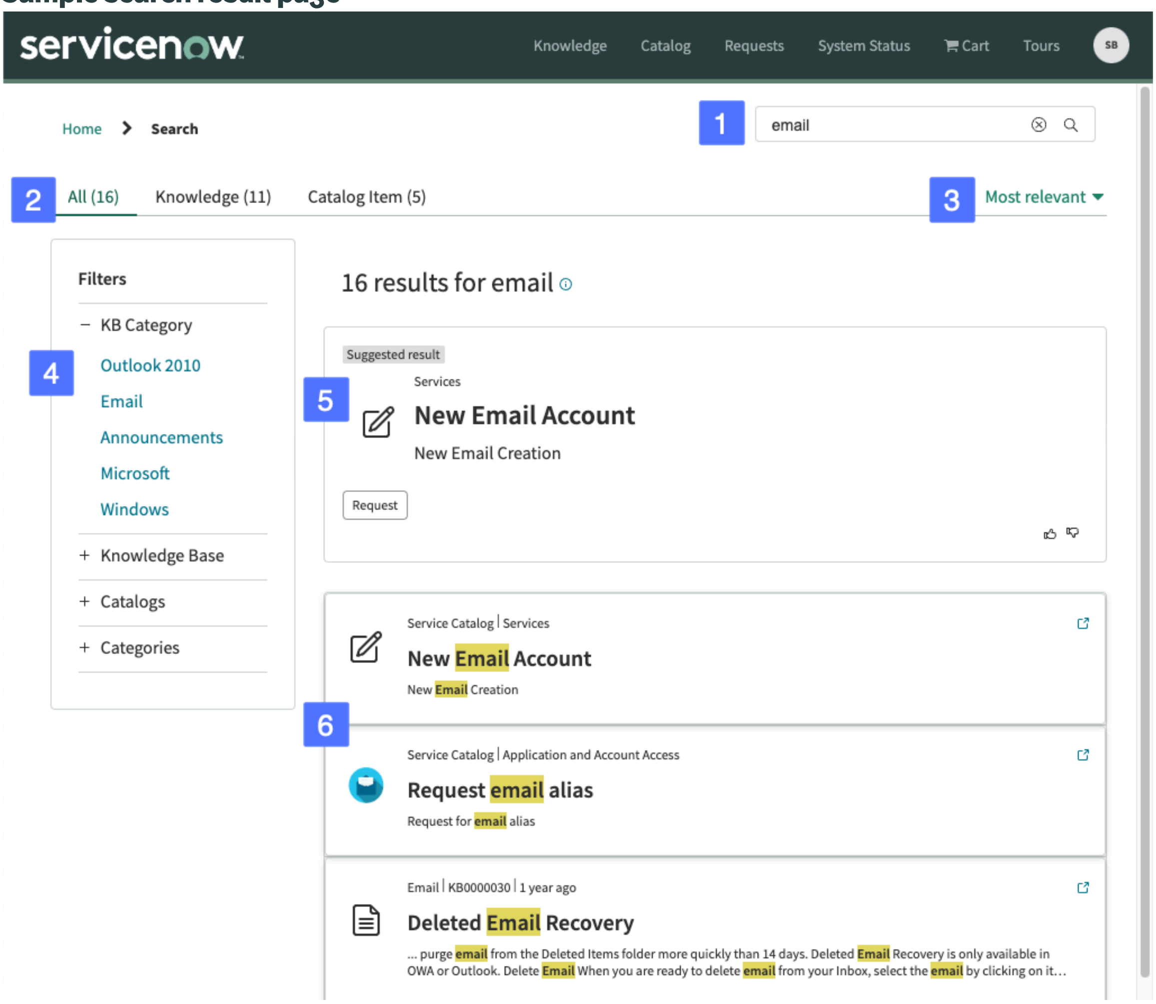

Search result page elements

Element	Description
1. Search input field	Displays the search query that produced the listed search results. You can modify this search query and search again to view different results.
2. Navigation tab filters	Displays static filter options for refining search results by search source. If you want to view only results from a single search source, select that search source's name. To remove the navigation tab filter, select **All**.
3. Search result sort order	Displays the currently selected sort order for search results. By default, AI Search displays the most relevant results first. To apply an alternative sort order, open this list and select one of the available search result sort options. **Note:** If no search result sort options are available, the sort menu doesn't display and results are automatically sorted by relevance.
4. Facet filters	Displays dynamic filter options for refining search results by field value. To expand or collapse a facet, select the **+** (plus) or **-** (minus) symbol by its name. To view only search results that contain a specific field value, select the field value option in the facet. Clear the field value option to remove the facet filter. **Note:** When faceting on a date field, the facet options are precomputed date ranges, such as `Past 3 months`, rather than exact date field values. Selecting an option displays only results that have date field values in the specified range. You can select and clear these date range options in the same way that you select and clear field value options for non-date facets. Some facets allow you to select multiple field values. Field values in multi select facets

Search result page elements (continued)

Element	Description
	display as options instead of links, and include a label showing how multiple selections are applied when filtering your search results. Labels include: • `Contains all selected`: Displays only search results that include all of the selected field values. This label indicates the Multi Select And facet type. • `Contains any selected`: Displays search results that include any one of the selected field values. This label indicates the Multi Select Or facet type. If your facet filter cannot be applied to search results from the navigation tab you've selected, AI Search ignores the facet filter.
5. Genius Result answer cards	Display the best answers for your search query along with actions you can take directly from each card. **Note:** Genius Result answer cards may not appear for all search queries. To provide feedback relating to a Genius Result answer, select one of its **Was this suggestion helpful?** links. If the answer is relevant to your search, select **Yes**; otherwise, select **No**.
6. Search results	Displays the icon, summary field values, title, and text snippet for each search result, with search query terms highlighted in title and text snippets. Select a result to view its source record in full. When a search result displays the Open in New Tab icon ( ), you can select it to view the source record in a new browser tab, preserving the state of your search in the current tab. **Note:** If you select an attachment search result, the system may display the attachment or may prompt you to download the attachment.

Auto-correction message

When AI Search auto-corrects a misspelled search query term, a Showing results for <corrected_term> message appears above the search results list.

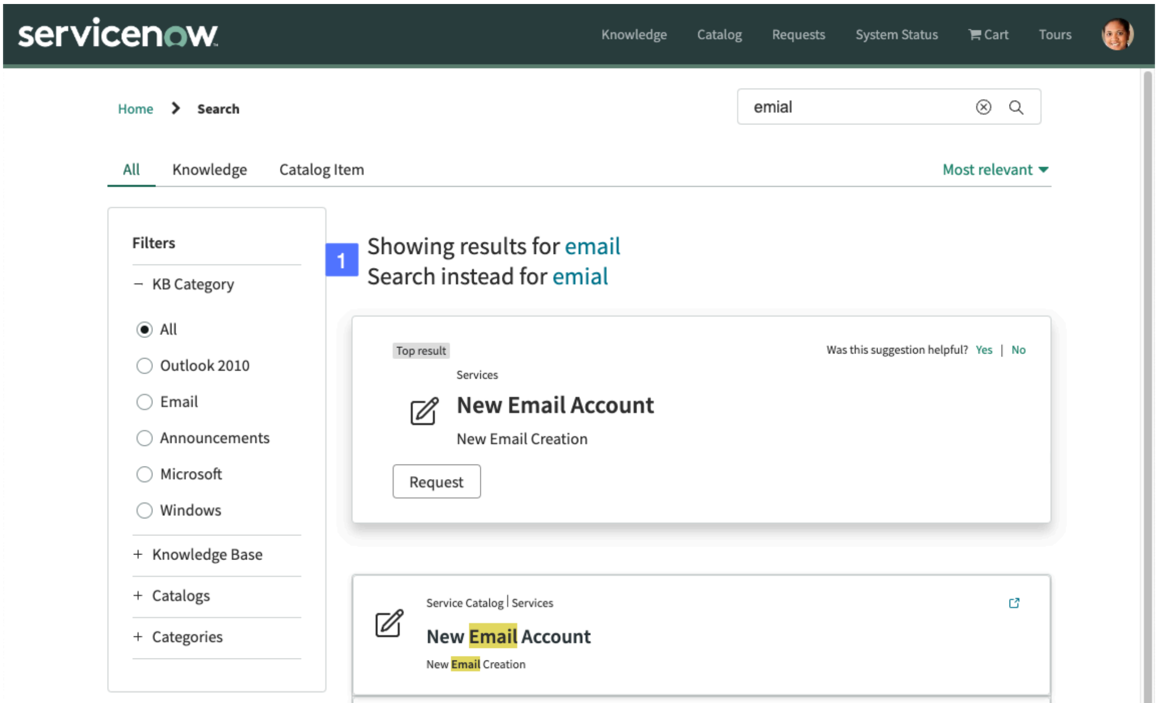

Auto-correction message elements

Element	Description
1. Typo handling auto-correction message and link to original search term	Displays the auto-corrected query term used for the search. To repeat the search using your original search query term, select the link in the **Search instead for <original_term>** line of the message.

Change the sort order for your search results

Choose the order in which AI Search displays results for your search. By default, the most relevant search results appear first. You can override this behavior and view search results ordered by their field values.

Before you begin

Role required: none

About this task

Sorting enables you to view your search results in an order based on their field values. You can select any of the sort options available in the search application.

Note:
If your search administrator has disabled sort options globally or for the search application, the sort menu doesn't appear and you can only see your search results in the default relevancy order.

Procedure

1. On the results page for your search, access the sort menu.
   - In Now Mobile, select **Filter** and select **Sort by**.
   - In other applications, select **Most relevant**.

   The sort menu displays the list of available sort options.
2. Select the sort option that you want to apply to your current search result set.

   Example

   As an example, to display the most recently updated search results first, select **Most recent**. The search results page refreshes to display your results in the desired order.

   > **Note:**
   > The new sort order persists when you select facets or navigation filters. It resets to the default relevancy order when you start a new search by modifying the search query terms.

Automatic search query resubmission

When a search query returns too few results, AI Search automatically modifies and resubmits it with the goal of returning at least the desired number of results.

Search administrators configure the minimum number of desired results per search query using the **Search Results Limit** field in search application configurations. For more details, see Create a search application configuration for AI Search.

AND-to-OR search query resubmission

When a search query with conjunctive terms or phrases returns insufficient results, AI Search resubmits it with the default Boolean operator changed from AND to OR and displays results for the modified query.

For example, if a query for `hotel "conference room" suite` matches no results, AI Search resubmits it as `hotel OR "conference room" OR suite`. This less-restrictive query may match more results than the original query.

> **Note:**
> If you set the **Boolean search operator to use when a search query includes multiple terms** (*`glide.ais.query.search_operator`*) system property to *`OR query`*, AI Search doesn't resubmit search queries in this way because the initial search already treats all query terms and phrases as though separated by the `OR` operator.

When automatically resubmitting a search query in this way, AI Search adds a minimum match count restriction to the OR expression. The modified query only returns results that match at least half (rounded up) of the terms separated by OR operators. This restriction improves query performance and provides more relevant results.

For example, suppose a search query for `laptop can't access wi-fi network` returns no results. AI Search automatically resubmits the query as `laptop OR can't OR access OR wi-fi OR network`. Only records that match at least three of the five query terms appear as results for the resubmitted search.

Conditions for automatic search query resubmission

Automatic search query resubmission ignores search queries that meet any of the following conditions.

- The original search query only includes a single term.
- The original search query includes 8 or more terms and returns at least one search result.

Provide feedback on Genius Results

Users can submit feedback for each Genius Result answer card, specifying whether it was helpful for their search or not. For cards from default Genius Result configurations, users can also select the reason they found an answer unhelpful. AI Search records feedback responses as search signals that you can use to assess the relevancy of your Genius Result configurations.

Before you begin

You must have performed a search that returned one or more Genius Result answer cards.

Role required: none

About this task

You can provide feedback for each Genius Result answer card in your search result set, marking it as helpful or unhelpful for your search. If you mark an answer card produced by a default Genius Result configuration as unhelpful, you can also select the reason it was unhelpful.

AI Search stores user feedback for Genius Result answers as signals in the Genius Result Event Action [sys_search_genius_result_event_action] table. For more information on this table, see Search signal tables.

Procedure

1. On the Genius Result answer card that you want to provide feedback for, find the **Helpful?** section.
2. Choose one of the following:
   - If you found the Genius Result answer helpful for your search, select **Yes**.
   - If you found the Genius Result answer unhelpful for your search, select **No**.
3. If you selected **No** on an answer card produced by a default Genius Result configuration, choose one of the following:
   - To indicate why you found the Genius Result answer unhelpful, select a reason from the list in the dialog box, then select **Submit feedback**.
   - To submit your feedback without a reason, select anywhere outside the dialog box.

Result

AI Search submits your feedback as a search signal.

Administering AI Search

Learn how AI Search features and behavior affect indexing and search. Gain insight into AI Search behavior by reviewing log messages, ingestion history for indexed sources, and session debugging messages.

AI Search and performance testing

Understand how the AI Search topology affects performance testing on instances created with System Clone.

> **Warning:**
>
> Be careful when testing AI Search performance in cloned non-production instances. The same dedicated AI Search nodes handle search query traffic from all of your instances. As a result, search queries on non-production instances may impact search query performance in your production instance.

Lemma and Unicode normalization

AI Search normalizes inflected words and Unicode glyphs during indexing and at search query time. Normalization improves search recall and enables users to find content with variant forms of their search query terms.

Normalization features are automatically enabled and aren't configurable.

Lemma normalization

Many languages include inflected forms of terms, such as plural nouns or verb tenses. AI Search normalizes inflected terms found in indexed content and search queries. Normalization enables matching based on a root form, such as the singular for a plural noun or the base form for a conjugated verb. This root form is called a lemma, and this process is referred to as lemma normalization.

For example, when a source record includes the conjugated verb `selling`, AI Search expands the indexed term to include the lemma form `sell` in addition to `selling`. When a user searches for the past-tense conjugated form `sold`, AI Search expands the search query term to include the lemma form `sell` as well as `sold`. Because the indexed term and the search query term include matching forms, the user's search returns the `selling` record as a result.

AI Search supports language-specific lemma normalization for Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese.

Decompounding

In addition to normalizing lemmas for German, Korean, and Swedish, AI Search indexes compound words and their individual component words. For example, when indexing a German record that contains the compound word `Humanressourcen`, AI Search indexes the component terms `Human` and `ressourcen` in addition to the compound term.

Unicode normalization

AI Search performs Unicode normalization on indexed terms and search query terms. This normalization makes alphabetical Unicode glyphs searchable using their nearest equivalent characters.

For example, when indexing a record containing the term `resumé`, AI Search expands the term to also include the non-accented form `resume`. This record appears as a search result when users search for either `resume` or `resumé`.

Unicode normalization includes NFKD (compatibility decomposition) and NFKC (compatibility composition) stages. For more information on these normalization forms, see the Unicode Standard Annex #15, https://www.unicode.org/reports/tr15/.

Interaction with other search features

The following table describes interactions between normalization and other search features.

Feature	Interaction with lemma and Unicode normalization
Genius Results	Search query terms added by lemma or Unicode normalization can't trigger Genius Result configurations with Term trigger conditions.
Result improvement rules	A search query term added by lemma or Unicode normalization can trigger a result improvement rule if it matches the rule's Query trigger.
Stop words	If a search query term is defined as a stop word, AI Search removes that term without normalizing it.
Synonyms	If a search query term is defined as a synonym, AI Search doesn't normalize it.
Typo handling	AI Search performs lemma and Unicode normalization on auto-corrected search query terms.

Internationalization support for AI Search

AI Search supports indexing and search in all languages offered by the ServiceNow AI Platform®. Search linguistic features are supported in Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese.

Internationalization support is automatically enabled and isn't configurable.

To view the full list of languages offered as ServiceNow AI Platform plugins and supported in AI Search, see Activate a language.

Note:
After you activate a new language plugin, you must reindex all indexed source content that you want to make searchable in the new language. For details on reindexing, see Perform a full table index or reindex for a single indexed source.

Language settings determine how AI Search separates the text of indexed content and search queries into individual terms. This process, called tokenization, is handled differently for each supported language, using language-specific settings. For example, most languages use spaces and punctuation to separate words and sentences, but when tokenizing Chinese or Japanese text, AI Search instead uses contextual interpretation to correctly identify word and sentence breaks. When tokenizing Japanese text, AI Search additionally recognizes the nakaguro (middle dot) as a word separator.

Note:
If you indexed content in Brazilian Portuguese, Dutch, Italian, Japanese, Korean, Portuguese, or Swedish prior to August 2024, you should reindex it to benefit from new tokenization improvements for these languages.

Indexing behavior in supported languages

When indexing content and metadata from a ServiceNow AI Platform source record or an external document, AI Search uses tokenization settings for the language of the record or document, as shown in the following table.

Record or Document	Tokenization Settings
Source record from the Task [task] table or one of its child tables	AI Search performs language identification and uses tokenization settings for the detected language to index the record's content and metadata. **Note:** Language identification only identifies Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese. Content in other languages is identified and treated as English.
Source record from a non-Task table	AI Search uses tokenization settings for the record's language to index its content and metadata. If the record has no language specified, the ServiceNow AI Platform treats it as being in the instance's default language. In an English instance, for example, AI Search indexes records without specified languages using tokenization settings for English.
External document	AI Search performs language identification and uses tokenization settings for the detected language to index the document's content and metadata. **Note:** Language identification only identifies Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese. Content in other languages is identified and treated as English.

Note:
When indexing content and metadata, AI Search recognizes regions of Japanese, Simplified Chinese, and Traditional Chinese text embedded within text in other languages. These text regions are indexed with the appropriate language tokenization settings regardless of the surrounding text's language. As an example, suppose you index an English-language Knowledge article that includes a paragraph of Simplified Chinese. AI Search indexes this paragraph's content as Simplified Chinese and the rest of the record's content as English.

Search query behavior in supported languages

When processing search query text, AI Search uses tokenization settings for the language of the current user's ServiceNow AI Platform session.

Note:
AI Search recognizes Japanese, Simplified Chinese, and Traditional Chinese terms in search queries. These terms are processed with the appropriate language tokenization settings regardless of the user session's language. As an example, if a user in a French user session searches for `remplacement ordinateur 笔#本##`, AI Search applies Simplified Chinese settings for the `笔#本##` term and French settings for the other search terms.

AI Search compares your search query terms with terms from indexed content and metadata, returning search results for indexed records or documents that contain matches. When your search terms are in the same language as the indexed terms, AI Search processes both sets of terms with the same tokenization settings, producing predictable matches and search results. If your search terms aren't in the same language as the indexed terms, AI Search processes the two sets of terms with different tokenization settings and matching may be unpredictable.

Language dependence for search features

The following search features are language-dependent and supported only for the listed languages.

Search feature language dependence

Feature	Language dependence and supported languages
Genius Results	AI Search only evaluates Genius Result configurations with NLU triggers if the linked NLU model has the same language as the search query. Supported languages: English.
Language identification and tokenization	During indexing, AI Search identifies supported languages in Task table records and external documents. Text processing for the indexed content uses tokenization settings for the identified language. Supported languages: Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean,

Search feature language dependence (continued)

Feature	Language dependence and supported languages
	Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese.
Lemma and Unicode normalization	AI Search performs language-specific lemma normalization for terms in indexed content and search queries. Supported languages: Brazilian Portuguese, Dutch, English, French, French - Canada, German, Italian, Japanese, Korean, Portuguese, Simplified Chinese, Spanish, Swedish, and Traditional Chinese. **Note:** For German, Korean, and Swedish, AI Search performs term decompounding in addition to lemma normalization. AI Search performs Unicode normalization for all terms in indexed content and search queries. For more information on normalization of lemmas and Unicode forms in indexed content and search queries, see Lemma and Unicode normalization.
Result improvement rules	AI Search only evaluates activation for result improvement rules that have the same language as the search query or that have **All Languages** specified. Supported languages: All languages activated in your instance. For the list of languages you can activate, see Activate a language.
Stop words	AI Search only considers stop words from dictionaries that have the same language as the search query. Supported languages: All languages activated in your instance. For the list of languages you can activate, see Activate a language.
Synonyms	AI Search only considers synonyms from dictionaries that have the same language as the search query.

Search feature language dependence (continued)

Feature	Language dependence and supported languages
	Supported languages: All languages activated in your instance. For the list of languages you can activate, see Activate a language.
Typo handling	AI Search derives a separate list of auto-correction terms for each supported language found in search source indexed content. Auto-correction only replaces search query terms with terms from the list that has the same language as the search query. Supported languages: Brazilian Portuguese, Dutch, English, French - Canada, French, German, Italian, Portuguese, Spanish, and Swedish. Typo handling isn't supported for Japanese, Korean, Simplified Chinese, or Traditional Chinese.

Searching translated content

Understand how AI Search searches translated fields and documents.

The ServiceNow AI Platform® includes two categories of translated content:

1. Translated fields, such as translations of Catalog Item record fields.
2. Translated documents (long text and HTML values), such as translations of Knowledge article records.

The default AI Search matching behavior differs for these two types of content.

Translated fields

When you search content that includes translated fields, AI Search matches and displays field translations in your ServiceNow AI Platform session language. If a field lacks a translation in the session language, AI Search instead falls back to matching against the English-language field value. For example, if your session language is Japanese, AI Search returns Japanese field translations where present on a Catalog Item record, but returns English values for Catalog Item fields that have no Japanese translation.

To learn more about how the ServiceNow AI Platform handles translated fields, see Translating individual UI strings.

If you insert or update field translation records in the Translated Text [sys_translated_text] table, reindex content from the relevant indexed source to make AI Search use the new translations. For details on reindexing content, see Perform a full table index or reindex for a single indexed source and Perform a full table index or reindex for multiple indexed sources.

Note:
AI Search doesn't support translated fields on the Knowledge [kb_knowledge] table. This behavior has the following effects.

- Search queries don't match translated versions of kb_knowledge translated fields.
- Search results don't display translated versions of kb_knowledge translated fields.
- Search application facets don't display translated versions of kb_knowledge translated fields.

Translated documents

When searching content with translated documents, by default AI Search only matches documents in the user's ServiceNow AI Platform session language. For example, a Japanese user search only matches Knowledge articles in Japanese, excluding all Knowledge articles in other languages.

Starting in the Utah release, you can configure a filter condition for globally searchable Knowledge articles. Articles that satisfy this condition can be matched by searches in any language. For information on this setting, see Configure globally searchable Knowledge articles.

Beginning in the Tokyo release, you can set a fallback language for each language enabled in your ServiceNow AI Platform instance. When a fallback language is configured for the user's session language, AI Search returns search results in both languages. For more details, see AI Search interaction with fallback languages.

You can configure a global fallback locale to make AI Search match documents in the instance locale's language as well as the user's session language. This configuration works best when you want to make all English-language documents globally searchable. To learn more about configuring a global fallback locale, see Enable a global fallback locale for translated content.

Assign Knowledge article search languages by country

Improve multilingual search recall by defining a set of Knowledge article search languages for each user country. User searches can match Knowledge articles in any of the search languages specified for their country as well as in the language associated with their ServiceNow AI Platform® session.

Before you begin

Role required: ais_admin

About this task

By default, user searches only match Knowledge articles in the language associated with the user's ServiceNow AI Platform session. As an example, a user searching from a Japanese session only sees search results for Japanese-language Knowledge articles.

Search administrators can expand the recall of Knowledge article searches for a country's users by specifying additional Knowledge article search languages for that country. For example, suppose you specify the Knowledge article search languages English and French for the country of Japan. With these country-to-language mappings in place, a user searching from a Japanese session can see search results for Knowledge articles in English and French as well as for articles in Japanese.

Procedure

1. Navigate to the AI Search Country To Search Language [ais_country_to_search_language] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `ais_country_to_search_language.list`.
   c. Press Enter.
2. For each Knowledge article search language that you want to map to a country, create a new record in the table.
   a. Select **New**.
   b. On the AI Search Country To Search Language form, fill in the fields.
      For a description of the field values, see AI Search Country To Search Language form.
   c. Select **Submit**.

   Note:
   You can define multiple AI Search Country To Search Language records for the same country, as long as each one has a different Knowledge article search language. If you attempt to create multiple records with the same country and the same Knowledge article search language, the system displays a warning message and the record insert operation fails.

Result

When users in a country with one or more Knowledge article search languages defined perform a search, they can match Knowledge articles in those search languages as well in the language associated with their ServiceNow AI Platform session.

Configure globally searchable Knowledge articles

Define an encoded query for Knowledge articles that you want to make searchable in any language.

Before you begin

Role required: ais_admin

About this task

By default, AI Search only matches translated content records, such as Knowledge articles, that are in the user's ServiceNow AI Platform® session language. For example, a user searching from a Spanish user session only matches Spanish Knowledge articles. For full details on this restriction, see Searching translated content.

Note:
If you've configured a global fallback locale, search can match translated content records in that locale's language as well as in the user's session language. For more details on this approach, see Enable a global fallback locale for translated content.

To enable users to search for Knowledge articles in any language, you can define an encoded query as the value of the *`glide.ais.global_searchable_filter.kb_knowledge`* system property. Records that match this query are globally searchable, meaning they bypass the normal language filters and can be matched regardless of the search's language. As an example, a user searching from a German user session can match English or Japanese Knowledge articles as long as they match the encoded query.

Search result languages with globally searchable articles

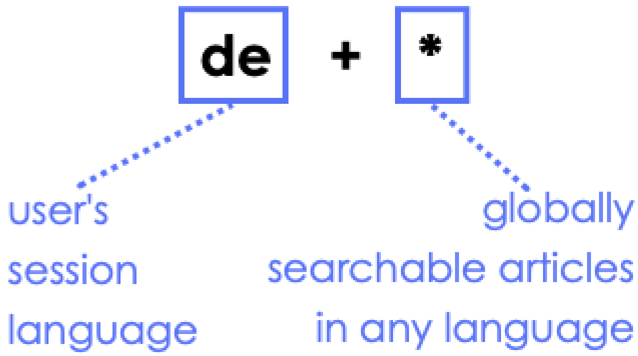

Note:
To generate Knowledge article query strings for your encoded query, you can navigate to **All** > *Knowledge* > **Articles** > **All** and follow the steps found in Generate an encoded query string through a filter. For descriptions of encoded query syntax and operators, see Encoded query strings and Operators available for filters and queries.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Open the *glide.ais.global_searchable_filter.kb_knowledge* system property record.
3. In the **Value** field, enter a filter expression specifying the Knowledge articles that you want to make globally searchable.
4. Select **Update**.

Result

Knowledge articles that match your encoded query can be returned as results for user searches in any language.

Example: Make Knowledge articles in the Email category globally searchable

Here is an example that defines the encoded query `kb_category=5681bf8bff0221009b20ffffffffff95` to make Knowledge articles in the Email category globally searchable.

Email category encoded query for Knowledge articles

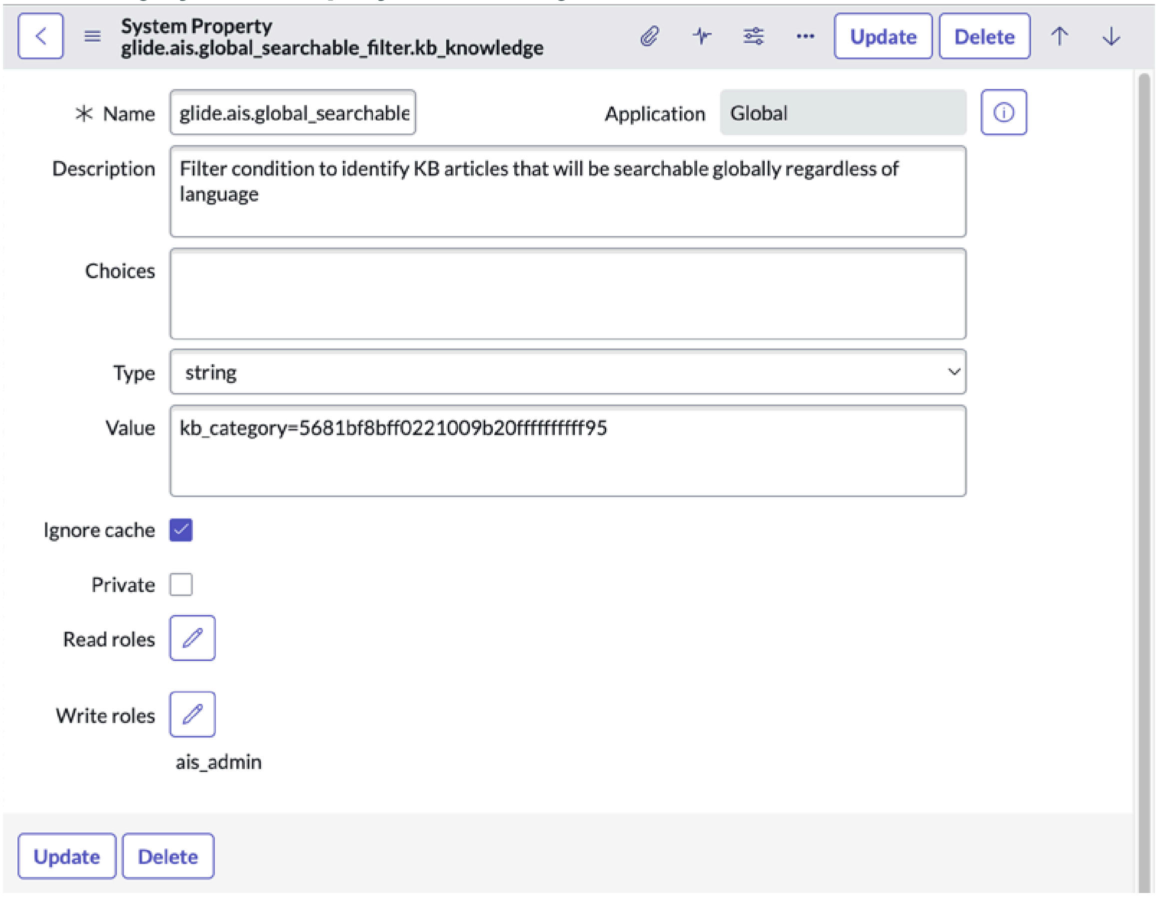

AI Search interaction with fallback languages

When a fallback language is configured for the user's session language, AI Search searches translated documents in both languages. Search uses the appropriate stop word, synonym, and typo handling dictionaries for the fallback language.

Starting in the Tokyo release, you can set a fallback language for each language enabled in your instance. The system uses this fallback language for elements of user interface text that aren't translated in the specified language. For details on configuring fallback languages, see Set a fallback language.

If you set a fallback language for the user's ServiceNow AI Platform® session language, AI Search returns results for translated content, such as Knowledge articles, in both the session language and your configured fallback language. When searching content in the fallback language, AI Search uses the stop word, synonym, and typo handling dictionaries for the fallback language.

As an example, suppose you configure Spanish as the fallback language for Mexican Spanish. With this configuration, users searching in Mexican Spanish sessions can find Knowledge articles in both Mexican Spanish and Spanish. AI Search uses the Mexican Spanish and Spanish dictionaries and result improvement rules when matching articles in Mexican Spanish. It uses the Spanish dictionaries and result improvement rules when matching articles in Spanish.

Search result languages with a fallback language

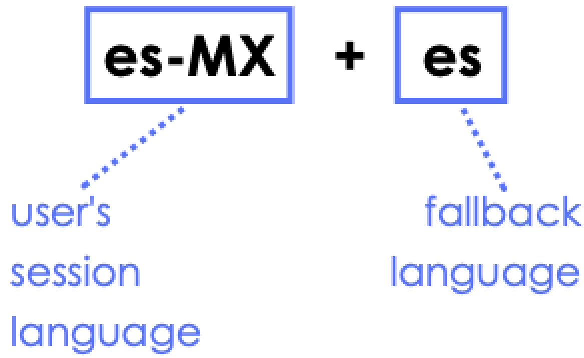

If you configure a global fallback locale in addition to a fallback language, AI Search returns matching results for translated content in the user's session language, the fallback language, and the global fallback locale's language. As an example, suppose you configure French as the fallback language for Canadian French, and also configure English as the global fallback locale. Users searching in Canadian French sessions can search Knowledge articles in Canadian French, French, and English.

Search result languages with a fallback language and a global fallback locale

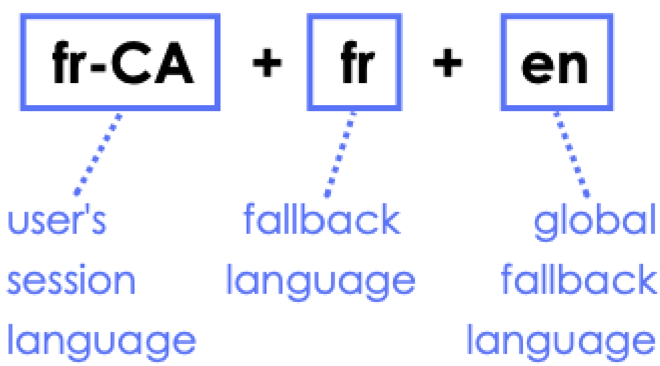

For details on configuring a global fallback locale, see Enable a global fallback locale for translated content.

Enable a global fallback locale for translated content

Configure AI Search to match translated documents in a global fallback locale's language as well as the user's session language.

Before you begin

Role required: ais_admin

About this task

By default, when searching content with translated documents, AI Search only matches indexed records in the user's ServiceNow AI Platform® session language. For example, a search in a Japanese user session only matches Knowledge articles in Japanese, excluding all articles in other languages.

You can configure AI Search to apply a global fallback locale to user searches. When you enable this option, AI Search matches records in the language from the ServiceNow AI Platform instance's locale (or another locale that you specify) as well as matching records in the user's

session language. As an example, if you set English as the global fallback locale, users searching in Japanese sessions can see Knowledge article results in both Japanese and English.

Search result languages with a global fallback locale

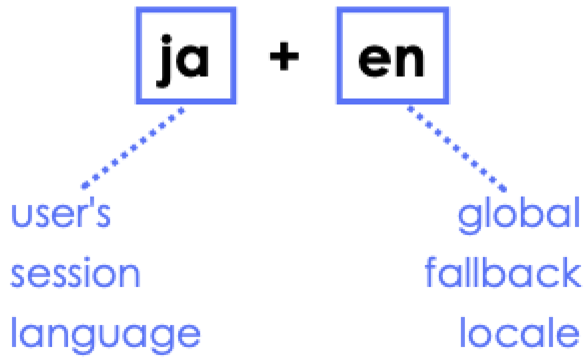

Note:
The global fallback locale works best when you want to make all English-language records globally searchable. To make other sets of records globally searchable, see Configure globally searchable Knowledge articles.

When you enable the global fallback locale option for search, AI Search determines the global fallback locale as follows:

1. If the value of the `glide.ais.translate.global_fallback_locale` system property is a valid locale string, such as `en-US` or `fr`, AI Search uses it as the global fallback locale.
2. Otherwise, AI Search uses the ServiceNow AI Platform instance's locale, as specified by the value of the `glide.system.locale` system property, as the global fallback locale.
3. If the `glide.system.locale` system property's value isn't a valid locale, AI Search uses English as the global fallback locale.

For more details on ServiceNow AI Platform locales and their associated languages, see Configuring locales.

Procedure

1. Navigate to **All** > **System Properties** > `AI Search`.
2. Select the option for the **Enable global language fallback** (`glide.ais.translate.enable_global_language_fallback`) system property.
3. Select **Save**.
4. **Optional:** Add the `glide.ais.translate.global_fallback_locale` system property and set its value to the locale that you want AI Search to use as the global fallback locale.
   For instructions on adding this system property, see Add a system property.

Result

AI Search searches translated document content in the global fallback locale's language as well as the user's session language.

Note:
When searching the Catalog Item table, AI Search ignores the instance locale. Catalog searches only match documents in the user's ServiceNow AI Platform session language.

Machine learning relevancy in AI Search

AI Search displays the most relevant search results for a query first. Machine learning automatically tunes search result relevancy scoring for search experiences based on aggregated user interactions.

Machine learning relevancy is automatically enabled and isn't configurable.

Relevancy models and scoring

AI Search uses a relevancy model to compute a relevancy score for each result returned by a search. Documents with higher relevancy scores appear first in the result set. A result's relevancy score is specific to the particular document, search terms, and user associated with the query.

Each search profile includes its own relevancy model. You can't view, modify, or delete this relevancy model.

Note:
AI Search doesn't apply relevancy ranking to `* * *` universal wildcard queries. Results from `* * *` queries appear in an unspecified order.

Search signals and machine learning relevancy tuning

AI Search UX components record signals associated with user searches. These search signals include data on how search users interact with the search input field, auto-complete suggestions, facet and navigation tab filters, Genius Result answer cards, and search results. To learn more about how search signals are recorded and stored, see Search signals.

Machine learning relevancy uses data from these search signals to intelligently tune relevancy models on a continual basis. Every 30 days, AI Search computes a new version of each relevancy model, iteratively modifying its parameters and regression testing it against aggregated search signal data for the search profile. When this tuning process is complete, AI Search compares the existing and new relevancy models to see which one produces better matches for user search behavior as recorded in the historical signal data.

If the new relevancy model produces better results with the signal data, AI Search uses its modified parameter values to perform A/B testing evaluations of live search traffic for the search profile. These evaluations test individual parameter changes to verify that they produce better search relevancy.

Note:
For details on the search query parameter evaluation framework used to perform A/B testing evaluations, see Search query parameter evaluation framework.

If the new model outperforms the original model in both the historical search-match comparison and the A/B testing, AI Search sets it as the active relevancy model for the search profile, overwriting the existing relevancy model. The updated relevancy model remains in use until the next tuning cycle begins.

These relevancy model tuning processes occur separately for each search profile. Changes made to the relevancy model in one search profile don't affect relevancy models in other search profiles.

Note:
When you upgrade to Washington DC from a previous release, the default relevancy scores for your search results may change. Relevancy models trained in the previous release should continue to produce the same result ordering. Models trained more than one release ago may revert to the default relevancy model.

Relevancy model for auto-complete suggestions

AI Search uses a dedicated relevancy model to rank records for display as auto-complete suggestions in the search field. This relevancy model scores records based on their freshness and on search query term matches in their **title** fields. The system doesn't train this auto-complete suggestion relevancy model. For details on configuring auto-complete suggestions, see Auto-complete suggestions in AI Search applications.

Viewing relevancy scores for search results

Search administrators can view the scores for search results in the Search Preview UI from the Advanced AI Search Management Tools ServiceNow® Store application. For details on using this feature to investigate search behavior, see Search Preview UI for AI Search.

Content security in AI Search

AI Search filters search query results and displays only records that the currently logged in user can access.

Content security is automatically enabled and isn't configurable.

Indexed security features

When indexing records from ServiceNow AI Platform® tables, AI Search preserves the following access control settings:

- Role-based field-level access control list rules (ACLs)
- Non-scripted conditional ACLs
- Scripted table-level ACLs
- Before Query business rules
- Domain separation
- User criteria for records indexed from the Knowledge [kb_knowledge] and Catalog Item [sc_cat_item] tables

AI Search doesn't support scripted field-level ACLs or conditional field-level ACLs.

Early and late binding security implementations for search queries

At query time, AI Search filters search results by evaluating user access to indexed tables and records. The filtering implementation used depends on the security features present on tables and records that match the search query.

Content security implementation	Description
Early binding	Default content security implementation that incorporates security filters into the search query. The system adds relevant security filters for the user to each search query. To appear in the search results, records must match the search query and pass all security filters. Early binding security supports non-scripted ACLs, Before Query business rules, and

Content security implementation	Description
	domain separation, plus user criteria for records indexed from the kb_knowledge and sc_cat_item tables. If a search matches records that have table-level scripted ACLs, the search query automatically falls back to use late binding security. **Note:** If an early binding filter produces an error, the search query automatically falls back to use late binding security.
Late binding	More computationally expensive alternate content security implementation that post-filters search query results. The system applies security filters to the search query just as early binding security does. After retrieving records that match the filtered query, it invokes the `GlideRecord.canRead()` method for each record. To appear in the search results, records must match the search query, pass all security filters, and return `true` for the API method. Late binding security supports all indexed security features. By default, late binding security serves as a fallback mechanism for early binding security. You can override this default behavior in the following ways: • To make AI Search use late binding security for all results from an indexed source, set the value of that indexed source's **Force Late Binding** field to `true`. • To make AI Search use late binding security for all search results, set the value of the **Force** `AI Search` **to use late binding for ACL evaluation** (`glide.ais.security.force_late_binding`) system property to `true`. If AI Search applies late binding security to any results from an indexed source used by a facet, the list of available filter selections for the facet changes. When this happens, the facet only displays field values found in the current page of search results. For details on facets, see Create a facet in an AI Search application configuration.

Content security implementation	Description
	When you enable search result counts for a search application, the counts indicate the number of results that matched the user's search. Late binding security can remove results after they're counted, causing the actual number of displayed results to be lower than the count. For more information on result counts, see Display search result counts on the results page for a search application.

Content security for external content

AI Search preserves access permissions when ingesting external documents. You can alias ServiceNow AI Platform users to the externally defined users and groups specified in these access permissions. AI Search constructs content security filters for external document records from these permissions and aliases.

To learn about using AI Search with external data sources, see Indexing and searching external content in AI Search. For full details on external content security, see External content security for AI Search.

Encrypted fields in AI Search

AI Search excludes source record fields encrypted with Encryption Support or Edge Encryption.

When dealing with an encrypted source record field, AI Search has the following limitations:

- Content indexing excludes the encrypted field. This exclusion means that search query terms won't match terms from the encrypted field's content.
- You can't create a facet using the encrypted field.
- If a search query attempts to filter on the encrypted field, such as when applying an access control list (ACL) filter, the search query fails. An `Encrypted Field is not supported` error message is logged.

For detailed descriptions of the encryption options available for source records, see Encryption.

Semantic vector search in AI Search

Semantic vector search allows the Now LLM Service to find results based on how similar they are in meaning to your search terms. Now Assist Q&A Genius Results and Now Assist in Virtual Agent use semantic vector search to improve recall with natural language interpretation of search queries.

By default, AI Search uses lexical search, meaning that it finds results for records that contain literal keyword matches for the terms in your search query. Starting in the Vancouver Patch 4 release, AI Search includes an alternate search mode called semantic vector search that's used in Now Assist Q&A and Now Assist Actions Genius Results.

Semantic vector search analyzes the meanings and context of your search terms and uses that information to find results with similar meanings. It improves search recall by interpreting natural language to more accurately reflect your search's intent.

As an example, suppose you index a source record including text `how to prevent phishing` and search for `avoiding scams`.

- In lexical search mode, AI Search returns no result for this record because your search query terms aren't literal matches for the terms included in the indexed content.
- In semantic vector search mode, however, AI Search might return a result for the record based on the following context-aware analysis:
  - Your search term `avoiding` is contextually similar in meaning to `prevent`
  - Your search term `scams` has meaning that overlaps with the meaning of `phishing`

AI Search orders results from semantic vector search based on how similar they are to your search. In the example, if you had a second indexed record including text `how to prevent scams`, it would be more similar to the `avoiding scams` search and its search result would appear before the `how to prevent phishing` result.

Unlike the default keyword-based search mode, semantic vector search does not rely on your synonym dictionaries to find term equivalences. AI Search identifies term similarities automatically when indexing source content and metadata for semantic vector search.

Semantic vector search is automatically enabled and doesn't include any configurable settings.

Benefits of semantic vector search

In comparison with the default keyword-based lexical search mode, semantic vector search provides the following benefits:

- Helps improve recall for Knowledge article, Catalog Item, and topic retrieval by matching results based on the underlying meaning of your search rather than matching on keywords.
- Provides a more natural and consistent search experience by returning similar results for varied search queries with the same underlying meaning.
- Reduces the need to create and maintain synonym dictionaries to anticipate similarities in searches.

Semantic vector search output

Semantic vector search overrides the normal AI Search term matching and relevancy ranking mechanisms. Results are ordered based on their computed similarity to your search query terms, with the most similar results appearing first.

Availability of semantic vector search

Semantic vector search is available in the following contexts.

- Now Assist in Virtual Agent uses semantic vector search for Catalog Item retrieval and live agent topic retrieval. For more details, see Now Assist in Virtual Agent.
- Now Assist Q&A Genius Results use semantic vector search along with legacy keyword search when looking for Knowledge articles that match your search query. They also use semantic vector search when looking for cached answers that match your query in the second-level cache. For more details, see Now Assist Q&A Genius Results and Caching for Now Assist Q&A Genius Results.
- Now Assist Actions Genius Results use semantic vector search along with legacy keyword search when looking for Catalog Items that match your search query. For more details, see Now Assist Actions Genius Results.

Semantic vector search only supports searches in English.

AI Search logging and debugging

Review AI Search logs and history to analyze performance and diagnose issues. View log messages and ingestion history, or use session debugging to capture detailed messages for a user session.

View AI Search ingestion log messages

Review AI Search ingestion and indexing events by viewing log messages.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Logs** > **Ingestion Warnings & Errors**.
   The Log [syslog] table appears with a **[Source] [starts with] [AIS-Datasource]** filter condition applied.
2. Use the list search, sort, filter, and group functions to locate messages of interest.

View all AI Search log messages

Review AI Search events by viewing log messages.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Logs** > **All**.
   The Log [syslog] table appears with a **[Message] [contains] [ais]** filter condition applied.
2. Use the list search, sort, filter, and group functions to locate messages of interest.

Review ingestion history for an indexed source

View processing results and performance statistics associated with an ingestion event for an indexed source.

Before you begin

Role required: ais_admin

About this task

The AI Search Indexed Source History [ais_ingest_datasource_stats] table contains records for indexed source ingestion events. Each record includes performance statistics and processing results for the ingestion event. You can review the ingestion event history record to analyze ingestion performance and diagnose issues.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **Index** > **Indexed Sources**.
2. Open the indexed source whose indexing history you want to review.
3. In the Indexing History related list, open the ingestion event record that you want to review.
4. On the Indexed Source History form, review the fields for the selected ingestion event.
   For a description of the field values, see Indexed Source History form.

Enable session debugging for AI Search

Capture and review detailed AI Search log messages for a user session.

Before you begin

Role required: admin

About this task

When you enable session debugging for AI Search in a user session, the system records debug log messages for that session. These messages include:

- Messages illustrating how AI Search interprets a user's search query
- Messages specifying the search source filters and content security filters applied to a search query
- Messages containing search query request and response objects in human-readable format

You can review these debug log messages to diagnose search issues or configure optimal search performance settings.

Once enabled, session debugging remains active for the user session until you disable it. For more information on session debugging, see Session debug.

Procedure

1. To enable debugging, navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug** `AI Search`.
   A message reports completion of the Debug AI Search session debug script.
2. To review debug logs, navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug Log**.
   Debug log messages appear in the **Session Log** window.
3. To disable debugging, navigate to **All** > **System Diagnostics** > **Session Debug** > **Disable All**.
   A message reports completion of the Disable All session debug script.

Diagnose search result access issues using the Search Preview UI

Use the Search Preview UI's impersonation tool and the AI Search High Security administrator elevated privilege role to determine whether filters prevent a search user from viewing a specific search result in the AI Search index.

Before you begin

Role required: ais_admin, ais_high_security_admin, and impersonator.

About this task

Use the Search Preview UI to diagnose filter configurations that prevent a search user from seeing a specific search result in the AI Search index. First, search as that user in the Search Preview UI to confirm whether the search result is visible. Next, end impersonation and repeat the search after elevating to the ais_high_security_admin to bypass all search source and content security filters. If the search result appears in the elevated search, examine its search source and content security filter settings to determine why the user can't view it.

> **Note:**
> Changing the value of the **Search as** field in the Search Preview UI's User impersonation tool terminates elevated privileges for the ais_high_security_admin role. To bypass filters after changing this field's value, you must re-elevate to the ais_high_security_admin role.

Procedure

1. Navigate to **All** > `AI Search` > **Preview** > **Search Preview**.
2. Use the **Search as** field in the Search Preview UI's User impersonation tool to issue the search query of interest as the affected user.

Note:
For more details on the **Search as** field, see Search preview admin tools.

Results appear, matching those seen by the affected user for the search query of interest. These results may be limited by search source or content security filters.

3. Elevate to the ais_high_security_admin role:

   a. Perform the appropriate action for your version of the UI:

UI version	Action
Next Experience UI	In the banner frame, select the icon for your account to open the user menu, then select **Elevate role**. 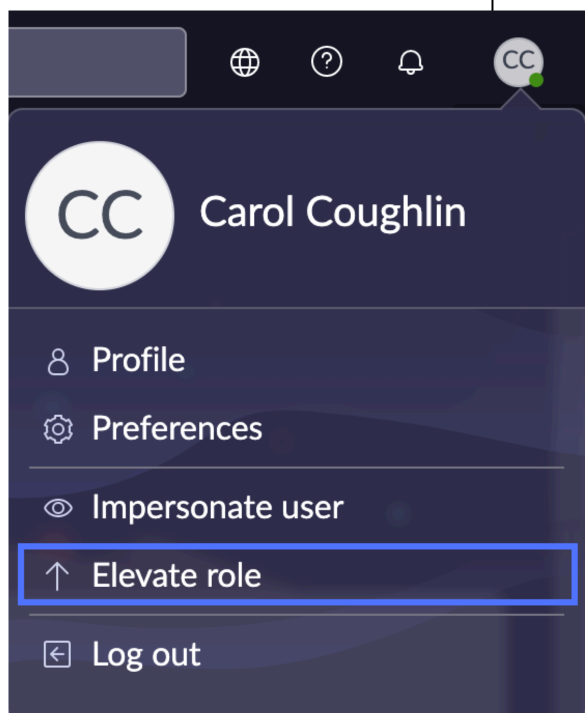
Core UI	In the banner frame, select your name to open the user menu, then select **Elevate Roles**. 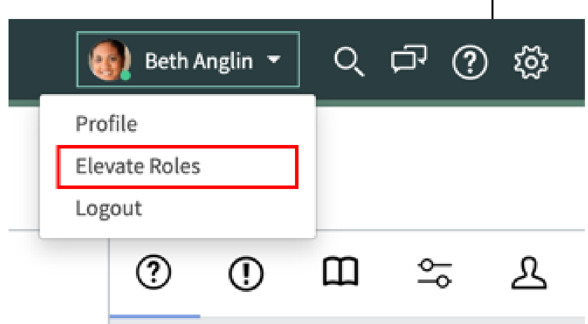

A dialog box appears, displaying a checklist of your available privileged roles.

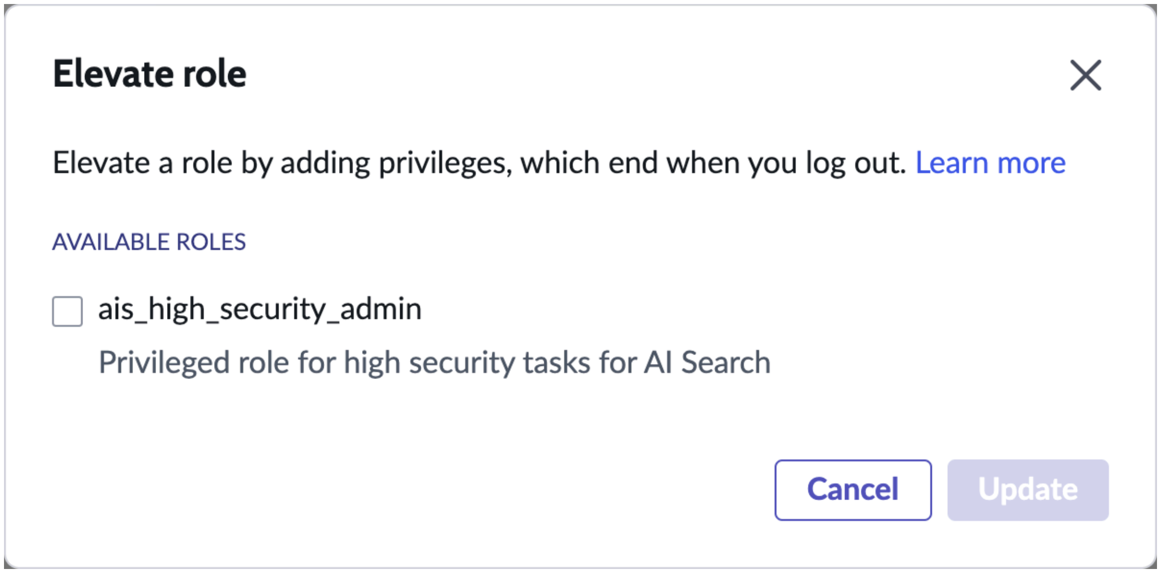

b. In the dialog box, select the **ais_high_security_admin** option, then select **Update** (in Next Experience UI) or **OK**.
The page reloads and an elevated role indicator appears next to your user name in the user menu. In Next Experience UI, this indicator displays the names of the active privileged roles.
In Core UI, the indicator displays the elevated role icon ().

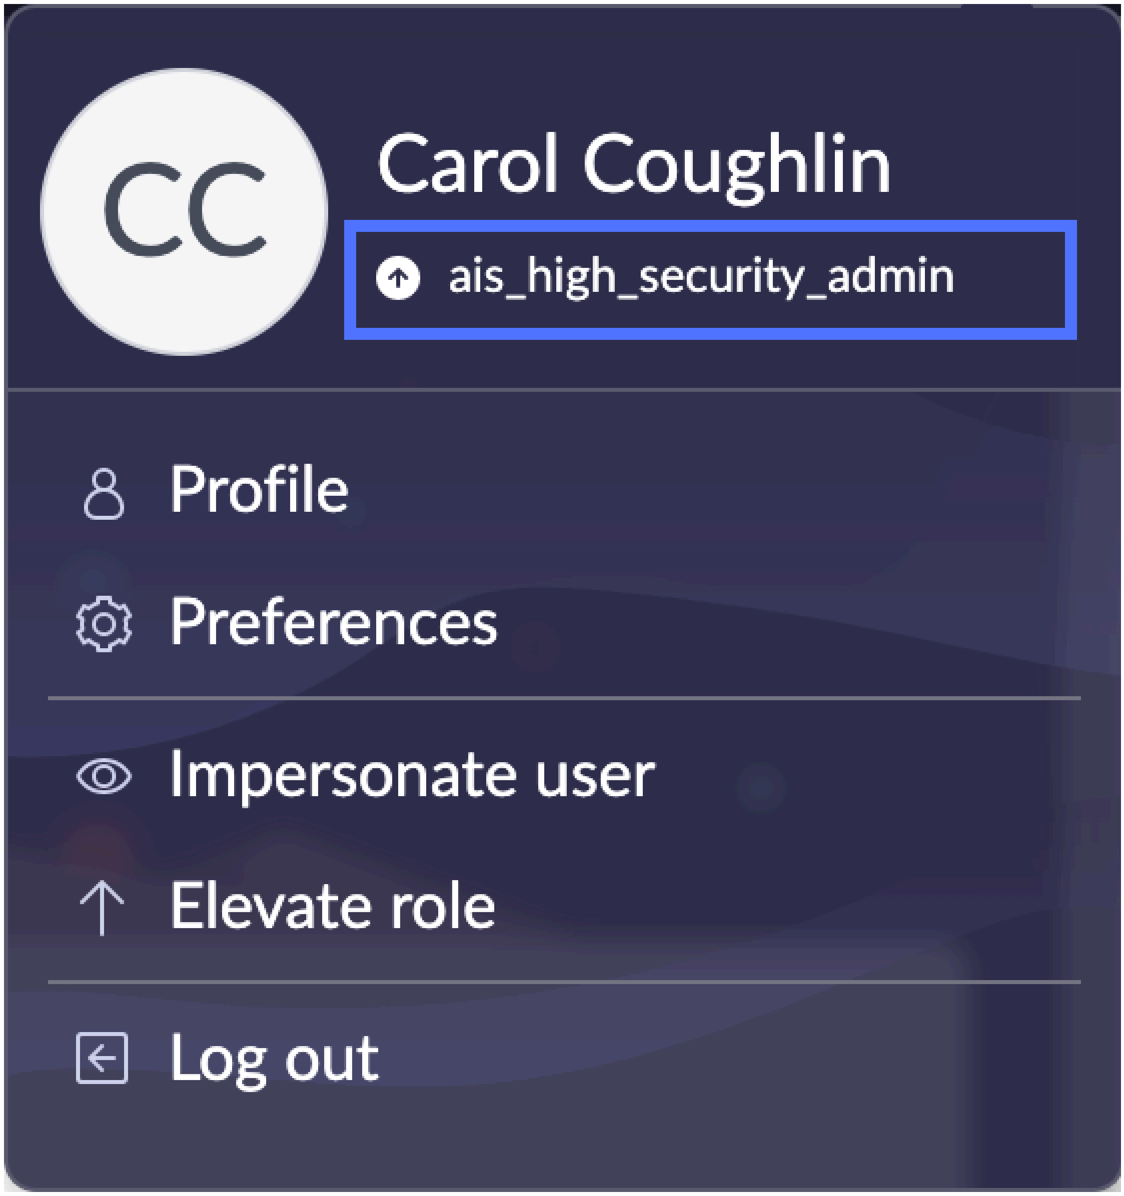

Note:
When the page reloads, any unsaved edits are lost.

4. Repeat the search of interest without impersonating the affected user.
   Search results appear with no search source or content security filters applied.

What to do next

Compare the two search result sets. If the search result in question appears only in the elevated search, a search source or content security filter prevents the affected user from viewing that result. If the search result doesn't appear in either search, it's not present in the AI Search index.

If the search result in question is indexed from an external document, you can debug the external content security access permission filter for the affected user to see whether it prevents access to the search result. For details on this procedure, see Debug external content security filters for a ServiceNow AI Platform user.

Debug search source and content security filters

Use the Session Log to view search source and content security filters applied to results for a search query. Administrators can use this approach to verify search source filters and content security settings.

Before you begin

Role required: admin

About this task

Enable session debugging for AI Search to view debug log messages for search source and content security filters. Each debug log message shows the filters that AI Search applies to determine which records appear in a search query's results.

Procedure

1. To enable session debugging for AI Search, navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug** `AI Search`.
2. Navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug Log**.
   The Session Log opens.
3. Use the Search Preview UI for AI Search to submit the search query of interest.

   Note:
   To debug content security filters for another user, specify their name in the Search Preview UI's user impersonation tool.

4. To see the search source filters applied for your search queries, filter the Session Log entries for `AIS:search sources`.

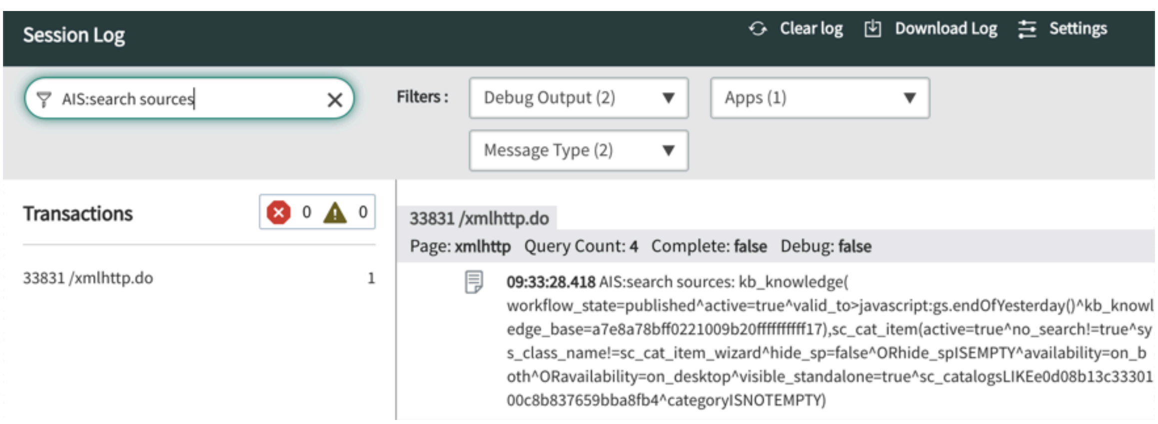

5. To see the content security filters applied when evaluating access to search results for the current (or impersonated) ServiceNow AI Platform® user, filter the Session Log entries for `security query`.

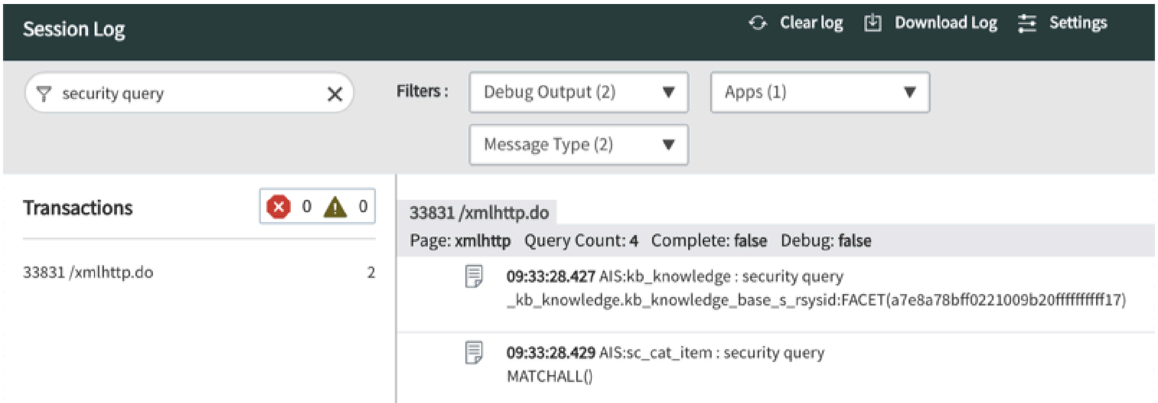

What to do next

If a search source's filters exclude records that should be searchable, modify those filters. For details, see Create a search source.

If content security filters prevent the user from seeing search results that should be accessible, correct the access control settings for the affected source records and then reindex content for the relevant indexed source. For details on content security settings, see Content security in AI Search. To reindex content, follow the steps from Perform a full table index or reindex for a single indexed source.

Debug external content security filters for a ServiceNow AI Platform® user

Use the Session Log to view the set of externally defined user and group access permission filters associated with the current ServiceNow AI Platform user. Administrators can use this approach to verify that the proper user mappings exist for the current user.

Before you begin

Role required: admin

About this task

Enable session debugging for AI Search to view debug log messages for external content security. Each debug log message shows the external content security filter that AI Search applies to determine whether the current ServiceNow AI Platform user can see a search result for an external document.

Procedure

1. To enable session debugging for AI Search, navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug** `AI Search`.
2. Navigate to **All** > **System Diagnostics** > **Session Debug** > **Debug Log**.
   The Session Log opens.
3. Use the Search Preview UI for AI Search to submit the external content search query of interest.

   Note:
   To debug external content security access permissions for another user, specify their name in the Search Preview UI's user impersonation tool.

4. In the Session Log, filter for `AIS:External Content Security Filter` to see the external content security filters applied when evaluating access to an external document search result for the current (or impersonated) ServiceNow AI Platform user.

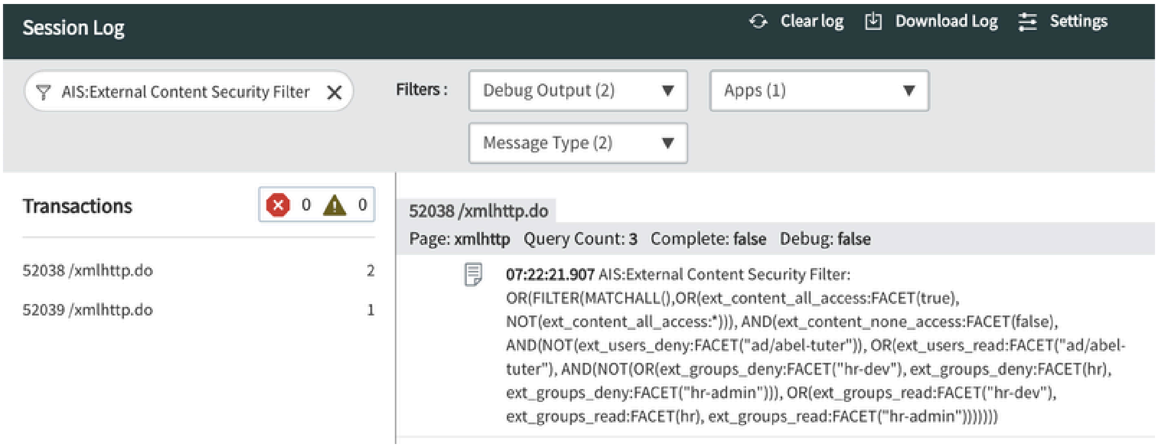

What to do next

If the external content security filters don't include the desired external user or group access permissions, a search administrator can import an updated and corrected user mapping for the relevant ServiceNow AI Platform user. For more details on defining and importing user mappings, see Mapping external users and groups to ServiceNow AI Platform users.

Review record counts for indexed sources

Understand where your indexed content originates by viewing record counts for your indexed sources in the AI Search Indexed Source Statistics table.

Before you begin

Role required: ais_admin

About this task

The AI Search Indexed Source Statistics [ais_datasource_stats] table contains entries showing the indexed record counts for each of your indexed sources. These entries are logged every 24 hours by the **AIS Collect Ingestion Stats** scheduled job. You can review the statistics entries to see how many records each indexed source contributes to the AI Search index.

Procedure

1. Navigate to the AI Search Indexed Source Statistic [ais_datasource_stats] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `ais_datasource_stats.list`.
   c. Press Enter.
2. If the Updated field doesn't appear, select the Personalize List icon ( ) and use the slushbucket to add the Updated field to the list view.
3. To display the most recent indexed record-count entries first, select the Updated field's column options icon (in Next Experience UI or), then select **Sort (z to a)**.
4. Review the indexed record-count entries for each of your indexed sources.

   Note:
   The entry for an indexed source indicates the number of indexed records from that source when the **AIS Collect Ingestion Stats** scheduled job ran at the time listed in the Updated field.

Search query parameter evaluation framework

AI Search performs A/B test evaluations that compare result relevancy for alternate sets of search query parameter values. These evaluations determine the models that AI Search uses for machine learning relevancy and for Q&A Genius Results.

The search query parameter evaluation framework enables A/B relevancy testing on live search traffic. AI Search considers results from this live testing along with its offline evaluation of aggregated search signal data when making the following determinations:

- Determining which relevancy model to publish for a search profile. For details on generation and publication of relevancy models, see Machine learning relevancy in AI Search.
- Determining which machine reading comprehension (MRC) model to use when validating potential Q&A Genius Result answers. For more information on the MRC model used for Q&A Genius Results, see Q&A Genius Results.

AI Search computes A/B testing evaluation results on a nightly basis.

The search query parameter evaluation framework is part of the Adaptive Telemetry suite of features.

Search query parameter evaluation framework tables

The search query parameter evaluation framework for AI Search uses tables to store results and metrics from A/B testing of live search traffic. AI Search uses the stored data when determining which relevancy model and Q&A Genius Result answer validation model to publish for a search application.

A/B Testing Evaluation table

The A/B Testing Evaluation [evaluation] table stores details for A/B testing evaluations performed on live search traffic.

Note:
The system automatically cleans this table, removing records for inactive evaluations that are more than two years old.

evaluation

Column	Description
Number	Autogenerated number to identify the evaluation.
Name	Name and description for the evaluation.
Source	Name of the table containing records to run the evaluation for.
Testing Scope	Scope of the evaluation. • Supported values: ◦ *ALL*: Run the evaluation against all records in the table specified by **Source**. ◦ *SELECTED*: Only run the evaluation against records specified in the **Selected Artifacts** field value. • Default value: *ALL*

evaluation (continued)

Column	Description
Artifact Provider	List of artifact filtering classes that provide records for testing.
Selected Artifacts	Comma-separated list of record sys_id values. When **Testing Scope** is set to **SELECTED**, the evaluation only runs against records with sys_id values specified in this list.
Minimum Signals for Evaluation	Minimum number of search signals the evaluation should collect before ending. • Type: integer • Default value: *0* (no minimum) • Maximum value: *30000* For details on how this condition interacts with the **Minimum Days for Evaluation** condition, see **End Criteria Operator**.
Minimum Days for Evaluation	Minimum number of days the evaluation should be active before ending. • Type: integer • Default value: *0* (no minimum) • Maximum value: *180* For details on how this condition interacts with the **Minimum Signals for Evaluation** condition, see **End Criteria Operator**.
End Criteria Operator	Logical operator to use for evaluation end conditions when both **Minimum Signals for Evaluation** and **Minimum Days for Evaluation** conditions exist. • Supported values: ○ *AND*: The evaluation remains active until the **Minimum Signals for Evaluation** and **Minimum Days for Evaluation** conditions are both satisfied. ○ *OR*: The evaluation remains active until either of the **Minimum Signals for Evaluation** or **Minimum Days for Evaluation** conditions is satisfied. • Default value: *OR*

A/B Testing Evaluation Execution table

The A/B Testing Evaluation Execution [evaluation_execution] table stores details of individual operations executed as part of A/B testing evaluations for live search traffic.

evaluation_execution

Column	Description
Number	Autogenerated number to identify the evaluation operation.
Evaluation	Reference to the record for the evaluation on the A/B Testing Evaluation [evaluation] table.
Artifact ID	Sys_id for the record analyzed by this evaluation operation.
State	State for the record evaluation execution. • Supported values: ◦ *Queued*: The evaluation operation is queued for execution. ◦ *In Progress*: The evaluation operation is underway. ◦ *Scoring*: The evaluation operation is in the scoring phase. ◦ *Complete*: The evaluation operation completed successfully. ◦ *Errored*: The evaluation operation failed with an error. ◦ *Canceled*: The evaluation operation was canceled by the system. ◦ *Suspended*: The evaluation operation was suspended by the system. • Default value: *Queued*
Start Date	Date and time when the evaluation operation started.
End Date	Date and time when the evaluation operation ended.

A/B Testing Evaluation Parameter table

The A/B Testing Evaluation Parameter [evaluation_parameter] table stores details of individual search query parameters considered in A/B testing evaluations for live search traffic.

evaluation_parameter

Column	Description
Evaluation	Reference to the record for the evaluation on the A/B Testing Evaluation [evaluation] table.

evaluation_parameter (continued)

Column	Description
Name	Name and description for the evaluation parameter.
Use Artifact's Assigned Parameter	• Type: true | false • Supported values: ○ *true*: Compare the specified *Parameter Values* to the currently assigned search query parameter value on the provided artifact. ○ *false*: Don't compare the specified *Parameter Values* to the currently assigned search query parameter value on the provided artifact • Default value: *true*
Parameter Values Traffic Percentage	Percentage of search queries to apply the specified *Parameter Values* to during the A/B evaluation. • Type: integer • Default value: *0* • Maximum value: *100*
Parameter Values	JSON list of values to compare to the current search query parameter value when *Use Artifact's Assigned Parameter* is true.
Parameter Type	Type for the search query parameter. The selected value determines how the *Parameter Values* are used during and after an A/B evaluation. Supported values: • *Search Context Parameters*: Merge and override search profile context parameters with the *Parameter Values* for search query requests. When the evaluation ends, use the search query context parameter value set with the winning score to override or merge the relevant search profile context parameters. • *Search QnA Genius Result Configuration*: Apply the *Parameter Values* to Q&A Genius Result configurations used in search query requests. When the evaluation ends, update the relevant search profile to include the

evaluation_parameter (continued)

Column	Description
	Genius Result configuration parameter value with the winning score. • *Search QnA Model*: Apply the *Parameter Values* to Q&A Genius Result models used in search query requests. When the evaluation ends, update the relevant search profile to include the Q&A Genius Result model parameter value with the winning score. • *Search Relevancy Model*: Apply the *Parameter Values* as search relevancy models to use in search query requests. When the evaluation ends, update the relevant search profile to include the relevancy model parameter value with the winning score. • *Search REST Parameters*: Merge and override search profile query parameters with the *Parameter Values* for search query requests. When the evaluation ends, use the search query parameter value set with the winning score to override or merge the relevant search profile parameters.
Score Calculation Type	Form of calculation used to compute search query parameter value scores and find the best value. Supported values: • *Average Click Position*: The best search query parameter value is the one with the lowest average click position score. **Note:** AI Search computes the average click position score by dividing the sum of all selected search result ranks by the number of searches. The highest-ranked search result has rank 1, the next-highest has rank 2, and so on. As an example, if you perform two searches, selecting the first result in one case and the second result in the other, your average click position score is (1 + 2) / 2 = 1.5. • *Genius Result Helpfulness*: The best search query parameter value is the one with the highest helpfulness score. This calculation takes into account whether a

evaluation_parameter (continued)

Column	Description
	relevant Genius Result answer appeared in your search results. **Note:** The helpfulness score is a metric that indicates whether Genius Result answers were helpful in the context of your search. • *Helpfulness*: The best search query parameter value is the one with the highest helpfulness score. Unlike *Genius Result Helpfulness*, this calculation doesn't take into account whether a relevant Genius Result answer appeared in your search results.
Signal Provider	The provider for search signals needed to score the search query parameter. *Search Event Signal Provider* is the only supported value.

A/B Testing Evaluation Parameter Result table

The A/B Testing Evaluation Parameter Result [evaluation_parameter_result] table stores calculation results for individual search query parameters considered in A/B testing evaluations for live search traffic.

evaluation_parameter_result

Column	Description
Evaluation Execution	Reference to the record for the evaluation execution on the A/B Testing Evaluation Execution [evaluation_execution] table.
Parameter Evaluation	Reference to the record for the search query parameter on the A/B Testing Evaluation Parameter [evaluation_parameter] table.
Best Value	Best value for the search query parameter, as determined by the *Winning Score*.
Winning Score	Numeric score for the search query parameter, determined using the *Score Calculation Type*.
Score Metadata	Metadata from the score computation for the search query parameter.

Exclude a search profile from the search query parameter evaluation framework

Exclude a search profile from A/B testing evaluations of live AI Search traffic. This procedure prevents AI Search from using A/B testing evaluation results when publishing the search profile's search result relevancy model and its Q&A Genius Result answer validation model.

Before you begin

Role required: ais_admin

About this task

The search query parameter evaluation framework performs A/B testing evaluations of search configuration settings using live search traffic. By default, AI Search evaluates configuration settings for all search profiles.

Search administrators can exclude individual search profiles from the search query parameter evaluation framework. Excluding a search profile from the framework prevents AI Search from performing A/B testing evaluations for live search traffic that uses the excluded search profile.

Note:
When you exclude a search profile from evaluations, AI Search no longer considers evaluation results when updating the machine learning relevancy and machine reading comprehension (MRC) models for that search profile. As a result, relevancy scoring settings and Q&A Genius Result answer filtering settings for the search profile may be less reflective of your search traffic. For more information on how AI Search uses A/B evaluation testing results when publishing these models, see Machine learning relevancy in AI Search and Q&A Genius Results.

Procedure

1. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
2. Open the record for the search profile that you want to exclude from A/B testing evaluations of live search traffic.
3. If the Search Profile form doesn't already display the **Exclude from evaluation** field, configure the form layout to make the field visible.
   For details on configuring a form layout to make fields visible, see Configuring the form layout.
4. Select the **Exclude from evaluation** option.
5. Select **Update**.

Result

AI Search no longer performs A/B testing evaluations for traffic that uses the excluded search profile. Machine learning relevancy no longer updates the relevancy model for the search profile.

AI Search reference

These reference topics provide additional information about AI Search installed components, system properties, features, and forms.

Plugins installed with AI Search

AI Search adds the following plugin.

AI Search plugins

Plugin	Description
AI Search plugin (com.glide.ais)	Enables core functionality for AI Search. Installed in the base system.

AI Search plugins (continued)

Plugin	Description
	Note: This plugin isn't visible in the **All** > **System Applications** > **All Available Applications** > **All** and **All** > **System Definition** > **Plugins** listings. To verify that AI Search is installed and provisioned for use on your instance, navigate to **All** > `AI Search` > `AI Search` **Status**.

Components installed with AI Search

Several types of components are installed with activation of the AI Search plugin, including scheduled jobs, tables, and user roles.

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

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.

Scheduled jobs installed

Scheduled job	Description
Aggregate Health Data for AIS Partition	Collects health check data for AI Search query nodes connected to your instance.
AI Search Expire Rules	Marks result improvement rules as expired when their end date is in the past.
AI Search Relevancy Data Processor	Processes search signal data as part of search profile relevancy model tuning by the machine learning relevancy feature.
AIS Collect Ingestion Stats	Collects ingestion history statistics for indexed sources.
AIS Derived Dictionary Auto Refresh	Refreshes derived typo handling dictionaries when their refresh interval has elapsed.
AIS Index Event Processor	Processes indexing events for indexed sources.
AIS Index Event Queue Size Monitor	Monitors the size of the indexing queue during indexing of source records.
AIS Index Purge Stale Document	Optimizes the search index by deleting stale documents.
AIS Validate Publish ID and Auto Publish	Publishes search profiles and typo handling, stop word, and synonym dictionaries.

Scheduled job	Description
ML Trigger for AIS Search Relevancy	Triggers automatic tuning of search profile relevancy models by the machine learning relevancy feature.

Tables installed

Table	Description
AI Search ACL Overrides [ais_acl_overrides]	List of rules used by content security.
AI Search Active Table Ingestion Tracker [ais_active_table_ingestion_tracker]	List of active content ingestion events for indexed source tables.
AI Search Child Table [ais_child_table]	List of child tables defined for indexed sources.
AI Search Configuration Attribute [ais_configuration_attribute]	List of indexed source attributes available for use in indexed sources.
AI Search Connection [ais_connection]	List of AI Search query nodes connected to your instance.
AI Search Dictionary [ais_dictionary]	List of typo handling, stop word, and synonym dictionaries.
AI Search Dictionary Term [ais_dictionary_term]	List of terms defined in typo handling, stop word, and synonym dictionaries.
AI Search External Content User Mapping Base [ais_external_search_user_mapping]	Empty list of external content user mappings. All custom external content user mapping tables extend this table. **Note:** This table is only installed after activation of the External Content for AI Search plugin (com.glide.ais.external_content).
AI Search External Content User Mapping History Record [ais_external_content_user_mapping_history]	List of import history events for external content user mapping tables. **Note:** This table is only installed after activation of the External Content for AI Search plugin (com.glide.ais.external_content).

Table	Description
AI Search External Content User Mapping Table Data Source Relationship [ais_external_search_user_mapping_data_source_m2m]	List of mappings between external content user mapping tables and indexed sources. **Note:** This table is only installed after activation of the External Content for AI Search plugin (com.glide.ais.external_content).
AI Search Genius Result Configuration [ais_genius_result_configuration]	List of Genius Result configurations.
AI Search Genius Result Configuration NLU Model Mapping [ais_genius_result_configuration_sys_nlu_model_m2m]	List of mappings between NLU models and Genius Result configurations.
AI Search Index Events stats [ais_index_queue_stats]	List of statistics entries for indexing events.
AI Search Indexed Batch History [ais_ingest_batch_stats]	List of ingestion history events for indexing batches.
AI Search Indexed Event History [ais_ingest_event_stats]	List of ingestion history events for indexer events
AI Search Indexed Retry History [ais_ingest_retry_stats]	List of ingestion history events produced when retrying indexing for a source table record.
AI Search Indexed Source [ais_datasource]	List of indexed sources.
AI Search Indexed Source Attribute [ais_datasource_attribute]	List of indexed source attributes configured for indexed sources.
AI Search Indexed Source Field Attribute [ais_datasource_field_attribute]	List of field settings configured for indexed sources.
AI Search Indexed Source History [ais_ingest_datasource_stats]	List of ingestion history events for indexed sources.
AI Search Indexed Source Statistics [ais_datasource_stats]	List of entries showing indexed record counts for indexed sources. For more information on these entries, see Review record counts for indexed sources.
AI Search Indexed Table History [ais_ingest_table_stats]	List of ingestion history events for individual source tables defined in indexed sources.

Table	Description
AI Search Ingestion Trigger Checkin [ais_ingestion_trigger_checkin]	List of ingestion trigger check-in times for AI Search nodes.
AI Search Partition Health [ais_partition_health]	List of health check events for AI Search query nodes connected to your instance.
AI Search Partition Health Response [ais_partition_health_response]	List of health check response events for AI Search query nodes connected to your instance.
AI Search Publish History [ais_publish_history]	List of publication history events for search profiles and typo handling, stop word, and synonym dictionaries.
AI Search Relevancy Auto Train Solution [ais_relevancy_autotrain_solution]	List of relevancy model tuning records used by machine learning relevancy.
AI Search Relevancy Solution Definition [ais_relevancy_solution_definition]	List of search profiles with relevancy models tuned by machine learning relevancy.
AI Search Relevancy Training Execution [ais_relevancy_training_execution]	List of relevancy model tuning tasks performed by machine learning relevancy.
AI Search Relevancy Training Staging [ais_relevancy_training_staging]	List of relevancy model staging tasks performed by machine learning relevancy.
AI Search Results Improvement Rule [ais_rule]	List of result improvement rules.
AI Search Results Improvement Rule Action [ais_rule_action]	List of actions defined for use in result improvement rules.
AI Search Results Improvement Rule Condition Builder Values [v_ais_result_improvement_rule_condition_builder_values]	List of static values defined for use in result improvement rule actions.
AI Search Retention Policy [ais_retention_policy]	List of retention policies available for use in indexed sources.
AI Search Search Profile [ais_search_profile]	List of search profiles.
AI Search Search Profile Dictionary Mapping [ais_search_profile_ais_dictionary_m2m]	List of mappings between search profiles and typo handling, stop word, and synonym dictionaries.

Table	Description
AI Search Search Profile Genius Result Configuration Mapping [ais_search_profile_ais_genius_result_configuration_m2m]	List of mappings between search profiles and Genius Result configurations.
AI Search Search Profile Relevancy Model [ais_search_profile_relevancy_model]	List of relevancy models created by machine learning relevancy.
AI Search Search Profile Search Source Mapping [ais_search_profile_ais_search_source_m2m]	List of mappings between search profiles and search sources.
AI Search Search Source [ais_search_source]	List of search sources.
AI Search Updatable Field Event [ais_updatable_field_event]	List of events for updatable fields on indexed records.
Raw User Mapping for External Content [ais_external_content_raw_external_user_mapping]	List of imported user mapping table entries. **Note:** This table is only installed after activation of the External Content for AI Search plugin (com.glide.ais.external_content).

Roles installed

AI Search includes the following roles.

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

AI Search administrator [ais_admin]

Manage configuration settings for the AI Search application.

- Creates, reads, updates, and deletes indexed sources
- Creates, reads, updates, and deletes search sources
- Creates and deletes mappings between search sources and search profiles
- Creates, reads, updates, and deletes search profiles
- Creates, reads, updates, and deletes synonym and stop word dictionaries
- Creates and deletes mappings between dictionaries and search profiles
- Reads and updates typo handling auto-correction dictionaries linked to search profiles
- Creates, reads, updates, and deletes Genius Result configurations
- Creates and deletes mappings between Genius Result configurations and search profiles
- Creates, reads, updates, and deletes result improvement rules
- Creates and deletes mappings between result improvement rules and search profiles

Contains roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

External Content for AI Search roles

External Content for AI Search includes the following roles.

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

AI Search external content API user [ais_external_content]

Access endpoints for AI Search external content APIs.

- Feeds external documents for AI Search to index as search results, or deletes records for external documents from the index, using the External Content Ingestion API.
- Imports user mappings for external content security using the AI Search External User Mapping API.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Assign this role to non-interactive users and integration users who require access to external content API endpoints. For more information on non-interactive users, see Non-interactive sessions. For details on integration users, see Mark service accounts as internal integration users.

Users with the admin role don't need this role to access external content API endpoints.

AI Search high security administrator [ais_high_security_admin]

Access High Security settings for AI Search.

- Bypasses all search filters from search sources and content security in the Search Preview UI, viewing all search query results available in the AI Search index
- Reads external content user mapping import history records

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

This is an elevated privilege role. Elevated privilege roles aren't assigned to users or groups, and must be used by elevation. For more information on elevated privilege roles, see Elevated privilege roles. To learn about elevation, see Elevate to a privileged role.

The Instance Security Center **Users with High Privilege Roles** user metric displays the count of users assigned this role. For more details, see User metrics.

AI Search system properties

Users with the admin role can configure the following system properties for AI Search by navigating to **All** > **System Properties** > `AI Search`.

AI Search system properties

Property	Description
Enable automatic switch over to a healthy `AI Search` **node in the near datacenter** (`glide.ais.ha.auto_switch_to_near_node`)	Option to enable automatic switch over to a healthy AI Search node in the near datacenter. • Type: true | false • Default value: `true` • Supported values: ◦ `true`: The system automatically switches over to use an AI Search node in the near datacenter when its read health score equals or exceeds the `glide.ais.monitor.health_threshold_green` threshold. ◦ `false`: The system doesn't automatically switch over to an AI Search node in the near datacenter. • Location: **System Properties** > `AI Search` (AI Search High Availability section)

AI Search system properties (continued)

Property	Description
Enable automatic failover to an alternate *AI Search* **node when the active node becomes unhealthy** (*glide.ais.ha.enable_auto_failover*)	Option to enable automatic failover to an alternate AI Search node when the active node becomes unhealthy. • Type: true | false • Default value: *true* • Supported values: ○ *true*: The system automatically fails over to an alternate AI Search node if the active node's write health score drops below the *glide.ais.monitor.health_threshold_yello* threshold. ○ *false*: The system doesn't automatically fail over an active partition based on its health score. • Location: **System Properties** > *AI Search* (AI Search High Availability section)
Ingestion batcher queue memory size in MB (*glide.ais.ingestion.batcher_queue_memory_size*)	Size in MB of batcher queue for indexing of source record content. • Type: integer • Default value: *20* • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion batcher queue length (records) (*glide.ais.ingestion.batcher_queue_size*)	Maximum number of records on batcher queue for indexing of source record content. • Type: integer • Default value: *1,000* • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion dictionary batcher queue length (records) (*glide.ais.ingestion.dictionary_term_batcher_queue_size*)	Maximum count of records on queue for addition of terms to a synonym, stop word, or typo handling dictionary. • Type: integer • Default value: *1,000* • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion failure retry timeout in seconds (*glide.ais.ingestion.failure_retry_timeout*)	Timeout period in seconds for indexing of content from a source record.

ServiceNow

AI Search system properties (continued)

Property	Description
	• Type: integer • Default value: *600* • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion feedback logger for error messages (*glide.ais.ingestion.feedback_logger_error*)	Destination for AI Search ingestion feedback messages with the ERROR level. • Type: string • Default value: *Sys Log* • Supported values: ◦ *None*: Don't log ingestion feedback messages with this level. ◦ *Sys Log*: Log ingestion feedback messages with this level to the . ◦ *File Log*: Log ingestion feedback messages with this level to a log file on the active AI Search node. • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion feedback logger for informational messages (*glide.ais.ingestion.feedback_logger_info*)	Destination for AI Search ingestion feedback messages with the INFO level. • Type: string • Default value: *None* • Supported values: ◦ *None*: Don't log ingestion feedback messages with this level. ◦ *Sys Log*: Log ingestion feedback messages with this level to the . ◦ *File Log*: Log ingestion feedback messages with this level to a log file on the active AI Search node. • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion feedback logger for minor error messages (*glide.ais.ingestion.feedback_logger_minor_error*)	Destination for AI Search ingestion feedback messages with the MINOR_ERROR level. • Type: string • Default value: *Sys Log* • Supported values:

AI Search system properties (continued)

Property	Description
	◦ *None*: Don't log ingestion feedback messages with this level. ◦ *Sys Log*: Log ingestion feedback messages with this level to the . ◦ *File Log*: Log ingestion feedback messages with this level to a log file on the active AI Search node. • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Ingestion feedback logger for warning messages (*glide.ais.ingestion.feedback_logger_warn*)	Destination for AI Search ingestion feedback messages with the WARN level. • Type: string • Default value: *Sys Log* • Supported values: ◦ *None*: Don't log ingestion feedback messages with this level. ◦ *Sys Log*: Log ingestion feedback messages with this level to the . ◦ *File Log*: Log ingestion feedback messages with this level to a log file on the active AI Search node. • Location: **System Properties** > *AI Search* (AI Search Ingestion Properties section)
Interval in seconds between AI Search node health status check requests (*glide.ais.monitor.health_status_check_interval*)	Interval in seconds between AI Search partition health status checks. • Type: integer • Default value: *60* • Location: **System Properties** > *AI Search* (AI Search Monitor Properties section)
Timeout period in seconds for AI Search node health status check request (*glide.ais.monitor.health_status_check_timeout*)	Timeout period in seconds for AI Search partition health status checks. • Type: integer • Default value: *30* • Location: **System Properties** > *AI Search* (AI Search Monitor Properties section)
Inclusive lower bound for GREEN (healthy) range of AI Search node health scores (*glide.ais.monitor.health_threshold_green*)	Inclusive lower percentage bound for green partition health score range. The inclusive upper bound for this health score range is 100.

AI Search system properties (continued)

Property	Description
	• Type: integer • Default value: *90* • Location: **System Properties** > *AI Search* (AI Search Monitor Properties section)
Inclusive lower bound for YELLOW range of AI Search node health scores (*glide.ais.monitor.health_threshold_yellow*)	Inclusive lower percentage bound for yellow partition health score range. The non-inclusive upper bound for this health score range is set by the *glide.ais.monitor.health_threshold_green* property value. The red partition health score range includes all scores below this property value. If the health score for an active partition drops into the red range, and if *glide.ais.ha.enable_auto_failover* is set to *true*, AI Search failover on the partition. • Type: integer • Default value: *50* • Location: **System Properties** > *AI Search* (AI Search Monitor Properties section)
Maximum request count per search query for AI Search node (*glide.ais.query.max_request_count_per_search*)	Maximum number of requests that an AI Search node can process per search query. • Type: integer • Default value: *10* • Location: **System Properties** > *AI Search* (AI Search Query Properties section)
Search query result count multiplier for AI Search node (*glide.ais.query.requested_rows_multiplier*)	Result-count multiplier for search queries that use the late-binding security implementation. The system multiplies the search query's minimum desired result count by this factor before requesting records from the AI Search node. Extra records replace any search results removed by the late-binding security implementation. For example, with the default value of *3*, a search query that uses the late binding security implementation and expects at least 10 results produces a request to the AI Search node for 30 matching records. The final result set includes the first 10 of these records that the user can access.

servicenow

AI Search system properties (continued)

Property	Description
	Lower values improve search query performance but may reduce the search result count below the requested minimum. • Type: integer • Default value: *3* • Location: **System Properties** > *AI Search* (AI Search Query Properties section)
Boolean search operator to use when a search query includes multiple terms (*glide.ais.query.search_operator*)	Boolean search operator to use when a search query includes multiple terms. When using the AND (conjunctive) operator, search only finds records that include all search query terms and phrases. When using the OR (disjunctive) operator, search finds records that include any search query term or phrase. • Type: string • Default value: *AND then OR 2+ key terms* • Supported values: ○ *AND then OR 2+ key terms*: Search using the AND operator. If the search query contains more than two search terms after stop word removal and returns less than one full page of results, automatically resubmit it using the OR operator. The resubmitted search query only finds records that contain at least half (rounded up) of the terms and phrases in the OR expression. **Note:** When you view the *glide.ais.query.search_operator* record on the System Property [sys_properties] table, this value appears as *AND_OR_ENHANCED*. ○ *AND then OR query*: Search using the AND operator. If the search query returns less than one full page of results, automatically resubmit it using the OR operator.

AI Search system properties (continued)

Property	Description
	The resubmitted search query only finds records that contain at least half (rounded up) of the terms and phrases in the OR expression. **Note:** When you view the *glide.ais.query.search_operator* record on the System Property [sys_properties] table, this value appears as *AND_OR*. ◦ *OR query*: Search using the OR operator. ◦ *AND query*: Search using the AND operator. **Note:** This value is not displayed on the AI Search system properties page, but can be set when editing the *glide.ais.query.search_operator* record on the System Property [sys_properties] table. • Location: **All** > **System Properties** > *AI Search* (AI Search Query Properties section)
Force *AI Search* **to use late binding for ACL evaluation** (*glide.ais.security.force_late_binding*)	Option to force use of the late binding security implementation for all search queries. For descriptions of the early binding and late binding security implementations, see Content security in AI Search. • Type: true | false • Default value: *false* • Supported values: ◦ *true*: The system uses the late binding security implementation for all search queries. ◦ *false*: The system uses the early binding security implementation for search queries. If search results include scripted ACLs, or if early

AI Search system properties (continued)

Property	Description
	binding encounters an error, the system switches to use the late binding security implementation. • Location: **System Properties** > *AI Search* (AI Search Security Properties section) **Note:** You can enable the late binding security implementation for all search results from a particular indexed source by setting the **Force Late Binding** field to *true* for the indexed source.
Enable global language fallback (*glide.ais.translate.enable_global_language_fallback*)	Option to enable matching of translated documents in the instance locale's language as well as the user's session language. For example, suppose a user searches from a Japanese user session on an instance with the default US English locale. With the default system property value of *false*, this search only matches Knowledge articles in Japanese. If the system property value is set to *true*, the same search matches articles in both Japanese and English. To learn more about translated document matching, see Searching translated content. For details on instance locales, see Configuring locales. • Type: true | false • Default value: *false* • Supported values: ◦ *true*: AI Search matches translated documents in the instance locale's language as well as the user's session language. ◦ *false*: AI Search only matches translated documents in the user's session language. • Location: **System Properties** > *AI Search* (AI Search Multi-Language Support section)
glide.mlpredictor.ml_x_snc_global_global_6e6d038834bf5010f877396127c8ffee.pr	Timeout period for ML predictions used in Genius Result processing. If Genius Result processing time exceeds this value for a search query, no Genius Result

AI Search system properties (continued)

Property	Description
	answers are displayed in the search query results. • Type: integer • Default value: *1000* • Location: System Properties [sys_properties] table
query.maxDictionaryTermLength	Character length limit for stop word and synonym dictionary terms. AI Search ignores dictionary terms that exceed this length limit. • Type: integer • Default value: *256* • Location: System Properties [sys_properties] table

ServiceNow AI Platform® tables excluded from AI Search indexing

AI Search excludes the following ServiceNow AI Platform tables when indexing content. You can't configure these tables, or tables derived from them, as AI Search indexed sources.

- AI Search External Search User Mapping Base [ais_external_search_user_mapping]
- Base Configuration Item [cmdb]
- Email [sys_email]
- Event [sysevent]
- Import Set Row [sys_import_set_row]
- Log [syslog]
- Queue [ecc_queue]
- Replicated Operation [sys_replication_queue]
- Sys Audit [sys_audit]

Variable types supported by AI Search indexing

AI Search can index searchable content and metadata from multiple variable types.

Variable types indexed from variables column

When indexing records from any source table, AI Search retrieves searchable content from some variable types when they're present in a source record column that has name **variables** and type *Variables*. Supported variable types for this context include:

- Email
- Multi-line text
- Single-line text
- URL
- Wide single-line text

Other variable types aren't supported for indexing.

For the full list of available Service Catalog variable types, see Types of Service Catalog variables.

AI Search index fields

List of fields available on records in the AI Search index. AI Search populates these fields with values from source record fields during indexing.

AI Search index fields

Field	Description
ai_search_parent_title	Title of the indexed record's parent record, if it has one. This field is automatically generated by AI Search from the parent record's **title** field value.
ai_search_teaser_text	Searchable text of the indexed record, displayed in search results. This field is automatically generated by AI Search from the indexed record's **text** field value.
ai_search_teaser_title	Title of the indexed record, displayed in search results. This field is automatically generated by AI Search from the indexed record's **title** field value.
datasource_id	Name of the indexed source that added the record to the AI Search index.
ext_record_id	Unique identifier for the external content source record that corresponds to the indexed record. Only populated for records from external content sources.
ext_groups_deny	Names of security principal groups that can't read the record. Only populated for records from external content sources. **Note:** Deny entries take precedence over read entries. If a group appears in this field and in *ext_groups_read*, members of that group can't read the record.
ext_groups_read	Names of security principal groups that can read the record. Only populated for records from external content sources.

AI Search index fields (continued)

Field	Description
	Note: Deny entries take precedence over read entries. If a group appears in this field and in *ext_groups_deny*, members of that group can't read the record.
ext_users_deny	Names of security principal users that can't read the record. Only populated for records from external content sources. **Note:** Deny entries take precedence over read entries. If a user appears in this field and in *ext_users_read*, that user can't read the record.
ext_users_read	Names of security principal users that can read the record. Only populated for records from external content sources. **Note:** Deny entries take precedence over read entries. If a user appears in this field and in *ext_users_deny*, that user can't read the record.
index_event_id	Sys_id of the indexing event that added the record to the AI Search index.
language	Language of the indexed record. AI Search uses the language specified for the record in the source table. If no language is specified in the source table, AI Search automatically detects the record's language.
record_class_name	The hierarchical class used to render the source record.
sys_id	Sys_id of the source record.
table	Name of the indexed record's source table. For records from external content sources, this index field indicates the name of the content schema table specified during document ingestion.
text	Searchable text for the indexed record. This index field may be populated directly from a

AI Search index fields (continued)

Field	Description
	source record field, such as **Description**, or may be extracted from binary content.
title	Title for the indexed record. This index field may be populated directly from a source record field, such as **Short description**, or may be extracted from binary content.
url	URL for the source record's form view. When you select the indexed record from a search result list, the system redirects you to this URL so you can view the source record. **Note:** AI Search automatically populates this index field from the source record. You cannot override the index field value using *map_ to* field settings.

Common index fields

The following AI Search index fields appear on all records indexed from ServiceNow AI Platform® source tables:

- ai_search_parent_title
- ai_search_teaser_text
- ai_search_teaser_title
- record_class_name
- sys_id
- table

[store-future: BEGIN review]

User context fields for boost result improvement rules

User context fields provide dynamically populated details about the current ServiceNow AI Platform® user and session. You can configure result improvement rules to boost search results that contain matches for user context field values.

For details on boosting search results that contain matches for user context field values, see the description for the Boost By User Context (dynamic) boost type in Boost search results using a result improvement rule.

User context field	Description
Browser (session.user_agent_browser)	Web browser name reported for the search user's ServiceNow AI Platform session.

User context field	Description
Browser Agent (session.user-agent)	User agent string for the web browser reported for the search user's ServiceNow AI Platform session.
City (location.city)	City field value from the Location [cmn_location] table record that's referenced by the Location field on the search user's record in the User [sys_user] table.
City (sys_user.city)	City field value from the search user's record in the User [sys_user] table.
Company (sys_user.company)	Company field value from the search user's record in the User [sys_user] table.
Country (location.country)	Country field value from the Location [cmn_location] table record that's referenced by the Location field on the search user's record in the User [sys_user] table.
Country (sys_user.country)	Country code field value from the search user's record in the User [sys_user] table.
Department (sys_user.department)	Department field value from the search user's record in the User [sys_user] table.
Device Form (session.user_agent_form)	Device form factor reported for the search user's ServiceNow AI Platform session.
Device Type (session.device_type)	Device type reported for the search user's ServiceNow AI Platform session.
Domain (sys_user.sys_domain)	Name field value from the Group [sys_user_group] table record that's referenced by the Domain field on the search user's record in the User [sys_user] table.
Home Office (sys_user.building)	Name field value from the Building [cmn_building] record that's referenced by the Building field on the search user's record in the User [sys_user] table.
Home Office City (sys_user.building.location.city)	City field value from the Location [cmn_location] record referenced by the Building [cmn_building] record that's referenced by the Building field on the search user's record in the User [sys_user] table.
Home Office Country (sys_user.building.location.country)	Country field value from the Location [cmn_location] record referenced by the Building [cmn_building] record that's referenced by the Building field on the search user's record in the User [sys_user] table.
Home Office State/Province (sys_user.building.location.state)	State / Province field value from the Location [cmn_location] record referenced by the Building [cmn_building] record that's referenced by the Building field on the search user's record in the User [sys_user] table.

User context field	Description
Job Title (sys_user.title)	Title field value from the search user's record in the User [sys_user] table.
Language (sys_user.language)	Language field value from the search user's record in the User [sys_user] table.
Location (sys_user.location)	Name field value from the Location [cmn_location] table record that's referenced by the Location field on the search user's record in the User [sys_user] table.
Mobile Device Type (session.mobile_device_type)	Mobile device type reported for the search user's ServiceNow AI Platform session.
Name (sys_user.username)	User ID field value from the search user's record in the User [sys_user] table.
OS (session.user_agent_os)	Operating system reported for the search user's ServiceNow AI Platform session.
OS Version (session.user_agent_version)	User agent version reported for the search user's ServiceNow AI Platform session.
Role (sys_user.roles)	Roles field value from the search user's record in the User [sys_user] table.
State (sys_user.state)	State / Province field value from the search user's record in the User [sys_user] table.
State (location.state)	State / Province field value from the Location [cmn_location] table record that's referenced by the Location field on the search user's record in the User [sys_user] table.

[End]

AI Search Country To Search Language form

The AI Search form contains information about relationships between user countries and Knowledge article search languages. Use this form when defining Knowledge article search languages for a country.

For details on defining Knowledge article search languages for a country, see Assign Knowledge article search languages by country.

AI Search Country To Search Language form

Field	Description
Country	Reference to the country from the Country [core_country] table that you want to define an additional Knowledge article search language for.
Search language	Reference to a language from the Language [sys_language] table that you want AI Search to use for Knowledge article searches by users in the selected country.

Note:
You can define multiple AI Search Country To Search Language records for the same country, as long as each one has a different Knowledge article search language. If you attempt to create multiple records with the same country and the same Knowledge article search language, the system displays a warning message and the record insert operation fails.

AI Search External Search User Mapping Table Data Source Relationship form

The AI Search External Search User Mapping Table Data Source Relationship form contains information about relationships between user mapping tables and indexed sources for external content. Use this form when linking a user mapping table to an indexed source for external content.

For details on linking a user mapping table to an indexed source for external content, see Mapping external users and groups to ServiceNow AI Platform users.

AI Search External Search User Mapping Table Data Source Relationship form

Field	Description
User Mapping Child Table	The name of the user mapping table that you want to link to the external content indexed source.
Index Source	The name of the external content indexed source that you want to link the user mapping table to.

AI Search Genius Result Configuration form

The AI Search Genius Result Configuration form contains information about Genius Result configurations defined for your search profiles in AI Search. Use this form when creating or modifying a Genius Result configuration.

For details on creating or editing a Genius Result configuration, see Create a Genius Result configuration.

AI Search Genius Result Configuration form

Field	Description
Name	Name for the Genius Result configuration.
Trigger condition	Condition that triggers evaluation of the Genius Result configuration. Supported values: • *Always*: The Genius Result configuration triggers for every search query. • **NLU**: The Genius Result configuration triggers when the search query's detected intent matches an NLU model intent. To use this trigger condition, you must link the NLU model and one or more of its intents to your Genius Result configuration. For

AI Search Genius Result Configuration form (continued)

Field	Description
	instructions, see Link an NLU model and intents to a Genius Result configuration. **Note:** For an NLU trigger condition to function, the NLU model must be published. You can verify NLU model status by navigating to **All** > `AI Search` > **Search Query Settings** > **Search Intents**. • *Term*: The Genius Result configuration triggers when the search query exactly matches a specified set of query terms. **Note:** Exact matching means that the search query must contain only the specified terms in the specified order. For example, if you define `conference room` as the trigger term, a search for `conference room` triggers the Genius Result condition but searches for `hotel conference room` or `room conference` don't trigger it.
Term	Term that triggers the Genius Result configuration when it exactly matches the search query. This field appears only when you select **Term** from the **Trigger condition** field.
Application	Application scope for the Genius Result configuration. This field is automatically set.
Active	Option to activate the Genius Result configuration. Inactive Genius Result configurations never trigger.
Processor script template	Sample server-side script function. Use the code displayed in this field as a basis for your `AI Search` **request processor** and `AI Search` **response processor** scripts. This field is automatically set.
AI Search request processor	Server-side script function that takes a `GeniusResultContext` object as its parameter. The function must return a `GeniusResultAnswer` object that contains a search query. Up to three results from this search query display as Genius Result answer cards.

AI Search Genius Result Configuration form (continued)

Field	Description
	Note: This field is required if **AI Search response processor** isn't populated. You can copy the code from the **Processor script template** to use as a starting point for your script function. When a search query triggers the Genius Result configuration, the system executes this script and then sends both the triggering search query and the scripted search query to the AI Search back end for computation of search results. To define a search query, use the `GeniusResultAnswer.setTable()`, `GeniusResultAnswer.addSearchPhrases()`, and `GeniusResultAnswer.setSearchPhrase()` API methods. For details of API methods usable in this script function, including examples of use, see the GeniusResultAnswer and GeniusResultContext scoped JavaScript API descriptions. **Note:** The `GeniusResultContext.getMatchingDocuments` API method returns no results when invoked in an **AI Search request processor** script, since the script is executed before AI Search computes results for the triggering search query. Don't use this method in your request processor scripts.
AI Search response processor	Server-side script function that takes a `GeniusResultContext` object as its parameter. The function must return a `GeniusResultAnswer` object that defines one or more maps of key-value pairs. You can display these key-value pairs on the Genius Result answer card by specifying them as **Return fields**. **Note:** This field is required if **AI Search request processor** isn't populated. You can copy the code from the **Processor script template** to use as a starting point for your script function.

AI Search Genius Result Configuration form (continued)

Field	Description
	When a search query triggers the Genius Result configuration, the system executes this script after AI Search computes search results. To define your map of key-value pairs, use the *`GeniusResultAnswer.addDataMap()`* API method. Each key-value pair represents a field name and value that you can display on the Genius Result answer card using the **Return fields** list. For details of API methods usable in this script function, including examples of use, see the GeniusResultAnswer and GeniusResultContext scoped JavaScript API descriptions.
Return fields	Comma-separated list of fields to display on the Genius Result answer card. This field is required if **AI Search response processor** is populated. Behavior depends on which processor script field is populated: • If **AI Search response processor** is populated, you must use this field to specify the names of fields defined in your *`GeniusResultAnswer`* response object that you want to display on the Genius Result answer card. • If **AI Search response processor** isn't populated, the Genius Result answer card automatically displays the common AI Search index fields: ◦ ai_search_parent_title ◦ ai_search_teaser_text ◦ ai_search_teaser_title ◦ record_class_name ◦ sys_id ◦ table You can optionally use this field to specify any additional AI Search index fields that you want to display for search results shown on the Genius Result answer cards. To view the list of AI Search index fields, see AI Search index fields.

AI Search Genius Result Configuration NLU Model Mapping form

The AI Search Genius Result Configuration NLU Model Mapping form contains information about relationships between Genius Result configurations and NLU models for AI Search. Use this form when linking Genius Result configurations to NLU models.

For details on linking Genius Result configurations to NLU models, see Link an NLU model and intents to a Genius Result configuration.

AI Search Genius Result Configuration NLU Model Mapping form

Field	Description
Genius Result Configuration	The Genius Result configuration to link the NLU model to. This field is automatically set to the selected Genius Result configuration.
NLU Model	Existing NLU model to use for the Genius Result configuration's NLU trigger condition.
State	Current state of the selected NLU model. This field is automatically set.
Intents	List of intents from the NLU model that can trigger the Genius Result configuration when detected in the search query.
Order	Order of evaluation for the NLU model in the context of the Genius Result configuration.
Application	Application scope for the NLU model link. This field is automatically set.

AI Search Results Improvement Rule form

The AI Search Results Improvement Rule form contains information about result improvement rules defined in your search profiles for AI Search. Use this form when creating or modifying a result improvement rule.

For details on creating or editing result improvement rules, see Create a result improvement rule.

AI Search Results Improvement Rule form

Field	Description
Label	Unique label for the result improvement rule.
Start Date	System time for the first date when the result improvement rule can activate. If left empty, this field is automatically set to the current date.

AI Search Results Improvement Rule form (continued)

Field	Description
End Date	System time for last date when the result improvement rule can activate. After this date, the rule state changes to `Expired`. If left empty, this field is automatically set to 30 days from the **Start Date**.
Activate on all queries	Option to make the result improvement rule activate for every search query.
Language	Language for the result improvement rule. AI Search only evaluates activation for result improvement rules that have the same language as the search query or that have **All Languages** specified for this field.
State	Current state for the result improvement rule. • **New**: The rule has never been published. • **Draft**: The rule is saved with unpublished changes. • **Published**: The rule is saved with all changes published. • **Expired**: The rule can't activate because its **End Date** is in the past. Publishing a search profile doesn't publish expired rules.
Marked for Deletion	Option indicating whether AI Search will delete the result improvement rule the next time you publish the search profile.
Apply only to Genius Results	Option to make the result improvement rule only operate on Genius Result answers when activated. **Note:** When this option is selected, the result improvement rule can only contain block actions.
Triggers	
Expand Synonyms For Query Conditions	Option for enabling synonym expansion for Query condition arguments. This field disappears if you select **Activate on all queries**. When expanding Query conditions, AI Search only considers active synonyms that are

AI Search Results Improvement Rule form (continued)

Field	Description
	defined in the same search profile and that have a language matching the language of the result improvement rule.
Triggers	List of conditions that must be met for the result improvement rule to activate. This field disappears if you select **Activate on all queries**. Select one of the following fields as the subject for your condition: • *`Keywords`*: Operates on one or more keywords found in the user's search query. For example, a **[Keywords] [are] [password]** condition is satisfied by searches for `How do I change my password` or `password reset`. • *`Query`*: Operates on the user's full search query. As an example, a **[Query] [is] [password reset]** condition is satisfied by a search for `password reset`, but not by a search for `password reset steps`. • *`User Context`*: Operates on a user context element from the search user's ServiceNow AI Platform® session. The system displays the source table and field after the user context element name. As an example, a **[User Context - Country (sys_user.country)] [is] [Japan]** condition is satisfied when the search user's User [sys_user] record has **country** set to Japan.

Child Table form

The Child Table form contains information about indexed source child tables defined for AI Search. Use this form when creating or modifying a child table of an indexed source.

For details on creating or modifying an indexed source child table, see Create an indexed source.

Child Table form

Field	Description
Indexed Source	Reference to the indexed source. This field is automatically set.
Table	Label and name for the child table that you want to enable content indexing for. You can only select from tables that extend the indexed source table.

Child Table form (continued)

Field	Description
Active	Option to activate indexing for the child table.
Application	Application scope for the child table indexing configuration. This field is automatically set.

Custom Matcher form

The Custom Matcher form includes information about global search custom matchers defined for AI Search. Use this form when creating a custom matcher to return exact-match results from a specific table and column for search terms matching a regular expression.

For details on creating a global search custom matcher, see Create a custom search matcher for global search.

Custom Matcher form

Field	Description
Name	Unique name for the custom matcher.
Regex	Java regular expression pattern that AI Search compares search queries to. When a search query matches this regular expression pattern, the custom matcher triggers and the search only returns results from the specified **Table** and **Column**. This field is automatically set when you select **Generate Regex**. **Note:** For details on Java regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class.
Table name	Table that you want to restrict search matches to when a user search satisfies the custom matcher's **Regex** regular expression pattern.
Application	Application scope for the custom matcher. This field is automatically set.
Column	Field from the specified **Table** that you want to restrict search matches to when a user search satisfies the custom matcher's **Regex** regular expression pattern.
Active	Option to activate the custom matcher. To deactivate the matcher, clear the **Active** option.

Dictionary and AI Search Dictionary forms

The Dictionary and AI Search Dictionary forms contain information about your AI Search typo handling, stop word, and synonym dictionaries. Use these forms when creating or modifying a typo handling, stop word, or synonym dictionary for use with AI Search.

For details on creating and modifying dictionaries, see Modify settings for a typo handling dictionary, Create a stop word dictionary linked to a search profile, and Create a synonym dictionary linked to a search profile.

Dictionary and AI Search Dictionary forms

Field	Description
Label	Unique label for the dictionary.
Dictionary name	Name for the dictionary. This field is automatically set based on the **Label** field value.
Language	Language for all terms defined in the dictionary. This field only displays when editing a stop word or synonym dictionary. AI Search automatically sets the language for typo handling dictionaries. **Note:** You can only link one stop word dictionary with a given language to a search profile.
Type	Dictionary type. This field is automatically set to *`Spell Check`*, *`Stop Word`*, or *`Synonym`*.
State	Current state of the dictionary. • **New**: The dictionary has never been published. • **Draft**: The dictionary is saved with unpublished changes. • **Published**: The dictionary is saved with all changes published.
Case sensitive	Option to enable case sensitivity for all terms in the dictionary. For more details on case sensitivity, see Typo handling, Stop words, or Synonyms.
Punctuation sensitive	Option to enable punctuation sensitivity for all terms in the dictionary. For more details on punctuation sensitivity, see Typo handling, Stop words, or Synonyms.
Derived	Option to derive dictionary entries from indexed content. This field and the **Derive Settings** section only display when editing a typo handling dictionary.

Dictionary and AI Search Dictionary forms (continued)

Field	Description
Derive Settings	
Indexed sources	List of indexed sources. The dictionary derives auto-correction terms from content indexed for these sources. This field only displays when editing a typo handling dictionary. For more information on indexed sources, see Indexed sources.
Auto refresh	Interval, in hours, after which derived typo handling entries automatically refresh. The minimum value is 24 hours. This field only displays when editing a typo handling dictionary.
Blocked Fields	Comma-separated list of fields. The dictionary excludes terms found in these fields when deriving typo handling auto-correction terms from indexed source records. This field only displays when editing a typo handling dictionary. You can specify fields using the following formats: • `*` (asterisk): Exclude terms found in all fields. This value disables derivation of typo-handling auto-corrections for the dictionary. • *`fieldname`*: Exclude terms found in the specified AI Search index field, such as **title** or **text**. For details on AI Search index fields, see AI Search index fields. • *`tablename.fieldname`*: Exclude terms found in the specified field on the specified ServiceNow AI Platform® table. • *`tablename.*`*: Exclude terms found in all fields in the specified ServiceNow AI Platform table.

Dictionary Term form

The Dictionary Term form contains information about individual typo handling, stop word, and synonym terms defined for AI Search. Use this form when creating or modifying a typo handling entry, a stop word, or a synonym in dictionaries linked to your search profiles.

For more details on creating or modifying dictionary terms, see Exclude a term from a typo handling dictionary, Create stop words, and Create synonyms.

Dictionary Term form

Field	Description
Dictionary	Name of the dictionary that includes the typo handling, stop word, or synonym term. This field is automatically set based on the dictionary that you selected.
Term	The term or phrase to include in the typo handling, stop word, or synonym dictionary.
Payload	List of equivalent terms or phrases in JSON array format. For example, `["wireless network", "WLANS"]`. This field only displays when editing a synonym term.
State	Current state for the dictionary term. • **New**: The dictionary term has never been published. • **Draft**: The dictionary term is saved with unpublished changes. • **Published**: The dictionary term is saved with all changes published. This field only displays when editing a synonym term.
Active	Option to activate the term for use in the dictionary. To deactivate the term, clear the **Active** option.
Marked for Deletion	Option indicating that AI Search will delete the term the next time you publish the search profile linked to the selected dictionary. This field only displays when editing a typo handling or stop word term, and is automatically set.
Application	Application scope for the dictionary term. This field only displays when editing a typo handling or stop word term, and is automatically set.

Facet form

The Facet form contains information about facets defined for search application configurations in AI Search. Use this form when creating or modifying a facet.

For details on creating or modifying a facet, see Create a facet in an AI Search application configuration.

Facet form

Field	Description
Name	Unique name for the facet.
Label	Label to display for the facet in the search application.
Facet Field	Comma-separated list of field names from AI Search indexed source tables. The facet displays search result refinement filters for values from one or more specified fields that are present in the current search result set. When you select a refinement filter after searching, AI Search restricts the search result set to show only results that contain the specified field value. **Note:** When you save or update a facet, AI Search verifies each facet field. To pass verification, the field must exist on at least one record that exists in a search source used by the search application. If a field doesn't satisfy this condition, the system displays an error message. Supports the following formats: **Table-specific field list** Enter a comma-separated list of qualified `[indexed_source_table_name].[field_name]` entries. AI Search displays search result refinement filters for values from the superset of the specified indexed source table fields. For example, enter `kb_knowledge.author` to display search result refinement filters for **Author** field values from the Knowledge [kb_knowledge] indexed source table. Or enter `kb_knowledge.author, sc_cat_item.owner, custom_table.author_field` to display search result refinement filters for values from all of the three specified indexed source table fields.

Facet form (continued)

Field	Description
	Note: All fields specified in a facet field list must be of the same type. The `[indexed_source_table_name]` qualifier must be the name of a table specified in an indexed source. Don't specify indexed source child table names for this qualifier. **Cross-table field** Enter an unqualified `[field_name]` entry. AI Search displays search result refinement filters for values from the specified cross-table field. For example, enter `language` to display search result refinement filters for **Language** field values from any indexed source table. **Note:** This format only supports the `language` and `tags` cross-table fields. Special cases include: **Faceting on parent table fields** If you specify an indexed source parent table field as a facet field, AI Search includes values from the same field on child tables from the indexed source. As an example, suppose you create an indexed source for a custom Parent [u_parent] table, which is extended by a custom Child [u_child] table. If you specify `u_parent.author` as a facet field, AI Search displays search result refinement filters for both `u_parent.author` and `u_child.author` field values. **Faceting on child table fields**

Facet form (continued)

Field	Description
	You can specify a field that's unique to an indexed source child table as a facet. Continuing the previous example, if your Child [u_child] table includes an **Editor** field that's not present on the Parent [u_parent] table, you can specify `u_parent.editor` as a facet to display search result refinement filters for **Editor** field values. **Note:** Be sure to specify the parent table name as the qualifier, not the child table name. **Faceting on dot-walk fields from referenced tables** If your indexed source includes a *`dot_walk_fields`* field setting for a dot-walk referenced table field, you can specify that field in a facet field entry using dot-separated reference field qualifiers. Add these qualifiers between the indexed source table name and the referenced table field name. For example, to facet Knowledge article search results based on the Department field values for their authors, you can define `author.department` as a *`dot_walk_fields`* field setting for the indexed source for the Knowledge [kb_knowledge] table and then specify `kb_knowledge.author.department` in your facet field entry. If you define `company.contact.department` as a *`dot_walk_fields`* field setting for the indexed source for the User [sys_user] table, you can specify `sys_user.company.contact.department` in your facet field entry.

Facet form (continued)

Field	Description
	Note: Each dot-walk reference level imposes a performance impact on indexing. Avoid using multi-level references unless necessary. To learn about *dot_walk_fields* field settings, see Enable indexing of referenced table fields for an indexed source. Facets work best when created from fields of the following types with relatively few unique values: • Choice • Reference • String • String (Full UTF-8) • True/False
Type	Type of the facet. The facet type determines how many field values you can select as facet refinement filters and the filter logic used by the facet. Supports the following values: **Single Select** In a Single Select facet, you can only select a single field value as a facet refinement filter. When you select a field value from a Single Select, AI Search filters the search results list, displaying only results that contain your chosen field value. A Single Select facet only allows you to select field values that are found in the current result set. **Multi Select And** In a Multi Select And facet, you can select multiple field values as facet refinement filters.

Facet form (continued)

Field	Description
	When you select a field value from a Multi Select And facet, AI Search filters the search results list, displaying only results that contain all of the field values you've selected from the facet. A Multi Select And facet only allows you to select field values that are found in the current result set. **Multi Select Or** In a Multi Select Or facet, you can select multiple field values as facet refinement filters. When you select a field value from a Multi Select Or facet, AI Search filters the search results list, displaying only results that contain any of the field values you've selected from the facet. A Multi Select Or facet initially only allows you to select field values that are found in the current result set. All field values that are available when you make your first selection remain available for selection. As an example, suppose you have a Multi Select Or facet for a custom **Color** field with blue, green, and red field values. If your search returns only green and red results, the **Color** facet only displays **green** and **red** field values. When you select the **red** field value, AI Search keeps the **green** field value available for selection because it was available before you made any **Color** selection.

Facet form (continued)

Field	Description
	Warning: Applying multiple filters from a Multi Select And or Multi Select Or facet makes your search query URL longer. Test your multi-select facet configurations to ensure that search query URLs with facet filters don't exceed the maximum URL length limit for your web browser. (Different web browsers have different URL length limits.)
Facet Value Limit	Maximum number of distinct field values to display for the facet.
Order	Order of appearance for the facet in the context of the search application configuration.
Active	Option to activate the facet for the application.

Field Setting form

The Field Setting form contains information about field-specific settings defined for your AI Search indexed sources. Use this form when creating or modifying a field setting for an indexed source.

For details on creating or editing field settings for indexed sources, see Field settings.

Field Setting form

Field	Description
Indexed Source	Indexed source that the field setting applies to.
Source	Specific table from the indexed source that the field setting applies to.
Attribute	Field setting to apply for the selected field on all records indexed from the selected source table. For the full set of available field settings, see List of AI Search indexed source field settings.
Application	Application scope for the field setting. This field is automatically set.
Field	Record field that the field setting operates on.
Value	Value for the specified field setting. For details on the values supported for each field setting, see List of AI Search indexed source field settings.

Indexed Source form

The Indexed Source form contains information about indexed sources defined for AI Search. Use this form when creating or modifying an indexed source.

For details on creating or modifying an indexed source, see Create an indexed source.

Indexed Source form

Title	Description
Name	Unique name for the indexed source.
Source	The ServiceNow AI Platform® table to index. AI Search excludes some ServiceNow AI Platform tables from indexing. You can't define indexed sources for these excluded tables or their derived tables. For a list of excluded tables, see ServiceNow AI Platform tables excluded from AI Search indexing. To index content from a remote table, create an external indexed source. For more details on configuring indexed sources for external content, see Indexing and searching external content in AI Search. Plugins and applications can define multiple indexed sources for the same ServiceNow AI Platform table, but only one of them can be active at a time. For more details, see `Active`.
Active	Option to activate the indexed source. AI Search only indexes content and metadata from active indexed sources. **Note:** Only one indexed source can be active at a time for a given ServiceNow AI Platform table. The system makes duplicate indexed sources for a table inactive by default. Before you can make one of these duplicate sources active, you must edit the currently active source and make it inactive. AI Search only indexes content and metadata from active indexed sources.
Application	Application scope for the indexed source. This field is automatically set.
Retention Policy	Policy specifying a retention time limit for records from the source table. AI Search only indexes source records that were last updated within the specified time period, and purges records from the index when their time since last update exceeds the limit.

Indexed Source form (continued)

Title	Description
	If this field is empty, AI Search indexes source records regardless of their last updated date and retains them in the index indefinitely. **Note:** This field is required when **Source** is set to the Task [task] table or any table that extends it.
Force Late Binding	Option to make AI Search use late binding security for all search results from the indexed source. This field isn't shown in the default Indexed Source form view. For instructions on configuring the form layout to show the field, see Configuring the form layout. For an explanation of late binding security and its effects on search, see Content security in AI Search.
Condition	
Filter	Optional filter conditions limiting the set of records to index from the source table. AI Search only indexes records that satisfy all defined filter conditions, and purges records from the index if they no longer satisfy these conditions.

Indexed Source Attribute form

The Indexed Source Attribute form contains details of attributes that affect content indexing behavior for all fields from an AI Search indexed source. Use this form when creating or modifying an indexed source attribute.

For details on creating or editing indexed source attributes, see Indexed source attributes.

Indexed Source Attribute form

Field	Description
Indexed Source	Indexed source the attribute applies to.
Attribute	Attribute to set for all records indexed from the selected source table. For the full set of indexed source attributes, see List of AI Search indexed source attributes.
Application	Application scope for the indexed source attribute. This field is automatically set.

Indexed Source Attribute form (continued)

Field	Description
Source	Specific table from the indexed source that the attribute applies to.
Value	Value for the specified attribute.

Indexed Source History form

The Indexed Source History form records details of content and metadata ingestion and indexing for an AI Search indexed source. Use this form when reviewing the ingestion history for an indexed source.

To learn how to review the ingestion history for an indexed source, see Review ingestion history for an indexed source.

Indexed Source History form

Field	Description
Indexed Source	Reference to the indexed source from which AI Search ingested records for indexing.
Ingestion Event ID	Sys_id of the ingestion event.
Partial Ingestion	Option indicating whether AI Search ingested all parent and child tables defined in the indexed source. Valid values: • `false`: AI Search ingested records from the indexed source table and all child tables configured for indexing. • `true`: AI Search only ingested records from a subset of the indexed source tables configured for indexing.
Number of Errors	Count of error messages logged for the ingestion event.
Number of Minor Errors	Count of minor error messages logged for the ingestion event.
Number of Warnings	Count of warning messages logged for the ingestion event.
Ingestion State	State for the ingestion event. Valid values: • `cancelled`: Ingestion and indexing didn't complete successfully. • `indexed`: Ingestion and indexing completed successfully.

Indexed Source History form (continued)

Field	Description
	• *indexing*: Ingestion and indexing are ongoing. • *not_started*: Ingestion and indexing are pending.
Start Time	System time at which the ingestion event began.
End Time	System time at which the ingestion event completed.
Content Preparation Duration	Total time that the ServiceNow AI Platform® spent retrieving source records for the ingestion event.
Indexing Duration	Total time that AI Search spent indexing source records for the ingestion event.
Records Processed	Count of source records processed for the ingestion event.
Attachments Processed	Count of source record attachments processed for the ingestion event.
Documents Processed per Second (DPS)	Average number of source records and attachments processed per second for the ingestion event.

Navigation Tab form

The Navigation Tab form contains information about navigation tabs defined for search application configurations in AI Search. Use this form when configuring display settings for a navigation tab.

For information on configuring navigation tab display settings for an AI Search application, see Configure navigation tabs in an AI Search application configuration.

Navigation Tab form

Field	Description
Name	Name for the navigation tab. This field is set by default to the name of the originating search source. Changing the name doesn't affect the navigation tab's filtering behavior.
Active	Option to enable the navigation tab. Inactive navigation tabs aren't displayed in the search application.
Order	Order of appearance for the navigation tab in the context of the search application configuration.

Navigation Tab form (continued)

Field	Description
Label	Label to display for the navigation tab in the search application.
Search Context Config	Search application configuration for the navigation tab.

Record Producer Configuration form

The Record Producer Configuration form contains information about your AI Search Assist widget settings for Service Portal. Use this form when creating or modifying an AI Search Assist record producer configuration.

Record Producer Configuration form

Field	Description
Title	Title of the record producer configuration. For example, `Search Results`.
Results per page	Number of search results to display per page.
Allow expand/collapse	Option to allow users to expand or collapse search results in this record producer configuration.
Active	Option to make this record producer configuration active for your instance.
Application	Application scope for this record producer configuration.

Rule - Action Mapping form

The Rule - Action Mapping form contains information about relationships between AI Search result improvement rules and boost actions. Use this form when linking boost actions to result improvement rules.

For information on configuring boost actions in result improvement rules, see Boost search results using a result improvement rule.

Rule - Action Mapping form

Field	Description
Label	Unique label for the boost action.
Type	Action type. This field is automatically set to **Boost Action**.
Application	Application scope for the boost action. This field is automatically set to the name of the current application.
Order	Order of evaluation for the boost action in the context of the result improvement rule.
Boost Action	

Rule - Action Mapping form (continued)

Field	Description
Boost Type	Type of boost the action applies to search results from the selected indexed source. Supported types include: • *`Boost All Documents In Indexed Source`*: Applies relevancy boost to every search result from the indexed source. • *`Boost By User Context (dynamic)`*: Applies relevancy boost to search results from the indexed source if they include a value for the selected field that matches a field value specified in the dynamically populated user context. • *`Boost By Field Match (static)`*: Applies relevancy boost to search results from the indexed source if they include a value for the selected field that matches a specified static value.
Indexed Source	Reference to an indexed source. The boost action only applies to search results for records from this indexed source.
Boost Weight	Integer value that AI Search uses to alter the relevancy scores for search results affected by the boost action. Boost weights can be positive or negative. Positive boost weights increase the relevancy scores for boosted search results. Each boost weight of 1 increases a boosted search result's relevancy score by 0.1% of its original value. For example, if you set a boost weight of 100, you increase each boosted search result's relevancy by +10%, so a search result with an original relevancy score of 60 gets a final relevancy score of 66. Setting a boost weight of 1000 adds +100%, doubling the relevancy scores for boosted search results. Negative boost weights reduce the relevancy scores for boosted search results. Each boost weight of −1 reduces a boosted result's relevancy score by 0.1% of its original value. As an example, if you set a boost weight of −150, you reduce a boosted search result's relevancy by 15%, so a search result with an original relevancy score of 100 gets a final relevancy score of 85. Any boost weight of −1000 or lower reduces a boosted search result's relevancy score to 0, causing it to appear at the end of the search result list.

Rule - Action Mapping form (continued)

Field	Description
	When a single search result is affected by multiple boost actions, AI Search adds the individual boost weights together to determine the total relevancy score modifier. For example, a search result boosted by three rules with boost weights 100, 150, and −25 has a total boost weight of 225, increasing its final relevancy score by 22.5% of its original value.
When	Name of a field from the selected indexed source. This field appears only when you select **Boost By User Context (dynamic)** or **Boost By Field Match (static)** as the boost type.
Matches the Searcher's	Dynamically populated user context element to compare with a search result's value for the field selected in the **When** field. This field appears only when you select **Boost By User Context (dynamic)** as the boost type. The choice list displays the source table and field for each user context element. As an example, the **Language** user context element is populated using the User [sys_user] table's **language** field value for the search user.
Contains	Static value to compare with a search result's value for the field selected in the **When** field. This field appears only when you select **Boost By Field Match (static)** as the boost type.

Search Application Configuration form

The Search Application Configuration form contains information about search application configurations defined for AI Search. Use this form when creating or modifying a search application configuration.

For details on creating or modifying a search application configuration for use with AI Search, see Create a search application configuration for AI Search.

Search Application Configuration form

Field	Description
Name	Name for the search application configuration. Type: string
Search Engine	Search engine to use in the ServiceNow AI Platform® search application that links to the search application configuration.

Search Application Configuration form (continued)

Field	Description
	Type: choice list Default value: AI Search Supported values: • *AI Search*: Use AI Search as the search engine for the application. Select this search engine if your search application has the **Enable** *AI Search* option enabled. • *Zing*: Use the Zing text indexing and search engine as the search engine for the application. Select this engine if your search application has the **Enable AI Search** option disabled, or if it lacks this option. **Note:** If your search engine selection conflicts with your search application's **Enable** *AI Search* setting, the system displays an error message. Before changing search engines, first make sure that your search application has the proper **Enable** *AI Search* setting.
Search Profile	Search profile to use for searches in the application. This field appears only when you select **AI Search** as the application's search engine. Type: reference to a record from the Search Profile [ais_search_profile] table
Search Results Limit	Maximum number of search results the application should display on each search results page. This limit also specifies the minimum number of search results each query should return. When a search query returns less than this number of results, AI Search attempts to broaden the result set by automatically resubmitting the query with modified parameters. For more details on query resubmission, see Automatic search query resubmission. Type: integer Default value: 10 Supported values: 0 or positive

Search Application Configuration form (continued)

Field	Description
Genius Results Limit	Maximum number of Genius Result answer cards the application should display on the search results page. This field appears only when you select **AI Search** as the application's search engine. Type: integer Supported values: 0 or positive Starting in Washington DC Patch 4, when **Genius Results Limit** is set to 2 or greater and a search returns multiple Genius Result answers, the answer cards display in a carousel. You can scroll the carousel and select the answer card you want to take action from. **Note:** Search applications display a maximum of 10 answer cards per search query. If you set this limit to 11 or more, the search application still displays a maximum of 10 answer cards.
Enable Typo Handling	Option to auto-correct search query terms to match terms found in the search profile's typo handling dictionary. This field appears only when you select **AI Search** as the application's search engine. For more details on typo handling auto-correction configuration and behavior, see Typo handling. Type: true | false Supported values: • `true`: Auto-correct search query terms in the search application. • `false`: Don't auto-correct search query terms in the search application.
Enable Exact Match	Option to use exact matching in the search application. When this option is selected, if a search result's **Number** field exactly matches the search, the application displays the record or document associated with that search result, bypassing the search results page. Exact

Search Application Configuration form (continued)

Field	Description
	matching also bypasses any custom matchers defined for the search application. **Note:** Exact match behavior is controlled by the *com.snc.agent_workspace.global_search.ty* system property. This property's value is a Java regular expression pattern. Only search queries that match the regular expression pattern are evaluated as possible exact record matches. For details on Java regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class. Type: true | false Default value: true Supported values: • *true*: Use exact matching in the search application. • *false*: Don't use exact matching in the search application.
Show parent for results that are attachments	Option to display attachment search results inline with their parent results on the search application's results page. If cleared, attachment search results display separately from their parent results. For more details on the effects of attachment grouping, see Group attachment search results with their parent results. Type: true | false Default value: true Supported values: • *true*: Group attachment search results inline with their parent results in the search application. • *false*: Display attachment search results separately from their parent results in the search application.
Show Search Results Count	Option to display the count of search results on the search application's results page.

Search Application Configuration form (continued)

Field	Description
	For more details on this option, see Display search result counts on the results page for a search application. Type: true | false Default value: true Supported values: • *true*: Display the count of search results on the search application's results page. • *false*: Don't display the count of search results on the search application's results page.

Search Profile form

The Search Profile form contains information about your search profiles defined for AI Search. Use this form when creating or modifying a search profile.

For details on creating or editing search profiles, see Create a search profile.

Search Profile form

Field	Description
Label	Label for the search profile.
Description	Text description for the search profile.
Application	Application scope for the search profile. This field is automatically set.
Updated	Last update date for the search profile. This field is automatically set.
Updated by	User who most recently updated the search profile. This field is automatically set.
Active	Option to activate the search profile.
State	State of the search profile. • **New**: The search profile has never been published. • **Draft**: The search profile is saved with unpublished changes. • **Published**: The search profile is saved with all changes published.

Search Profile - Dictionary Mapping form

The Search Profile - Dictionary Mapping form contains information about mappings between search profiles and stop word and synonym dictionaries. Use this form when linking stop word and synonym dictionaries to search profiles.

For details on linking stop word and synonym dictionaries to your search profiles, see Link a stop word dictionary to a search profile and Link a synonym dictionary to a search profile.

AI Search Profile Dictionary M2M form

Field	Description
Dictionary	The stop word or synonym dictionary that you want to link to the selected search profile.
Application	Application scope for the dictionary mapping record. This field is automatically set.
Profile	The search profile that you want to link the selected stop word or synonym dictionary to. This field is automatically set to the profile that you edited.

Search Profile - Genius Result Mapping form

The Search Profile - Genius Result Mapping form contains information about relationships between Genius Result configurations and search profiles in AI Search. Use this form when linking Genius Result configurations to search profiles.

For details on linking Genius Result configurations to your search profiles, see Link a Genius Result configuration to a search profile.

Search Profile - Genius Result Mapping form

Field	Description
Profile	The search profile to link the Genius Result configuration to. This field is automatically set to the profile that you edited.
Genius Result Configuration	The Genius Result configuration to link to the specified search profile.
Order	Order of evaluation for the Genius Result configuration in the context of the search profile.
Application	Application scope for the Genius Result configuration mapping record. This field is automatically set.

Search Profile - Search Source Mapping form

The Search Profile - Search Source Mapping form contains information about relationships between search sources and search profiles in AI Search. Use this form when linking search sources to search profiles.

For details on linking search sources to your search profiles, see Link a search source to a search profile.

Search Profile - Search Source Mapping form

Field	Description
Profile	The search profile that you want to link the search source to. This field is automatically set to the profile that you edited.
Search source	The search source that you want to link to the specified search profile.
Exclude in Regular Search	Option to exclude content from this search source when generating standard search results for searches that use the selected search profile.
Application	Application scope for the search source mapping record. This field is automatically set.
Exclude in Now Assist Genius Results	Option to exclude content from this search source when generating Genius Result answers for searches that use the selected search profile. This option applies to all Genius Results, not just Now Assist Genius Results.
Exclude Attachments in Now Assist Genius Results	Option to exclude attachments from this search source when generating Genius Result answers for searches that use the selected search profile. If `Exclude in Now Assist Genius Results` is not selected, selecting this option has no effect. This option applies to all Genius Results, not just Now Assist Genius Results. **Note:** Selecting this option has no effect on search sources created from external indexed sources (those that index external content). It only affects search sources created from internal indexed sources (those that index content from ServiceNow AI Platform® tables).

Search Scripted Post-processors form

The Search Scripted Post-processors form contains information about post-processing scripts defined in search application configurations for AI Search. Use this form when creating or modifying search scripted post-processors.

For details on creating or editing search scripted post-processors, see Using search scripted post-processors in AI Search application configurations.

Search Scripted Post-processors form

Field	Description
Active	Option to activate the search script post-processor.
Name	Name for the search script post-processor.

Search Scripted Post-processors form (continued)

Field	Description
Application	Application scope for the search script post-processor. This field is automatically set.
Type	Type of record the search script post-processor operates on. Valid values: • *Genius Result*: Script operates on a Genius Result answer record. • *Search Result*: Script operates on a search result record.
Post processing script	Server-side script function that takes a record object as its parameter. When the search scripted post-processor is linked to a search application configuration, search results reflect changes made to the record object in this script function. As an example, this post-processing script function operates on search results from the kb_knowledge table, transforming sys_updated_on date field values into a human-readable format:

```
function process(record) {

    var transformableTables =
 ['kb_knowledge'];
    var
 transformableDateFields =
 ['sys_updated_on'];

    var table =
 record.getTable();

    // Terminate early
 if possible, to avoid
 unnecessary overhead
    if
 (transformableTables.indexOf(
table) === -1)
        return;

    function
 transformToHumanFriendlyDate(
rawDate) {
        // Time gaps in ms
        var MINUTE_GAP =
 60000;
        var HOUR_GAP =
 3600000;
```

Search Scripted Post-processors form (continued)

Field	Description
	`` var DAY_GAP = 86400000; var WEEK_GAP = 604800000; var MONTH_GAP = 2592000000; var YEAR_GAP = 31536000000; var gap = new Date().getTime() - rawDate; if (gap < MINUTE_GAP) { return gs.getMessage('Just Now'); } else if (gap > MINUTE_GAP && gap < 2 * MINUTE_GAP) { return gs.getMessage('1 minute ago'); } else if (gap < HOUR_GAP) { return gs.getMessage('{0} minutes ago', '' + Math.floor(gap / MINUTE_GAP)); } else if (gap > HOUR_GAP && gap < 2 * HOUR_GAP) { return gs.getMessage('1 hour ago'); } else if (gap < DAY_GAP) { return gs.getMessage('{0} hours ago', '' + Math.floor(gap / HOUR_GAP)); } else if (gap > DAY_GAP && gap < 2 * DAY_GAP) { return gs.getMessage('1 day ago'); } else if (gap < WEEK_GAP) { return gs.getMessage('{0} days ago', '' + Math.floor(gap / DAY_GAP)); } else if (gap > WEEK_GAP && gap < 2 * WEEK_GAP) { ```

Search Scripted Post-processors form (continued)

Field	Description
	`` return gs.getMessage('1 week ago'); } else if (gap < MONTH_GAP) { return gs.getMessage('{0} weeks ago', '' + Math.floor(gap / WEEK_GAP)); } else if (gap > MONTH_GAP && gap < 2 * MONTH_GAP) { return gs.getMessage('1 month ago'); } else if (gap < YEAR_GAP) { return gs.getMessage('{0} months ago', '' + Math.floor(gap / MONTH_GAP)); } else if (gap > YEAR_GAP && gap < 2 * YEAR_GAP) { return gs.getMessage('1 year ago'); } else { return gs.getMessage('{0} years ago', '' + Math.floor(gap / YEAR_GAP)); } } for (var i = 0; i < transformableDateFields.leng th; i++) { var fieldName = transformableDateFields[i]; var updatedDateField = record.getField(fieldName); if (updatedDateField !== null) { var updatedDateRaw = updatedDateField.getValue(); if (updatedDateRaw != null) { var humanFriendlyDate = transformToHumanFriendlyDate(updatedDateRaw); updatedDateField.setDisplayVa lue(humanFriendlyDate); ```

Search Scripted Post-processors form (continued)

Field	Description
	`}` `}` `}` `}` **Note:** The UI displays an overview of API methods usable in the post-processing script function. You can copy code from an existing post-processor script to use as a starting point for your own script function. To access existing post-processor scripts, open the Search Scripted Post-processors [sys_search_scripted_processor] table in list view by following these steps: 1. Select **All**. 2. In the Filter field, enter `sys_search_scripted_processor.list`. 3. Press Enter.

Search Source form

The Search Source form contains information about search sources that you can make available in AI Search user search experiences. Use this form when creating or modifying a search source.

For details on creating or editing search sources, see Create a search source.

AI Search Source form

Field	Description
Indexed source	Reference to an indexed source. For more details on indexed sources, see Indexed sources.
Conditions	Set of filter conditions that documents from the indexed source must meet to be included in the search source. For more details on filter conditions, see Filters.

Sort Option form

The Sort Option form contains information about search result sort options defined in your AI Search application configurations. Use this form when creating or modifying a search result sort option for a search application configuration.

For details on creating or editing search result sort options, see Search result sort options in AI Search application configurations.

Sort Option form

Field	Description
Label	Label displayed for the search result sort option in the search application's sort menu.
Sort Field	Name of a source table field or common index field that the sort option uses to sort search results. The sort field can be a date, number, or string field. **Note:** If you specify a sort field of an unsupported type, an error message appears when you attempt to save the sort option. Supports the following formats: **Qualified `[indexed_source_table_name].[field_name]` entry** Sort search results on their values from the specified table's field. For example, enter `kb_knowledge.author` to sort search results by their Author field values. **Note:** The parent search application configuration's linked search profile must include a search source derived from the specified indexed source table. If this condition isn't met, an error message appears when you attempt to save the sort option. **Unqualified common index field name** Sort search results on their values from the specified common AI Search index field. For example, enter `title` to sort search results based on

Sort Option form (continued)

Field	Description
	their titles, or `language` to sort search results based on their language. The `[indexed_source_table_name]` qualifier must be the name of a table specified in an indexed source. Don't specify indexed source child table names for this qualifier. Special cases include: **Sorting on parent table fields** If you specify an indexed source parent table's field in a sort option, AI Search includes values from the same field on child tables from the indexed source. For example, suppose you create an indexed source for custom table Parent [u_parent] which is extended by custom table Child [u_child], and then specify `u_parent.author` as the field for a sort option. When a user applies the sort option, AI Search uses both `u_parent.author` and `u_child.author` field values to sort the search result list. **Sorting on child table fields** You can specify a field that's unique to an indexed source's child table as a sort option field. Continuing the previous example, if your Child [u_child] table includes an **Editor** field that's not present on the Parent [u_parent] table, you can define `u_parent.editor` as a sort option field to sort search result s using this field's values. **Note:** Be sure to specify the parent table's name as the qualifier, not the child table's name.
Application	Application scope for the search result sort option. This field is automatically set.

Sort Option form (continued)

Field	Description
Ascending	Option to make the search result sort option display results in ascending order (a to z). Clear this option to make the search result sort option display results in descending order (z to a).
Active	Option to activate the search result sort option for use in the search application.

Suggestion Reader Group form

The Suggestion Reader Group form contains information about auto-complete suggestion types defined for AI Search. Use this form when modifying settings for auto-complete suggestion types in AI Search applications.

For details on configuring auto-complete suggestions in AI Search applications, see Auto-complete suggestions in AI Search applications.

Suggestion Reader Group form

Field	Description
Suggestion Type	Application-specific name for the auto-complete suggestion type. This field is automatically set.
Section Header	Header displayed in the application search field for auto-complete suggestions of this type.
Application	Application scope for the auto-complete suggestion type. This field is automatically set.
Order	Order of appearance in the application search field for the section containing auto-complete suggestions of this type. The system uses this field to determine display order when a search application configuration includes multiple suggestion types with the same **Click-in Mode** setting.
Suggestion Limit	Maximum number of auto-complete suggestions of this type that the application should display in the search field.
Click-in Mode	Option to control the display of the section containing auto-complete suggestions of this type.

Suggestion Reader Group form (continued)

Field	Description
	• If selected, auto-complete suggestions of this type are displayed when the user first selects the empty application search field. • If deselected, auto-complete suggestions of this type are displayed as the user enters search query terms into the application search field. **Note:** A search application configuration can't link to multiple search suggestion types that include the same reader and have the same click-in mode. If you try to link a second suggestion type with the same reader and click-in mode to a search application configuration, the system displays an error message.

Extending AI Search with ServiceNow Store applications and integrations

Applications and integrations from the ServiceNow® Store extend AI Search functionality and provide insight into search usage and behavior.

AI Search spoke

This Integration Hub spoke enables storage of document data from third-party applications in the AI Search index using Flow Designer. Users can access the indexed data through portals that have AI Search enabled.

Platform Analytics Solution for Advanced AI Search Management Tools

This Platform Analytics Solution helps AI Search administrators analyze search query traffic trends, review indexed record and configuration settings, and preview search query results.

AI Search for Next Experience

The AI Search for Next Experience ServiceNow® Store application enables administrators to configure AI Search as the search engine for Unified Navigation searches in Next Experience, including global search and configurable workspace search.

Now Assist in AI Search

The Now Assist in AI Search ServiceNow® Store application combines the power of search with the Now LLM Service agentic AI model to answer questions in user searches with actionable AI-generated summaries of relevant Knowledge articles.

Platform Analytics Solution for Advanced AI Search Management Tools

ServiceNow® Platform Analytics Solutions contain prepackaged Performance Analytics and Reporting content for use with other ServiceNow AI Platform® products. This Platform Analytics Solution helps AI Search administrators analyze search query traffic trends, review indexed record and configuration settings, and preview search query results.

Required roles

The following roles are associated with this solution:

- ServiceNow AI Platform administrator (admin): Install and activate this Platform Analytics Solution and make any necessary changes to system properties.
- AI Search administrator (ais_admin): Review the AI Search dashboards. Share the dashboards with appropriate stakeholders. Test search queries using the Search Preview UI module.
- User Experience Analytics administrator (analytics_admin), User Experience Analytics viewer (analytics_viewer), User Portal Analytics administrator (portal_analytics_admin), Portal Analytics viewer(portal_analytics_viewer), Web Analytics administrator (web_analytics_admin), or Web Analytics viewer(web_analytics_viewer): View the AI Search Analytics dashboard metrics and reports.
- Impersonator (impersonator): View NLU intents detected in queries, enable debugging, or search as another user in the Search Preview UI module.

Related topics

Analytics and Reporting Solutions

Activate the Platform Analytics Solution for Advanced AI Search Management Tools

Activate the content pack ServiceNow® Store application for the Platform Analytics Solution.

Before you begin

The Usage Insights API application from the ServiceNow Store must be activated in your instance.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **Available To Obtain From Store**.
2. Browse for Advanced AI Search Management Tools and verify that you have the necessary valid entitlements.

   You need a Now Support account and permission to request applications for the instances under consideration. The Now Support account and the instance must be part of the same company. For more information, see Install a ServiceNow Product on the ServiceNow Store.
3. Search for `Advanced AI Search Management Tools`.
4. Install the application.

Result

The following application menus and modules appear in the application navigator:

Application	Application navigator path
AI Search Profile dashboard	`AI Search` > `AI Search` **Analytics** > **Search Profile Analytics**
AI Search dashboard	`AI Search` > `AI Search` **Analytics** > **Search Index Analytics**
Search preview UI	`AI Search` > **Preview** > **Search Preview**

Note:
Immediately after installation, the AI Search dashboards will be empty. The dashboards may take up to an hour to populate with results.

AI Search Profile dashboard

The AI Search Profile dashboard summarizes indexed record counts and search query traffic associated with a search profile defined in AI Search. Interactive filters enable users to choose a search profile and select the time frame for analysis of search query traffic.

AI Search Profile dashboard

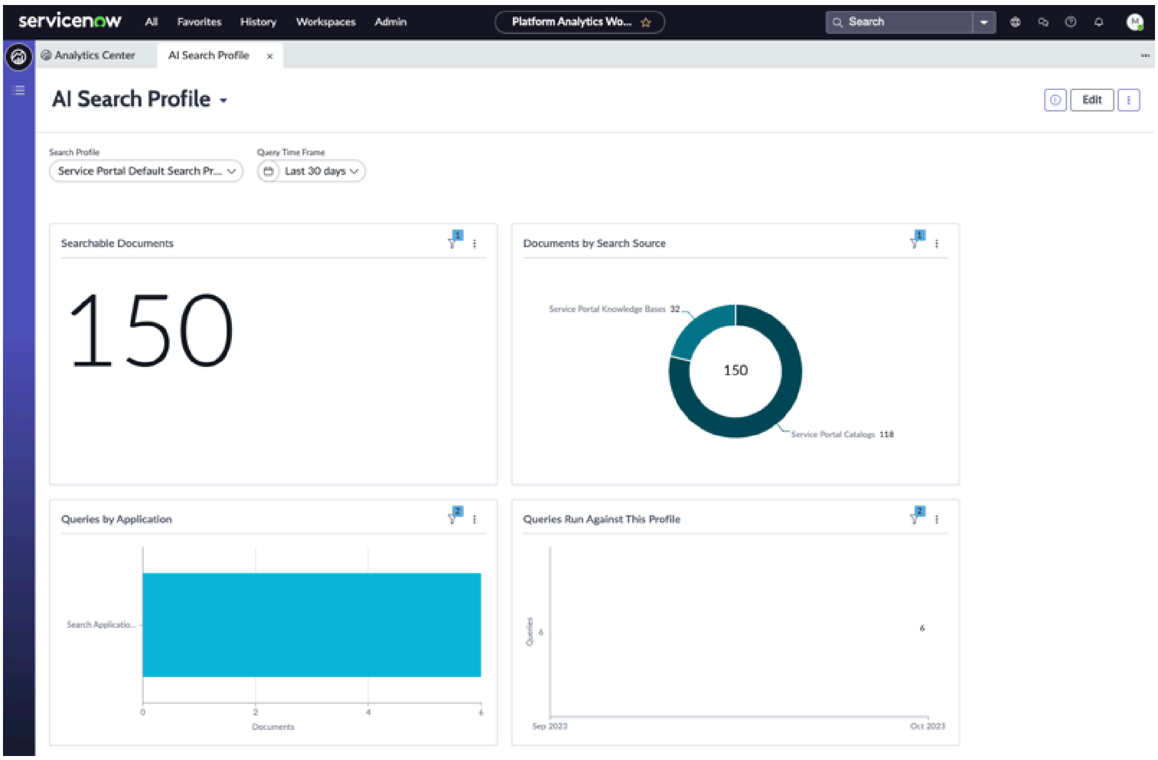

To access the dashboard, navigate to **All** > `AI Search` > `AI Search` **Analytics** > **Search Profile Analytics**.

Note:
If the dashboard displays a `Read operation on table '<name>' from scope 'Advanced AI Search Management Tools' was denied` informational message, ask your administrator to perform the steps described in Create a cross-scope access privilege for the AI Search dashboards for the listed table.

Required ServiceNow AI Platform® roles

To view or edit the dashboard, you must have the ais_admin role.

You may need additional roles to view and select search profiles for some AI Search applications or portals. For example, the **ESC Portal Default Search Profile** for the ESC portal is only visible in the search profile list if you have Employee Center administrator [sn_hr_sp.esc_admin] role.

Use cases

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

User	Dashboard use
AI Search administrator	Reviews indexed record counts and follows search query traffic trends for search profiles defined in AI Search.

Data visualizations

Title	Type	Source table	Description
Searchable Documents	Single Score ( )	sn_ais_admin_tools_ai_search_dashboard_documents_by_sea	Shows the number of indexed records that users can find when searching with the selected search profile.
Documents by Search Source	Donut ( )	sn_ais_admin_tools_ai_search_dashboard_documents_by_sea	Shows the number of indexed records accessible from each search source linked to the selected search profile.
Queries by Application	Bar ( )	sys_search_event	Shows the number of search queries that used the selected search profile in the selected query time frame, grouped by search application.
Queries Run Against This Profile	Line ( )	sys_search_event	Shows the number of search queries that used the selected search profile in the selected query time frame, grouped by month.

AI Search dashboard

The AI Search dashboard summarizes AI Search indexed documents, configuration settings in use, and search query traffic. An interactive filter enables users to select the time frame for analysis of search query traffic.

AI Search dashboard - AI Search Index tab

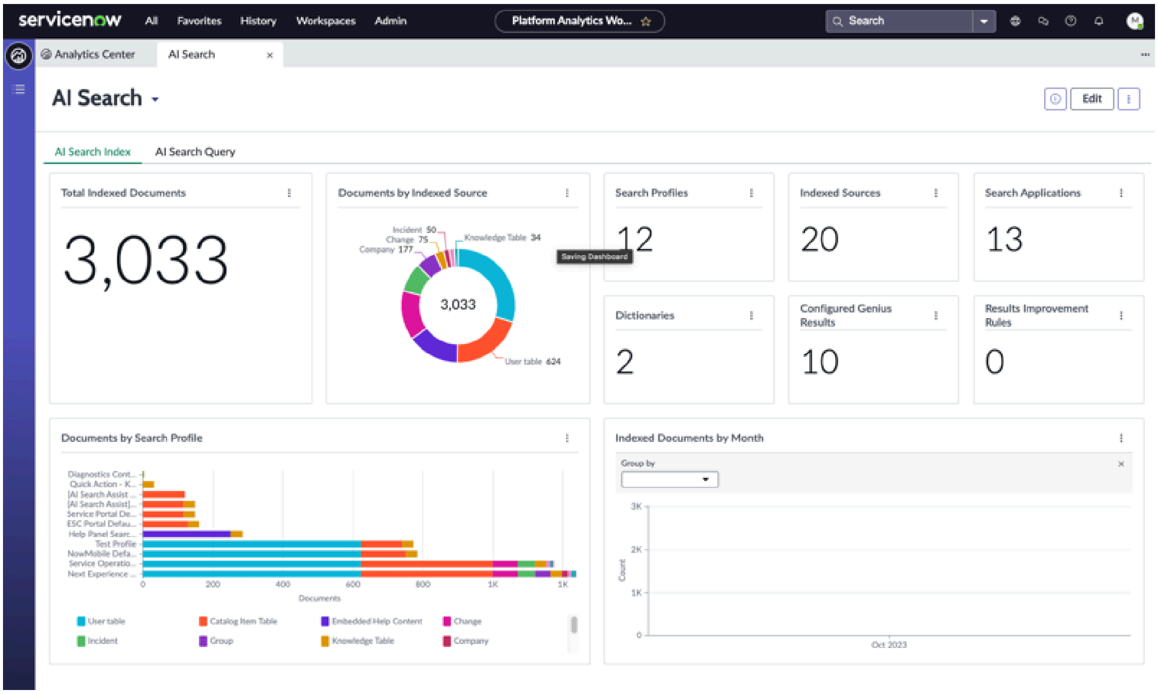

AI Search dashboard - AI Search Query tab

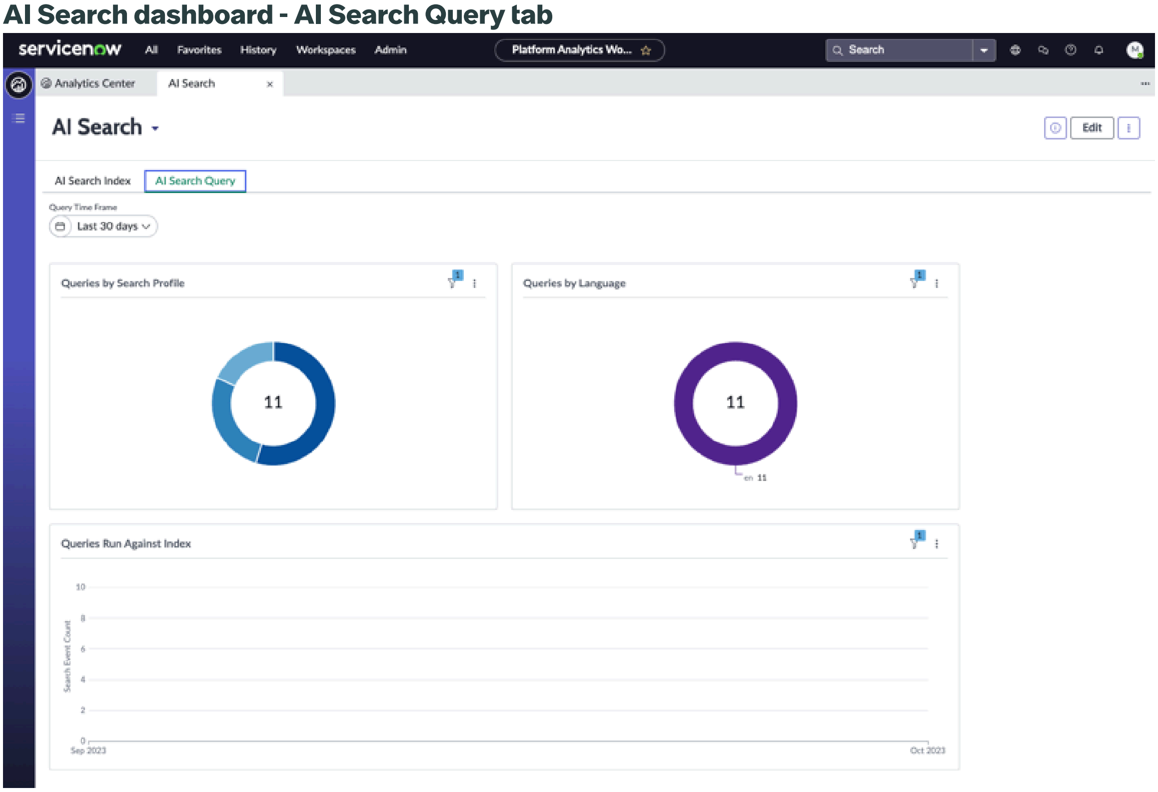

To access the dashboard, navigate to **All** > `AI Search` > `AI Search` **Analytics** > **Search Index Analytics**.

Note:
If the dashboard displays a `Read operation on table '<name>' from scope 'Advanced AI Search Management Tools' was denied` informational message, ask your administrator to perform the steps described in Create a cross-scope access privilege for the AI Search dashboards for the listed table.

Required ServiceNow AI Platform® roles

To view or edit the dashboard, you must have the ais_admin role.

Use cases

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

User	Dashboard use
AI Search administrator	Reviews indexed record counts and follows search query traffic trends for AI Search.

Data visualizations

AI Search Index tab

Title	Type	Source table	Description
Total Indexed Documents	Single Score ( )	sn_ais_admin_tools_st_ai_search_index_analytics_latest_inde	Shows the total number of records indexed by AI Search. **Note:** Because search source filters can exclude indexed records, this number may exceed the total number of searchable records in the index.
Documents by Indexed Source	Donut ( )	sn_ais_admin_tools_ai_search_dashboard_total_indexed_docu	Shows the number of records indexed by AI Search, grouped by indexed source. To drill down and view the number of records for an indexed source grouped by source table, select the graph segment for that indexed source. Refresh the report to return to the top-level view.

AI Search Index tab (continued)

Title	Type	Source table	Description
			Note: Because search source filters can exclude indexed records, the number of records reported for an indexed source may exceed the total number of searchable records from that indexed source.
Search Profiles	Single Score ( )	ais_search_profile	Shows the number of search profiles defined in AI Search.
Indexed Sources	Single Score ( )	ais_datasource	Shows the number of indexed sources defined in the AI Search index.
Search Applications	Single Score ( )	sys_search_context_config	Shows the number of search application configurations defined that use the AI Search search engine.
Dictionaries	Single Score ( )	ais_dictionary	Shows the number of synonym and stop word dictionaries defined in AI Search.
Configured Genius Results	Single Score ( )	ais_genius_result_configuration	Shows the number of Genius Result configurations defined in AI Search.
Results Improvement Rules	Single Score ( )	ais_rule	Shows the number of result improvement rules defined in AI Search.
Documents by Search Profile	Bar ( )	sn_ais_admin_tools_ai_search_dashboard_documents_by_sea	Shows the number of indexed records

AI Search Index tab (continued)

Title	Type	Source table	Description
			searchable through each search profile defined in AI Search.
Indexed Documents by Month	Bar ( )	sn_ais_admin_tools_ai_search_dashboard_documents_by_mon	Shows the number of newly indexed records grouped by month.

AI Search Query tab

Title	Type	Source table	Description
Queries by Search Profile	Donut ( )	sys_search_event	Shows the number of search queries in the selected time frame, grouped by search profile used.
Queries by Language	Donut ( )	sys_search_event	Shows the number of search queries in the selected time frame, grouped by query language (ServiceNow AI Platform language context).
Queries Run Against Index	Line ( )	sys_search_event	Shows the number of search queries in the selected time frame, grouped by month.

Create a cross-scope access privilege for the AI Search dashboards

Enable the AI Search dashboards to read records from tables that cannot be accessed from the Advanced AI Search Management Tools application scope.

Before you begin

The Platform Analytics Solution for Advanced AI Search Management Tools must be activated on your instance. For details on activating this solution, see Activate the Platform Analytics Solution for Advanced AI Search Management Tools.

Role required: admin

About this task

The AI Search Profile and AI Search dashboards belong to the Advanced AI Search Management Tools application scope.

If an indexed source or search source table can't be accessed from this scope, the dashboards display the following informational message:

```
Read operation on table '<name>' from scope 'Advanced AI Search
Management Tools' was denied. The application 'Advanced AI Search
Management Tools' must declare a cross scope access privilege.
```

```
Please contact the application author to update their privilege
requests.
```

To resolve this issue, define a cross-scope table access privilege for the Advanced AI Search Management Tools application.

For more information on cross-scope privileges, see Cross-scope privilege record.

Procedure

1. Set your current application scope context to Advanced AI Search Management Tools using the application picker.
2. Navigate to **All** > **System Applications** > **Application Cross-Scope Access**.
3. Select **New**.
4. On the Cross scope privilege form, enter the following field values.

Field	Value
Target Scope	<table scope>
Target Name	<table name>
Target Type	Table
Operation	Read
Status	Allowed

5. Select **Submit**.

Result

The new cross-scope privilege appears in the Cross scope privilege table.

AI Search Analytics dashboard

The AI Search Analytics dashboard displays key performance metrics, trends, and reports for AI Search usage on its main page. Its Queries page displays additional information on search queries. Interactive filters enable analysts to view performance analytics for individual search applications over the last 180 days and to select the time frame to analyze.

AI Search Analytics dashboard - main page, Overview section

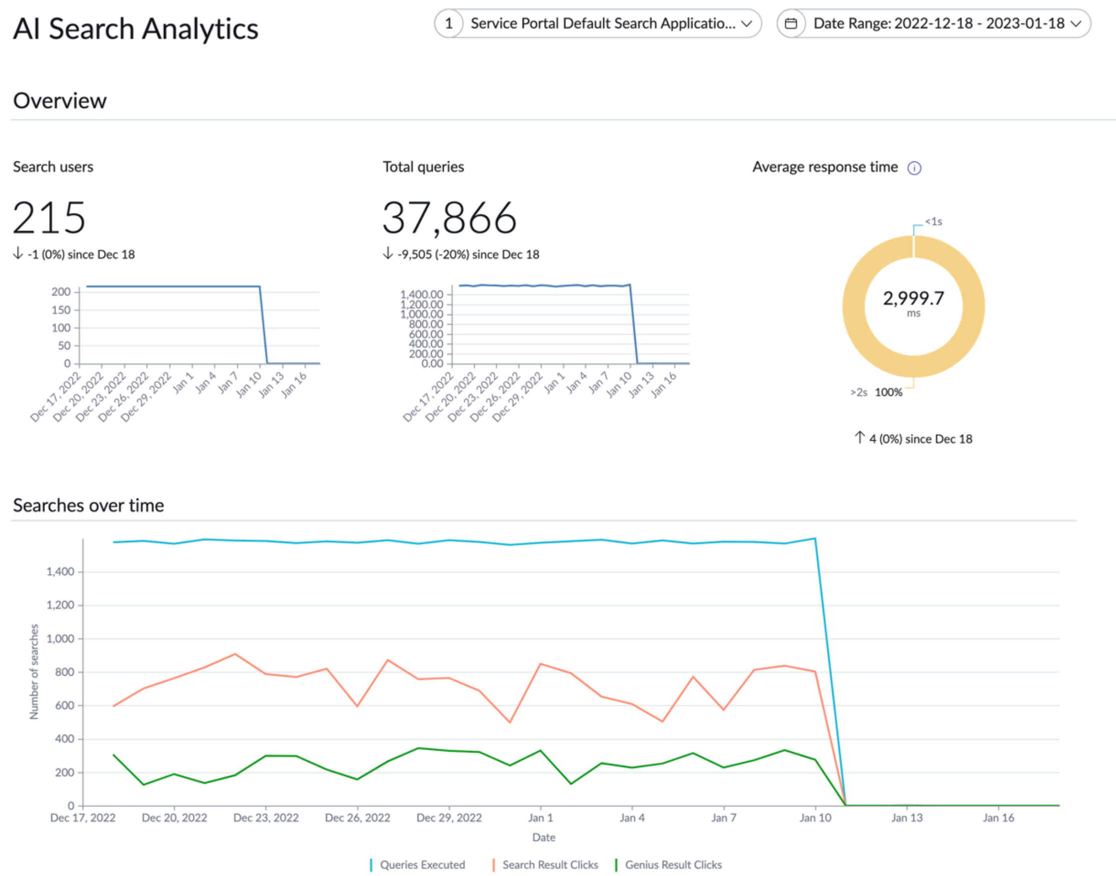

AI Search Analytics dashboard - main page, Search results section

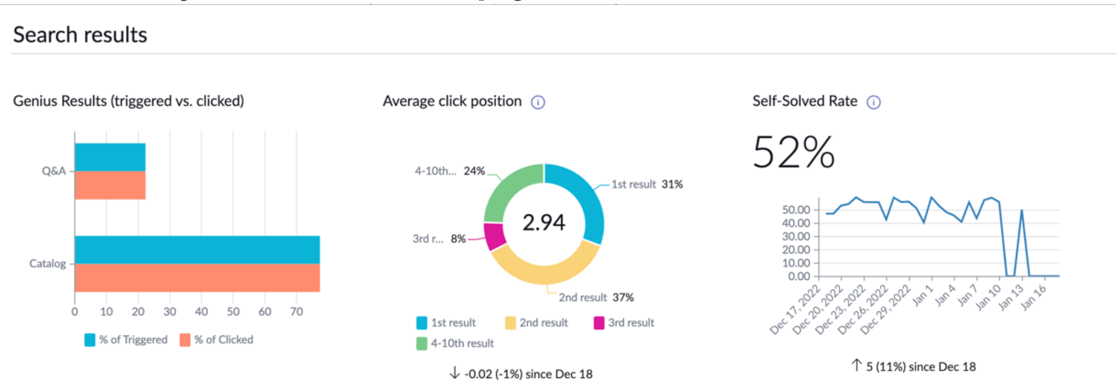

AI Search Analytics dashboard - main page, Queries section

Queries

View all

Top queries	% of total queries
outlook	12.6
headset	12.5
ipad	12.5
backpack	12.5
bluetooth	12.5

Queries with no clicks	Queries
iphone	1,735
outlook	1,729
what is spam	1,571
macbook	1,253
headset	1,232

Queries with no results	Queries
backpack	4,728
bluetooth	4,727

AI Search Analytics dashboard - main page, Top clicked results section

Top clicked results

Title	Data source	% of views	Average click position
How to configure VPN for Apple Devices	Knowledge	8.7	1
(empty)	Knowledge	8.5	2
Logitech USB Headset for PC & Mac	Catalog Item	8.4	2
What is a cookie?	Knowledge	8.4	2
Plantronics Wideband Headset	Catalog Item	8.3	1

AI Search Analytics dashboard - Queries page, Queries over time report

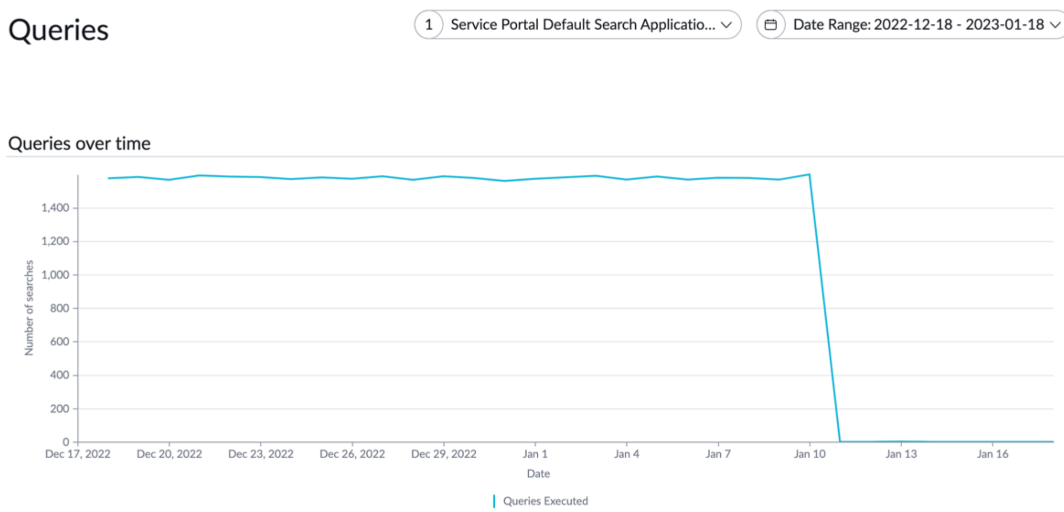

AI Search Analytics dashboard - Queries page, Query Analysis section

Query Analysis

Top queries | Queries with no clicks | Queries with no results

Top queries	Language	% of total queries	Average click position
outlook	EN	12.6	5.5
headset	EN	12.5	1.5
ipad	EN	12.5	2
backpack	EN	12.5	
bluetooth	EN	12.5	
iphone	EN	12.5	1.5
what is spam	EN	12.5	1.5
macbook	EN	12.4	5.5
email	EN	0	
apple	EN	0	
phone	EN	0	4

To access the AI Search Analytics dashboard, navigate to **All** > **User Experience Analytics** > **Dashboard**, select the SearchDashboard record, then select the Search Analytics icon ( ) or link in the modules menu. To learn more about the Usage Insights dashboard, see Overview of the dashboard.

Note:
If the dashboard doesn't load, follow the steps from Configure Service Portal to send analytics data to enable it.

To access the Queries page, select **View all** in the **Queries** section header of the main page.

You can improve the visual experience by using the dark theme. For more information, see Working with themes in Next Experience.

Required ServiceNow AI Platform roles

To view the AI Search Analytics dashboard, you must have the ais_admin role and one of the following roles:

- analytics_admin
- analytics_viewer
- portal_analytics_admin
- portal_analytics_viewer
- web_analytics_admin
- web_analytics_viewer

Use cases

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

User	Dashboard use
Any of the following: • Portal Analytics administrator or viewer • User Experience Analytics administrator or viewer • Web Analytics administrator or viewer	• View and monitor search usage and search performance for a specified search application and date range • Review search trends to better understand the needs of search application end users • Gain insights on how to improve search experience and knowledge coverage by identifying common search queries that yield no results or low-quality results

Interactive filters

Filter	Description
Search application	Specifies the search application for which you want to see search analytics metrics, trends, and reports. You can search for a search application or select its name in the list. Select **Apply** to make your filter selection take effect. **Note:** You can only select a single search application in this filter. The filter list includes search applications that are configured to use AI Search and that are used in any of the following contexts:

Filter	Description
	• A service portal • A Next Experience workspace search • Next Experience global search • Mobile Platform search The dashboard reports metrics for all search queries in the selected search application that satisfy your **Date range** constraint. These may include user searches submitted directly to the search application and searches submitted through interaction with other ServiceNow AI Platform features, such as Virtual Agent.
Date range	Specifies the range of dates for which you want to see search analytics metrics, trends, and reports. By default, the dashboard displays data from the last 30 days. You can choose one of the preset date ranges or set your own start and end dates. Select **Apply** to make your filter selection take effect. The dashboard retains data from the last 180 days. If you select a date from outside of this range, the system displays an error message and the dashboard reverts to the last valid date range selection. If you haven't selected any valid date range, the dashboard reverts to its default date range. **Note:** When you select a custom date range, the dashboard may take longer to load than when you select a preset date range. The dashboard reports metrics for all search queries in the selected date range that satisfy your **Search application** constraint. These may include user searches submitted directly to the search application and searches submitted through interaction with other ServiceNow AI Platform features, such as Virtual Agent.

Metrics, trends, and reports

The dashboard includes a main page and a Queries page accessible from the **View all** link in the **Queries** section.

Note:
To improve performance, the AI Search Analytics dashboard caches results from its source tables for 60 seconds. You may not see search queries from the last 60 seconds reflected in the dashboard's metrics, trends, and visualizations.

Main page

Title	Visualization type	Source tables	Description
Overview section			
Search users	Line ( )	• Metric and trend: AI Search Analytics Dashboard Unique Search Users [sn_ais_admin_tools_ai_search_analytics_dashboard_unique • Report: AI Search Analytics Dashboard Daily Search Users [sn_ais_admin_tools_ai_search_analytics_dashboard_daily_s	• Metric indicates the total count of unique users who submitted search queries in the selected application and date range. • Trend describes how the number of unique search users changed in the selected application and date range. • Report displays the number of unique search users in the selected application for dates in the selected range. **Note:** Searches performed while impersonating another user are ignored for this metric, trend, and report.
Total queries	Line ( )	AI Search Analytics Dashboard Total Searches and Clicks by Date [sn_ais_admin_tools_ai_search_analytics_dashboard_total_sea	• Metric indicates the total count of search queries submitted in the selected application and date range. • Trend describes how the count of search queries submitted changed in the selected application and date range.

Main page (continued)

Title	Visualization type	Source tables	Description
			• Report displays the count of search queries submitted in the selected application for dates in the selected range.
Average response time	Donut ( )	• Metric and trend: AI Search Analytics Dashboard Average Search Response Time [sn_ais_admin_tools_ai_search_analytics_dashboard_averag • Report: AI Search Analytics Dashboard Daily Average Response Time [sn_ais_admin_tools_ai_search_analytics_dashboard_daily_a	• Metric indicates the average search query response time for the selected application and date range. • Trend describes how the average search query response time changed in the selected application and date range. • Report displays the distribution of search query response times in the selected application for dates in the selected range. Response times are categorized as <1 second, 1–2 seconds, or >2 seconds.
Searches over time section			
Searches over time	Line ( )	AI Search Analytics Dashboard Total Searches and Clicks by Date [sn_ais_admin_tools_ai_search_analytics_dashboard_total_sea	The count of search queries executed, search result clicks, and Genius Result clicks in the selected application for dates in the selected range.
Search results section			
Genius Results (triggered vs. clicked)	Horizontal Bar ( )	AI Search Analytics Dashboard Trending Genius Result [sn_ais_admin_tools_ai_search_analytics_dashboard_trending_	The most frequently triggered and selected Genius Result configurations

Main page (continued)

Title	Visualization type	Source tables	Description
			in the selected application and date range.
Average click position	Donut ( )	• Metric and trend: AI Search Analytics Dashboard Average Click Position [sn_ais_admin_tools_ai_search_analytics_dashboard_averag • Report: AI Search Analytics Dashboard Daily Average Click Position [sn_ais_admin_tools_ai_ shboard_daily_a	• Metric indicates the average position of user-selected search query results for the selected application and date range. **Note:** Positions indicate result numbers. For example, a position of 2 means that the user selected the second result returned for a search query. • Trend describes how the average position of user-selected search query results changed in the selected application and date range. • Report displays the distribution of positions for user-selected search query results in the selected application for dates in the selected range.
Self-Solved Rate	Line ( )	AI Search Analytics Dashboard Total Searches and Clicks by Date [sn_ais_admin_tools_ai_search_analytics_dashboard_total_sea	• Metric indicates the percentage of search queries that produced a search result click for the selected application and date range.

Main page (continued)

Title	Visualization type	Source tables	Description
			• Trend indicates how the percentage of search queries that produced a search result click changed in the selected application and date range. • Report indicates the percentage of search queries that produced a search result click in the selected application for dates in the selected range.
Queries section			
Top queries	Table ( )	AI Search Analytics Dashboard Trending Search Terms [sn_ais_admin_tools_ai_search_analytics_dashboard_trending_	The most frequently submitted search queries in the selected application and date range, by percentage of total search queries for the application and date range. To access the Queries dashboard and see more details and entries for this report, select **View all** in the **Queries** section header.
Queries with no clicks	Table ( )	AI Search Analytics Dashboard Queries with no Clicks [sn_ais_admin_tools_ai_search_analytics_dashboard_queries_	The most frequently submitted search queries in the selected application and date range for which users didn't select any search result, by count. To access the Queries dashboard and see more details and entries for this report,

Main page (continued)

Title	Visualization type	Source tables	Description
			select **View all** in the **Queries** section header.
Queries with no results	Table ( )	AI Search Analytics Dashboard Queries with no Results [sn_ais_admin_tools_ai_search_analytics_dashboard_queries_w	The most frequently submitted search queries in the selected application and date range that produced no results, by count. To access the Queries dashboard and see more details and entries for this report, select **View all** in the **Queries** section header.
Top clicked results section			
Top clicked results	Table ( )	AI Search Analytics Dashboard Trending Search Content [sn_ais_admin_tools_ai_search_analytics_dashboard_trending_	The most frequently clicked search results in the application and date range, by percentage of views that resulted in a click.

Queries page

Title	Report type	Source tables	Description
Queries over time	Line ( )	AI Search Analytics Dashboard Total Searches and Clicks by Date [sn_ais_admin_tools_ai_search_analytics_dashboard_total_sea	The count of search queries executed in the selected application for dates in the selected range.
Query Analysis section			
Top queries	Table ( )	AI Search Analytics Dashboard Trending Search Terms [sn_ais_admin_tools_ai_search_analytics_dashboard_trending_	The most frequently submitted search queries in the selected application and date range, by percentage of total search queries for the application and date range. To view additional details on a search

Queries page (continued)

Title	Report type	Source tables	Description
			query, including its average response time and top results, select the query entry in the Top queries column. To display additional report entries, select **Load More** at the bottom of the table.
Queries with no clicks	Table ()	AI Search Analytics Dashboard Queries with no Clicks [sn_ais_admin_tools_ai_search_analytics_dashboard_queries_	The most frequently submitted search queries in the selected application and date range for which users didn't select any search result, by count. To display additional report entries, select **Load More** at the bottom of the table.
Queries with no results	Table ()	AI Search Analytics Dashboard Queries with no Results [sn_ais_admin_tools_ai_search_analytics_dashboard_queries_	The most frequently submitted search queries in the selected application and date range that produced no results, by count. To display additional report entries, select **Load More** at the bottom of the table.

Customize the banner logo for the AI Search Analytics dashboard

Add your own logo to the banner for the AI Search Analytics dashboard.

Before you begin

The Platform Analytics Solution for Advanced AI Search Management Tools must be activated on your instance. For details on activating this solution, see Activate the Platform Analytics Solution for Advanced AI Search Management Tools.

Role required: admin

About this task

The AI Search Analytics dashboard displays the ServiceNow logo in the banner. You can modify the UX theme assets for the dashboard to display your own logo instead.

Procedure

1. Switch to the Advanced AI Search Management Tools application.
   For details on switching applications, see Switch between applications.
2. To replace the default **header_for_dashboard** UX theme asset with a version containing your logo, perform the following steps.

   a. Navigate to the UX Theme Asset [sys_ux_theme_asset] table's list view by selecting **All**, entering `sys_ux_theme_asset.list` in the **Filter** field, and pressing Enter.

   b. Open the **header_for_dashboard** record.

   c. Select **Delete**.

   d. Select **New**.

   e. On the UX Theme Asset form, enter the following field values.

Field	Value
Category	Image
Name	header_for_dashboard

   f. In the form header, select the Manage Attachments icon ( ) and attach your logo image to the new record.
   To learn more about managing record attachments, see Manage attachments.

   g. Select **Submit**.

3. To link your new UX theme asset to the banner for the AI Search Analytics dashboard, perform the following steps.

   a. Navigate to the UX Theme [sys_ux_theme] table's list view by selecting **All**, entering `sys_ux_theme.list` in the **Filter** field, and pressing Enter.

   b. Open the **SearchAnalytics Dashboard Theme** record.

   c. Select the **Legacy: Experience Theme** tab.

   d. In the UX Theme Assets list, edit the existing asset record.

   Note:
   If you don't see the UX Theme Assets list, configure the form layout, adding **UX Theme Assets->Theme** to the list of selected fields. When prompted, switch to editing the form section in the Global application. For details on this procedure, see Configuring the form layout.

e. On the UX Theme Assets form, in the **Asset** field, enter the value `header_for_dashboard`.

f. Select **Update**.

Result

Your custom logo image appears in the banner for the AI Search Analytics dashboard.

Configure Service Portal to send analytics data

Enable loading of the AI Search Analytics dashboard by configuring Service Portal to send analytics data.

Before you begin

The Platform Analytics Solution for Advanced AI Search Management Tools must be activated on your instance. For details on activating this solution, see Activate the Platform Analytics Solution for Advanced AI Search Management Tools.

Role required: admin

About this task

If the AI Search Analytics dashboard doesn't load, you must configure tracking of Usage Insights data in the Service Portal. This configuration is a prerequisite for the dashboard.

Procedure

1. Navigate to **All** > **Service Portal** > **Portals**.
2. Edit the **Service Portal** record.
3. Select **Create Analytics Settings**.
4. On the Usage Insights Settings form, select the **Enable Unauthenticated User Tracking** option, then select **Update**.

Result

The AI Search Analytics dashboard should now load correctly.

Search Preview UI for AI Search

The Search Preview UI enables you to test AI Search queries using settings from a selected search profile. Administrator tools help you review search query performance data and feedback, debug queries, and override default query settings for testing purposes.

Access the Search Preview UI by navigating to *`AI Search`* > **Preview** > **Search Preview**.

Required ServiceNow AI Platform® roles

- The ais_admin role is required to view and use the Search Preview UI.
- The impersonator role is required for some AI Search administrator tools.
- The ais_high_security_admin role is required to bypass search source and content security filtering. For details on this procedure, see Diagnose search result access issues using the Search Preview UI.

Search Preview UI components

Search Preview UI

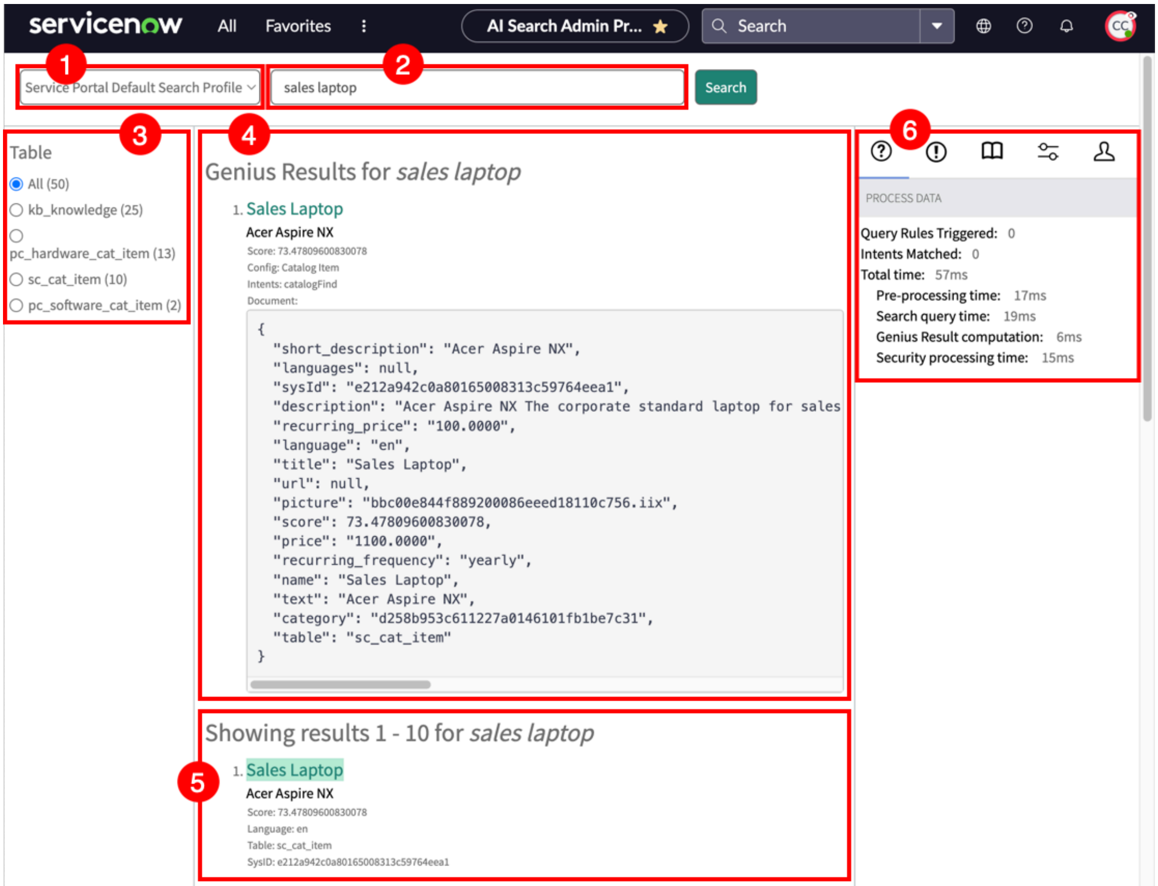

Search Preview UI components

Component	Description
1. Search profile list	Select the search profile to apply when previewing search queries. The list only includes published search profiles that have one or more linked search sources. **Note:** If a search profile has unpublished changes, the list displays a warning icon (!) and shows [DRAFT] after the profile's name. Publish the search profile to see these changes reflected in the search query preview.
2. Search input field	Enter search query terms and operators in this field. To preview the query's results using settings from the selected search profile, select **Search**.

Search Preview UI components (continued)

Component	Description
3. Table facet filter	View the distribution of search query results by source table. To filter the query results by table, select a table's name. To remove the table facet filter, select **All**.
4. Genius Result answers	Review the key-value pairs for each Genius Result answer the search query produces when executed with the selected search profile. **Note:** Genius Result answers may not appear for all search queries.
5. Search query results	Review the results the search query produces when executed with the selected search profile. To page through results, select **Next** or **Previous**.
6. AI Search administrator tools	Select the icon for the administrator tool that you want to use: • Process data tool (): View summary information on the search query's triggered result improvement rules, matched NLU model intents, and processing time. • Alerts tool (): Review alert feedback messages produced during processing of the search query. • Dictionary and query feedback tool (): Review dictionary and query feedback messages and NLU model intent data for the search query. Enable debugging. • Context settings tool (): Specify user context values for the search query to trigger result improvement rules. • User impersonation tool (): Change the search query's user or locale to test access settings and view translated documents. For more detailed descriptions of each administrator tool's output, controls, and fields, see the following reference section.

Search preview admin tools

Understand the output, controls, and fields provided by admin tools in the search preview UI.

Process data tool

Entry	Description
Query Rules Triggered	Count of result improvement rules triggered by the search query.
Intents Matched	Count and names of NLU model intents detected in the search query.
Total time	Total time spent processing the search query. Includes pre-processing, search query, Genius Result computation, and security processing times.
Pre-processing time	Time spent preparing the search query for execution. Includes time spent looking up the user's access permissions for content security.
Search query time	Time spent retrieving records that match the search query and generating the result set.
Genius Result computation	Time spent computing Genius Result answers for the search query.
Security processing time	Time spent filtering inaccessible records out of the search query's result set based on the user's access permissions.

Alerts tool

This tool displays a list of alert feedback messages produced while processing the search query.

If the search query produced no alert feedback messages, the tool instead displays `No Alert Feedback`.

Dictionary and query feedback tool

This tool displays separate panes for dictionary feedback entries and query feedback controls.

Dictionary feedback

Entry	Description
Stop words	List of feedback messages relating to stop word removal for the search query.
Spellcheck	List of feedback messages relating to typo handling auto-correction for the search query.
Synonyms	List of feedback messages relating to synonym expansion for the search query.

Query feedback

Control	Description
Show all feedback	Display a JSON array of feedback messages produced for the search query.
Show all intent data	Display a JSON array of NLU model intents detected in the search query. **Note:** To use this tool, you must have the impersonator role.
Debug	Enable this option to set *sysparm_debug* to *true* when submitting search preview queries. **Note:** To use this tool, you must have the impersonator role.

Context settings tool

With these tools, you can specify user context values to test user context triggers for result improvement rules.

Field	Description
User Domain	To set a domain in the user context for search preview, enter the name of a group from the Group [sys_user_group] table.
User Company	To set a company in the user context for search preview, enter the name of a company from the Company [core_company] table.
User Office	To set an office in the user context for search preview, enter the name of a building from the Building [cmn_building] table.
User Department	To set a department in the user context for search preview, enter the name of a department from the Department [cmn_department] table.
User Title	To set a title in the user context for search preview, select it from the list. This list includes all Title values found in the User [sys_user] table.

Field	Description
User Location	To set a location in the user context for search preview, enter the name of a location from the Location [cmn_location] table.
User Role	To set a role in the user context for search preview, enter the name of a role from the Role [sys_user_role] table.
User Country	To set a country in the user context for search preview, select it from the list. This list includes all countries found in user locations from the User [sys_user] table.
User Language	To set a language in the user context for search preview, select it from the list of supported languages.
User Office Country	To set an office country in the user context for search preview, select it from the list. This list includes all countries found in user locations from the User [sys_user] table.
User Office City	To set an office city in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.
User Office State	To set an office state in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.
User OS	To set an operating system in the user context for search preview, enter its name in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.
User Device Type	To set a user device type in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.

Field	Description
User OS Version Type	To set an operating system version type in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.
User Browser Agent	To set a browser agent string in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.
User Browser	To set a browser name in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.
User Device Form	To set a device form in the user context for search preview, enter it in this field. You can enable session debugging for AI Search to review values submitted for this user context field in user queries.

User impersonation tool

Field	Description
Search as	To submit search preview queries as another user, enter the user's name from the User [sys_users] table. You must have the impersonator role. Use this control to see how content security affects search query results for specific user accounts. Searches performed while impersonating another user don't affect the **Search users** metric, trend, and report on the AI Search Analytics dashboard. **Note:** Changing the value of this field terminates elevated privileges for the ais_high_security_admin role. To bypass filters after changing this field's value, you must re-elevate to the ais_high_security_admin role.

Field	Description
Query Locale	To submit search preview queries using a specific supported language, select the language from the list. Use this control to observe differences in search query results for different languages.

Advanced AI Search Management Tools reference

This reference topic describes system components installed with the Advanced AI Search Management Tools application.

Plugins installed with Advanced AI Search Management Tools

Plugin	Description
Advanced AI Search Management Tools (sn_ais_admin_tools)	Includes AI Search dashboards and Search Preview tool.

Scheduled jobs installed with Advanced AI Search Management Tools

Scheduled job	Description
Load AI Search Dashboard Documents by Month	Populates the Indexed Documents by Month data visualization on the AI Search dashboard. Runs every minute.
Load AI Search Dashboard Documents by Profile	Populates the Documents by Search Profile data visualization on the AI Search dashboard. Runs every hour.
Load AI Search Dashboard Documents by Search Source	Populates the Documents by Search Source data visualization on the AI Search Profile dashboard. Runs every hour.
Load AI Search Dashboard Total Indexed Documents	Populates the Total Indexed Documents data visualization on the AI Search dashboard. Runs every hour.
Update Search Analytics Data	Pre-calculates metrics to improve load time for the AI Search Analytics dashboard. Runs every hour. **Note:** If you pause or disable this scheduled job, the AI Search Analytics dashboard takes longer to load and displays the metrics calculated by the last scheduled job run. You can specify a new date range for the dashboard to force recalculation of metrics.

List of all components installed with Advanced AI Search Management Tools

To view the complete list of components installed with Advanced AI Search Management Tools, follow the steps described in Find components installed with an application. The application's package name is **Advanced** `AI Search` **Management Tools**.

AI Search for Next Experience

This ServiceNow® Store application enables administrators to configure AI Search as the search engine for Unified Navigation searches in Next Experience, including global search and configurable workspace search.

AI Search for Next Experience is automatically installed and enabled in the Washington DC release, so you don't need to perform the initial migration and enablement configuration steps. To add new Unified Navigation workspace search configurations later, see Add a new Unified Navigation workspace search configuration to AI Search for Next Experience.

To learn about how global search and configurable workspace search work when AI Search for Next Experience is enabled, see Using AI Search for Next Experience.

Install AI Search for Next Experience

Install the AI Search for Next Experience application from the ServiceNow® Store.

Before you begin

If your instance was zBooted in Washington DC, the AI Search for Next Experience application is already installed. You only need to install the application if you upgraded your instance from Tokyo or a previous release.

- 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.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the AI Search for Next Experience application (sn_aisearch_global) 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.

   Visit the ServiceNow Store website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.
3. Select **Install**.

What to do next

With the plugin installed, you can walk through Guided Setup tasks to enable AI Search as the search engine for the Unified Navigation search field in Next Experience. To get started with Guided Setup, see Verify that your instance meets the prerequisites for AI Search for Next Experience.

Configuring AI Search for Next Experience

Use a Guided Setup module to step through configuration of search application settings and enablement of AI Search for Next Experience on your instance.

The `Guided Setup` **for Zing to** `AI Search` **Migration** module walks you through the following configuration steps:

1. Verifying that your instance satisfies the prerequisites for AI Search for Next Experience.
2. Migrating your existing Zing search application configurations for use with AI Search.
3. Enabling AI Search for Next Experience on your instance. This step makes AI Search the search engine for global and workspace searches in the Unified Navigation search field.

Once you've enabled AI Search for Next Experience, you can use the Guided Setup module at any time to perform either of these tasks:

- Enable AI Search for a newly installed workspace by adding its search application configuration into AI Search for Next Experience.
- Revert to using Zing as the search engine for the Unified Navigation search field by disabling AI Search for Next Experience.

Verify that your instance meets the prerequisites for AI Search for Next Experience

Walk through Guided Setup tasks to confirm that your instance meets the requirements for AI Search for Next Experience.

Before you begin

The AI Search for Next Experience ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search for Next Experience.

Role required: ais_admin

About this task

The **Prerequisites** Guided Setup tasks walk you through verifying that your instance can use AI Search for Next Experience. Complete these tasks before you begin migrating your Zing search application configurations.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **for** `Next Experience` > `Guided Setup` **for Zing to** `AI Search` **Migration**.
2. Select **Get Started**.
3. In the **Prerequisites** category, select **Get Started**.
4. Confirm that AI Search is provisioned and enabled on your instance.

    a. In the **Confirm that** `AI Search` **is enabled** task, select **Configure**.
    A modal window reports whether AI Search is provisioned and enabled on your instance.

    b. If the window reports `AI Search` **is ready**, select **Close**, then select **Mark as Complete** for the **Confirm that** `AI Search` **is enabled** task.

    c. If the window reports any other status, select **Close**, then repeat this task after AI Search is provisioned and enabled on your instance.

5. Confirm that Next Experience is enabled on your instance.

a. In the **Confirm that** *`Next Experience`* **is Enabled** task, select **Configure**.
A modal window reports whether Next Experience is enabled on your instance.

b. If the system reports *`Next Experience`* **is enabled**, select **Close**, then select **Mark as Complete** for the **Confirm that** *`Next Experience`* **is Enabled** task.

c. If the system reports any other enablement status, ask your administrator to enable Next Experience on your instance, then repeat this task after it's enabled.

For details on enabling Next Experience, see Enable Next Experience.

6. Index content and metadata for AI Search indexed sources used in global search.

a. In the **Index searchable content and metadata for global search sources** task, select **Configure**.
The AI Search Indexed Sources list opens in a new browser tab.

b. Select all indexed sources marked with the alert triangle icon (⚠), then index their content by selecting **Actions on selected rows...** > **Index All Tables**.

For more information on indexing tables from multiple indexed sources, see Perform a full table index or reindex for multiple indexed sources.

c. When indexing is complete, close the browser tab and return to the **Prerequisites** Guided Setup category tab.

d. Select **Mark as Complete** for the **Index searchable content and metadata for global search sources** task.

7. When all three **Prerequisites** tasks are marked as complete, select the Back icon (<) to return to the top-level *`Guided Setup`* **for Zing to** *`AI Search`* **Migration** page.

What to do next

Now that you've satisfied the prerequisites, it's time to migrate your existing Zing search application configurations for use with AI Search. For details on this procedure, see Migrate Zing search application configurations to AI Search.

Migrate Zing search application configurations to AI Search

Walk through a Guided Setup task to migrate your existing Zing search application configurations to AI Search application configurations.

Before you begin

The AI Search for Next Experience ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search for Next Experience.

You must have already completed the **Prerequisites** Guided Setup tasks. For details on completing these tasks, see Verify that your instance meets the prerequisites for AI Search for Next Experience.

Role required: ais_admin

About this task

The base system includes Zing search application configurations for global search and some workspaces. To use AI Search for Next Experience, you need AI Search configurations for these search applications.

You can create the new AI Search application configurations you need by migrating settings from your existing Zing search application configurations. The **Convert Zing search applications** Guided Setup task walks you through this procedure.

> **Note:**
> When you migrate a Zing search application configuration, the system creates a new search application configuration for use with AI Search. This procedure doesn't delete or change your original Zing search application configuration.

Migration of a Zing search application configuration to AI Search includes these steps:

1. The migration tool analyzes your selected Zing search application configuration and determines what record changes in ServiceNow AI Platform® tables are needed to replicate its settings in a new AI Search application configuration.
2. You accept the proposed migration changes for review, and then commit them to make the required record changes in ServiceNow AI Platform tables.
3. You index content and metadata from any new indexed sources that the migration job created.

For more information on AI Search application configurations, see Search application configurations. To learn more about indexing content and metadata from AI Search indexed sources, see Indexed sources.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **for** `Next Experience` > `Guided Setup` **for Zing to** `AI Search` **Migration**.
2. In the **Convert Zing search applications** category, select **Get Started**.
3. In the **Migrate settings for a** `Next Experience` **search application** task, open the Zing to AI Search Migration Job form by selecting **Configure**.
   A modal window shows a list of Zing search application configurations and their migration status. Custom search application names are only translated to your session language if you've created appropriate Message [sys_ui_message] table records for them. For details on translating message strings by creating Message table records, see Translate a client script message.
4. Close the modal list of Zing search application configurations.

   > **Note:**
   > You can reopen the list of Zing search application configurations at any time by selecting the **View Search Application Migration Status** related link.

5. For each Zing search application configuration with **Needs migration** status, migrate the search application configuration to AI Search.

a. In the **Search Application Configuration** field, use the Lookup using list icon ( ) to select the Zing search application configuration that you want to migrate to AI Search.

Note:
Only Zing search application configurations with **Needs migration** status are available for selection. Application configurations with **Migrated** status already have corresponding AI Search application configurations.

b. **Optional:** To have the system send email notifications to users when indexing of content for the migration job starts and completes, add those users to the watch list.

Note:
The system automatically adds the current user to the watch list when a migration job is created or updated.

c. Select **Start Migration**.
The migration tool converts the Zing search application configuration for use with AI Search, and the migration job's state changes from **Pending Conversion** to **Configuration Converted**.

d. Select **Accept All Changes**.
The system displays proposed changes to system tables in the Changes to Review related list.

e. Select **Commit Changes**.
The system commits the changes to system tables shown in the Changes to Review related list, and the migration job's state changes to **Configuration Committed**.

f. Select **Index Searchable Content**.
AI Search indexes content and metadata for new indexed sources created by the migration job. When indexing completes, the job's state changes to **Indexing Complete** and the modal window with the list of Zing search application configurations reopens.

6. When all Zing search application configurations in the modal list show **Migrated** status, close the browser tab and return to the **Convert Zing search applications** Guided Setup category tab, then select the Back icon ( ) to return to the top-level `Guided Setup` **for Zing to** `AI Search` **Migration** page.

What to do next

Now that you've got AI Search application configurations for global search and your workspaces, it's time to enable AI Search for Next Experience on your instance. For details on this procedure, see Enable AI Search for Next Experience.

Enable AI Search for Next Experience

Walk through a Guided Setup task to configure AI Search as the search engine for global and workspace searches from the Unified Navigation search field.

Before you begin

The AI Search for Next Experience ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search for Next Experience.

You must have already completed the **Convert Zing search applications** Guided Setup task. For details on completing this task, see Migrate Zing search application configurations to AI Search.

Role required: ais_admin

About this task

The **Activate** *AI Search* **for all** *Next Experience* **search applications** Guided Setup task walks you through enabling the AI Search for Next Experience application.

Once you enable the application, Next Experience search applications use AI Search instead of Zing. Each search application uses the specific settings defined in its search application configuration.

Procedure

1. Navigate to **All** > *AI Search* > *AI Search* **for** *Next Experience* > *Guided Setup* **for Zing to** *AI Search* **Migration**.
2. In the **Complete migration to** *AI Search* category, select **Get Started**.
3. In the **Activate** *AI Search* **for all** *Next Experience* **search applications** task, select **Configure**.
   A modal window reports whether AI Search for Next Experience is enabled on your instance.
4. If the system reports *AI Search* **for** *Next Experience* **is disabled**, select **Enable** *AI Search* **for** *Next Experience*.

   Note:
   If AI Search for Next Experience is already enabled, ensure that all migrated search application configurations are activated by disabling the application and then re-enabling it.

5. When the system reports *AI Search* **for** *Next Experience* **is enabled**, select **Close**, then select **Mark as Complete** for the **Activate** *AI Search* **for all** *Next Experience* **search applications** task.

Result

The system uses AI Search as the search engine for global search and Next Experience workspace application searches issued from the Unified Navigation search field.

What to do next

Notify all users of your instance about the following details.

- Users need to log out and back in before they can use AI Search for Next Experience.
- Once you enable AI Search for Next Experience, only workspace applications that use AI Search as their search engine appear in the Unified Navigation search context menu. To search from workspace applications that use the Zing text indexing and search engine, you must navigate to the workspace instead of using the Unified Navigation search context menu.

Add a new Unified Navigation workspace search configuration to AI Search for Next Experience

Enable Unified Navigation search for newly installed workspaces in AI Search for Next Experience.

Before you begin

The AI Search for Next Experience ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search for Next Experience.

You must have already enabled AI Search for Next Experience on your instance. For details on this procedure, see Enable AI Search for Next Experience.

Role required: ais_admin

About this task

When a new workspace application that includes a search application configuration is installed on your instance, you can enable search for the workspace in AI Search for Next Experience by following this procedure.

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **for** `Next Experience` > `Guided Setup` **for Zing to** `AI Search` **Migration**.
2. In the **Post-Migration Actions** category, select **Get Started**.
3. In the **When a new** `Next Experience` **workspace has been added to your instance** task, open the Zing to AI Search Migration Job form by selecting the **migration tool** link.
   A modal window shows a list of Zing search application configurations and their migration status.
4. If the new workspace's search application configuration shows **Migrated** status, reindex its associated content and publish its search profiles and dictionaries.

   a. Repeat the **Index searchable content and metadata for global search sources** task in the **Prerequisites** Guided Setup category.
      For details on this procedure, see Verify that your instance meets the prerequisites for AI Search for Next Experience.

   b. Publish all search profiles and dictionaries provided by the new workspace.
      To learn how to publish search profiles, see Publish a search profile. For details on publishing stop word and synonym dictionaries, see Create stop words and Create synonyms.

5. If the new workspace's search application configuration shows **Needs migration** status, migrate it for use with AI Search.

   a. Close the modal window.

      Note:
      You can reopen the list of Zing search application configurations at any time by selecting the **View Search Application Migration Status** related link.

   b. In the **Search Application Configuration** field, use the Lookup using list icon ( ) to select the new Zing search application configuration.

      Note:
      Only Zing search application configurations with **Needs migration** status are available for selection. Application configurations with **Migrated** status already have corresponding AI Search application configurations.

   c. **Optional:** To have the system send email notifications to users when indexing of content for the migration job starts and completes, add those users to the watch list.

      Note:
      The system automatically adds the current user to the watch list when a migration job is created or updated.

   d. Select **Start Migration**.

The migration tool converts the Zing search application configuration for use with AI Search, and the migration job's state changes from **Pending Conversion** to **Configuration Converted**.

e. Select **Accept All Changes**.
The system displays proposed changes to system tables in the Changes to Review related list.

f. Select **Commit Changes**.
The system commits the changes to system tables shown in the Changes to Review related list, and the migration job's state changes to **Configuration Committed**.

g. Select **Index Searchable Content**.
AI Search indexes content and metadata for new indexed sources created by the migration job. When indexing completes, the job's state changes to **Indexing Complete** and the modal window with the list of Zing search application configurations reopens.

6. Refresh the list of search application configurations included in AI Search for Next Experience.

a. Navigate to **All** > *AI Search* > *AI Search* **for** *Next Experience* > *Guided Setup* **for Zing to** *AI Search* **Migration**.

b. In the **Complete migration to** *AI Search* category, select **Edit**.

c. In the **Activate** *AI Search* **for all** *Next Experience* **search applications** task, select **Configure**.
A modal window reports whether AI Search for Next Experience is enabled on your instance.

d. Select **Refresh** *AI Search* **for** *Next Experience*.

Result

AI Search for Next Experience updates to include the new workspace application's search application configurations. The system uses AI Search as the search engine for workspace searches from the Unified Navigation search field.

What to do next

Notify all users of your instance that they need to log out and back in before they can use the new workspace search in AI Search for Next Experience.

Revert to Zing as the Unified Navigation search engine

Revert to using Zing text indexing and search engine for the Unified Navigation search field in Next Experience.

Before you begin

The AI Search for Next Experience ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search for Next Experience.

Role required: ais_admin

About this task

You can undo the configuration changes you made to enable AI Search for Next Experience. This may be for either of the following reasons:

- You prefer to use Zing text indexing and search engine for Unified Navigation in Next Experience.
- You plan to disable Next Experience and revert to the legacy UI. (The legacy UI only supports the Zing engine for global search and workspace search.)

Procedure

1. Navigate to **All** > `AI Search` > `AI Search` **for** `Next Experience` > `Guided Setup` **for Zing to** `AI Search` **Migration**.
2. In the **Complete migration to** `AI Search` category, select **Edit**.
3. In the **Activate** `AI Search` **for all** `Next Experience` **search applications** task, select **Configure**.
   A modal window reports the enablement status of AI Search for Next Experience on your instance.
4. If the system reports `AI Search` **for** `Next Experience` **is enabled**, select **Disable** `AI Search` **for** `Next Experience`.
5. When the system reports `AI Search` **for** `Next Experience` **is disabled**, select **Close**, then select **Mark as Incomplete** for the **Activate** `AI Search` **for all** `Next Experience` **search applications** task.

Result

The system reverts to using Zing as the search engine for global search and Next Experience workspace application searches issued from the Unified Navigation search field.

What to do next

Notify all users of your instance that they need to log out and back in before they can use Zing search in the Unified Navigation search field.

Repair AI Search for Next Experience after activating a new ServiceNow AI Platform® language

Install translated search UI element labels for a newly activated ServiceNow AI Platform language by repairing the AI Search for Next Experience application.

Before you begin

You must have activated a new ServiceNow AI Platform language plugin after AI Search for Next Experience was installed in your instance. For details on activating language plugins, see Activate a language.

Role required: admin

About this task

At installation time, the AI Search for Next Experience application includes search UI element translations for all languages currently active in your ServiceNow AI Platform instance. These translations include language-specific labels for the Unified Navigation search field, navigation tabs, and other elements of the global search and configurable workspace search user interfaces.

If you activate a new ServiceNow AI Platform language plugin in an instance that already has AI Search for Next Experience installed, the application doesn't automatically update to include UI element translations for the newly added language. To install these translations, you can repair the AI Search for Next Experience application.

For more details on repairing applications, see Repair a ServiceNow application.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Locate the AI Search for Next Experience and select its menu icon (⋮).
3. In the Activate Plugin dialog box, select **Repair**.

Result

The AI Search for Next Experience application and its dependencies are reinstalled. After reinstallation, the application includes search UI element translations for all languages activated in your ServiceNow AI Platform instance.

Enable search-based auto-complete suggestions in AI Search for Next Experience

Run a scheduled job to make search-based auto-complete suggestions available in AI Search for Next Experience.

Before you begin

The AI Search for Next Experience ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search for Next Experience.

You only need to perform this task if you installed AI Search for Next Experience 3.x on a San Diego instance that you then upgraded to Tokyo or a later release.

Role required: admin

About this task

If you install AI Search for Next Experience 3.x on a San Diego instance and then upgrade to Tokyo or a later release, the new search-based auto-complete suggestion types introduced in Tokyo aren't automatically enabled in the **[AIS]** *`Next Experience`* **Search Configuration** search application configuration used by global search.

To enable the new search-based auto-complete suggestion types in AI Search for Next Experience and remove the legacy suggestion types that they replace, administrators can run a scheduled job.

Running the scheduled job adds these auto-complete suggestion types to the **[AIS]** *`Next Experience`* **Search Configuration** search application configuration, making them available for global search:

- Search-based Suggested Results Reader Group for [AIS] Next Experience Search Configuration
- User History Reader Group for [AIS] Next Experience Search Configuration

The scheduled job removes these two legacy auto-complete suggestion types from the **[AIS]** *`Next Experience`* **Search Configuration** search application configuration, making them unavailable for global search:

- Suggested Query Reader Group for [AIS] Next Experience Search Configuration
- Suggested Result Reader Group for [AIS] Next Experience Search Configuration

To learn more about auto-complete suggestion types in AIS applications, see Auto-complete suggestions in AI Search applications.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Open the **Enable search-based suggestions** scheduled job record.
3. Select **Execute Now**.

Result

The scheduled job adds these auto-complete suggestion types to the **[AIS]** *Next Experience* **Search Configuration** search application configuration:

- Search-based Suggested Results Reader Group for [AIS] Next Experience Search Configuration
- User History Reader Group for [AIS] Next Experience Search Configuration

The scheduled job removes these auto-complete suggestion types from the **[AIS]** *Next Experience* **Search Configuration** search application configuration:

- Suggested Query Reader Group for [AIS] Next Experience Search Configuration
- Suggested Result Reader Group for [AIS] Next Experience Search Configuration

Global search uses the updated set of auto-complete suggestion types.

Using AI Search for Next Experience

Global search enables you to search multiple record types at once from the Next Experience Unified Navigation search field. You can switch between global search results and results from workspace applications that you have access to.

Search in the Unified Navigation search field

To use global search, enter your search query in the Unified Navigation search field, then select **View all results for "<search-query>"** or press Enter.

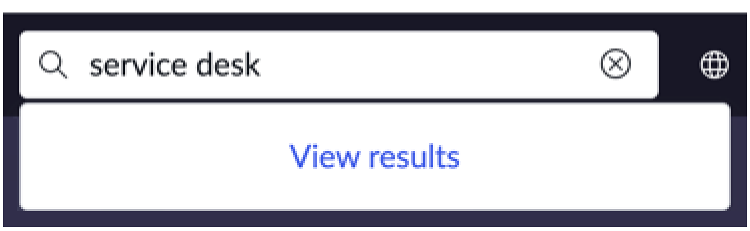

The search results page displays a list of previews for records that match your search query. You can open any search result record by selecting its preview in the list.

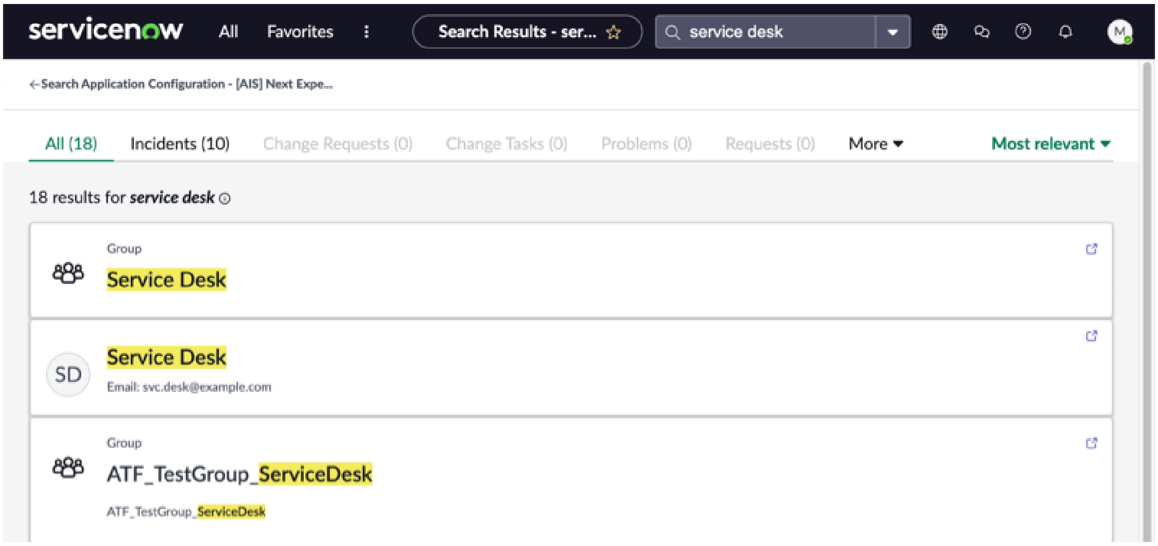

Note:
When you search with a search query, the system only displays results for records included in the AI Search index. Search queries won't match records that have never been indexed or records that have been purged from the index by a retention policy. AI Search administrators can configure indexing and retention policy settings. For details on these settings, see Indexing content from indexed sources and Indexed source retention policies and filter conditions.

AI Search administrators can customize the index fields displayed on search result previews. For details, see the customization section found in this topic.

Narrow your search by source on the search results page

You can narrow your search to display only results from a particular search source. When you select one of the search results page's navigation tabs, the page shows only those matching records that come from the selected source. To return to viewing all matching records, select the **All** navigation tab.

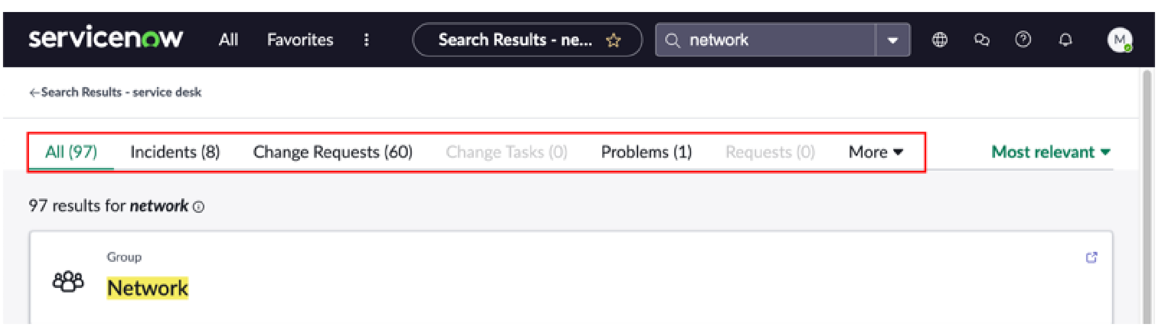

AI Search administrators can configure display settings for navigation tabs. For details, see Configure navigation tabs in an AI Search application configuration.

Narrow your search by field value on the search results page

After you select a navigation tab, you may see a Filters list. Each facet filter in this list indicates a field from your selected search source. The filter displays the field's values that are found on one or more of your currently displayed search result records. When you select a facet filter value, your search narrows to show only those matching records that include the specified field value.

For example, suppose you select the **Incident** navigation tab to view only Incident records that match your search query. The Filters list appears, displaying facet filters for Incident fields such

as **Priority**, **State**, **Category**, and **Assignment group**. Selecting the **In Progress** value from the **State** facet filter narrows your search to show only matching records for Incidents that are in progress.

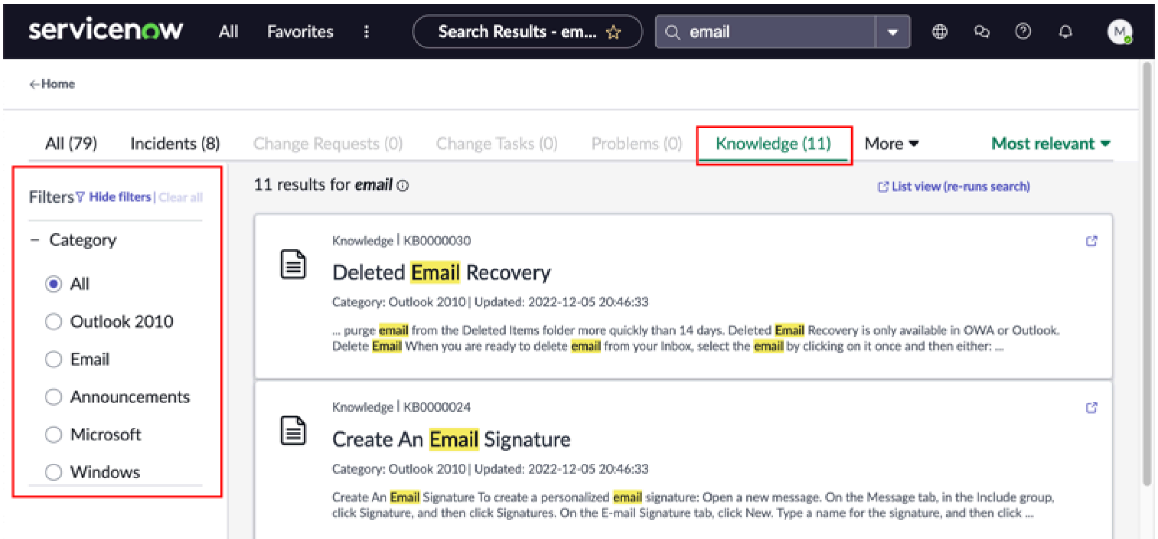

You can use cross-table facets to filter search results using field values found on multiple source tables. Cross-table facet filters only appear in the Filters list on the **All** navigation tab.

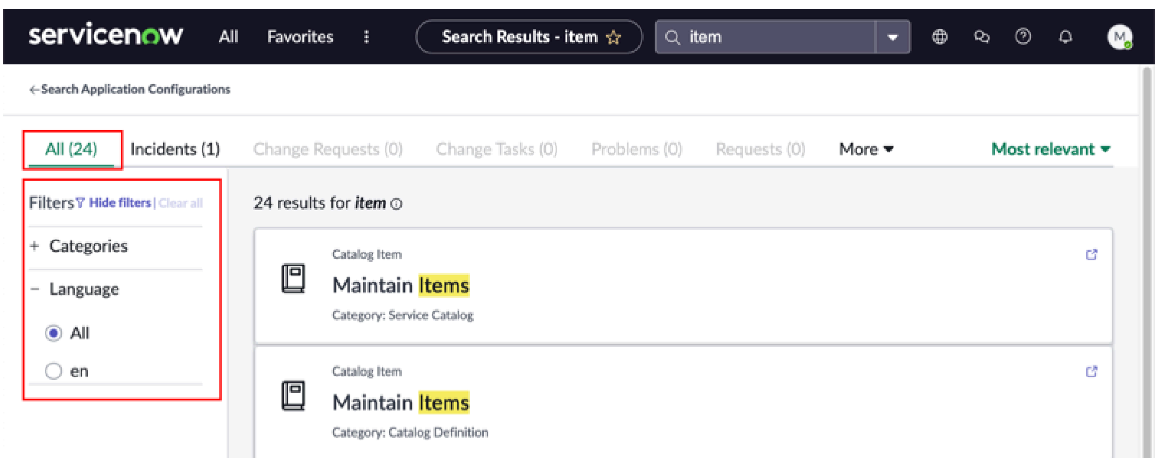

To remove an applied facet filter, select **Clear** by the facet field name, or **Clear all** to remove all applied filters at once. You can hide the Filters list by selecting **Hide filters**.

AI Search administrators can define single-table and cross-table facets. For details, see Create a facet in an AI Search application configuration.

Sort your search results

By default, AI Search displays the most relevant search results first. To change the display order for records on the search results page, select a sort option from the sort menu.

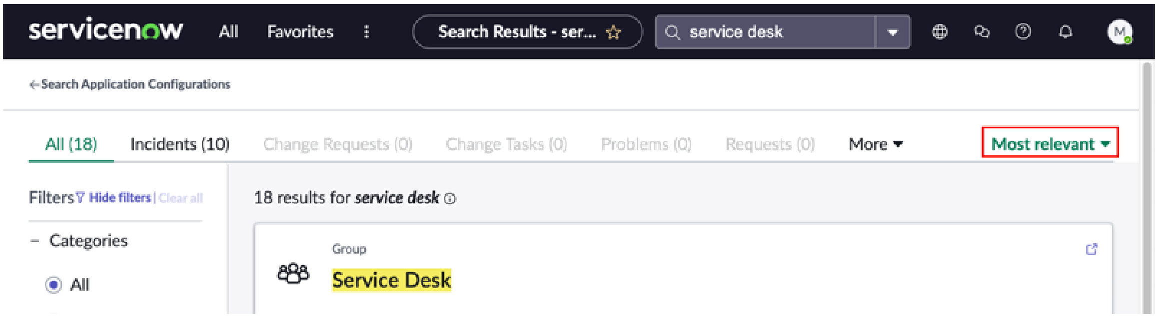

AI Search administrators can configure search result sort options. For details, see Search result sort options in AI Search application configurations.

Return to your starting point from the search results page

To return to the page where you initiated your search, select the link with a left arrow that appears above the navigation tabs on the search results page. The title of this link indicates the page it returns you to. For example, if you searched from the Home page, the link text would be **Home**, but if you searched from the Dashboards page, the link text would be **Dashboards Overview**.

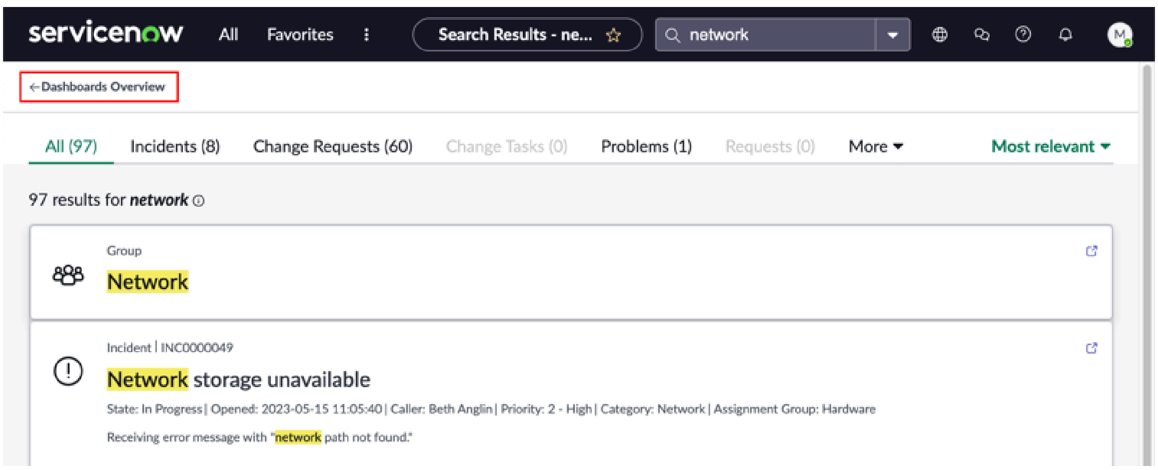

Search suggestions in the Unified Navigation search field

AI Search displays auto-complete suggestions to help you get directly to the results you need. Suggestions may include recent search queries, recently viewed search results, or search queries or results that match the terms you're typing into the search field.

Some suggestion types appear when you initially select the search field.

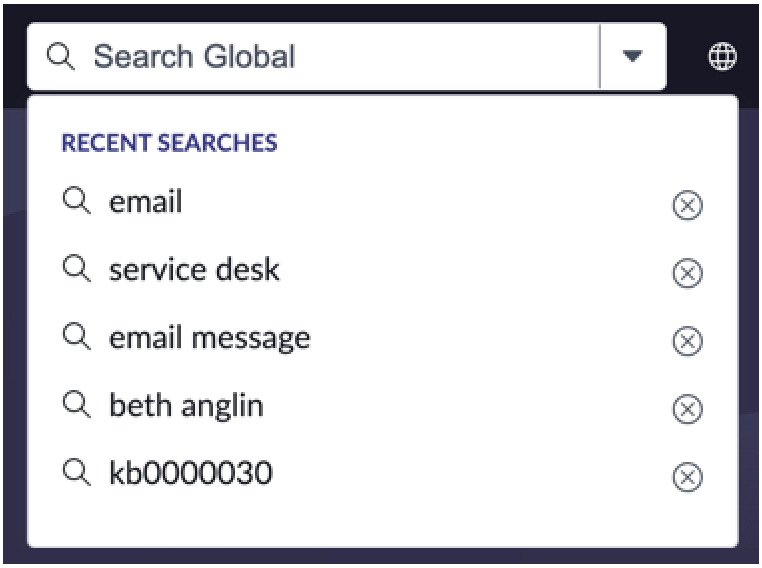

Other suggestion types appear when you begin typing your search query into the search field.

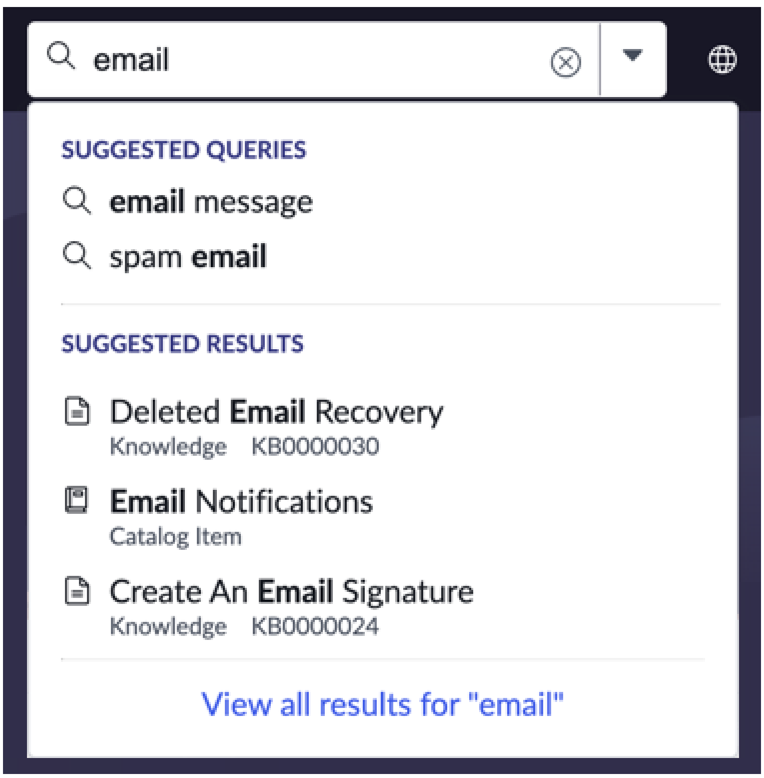

AI Search administrators can configure auto-complete suggestions. For details, see Auto-complete suggestions in AI Search applications.

Search for a specific record using Exact Match

Enter a record number into the Unified Navigation search field but don't press Enter or select **View results**. The search field displays a preview of the record with an **Exact Match** tag. Select the preview to go directly to the matching record, bypassing the search results page.

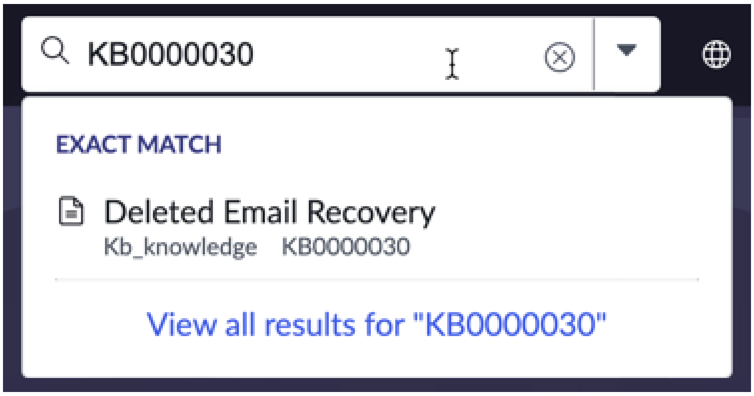

Note:
When you search for a record by number, the system displays the matching record regardless of whether it's included in the AI Search index or not.

Exact matching returns results from tables with a prefix defined in the Number [sys_number] table. If the search includes a single term that starts with the table's prefix, AI Search returns an exact match result from the table if one exists. As an example, searching for the single term `INC3263827` returns an exact match from the Incident [incident] table because the search term starts with the INC prefix defined for that table.

If more than one record has a **Number** field value that exactly matches your search, AI Search displays the first exact match. An informational message reports the total number of exact matches and provides links to the other exact matches.

Search administrators can override the default exact match behavior by configuring custom search matchers. For details on this procedure, see Create a custom search matcher for global search.

View results for your search in an available workspace application

If you have access to search in workspace applications, an arrow appears at the right of the Unified Navigation search field after you search that enables you to switch between the results in global search and in your workspace applications.

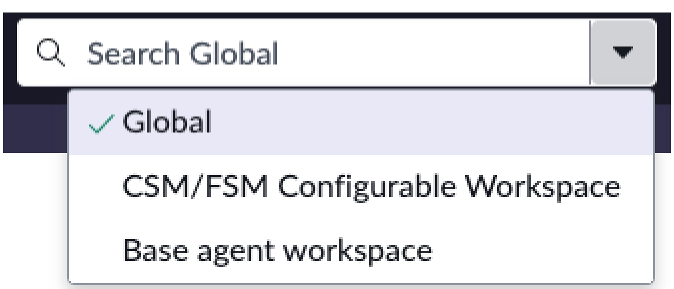

For example, if the context menu in the illustration was available after you performed a global search, you could select **CSM/FSM Configurable Workspace** from the context menu to view results for the search in CSM/FSM Configurable Workspace.

Exact matches open in the selected workspace application. For example, if you selected **CSM/FSM Configurable Workspace**, entered a record number, and selected the record preview in the search results, the record would open in CSM/FSM Configurable Workspace.

Customize display fields for your search result previews

Search result previews display a default set of AI Search index fields. If you have the ais_admin role, you can customize the set of index fields a search result preview displays by modifying its EVAM view configuration. For information on EVAM view configurations, see Entity View Action Mapping.

When customizing a search result preview's EVAM view configuration, you can reference the AI Search index fields that exist on the previewed search result. To add a new AI Search index field to search results from an indexed source, define a *map_to* field setting. This field setting populates the index field on each affected search result with the value of a field you specify from the search result's source record or document. For an overview of mapping source fields to AI Search index fields, see Field settings. To create a new *map_to* field setting for one of your indexed sources, see Create a field setting for an indexed source.

AI Search for Next Experience reference

Reference information for the AI Search for Next Experience application.

Plugins installed with AI Search for Next Experience

Plugin	Description
AI Search for Next Experience (sn_aisearch_global)	Includes system and search configuration settings required to enable AI Search as the search engine for the Unified Navigation search field in Next Experience. • Search profile: Now Experience Search Configuration • Search application configuration: [AIS] Now Experience Search Configuration To learn more about search profiles and search application configurations, see Search profiles and Search application configurations.
AI Search Index Sources (com.glide.ais.index_sources)	Includes indexed source definitions for ServiceNow AI Platform® tables beyond those provided with AI Search: • Catalog Task [sc_task] • Change [change_request] • Change Task [change_task] • Company [core_company] • Group [sys_user_group] • Incident [incident] • Location [cmn_location] • Problem [problem] • Request [sc_request] • Requested Item [sc_req_item]

Plugin	Description
	For more details on indexed sources, see Indexed sources.

System properties installed with AI Search for Next Experience

System property	Description
glide.ui.polaris.ais_ready	Option to enable AI Search as the search engine for the Unified Navigation search field in Next Experience. • Type: true | false • Default value: *false* • Supported values: ◦ *true*: The system uses AI Search as the search engine for the Unified Navigation search field in Next Experience. ◦ *false*: The system uses Zing indexing and text search as the search engine for the Unified Navigation search field in Next Experience. • Location: System Property [sys_properties] table

Search profiles installed with AI Search for Next Experience

Search profile	Description
Next Experience Search Configuration	Contains search experience settings for AI Search for Next Experience: • Search sources • Stop word dictionaries • Synonym dictionaries • Typo handling dictionaries • Result improvement rules • Genius Results For details on configuring these settings, see Search profiles. • Application scope: AI Search For Next Experience • Location: **All** > *AI Search* > **Search Experience** > **Search Profiles**

Search application configurations installed with AI Search for Next Experience

Search application configuration	Description
[AIS] Next Experience Search Configuration	Contains search result display and refinement settings for AI Search for Next Experience: • Navigation tabs • Facets • Auto-complete suggestions • Search scripted post-processors • Search result sort options For details on configuring these settings, see Search application configurations. • Application scope: AI Search For Next Experience • Location: **All** > `AI Search` > **Search Experience** > **Search Applications** **Note:** If you install AI Search for Next Experience 3.x on a San Diego instance and then upgrade to Tokyo or a later release, an informational message appears when you view the search application configuration record, pointing out that search-based auto-complete suggestion types aren't enabled for global search. To enable these auto-complete suggestion types for global search, see Enable search-based auto-complete suggestions in AI Search for Next Experience.

Entity View Action Mapper (EVAM) configurations installed with AI Search for Next Experience

Configuration	Description
AI Search Global | Catalog Genius Results	Contains view configuration settings for the Catalog Item Genius Result configuration in AI Search for Next Experience. To learn more about this Genius Result configuration, see Catalog Item Genius Results.

Configuration	Description
	• Application scope: AI Search For Next Experience • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**
AI Search Global | Q&A Genius Result	Contains view configuration settings for the Q&A Genius Result configuration in AI Search for Next Experience. To learn more about this Genius Result configuration, see Q&A Genius Results. • Application scope: AI Search For Next Experience • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**
AI Search Global | User Genius Results	Contains view configuration settings for the People Genius Result configuration in AI Search for Next Experience. To learn more about this Genius Result configuration, see People Genius Results. • Application scope: AI Search For Next Experience • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**

Entity View Action Mapper (EVAM) templates installed with AI Search for Next Experience

Template	Description
AI Search Global | Catalog Genius Results Template	Contains component, static value, field mapping, and action mapping settings for the Catalog Item Genius Result configuration in AI Search for Next Experience. To learn more about this Genius Result configuration, see Catalog Item Genius Results. • Application scope: AI Search For Next Experience • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Templates**
AI Search Global | Q&A Genius Result Template	Contains component, static value, field mapping, and action mapping settings for the Q&A Genius Result configuration in AI Search

Template	Description
	for Next Experience. To learn more about this Genius Result configuration, see Q&A Genius Results. • Application scope: AI Search For Next Experience • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Templates**
AI Search Global | User Genius Results Template	Contains component, static value, field mapping, and action mapping settings for the People Genius Result configuration in AI Search for Next Experience. To learn more about this Genius Result configuration, see People Genius Results. • Application scope: AI Search For Next Experience • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Templates**

List of all components installed with AI Search for Next Experience

To view the complete list of components installed with AI Search for Next Experience, follow the steps described in Find components installed with an application. The application's package name is `AI Search` **For** `Next Experience`.

AI Search for Next Experience default navigation tabs for global search

List of navigation tabs that appear by default in AI Search for Next Experience global search results.

When you activate AI Search for Next Experience, global search results display in navigation tabs. Each navigation tab represents a search source that's linked to the Next Experience Search Configuration search profile. This search profile is the one used by the [AIS] Next Experience Search Configuration search application.

The following table lists the navigation tabs that appear by default when you activate AI Search for Next Experience.

Navigation tab	Description
All	Global search results from all search sources linked to the Next Experience Search Configuration search profile.
Incidents	Global search results from the Tasks - Incidents search source. • Indexed source: Incident • Source table: Incident [incident] Incident [incident]

Navigation tab	Description
Change Requests	Global search results from the Tasks - Change Requests search source. • Indexed source: Change • Source table: Change Request [change_request]
Change Tasks	Global search results from the Tasks - Change Tasks search source. • Indexed source: Change Task • Source table: Change Task [change_task]
Problems	Global search results from the Tasks - Problems search source. • Indexed source: Problem • Source table: Problem [problem]
Requests	Global search results from the Tasks - Requests search source. • Indexed source: Request • Source table: Request [sc_request]
Catalog Tasks	Global search results from the Tasks - Catalog Tasks search source. • Indexed source: Catalog Task • Source table: Catalog Task [sc_task]
Users	Global search results from the People - Users search source. • Indexed source: User table • Source table: User [sys_user]
Requested Items	Global search results from the Tasks - Requested Items search source. • Indexed source: Requested Item • Source table: Requested Item [sc_req_item]
Groups	Global search results from the People - Groups search source. • Indexed source: Group • Source table: Group [sys_user_group]

Navigation tab	Description
Knowledge	Global search results from the Knowledge & Catalog - Knowledge search source. • Indexed source: Knowledge Table • Source table: Knowledge [kb_knowledge]
Catalog Items	Global search results from the Knowledge & Catalog - Catalog Items search source. • Indexed source: Catalog Item Table • Source table: Catalog Item [sc_cat_item]
Companies	Global search results from the People - Companies search source. • Indexed source: Company • Source table: Company [core_company]

You can change the search sources linked to the Next Experience Search Configuration search profile. To learn more about creating search sources and linking them to search profiles, see Search sources.

For details on configuring labels and display order for navigation tabs, see Configure navigation tabs in an AI Search application configuration.

Now Assist in AI Search

The Now Assist in AI Search ServiceNow® Store application combines the power of search with the Now LLM Service agentic AI model to provide actionable AI-generated or AI-selected answers in user searches. By constraining the context passed to the Now LLM Service, Now Assist in AI Search can increase the likelihood that Now LLM Service responses are grounded in indexed content.

Now Assist in AI Search overview

Now Assist in AI Search includes the following Genius Result configurations:

Now Assist Q&A Genius Results

Now Assist Q&A Genius Results use the Now LLM Service to generate actionable Q&A Genius Result answers from the most relevant Knowledge article results.

External Content Q&A Genius Results

External Content Q&A Genius Results generate actionable Q&A Genius Result answers from documents in your Microsoft SharePoint Online instance.

Now Assist Actions Genius Results

Now Assist Actions Genius Results use the Now LLM Service to select, filter, and display the most relevant Catalog Item and Virtual Agent topic results as answers.

All Now Assist Genius Result answers augment the user's search results, displaying as actionable answer cards. For reference, each answer card includes links to its source records or external documents.

Get started

Install	Configure	Reference
Activate the Now Assist in AI Search plugin on your instance	Configure Now Assist in AI Search Genius Results for use in Service Portal, Virtual Agent, Employee Center, and global search	View components installed with Now Assist in AI Search

Language support

Now Assist in AI Search natively supports searches, content, and answers in English. Additional languages are supported through Dynamic Translation. For details on translation functionality and supported languages, see Dynamic Translation for Now Assist Q&A Genius Results.

Helpful resources

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

ServiceNow Community

ServiceNow Community

Customer Success Center

Impact

Developer

developer.servicenow.com

Impact

http://impact.servicenow.com

ServiceNow University

ServiceNow University

NowCreate

Best Practices

Partner

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

ServiceNow

https://www.servicenow.com

ServiceNow Store

https://store.servicenow.com/

servicenow

Support

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

AI limitations

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

Data processing

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

Data collection

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

For more information, see the Now Assist documentation.

Install Now Assist in AI Search

As an administrator, you can install the Now Assist in AI Search application (sn_ais_assist) from the Now Assist Admin module.

Before you begin

Make sure that your instance is upgraded to Washington DC Patch 3 or later. To learn more about family releases and patches, see Available patches and hotfixes.

Role required: admin

About this task

The Now Assist in AI Search plugin is automatically installed when you install any of these feature plugins from Now Assist Admin:

- 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)

Note:
When you upgrade to the latest version of Now Assist in AI Search, the system reindexes content from the Catalog Item Table and Knowledge Table indexed sources. While these reindexing tasks are ongoing, searches may not return Catalog and Q&A Genius Result answers. Once reindexing completes, Catalog and Q&A Genius Result answers should appear normally.

Procedure

1. In Now Assist Admin, install one or more of the prerequisite feature plugins found in "About this task".
   For Now Assist feature plugin installation instructions, see Install Now Assist plugins.
2. Verify that Now Assist in AI Search is installed:
   a. Navigate to **All** > *Now Assist Admin* > **Features**.
   b. In the workflow list, select **Platform**.
   c. Verify that the **Conversational experience** feature card displays, and that the *Now Assist* **Q&A Genius Results** and *Now Assist* **Actions** skills appear in the **All Conversational experience skills** listing.

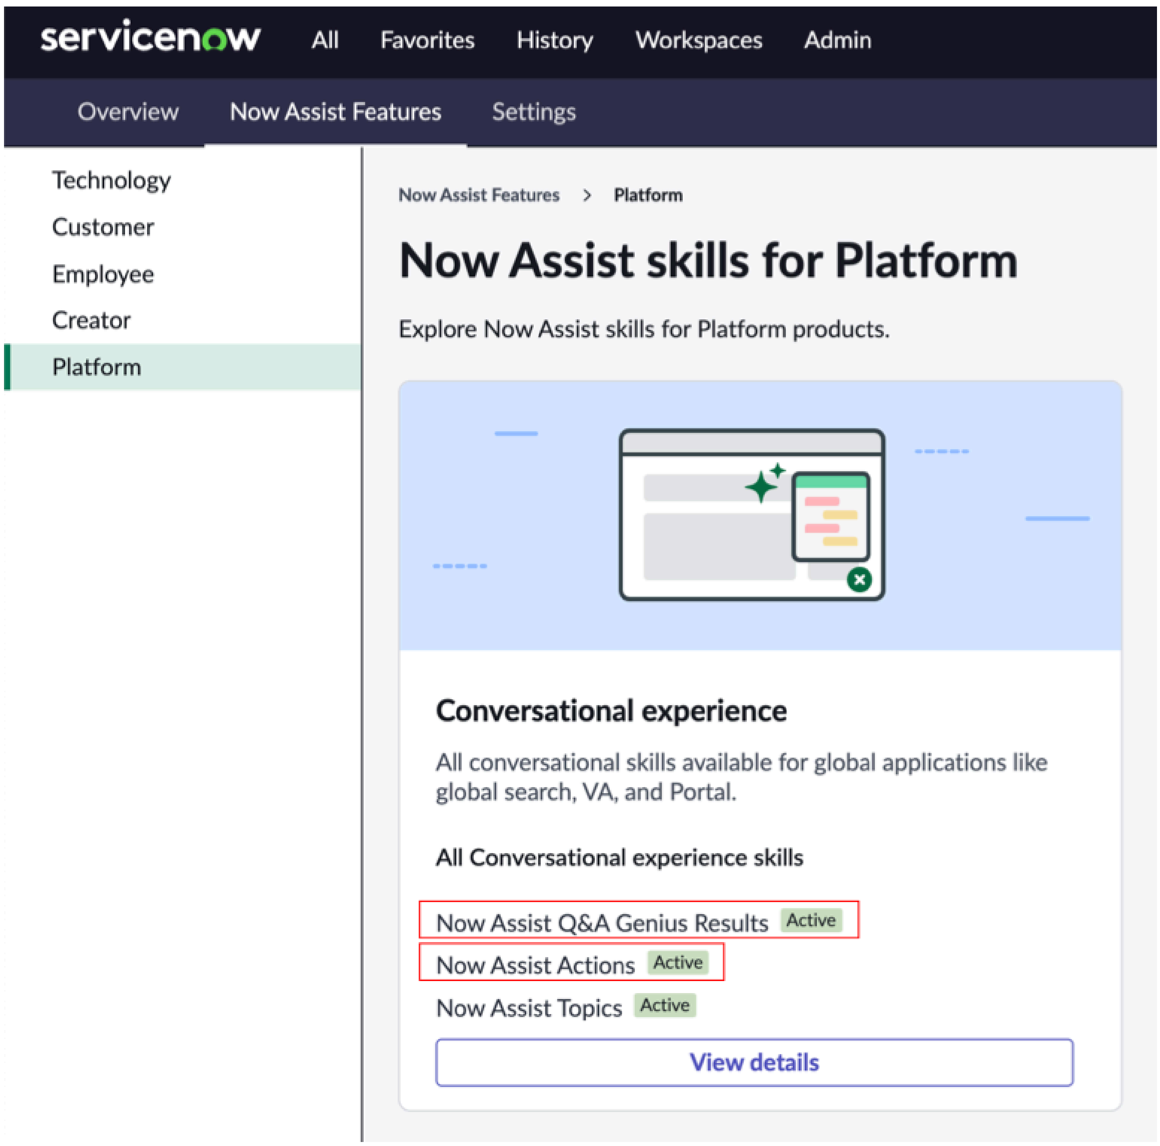

What to do next

With the plugin installed, search administrators can enable Now Assist Genius Results in the following contexts.

- Enable all available Now Assist Genius Result configurations in individual AI Search portals. For details, see Enable Now Assist Genius Results in AI Search portals.
- Enable Now Assist Q&A Genius Results in search profiles for AI Search applications. For steps, see Enabling Now Assist Q&A Genius Results.
- Enable Now Assist Q&A Genius Results in global search using the AI Search for Next Experience application. For steps, see Enabling Now Assist Q&A Genius Results.

To learn more about configuration settings for the plugin, see Configuring Now Assist in AI Search.

Review available versions of Now Assist in AI Search

View all versions of the Now Assist in AI Search application on the ServiceNow Store. Use this information to find the latest version of the application that's compatible with your instance's current ServiceNow AI Platform® family release.

Before you begin

You must have access to the ServiceNow Store site.

Role required: none

Procedure

1. In a web browser, navigate to the ServiceNow Store site: https://store.servicenow.com/sn_appstore_store.do#!/store/home
2. Use search or navigation to find the listing for any one of the following Now Assist feature applications:
   - 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)
3. Open the application listing.
   The listing shows the **Product Details** view for the selected application.
4. Select the **Dependencies** tab, then select `Now Assist` **in** `AI Search` from the **App Dependency** listing.
   The listing shows the **Product Details** view for the Now Assist in AI Search application.
5. In the **Version** section, select **Other App Versions**.
   The Version History window shows the list of all available versions of the Now Assist in AI Search application. Each entry shows a version's number, family release compatibility, release type, and release date.
6. To view the listing for another version of the Now Assist in AI Search application, expand the entry for that version by selecting its arrow icon (⌄), then select **View Details**.
   The listing shows the **Product Details** view for the selected version of the Now Assist in AI Search application.

Configuring Now Assist in AI Search

As an administrator or search administrator, you can configure settings to control the behavior of Now Assist in AI Search.

Enable Now Assist Genius Results in AI Search portals

Specify the Now Assist Genius Result types you want to make available in each of your AI Search portals.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

Role required: ais_admin

About this task

Now Assist in AI Search provides multiple AI-powered Genius Result configurations. As a search administrator, you can enable individual Now Assist Genius Result configurations in AI Search portal applications such as Service Portal and Employee Center.

To learn more about how Genius Results work, see Genius Results. For details on the Now Assist Genius Result configurations, see Now Assist Q&A Genius Results and Now Assist Actions Genius Results.

Procedure

1. Navigate to **All** > *AI Search* > *Now Assist* **in** *AI Search* **Setup**.
   The Set up Now Assist in AI Search panel displays a list of search profiles associated with AI Search portals and a Prerequisites section for each Now Assist Genius Result configuration.
2. **Optional:** If a Now Assist Genius Result enablement option isn't selectable, follow the guidance in its Prerequisites entry to make it selectable.

   Note:
   After following the Prerequisites guidance, you may need to refresh the Set up Now Assist in AI Search panel to make the Now Assist Genius Result enablement option selectable.

3. In each search profile row, select the options for the individual Now Assist Genius Results you want to enable for the associated AI Search portal.

   As an example, to enable Now Assist Q&A Genius Results in Service Portal, select the *Now Assist* **Q&A** option in the **Service Portal Default Search Profile** row.

   Note:
   You can disable a Now Assist Genius Result for an AI Search portal by clearing its option.

4. Select **Apply Changes**.

Result

AI Search enables your selected Now Assist Genius Result configurations for the specified AI Search portals.

Now Assist Q&A Genius Results

Now Assist Q&A Genius Results use the Now LLM Service to generate concise, actionable answers from Knowledge article results in Service Portal, Virtual Agent, Employee Center, and global searches.

Now Assist Q&A Genius Results overview

Now Assist Q&A Genius Results send the most relevant Knowledge articles from your search to the Now LLM Service, which generates answer snippets from the articles' HTML fields.

Each Now Assist Q&A Genius Result answer card displays up to three generated answer snippets. For reference, the answer card also includes a link you can select to view the source Knowledge articles.

The following example shows a Now Assist Q&A Genius Result answer card containing a snippet summarizing a Knowledge article. Select the answer card's **View article** action link to view the full Knowledge article.

Sample Now Assist Q&A answer card

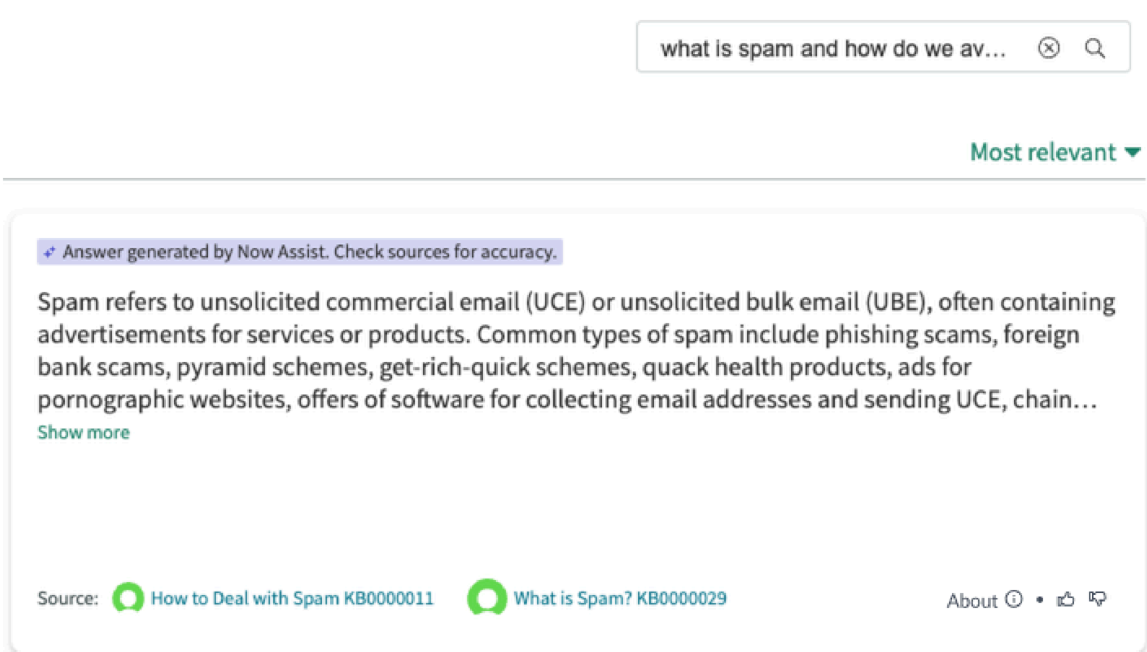

Note:
Because the Now Assist Q&A Genius Result answer is automatically generated, it's a good idea to review it for accuracy. You can provide feedback on the answer by selecting the thumbs-up icon () if the generated answer is accurate, or the thumbs-down icon () if it's not. Your feedback helps ServiceNow improve future versions of this Genius Result configuration.

Now Assist Q&A Genius Results use semantic vector search and legacy keyword search to find Knowledge articles that best match the meaning and intent of your search query. For more details on semantic vector search, see Semantic vector search in AI Search.

The Now Assist Q&A Genius Result configuration replaces the original Q&A Genius Result configuration from the base system. The base system's Q&A Genius Result configuration extracts answers from Knowledge articles using internal routines instead of using the Now LLM Service. To learn more about the base system's Q&A Genius Result configuration, see Q&A Genius Results.

Enabling Now Assist Q&A Genius Results

You can enable Now Assist Q&A Genius Results in your AI Search portals and mobile applications using the Now Assist in AI Search Setup module. For details on this procedure, see Enable Now Assist Genius Results in AI Search portals.

To use Now Assist Q&A Genius Results in global search, you can enable the Now Assist Q&A Genius Result configuration in the AI Search for Next Experience application. For details on this procedu re, see Enabling Now Assist Q&A Genius Results.

Limitations

By default, Now Assist Q&A Genius Results only support English-language searches. Administrators can enable support for other languages by activating Dynamic Translation. To learn more about how content and answers are translated, see Dynamic Translation for Now Assist Q&A Genius Results. For more details on Dynamic Translation, see Dynamic Translation.

Knowledge articles that are boosted or promoted by result improvement rules are more likely to appear as Now Assist Q&A Genius Results, but aren't guaranteed to appear.

Note:
The Knowledge search property settings don't affect Now Assist Q&A Genius Results. For more information on these settings, see Knowledge search properties.

Answer snippet creation for Now Assist Q&A Genius Results

AI Search uses the Now LLM Service to create each Now Assist Q&A answer snippet from a Knowledge article record's HTML fields. Each Now Assist Q&A answer card can include up to three snippets. These snippets may all be generated from the same source article or from different source articles.

Now LLM Service automatically determines which elements of a Knowledge article's text to include in a Now Assist Q&A answer snippet. You can't configure the criteria for this behavior.

The Now LLM Service summarizes and abstracts content from the Knowledge articles' text fields. Answer snippets displayed on Now Assist Q&A Genius Result answer cards may not exist word for word in the source records.

Both AI Search content retrieval and the Now LLM Service are continually improving, so Now Assist Q&A results for specific queries may vary over time. Because results from the Now LLM Service are non-deterministic, you should expect a higher answer variability compared to the base system's Q&A Genius Results.

Interaction with other search features

The following table describes the interactions between Now Assist Q&A Genius Results and other search features.

Feature	Interaction with Now Assist Q&A Genius Results
Result improvement rules	When computing Now Assist Q&A Genius Result answers for a search query, AI Search applies result improvement rules normally. The effects depend on the result improvement rule's action, as follows: • *block*: Now Assist Q&A Genius Results don't generate answers from blocked records. • *boost* or *promote*: Boosted and promoted records are more likely to be used when generating Now Assist Q&A Genius Result answers.

Feature	Interaction with Now Assist Q&A Genius Results
Stop words	Now Assist Q&A Genius Results use a blend of semantic vector search, which doesn't support stop words, and keyword-based search. AI Search only removes stop words from keyword-based searches, so Now Assist Q&A Genius Result answers may not reflect your stop words settings.
Synonyms	Now Assist Q&A Genius Results use a blend of semantic vector search, which doesn't support synonyms, and keyword-based search. AI Search expands synonyms in keyword-based searches, so your synonyms are likely to improve the relevancy of Now Assist Q&A Genius Result answers.
Typo handling	When computing Now Assist Q&A Genius Result answers for a search query, AI Search corrects misspelled terms in the query.

Enabling Now Assist Q&A Genius Results

As a search administrator, you can use the Now Assist Q&A Genius Results skill in AI Search portals and mobile applications by enabling Now Assist Q&A Genius Results in search profiles. You can also use the skill in global search by enabling Now Assist Q&A Genius Results in the AI Search for Next Experience application.

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

For details on how to enable Now Assist Q&A Genius Results and other Now Assist Genius Results in your AI Search portals and mobile applications, see Enable Now Assist Genius Results in AI Search portals.

Enable Now Assist Q&A Genius Results in global search

Enable Now Assist Q&A Genius Results in global search using the AI Search for Next Experience application.

Before you begin

The Next Experience UI must be activated on your instance. To learn about activating Next Experience UI, see Considerations for activating Next Experience.

The AI Search for Next Experience application must be installed and enabled on your instance. For more information about the application, see AI Search for Next Experience.

The Now Assist for Search application must be installed on your instance. For details on installing this application from the Now Assist Admin console, see Install Now Assist in AI Search.

Role required: ais_admin

About this task

As a search administrator, you can enable agentic AI model answer generation for Knowledge article results in global searches by adding the Now Assist Q&A Genius Result configuration to the Next Experience Search Configuration search profile used by AI Search for Next Experience.

Note:
Linking the Now Assist Q&A Genius Result configuration to the Next Experience Search Configuration search profile changes the status of the Now Assist Q&A Genius Results skill to **Active** in the Now Assist Admin console. The skill remains active as long as any search profile includes a link to the Now Assist Q&A Genius Result configuration.

Procedure

1. In the Unified Navigation, select the application scope and update set icon ( ), and then set the application scope to `AI Search` **for** `Next Experience`.
2. If you have version 1.x or 2.x of the AI Search for Next Experience application installed on a Tokyo or Utah instance, add the Genius Result data source to the AI Search For Next Search EVAM definition.
   If you have version 3.x or a later version of the AI Search for Next Experience installed, or are on Vancouver or a later release, go to the next step.
   a. Navigate to **All** > **Entity View Action Mapper (**`EVAM`**)** > `EVAM` **Definitions**.
   b. Open the AI Search For Next Search definition record by selecting it from the list.
   c. In the EVAM Datasource M2Ms related list, select **Link Existing**.
   d. On the EVAM Datasource M2M form, populate the **Datasource** field with a reference to the Genius Result Data Source record from the EVAM Datasources table by using the Lookup using list icon ( ).
   e. Select **Submit**.
   The Genius Result Data Source reference appears in the EVAM Datasource M2Ms related list.
3. Add the Now Assist Q&A Genius Result configuration to the search profile used by AI Search for Next Experience.
   a. Navigate to **All** > `AI Search` > **Search Experience** > **Search Profiles**.
   b. Open the Next Experience Search Configuration search profile record by selecting it from the list.
   c. In the Genius Results related list, select **Link Existing**.
   d. On the Search Profile - Genius Result Mapping form, fill in the following field values.

Field	Value
Genius Result Configuration	Enter or select `Now Assist` **Q&A**.

Field	Value
Order	Specify the evaluation order for the Now Assist Q&A Genius Result configuration in the search profile. AI Search evaluates active Genius Result configurations for a search profile in order, from lowest **Order** field value to highest. For more details on the evaluation order for Genius Results, see Set the evaluation order for Genius Result configurations in a search profile.

e. Select **Submit**.

4. Select the existing related list row for the Q&A Genius Result configuration, and then select **Unlink Selected**.
   The Now Assist Q&A Genius Result configuration supersedes the original Q&A Genius Result configuration. You don't need both configurations in the same search profile.

Result

The Now Assist Q&A Genius Result configuration is enabled for global searches.

In the Now Assist Admin console, the status for the Now Assist Q&A Genius Results skill is **Active**.

Define a query filter for Now Assist Q&A Genius Results

Define a Java regular expression pattern that a search must match to be eligible for triggering Now Assist Q&A Genius Results. Searches that don't match this pattern don't return Genius Result answers from Now Assist Q&A Genius Results.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

Make sure that you're familiar with Java's regular expression pattern syntax. To learn about the regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class.

Role required: admin

About this task

Now Assist Q&A Genius Results send eligible Knowledge article search results to the Now LLM Service for Q&A Genius Result answer generation.

By default, the system only sends Knowledge article results to the Now LLM Service when the search satisfies all of these conditions.

- The search application is Service Portal, Virtual Agent, Employee Center, or global search.
- The user's session language is English.
- The search contains two or more terms.
- The search matches the Java regular expression pattern defined by the *sn_ais_assist.u_question_regex* system property's value. This system property has no value set in the base system, so by default all search queries satisfy this condition.

You can customize the pattern match condition by defining your own Java regular expression pattern for use by Now Assist Q&A Genius Result configurations.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Open the record for the *`sn_ais_assist.u_question_regex`* system property by selecting it from the list.
   If the system prompts you to change the application scope so that you can edit the record, select the provided link.
3. In the system property's **Value** field, enter a Java regular expression pattern that you want Q&A search queries to match.
   The Java regular expression match is case-insensitive. As an example, to restrict Q&A answer expansion to search queries that start with a question keyword, you can set the property's value to `^(what|where|when|why|who|whom|which|how)\b`.
4. Select **Update**.

Result

Now Assist Q&A Genius Results don't send Knowledge article search results to the Now LLM Service for Q&A answer generation unless the search query matches your Java regular expression pattern.

Change the minimum search term count for Now Assist Q&A Genius Results

Specify the minimum number of terms that a search must contain to be eligible for triggering Now Assist Q&A Q&A Genius Results. Searches with fewer terms don't return Now Assist Q&A Genius Result answers.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

Role required: ais_admin

About this task

Now Assist Q&A Genius Results send eligible Knowledge article search results to the Now LLM Service for Q&A Genius Result answer generation.

By default, the system only sends Knowledge article results to the Now LLM Service when the search satisfies all of these conditions.

- The search application is Service Portal, Virtual Agent, Employee Center, or global search.
- The user's session language is English.
- The search contains two or more terms.
- The search matches the Java regular expression pattern defined by the *`sn_ais_assist.u_question_regex`* system property's value. This system property has no value set in the base system, so by default all search queries satisfy this condition.

You can customize the term count condition by altering the minimum term count required by Now Assist Q&A Genius Result configurations.

Procedure

1. Navigate to **All** > `AI Search` > **Search Query Settings** > **Genius Results**.
2. Open the record for the Q&A - Now Assist Genius Result configuration by selecting it from the list.
3. In the `AI Search` **response processor** field, locate the following code block:

```
var searchPhrase = context.getOriginalSearchPhrase();
if (searchPhrase.split(" ").length < 2) {
    return answer;
}
```

4. Change the inequality statement's numeric argument from 2 to another positive integer.
   As an example, if you only want the Genius Result configuration to operate on searches with three or more search terms, change 2 to `3`. If you want the Genius Result configuration to operate on searches with any number of search terms, change 2 to `1`.
5. Select **Update**.

Result

The modified Genius Result configuration doesn't send Knowledge article search results to the Now LLM Service for Q&A answer generation unless the search query has enough search terms to satisfy your chosen limit.

[store-future: BEGIN review]

Activate user context matching for Now Assist Q&A Genius Results

Configure the Now LLM Service to prefer Knowledge articles that match the search user's location and company name when selecting articles for Now Assist Q&A answer generation.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

Role required: admin

About this task

By default, the Now LLM Service does not include user context features, such as the search user's location or company, when determining which Knowledge article to select for Now Assist Q&A Genius Result answer generation. As a result, a user in France who searches for `employee benefits` might see a Now Assist Q&A Genius Result answer card with an answer generated from a Knowledge article for US employee benefits.

Administrators can configure the Now LLM Service to consider the location and company name user context features when selecting Knowledge articles for Now Assist Q&A answer generation. When this capability is activated, the Now LLM Service selects Knowledge articles that match the user context's location or company name in preference to articles that don't match those features. For example, when a user in France searches for `employee benefits`, the Now LLM Service selects Knowledge articles that match `employee benefits` and `France` in preference to articles that just match `employee benefits`.

Warning:
Activating user context matching causes Now Assist Q&A Genius Results to generate more `No answer` responses. It's a good idea to test this capability in a non-production environment before you activate it in production.

Now Assist Q&A Genius Result user context matching activates when all of the following conditions are satisfied:

- The *glide.ais.query.disable_augmentation_for_boost* system property's value is set to *false*.
- One or more of the Knowledge articles sent to the Now LLM Service for consideration has its relevancy score boosted by a result improvement rule that satisfies the following conditions:
  - The result improvement rule's boost action applies a boost value greater than 1.0.
  - The result improvement rule's boost action operates on a field whose name contains any of these words: `city`, `company`, `country`, `department`, `state`.
- The active search application's search profile includes a link to the User Context Synonym Dictionary. This dictionary defines synonym terms for geographical locations found in user context fields, such as countries, provinces, states, and cities. You can add your own user context synonym terms to the User Context Synonym Dictionary.

To learn more about result improvement rules and boost actions, see Result improvement rules and Boost search results using a result improvement rule. For details on synonyms and creation and modification of synonym terms, see Synonyms and Create synonyms.

Procedure

1. In the **Filter** field, enter `sys_properties.list` and press Enter.
2. Search for the *glide.ais.query.disable_augmentation_for_boost* system property.
   - If the system property already exists, set its value to *false* and select **Update**.
   - If the system property doesn't exist, go to the next step.
3. Select **New**.
4. On the System Property form, enter the following field values.

Field	Value
Name	glide.ais.query.disable_augmentation_for_boost
Type	true | false
Value	false

   For more details on System Property record fields, see Add a system property.
5. Select **Submit**.
   The system property's value is set to *false*.
6. For each AI Search application that you want to activate Now Assist Q&A Genius Result user context matching for, link the User Context Synonym Dictionary to the search application's search profile.

   a. Navigate to **All** > *AI Search* > **Search Experience** > **Search Applications**.
      The Search Application Configuration [sys_search_context_config] table displays in list view.

   b. Find the record for your chosen AI Search application and select the search profile record link in its **Search Profile** column.

Example
For example, if you want to activate Now Assist Q&A Genius Result user context matching for Service Portal, find the *`Service Portal`* **Default Search Application** and select the *`Service Portal`* **Default Search Profile** link in the **Search Profile** column.

c. In the Search Profile form's Synonyms related list, select **Link Existing**.

d. On the Search Profile - Dictionary Mapping form, fill in the **Dictionary** field.

Field	Value
Dictionary	user_context_synonym_dictionary

For a description of the field values, see Search Profile - Dictionary Mapping form.

e. Select **Submit**.
The User Context Synonym Dictionary appears in the search profile's Synonyms related list.

f. Select **Publish**.
For more details on publishing search profiles, see Publish a search profile.

The User Context Synonym Dictionary is linked to the chosen search profile.

[End]

Dynamic Translation for Now Assist Q&A Genius Results

Dynamic Translation improves the international search experience for Knowledge article content. When Dynamic Translation is activated, AI Search can generate Now Assist Q&A Genius Result answers from non-English Knowledge articles. Dynamic Translation also enables AI Search to translate Now Assist Q&A Genius Result answers into the search user's language.

When computing Now Assist Q&A Genius Result answers, the Now LLM Service only accepts English-language Knowledge article content and only generates answers in English.

Dynamic Translation expands the range of languages supported by Now Assist Q&A Genius Results by adding the following features:

Feature	Description
Content translation	The Now LLM Service automatically translates Knowledge article content from non-English languages into English before using that content to generate Now Assist Q&A Genius Result answers. **Note:** Translation is based on the language specified for the Knowledge article record. If the record has no language specified, the Now LLM Service analyzes the record's content to determine its language.

Feature	Description
Answer translation	When you search in a non-English language, the Now LLM Service automatically translates answers generated by Now Assist Q&A Genius Results from English into your search language.

To read more about the capabilities of Dynamic Translation and learn how to activate it, see Dynamic Translation.

Languages supported

Dynamic Translation works with all languages supported by the ServiceNow AI Platform®.

Interaction with Now Assist Q&A Genius Result cache

The query-time caches for Now Assist Q&A Genius Results only support English-language searches, articles, and answers. When searching in a non-English language or matching a non-English Knowledge article, Now Assist Q&A Genius Result behavior depends on the cache mode:

Cache mode	Description
`on` or `online`	AI Search bypasses the caches and queries the Now LLM Service to generate an answer for the Now Assist Q&A Genius Result.
`offline`	AI Search bypasses the caches and returns no answer for the Now Assist Q&A Genius Result.

For more details on the query-time caches for Now Assist Q&A Genius Results, see Caching for Now Assist Q&A Genius Results.

Caching for Now Assist Q&A Genius Results

AI Search provides two query-time caches to improve search performance for Now Assist Q&A Genius Results. Caching enables AI Search to return previously generated Now Assist Q&A Genius Result answers without submitting Knowledge articles to the Now LLM Service for answer generation.

Now Assist Q&A Genius Results caching overview

When a search query triggers a Now Assist Q&A Genius Result configuration, AI Search uses the search query and the sys_id of the search's top-ranked Knowledge article search result to look for cached article summaries. If no article summary is found in the active caches, AI Search submits the search query and article sys_id to the Now LLM Service for answer generation.

Note:
By default, AI Search only uses the first-level cache and the Now LLM Service when finding Now Assist Q&A Genius Result answers. To enable the second-level cache or disable use of the Now LLM Service, administrators can change the operational mode for the caches as described in this topic.

The query-time caches only support English-language searches, Knowledge articles, and answers.

Cache levels

The cache includes two levels which work in different ways to improve search performance for Now Assist Q&A Genius Results.

First-level cache

The first-level cache comprises a list of key-value pairs stored in memory for fast access. Each cache entry has a key that includes a search query and the sys_id of a Knowledge article returned by that query. The cache entry's value includes the summary generated by the Now LLM Service for the specified search query and Knowledge article.

When checking the first-level cache, AI Search compares your search query and the sys_id of your Knowledge article search result to the cache entry keys. If it finds a matching key, it returns the article summary from the corresponding cache entry value. Otherwise, it goes on to check the second-level cache.

Note:
The first-level cache only yields a result when your search query is an exact lexical match for the cached search query. For example, if you search for `avoiding scams` you won't get a result for a cached entry with search query `how to prevent scams` because the two search queries don't contain the same terms.

Second-level cache

The second-level cache comprises a table that is configured as an AI Search indexed source. Each record on this table is a cache entry, and includes a search query, the sys_id for an associated Knowledge article search result, the summary generated for that query and Knowledge article, and other fields such as pinned, sys_updated_on, and run_as. AI Search updates the index for this table whenever its records are created, updated, or deleted. This index update operation can take up to a minute.

When checking the second-level cache, AI Search queries the indexed table, looking for an entry that matches your search query and the sys_id of your Knowledge article search result. If it finds a matching entry, it returns the article summary stored in the indexed table. Otherwise, it goes on to submit your search query and Knowledge article search result to the Now LLM Service.

Note:
Unlike the first-level cache, the second-level cache compares search query meanings using semantic vector search, so you may get a cache result even if your search query isn't an exact lexical match for the cached search query. For example, if you search for `avoiding scams` you might get the result for a cached entry with search query `how to prevent scams` because the meanings of the two search queries are similar. For more information on semantic vector search, see Semantic vector search in AI Search.

Benefits of caching

Caching for Now Assist Q&A Genius Results provides the following benefits:

- Decreases average response time for common Now Assist Q&A Genius Result answers
- Lowers Now Assist entitlement consumption by reducing the number of search query results sent to the Now LLM Service for Now Assist Q&A answer extraction

- Increases the likelihood of returning a Now Assist Q&A Genius Result answer
- Improves search consistency by returning the same Now Assist Q&A Genius Result answer for similar searches

Content Security for cached queries

Because AI Search applies Content Security restrictions to your search before it matches Knowledge articles and checks the caches, neither cache returns hits for Knowledge articles that you don't have access to. For full details on AI Search's Content Security model, see Content security in AI Search.

Cache modes

The Now Assist Q&A Genius Result answer caches support the following operational modes:

- *off*: Use the first-level cache and the Now LLM Service to find Now Assist Q&A Genius Result answers.

  AI Search looks in the first-level cache for Now Assist Q&A Genius Result answers that exactly match your search query and Knowledge article result. If it doesn't find a matching answer, it sends your query and Knowledge article sys_id to the Now LLM Service for answer generation.

  Now Assist Q&A Genius Result answers generated by the Now LLM Service populate the first-level cache.

  When using Dynamic Translation, AI Search bypasses the caches and queries the Now LLM Service to generate an answer for the Now Assist Q&A Genius Result.

- *offline*: Use the first-level and second-level caches to find Now Assist Q&A Genius Result answers. Don't submit queries to the Now LLM Service.

  AI Search looks in the first-level cache for Now Assist Q&A Genius Result answers that exactly match your search query and Knowledge article result. If it doesn't find a matching answer, it uses semantic vector search to look for answers that match the meaning of your query in the second-level cache. If no cached answers match your query and Knowledge article result, AI Search returns no answer for the Now Assist Q&A Genius Result.

  Now Assist Q&A Genius Result answers found in the second-level cache populate the first-level cache.

  When using Dynamic Translation, AI Search bypasses the caches and returns no answer for the Now Assist Q&A Genius Result.

- *online*: Use the first-level and second-level caches and the Now LLM Service to find Now Assist Q&A Genius Result answers.

  AI Search looks in the first-level cache for Now Assist Q&A Genius Result answers that exactly match your search query and Knowledge article result. If it doesn't find a matching answer, it uses semantic vector search to look for answers that match the meaning of your query in the second-level cache. If no cached answers match your query and Knowledge article result, AI Search submits the query and article sys_id to the Now LLM Service for answer generation.

Now Assist Q&A Genius Result answers generated by the Now LLM Service populate the first-level and second-level caches.

When using Dynamic Translation, AI Search bypasses the caches and queries the Now LLM Service to generate an answer for the Now Assist Q&A Genius Result.

The default operational mode is `off`.

Administrators can change the operational mode for the Now Assist Q&A Genius Result answer caches by setting the value for the `sn_ais_assist.semantic_cache_mode` system property to `off`, `offline`, or `online`. For details on system property settings, see Add a system property.

For more details on using Dynamic Translation with Now Assist Genius Results, see Dynamic Translation for Now Assist Q&A Genius Results.

Scheduled job for cache management

The `Update Semantic Cache` scheduled job runs daily to perform the following tasks:

- Populate the second-level cache with results for the most frequently submitted queries found in the Search Event [sys_search_event] search signal table. For more information on this table, see Search signal tables.
- Purge all unpinned second-level cache entries that have not been used in the past seven days. Search administrators can pin results in the second-level cache table to prevent them from being purged. For more details on this procedure, see Pin cached answers for Now Assist Q&A Genius Results.

Pin cached answers for Now Assist Q&A Genius Results

Improve performance for Now Assist Q&A Genius Results by pinning frequently used answers in the second-level cache. Pinning an answer exempts it from the cache's purge mechanism.

Before you begin

Role required: ais_admin

About this task

Search administrators can pin entries in the second-level cache for Now Assist Q&A Genius Result answers. The `Update Semantic Cache` scheduled job ignores pinned entries when purging the second-level cache.

Pinning frequently used entries helps improve search performance by enabling AI Search to return previously generated Now Assist Q&A Genius Result answers without submitting Knowledge articles to the Now LLM Service for answer generation.

To learn more about the second-level Now Assist Q&A Genius Result answer cache and its usage, see Caching for Now Assist Q&A Genius Results.

Procedure

1. Navigate to the Now Assist in AI Search Semantic Cache [sn_ais_assist_semantic_cache] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sn_ais_assist_semantic_cache.list`.
   c. Press Enter.

Each record in this table represents a second-level cache entry. The table is populated with Now Assist Q&A Genius Result answers previously generated by the Now LLM Service.

2. Update each cache entry record that you want to pin.
   a. Open the cache entry record by selecting it from the list view.
   b. Set the value of the record's **Pinned** field to **true**.
   c. Select **Update**.

Result

The *Update Semantic Cache* scheduled job ignores your pinned entries when purging the second-level cache.

External Content Q&A Genius Results

External Content Q&A Genius Results display concise, actionable AI-generated answers derived from documents in your Microsoft SharePoint Online instance.

External Content Q&A Genius Results overview

The External Content Q&A Genius Result configuration searches for documents relevant to your search using an Integration Hub subflow. If it finds sufficiently relevant documents, it sends them to the Now LLM Service, which generates answers from their content and metadata.

Each External Content Q&A Genius Result answer card includes a generated answer. For reference, the answer card also includes a link you can select to download the source documents.

Note:
Because the External Content Q&A Genius Result answer is automatically generated, it's a good idea to review it for accuracy. You can provide feedback on the answer by selecting the thumbs-up icon () if the generated answer is accurate, or the thumbs-down icon () if it's not. Your feedback helps ServiceNow improve future versions of this Genius Result configuration.

When searching for relevant documents, the system uses your Microsoft SharePoint Online credentials and user profile, so your searches generate External Content Q&A Genius Result answers from documents that you can access.

Enabling External Content Q&A Genius Results

To make External Content Q&A Genius Results available for use in AI Search applications, administrators first need to perform the following steps:

1. Ensure that the ServiceNow IntegrationHub Enterprise Pack Installer [com.glide.hub.integrations.enterprise] plugin is installed on your instance. If this plugin isn't installed, you can request it. For details on requesting IntegrationHub plugins, see Request an Integration Hub plugin.
2. Create and configure an OAuth application in the Microsoft Azure portal to verify and authorize connection requests from your ServiceNow AI Platform® instance. To learn about this procedure, see .
3. Add the AllSites.FullControl SharePoint delegated permission (and grant admin consent) to your new OAuth application in the Microsoft Azure portal. For details on this procedure, see

Configure SharePoint API permissions for your External Content Q&A Genius Results OAuth 2.0 application.

4. Define OAuth configuration settings for External Content Q&A Genius Results. For full details on the required OAuth settings, see Configure OAuth settings for External Content Q&A Genius Results.

Next, search administrators need to enable External Content Q&A Genius Results in one or more AI Search portals. To learn about this procedure, see Enable Now Assist Genius Results in AI Search portals.

Once all these steps are completed, AI Search portal users can authorize External Content Q&A Genius Results for their searches by authenticating with Microsoft SharePoint Online. For details on this authentication procedure, see Sign in to Microsoft SharePoint Online to authorize External Content Q&A Genius Results.

Limitations

By default, External Content Q&A Genius Results only support English-language searches. Administrators can enable support for other languages by activating Dynamic Translation. To learn more about how content and answers are translated, see Dynamic Translation for Now Assist Q&A Genius Results. For more details on Dynamic Translation, see Dynamic Translation.

External Content Q&A Genius Results only generate answers from documents found in your Microsoft SharePoint Online instance. They don't generate answers from sites found in your instance.

Interaction with other search features

The following table describes the interactions between External Content Q&A Genius Results and other search features.

Feature	Interaction with External Content Q&A Genius Results Genius Results
Result improvement rules	When computing External Content Q&A Genius Result answers, AI Search doesn't apply your result improvement rules because they don't operate on documents from your Microsoft SharePoint Online instance.
Stop words	When computing External Content Q&A Genius Result answers, AI Search doesn't remove stop words from your search.
Synonyms	When computing External Content Q&A Genius Result answers, AI Search doesn't expand synonyms found in your search.
Typo handling	When computing External Content Q&A Genius Result answers, AI Search doesn't correct misspelled search terms.

Configure SharePoint API permissions for your External Content Q&A Genius Results OAuth 2.0 application

Add the AllSites.FullControl SharePoint API delegated permission to your External Q&A Genius Results OAuth 2.0 application in Microsoft Azure portal and grant admin consent to allow the application to access this permission. The OAuth 2.0 application for External Q&A Genius Results requires the delegated permission to search your Microsoft SharePoint Online sites.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

You need to have an OAuth 2.0 application configured in Microsoft Azure portal for External Content Q&A Genius Results. For details on this procedure, see .

You must have permissions to change the OAuth 2.0 application's settings in Microsoft Azure portal.

Role required: admin

Procedure

1. Access your OAuth 2.0 application's settings in Microsoft Azure portal.

   a. Log in to https://portal.azure.com/ .

   b. In the **Azure services** section, select **App registrations**.

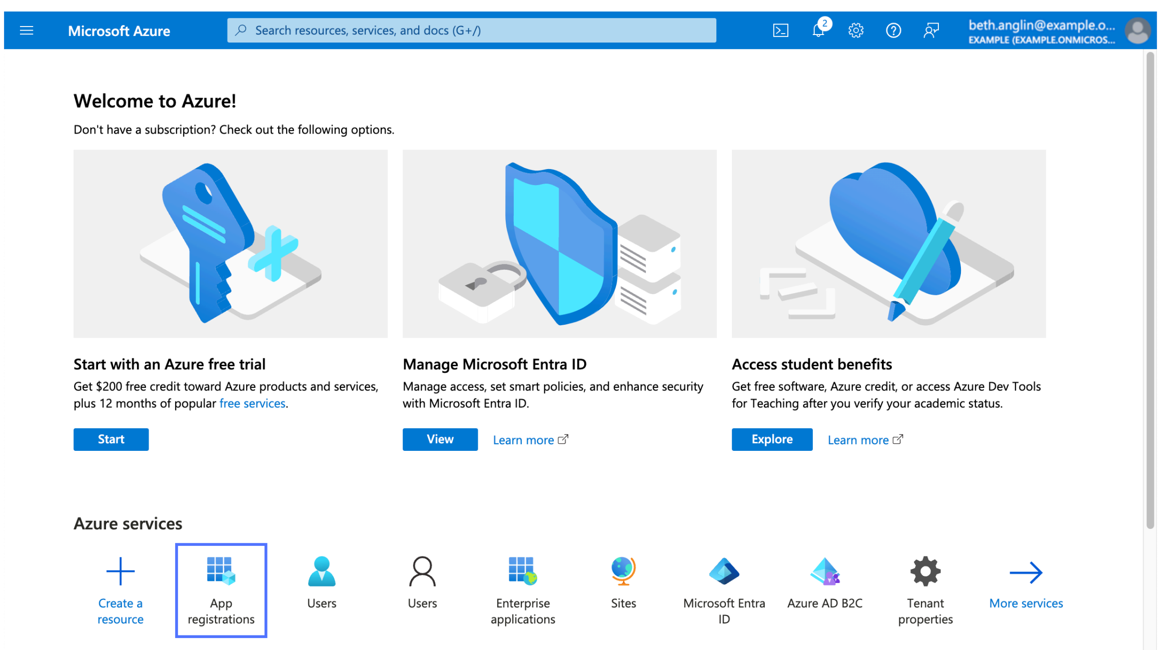

   c. Select **All applications** or **Owned applications**.

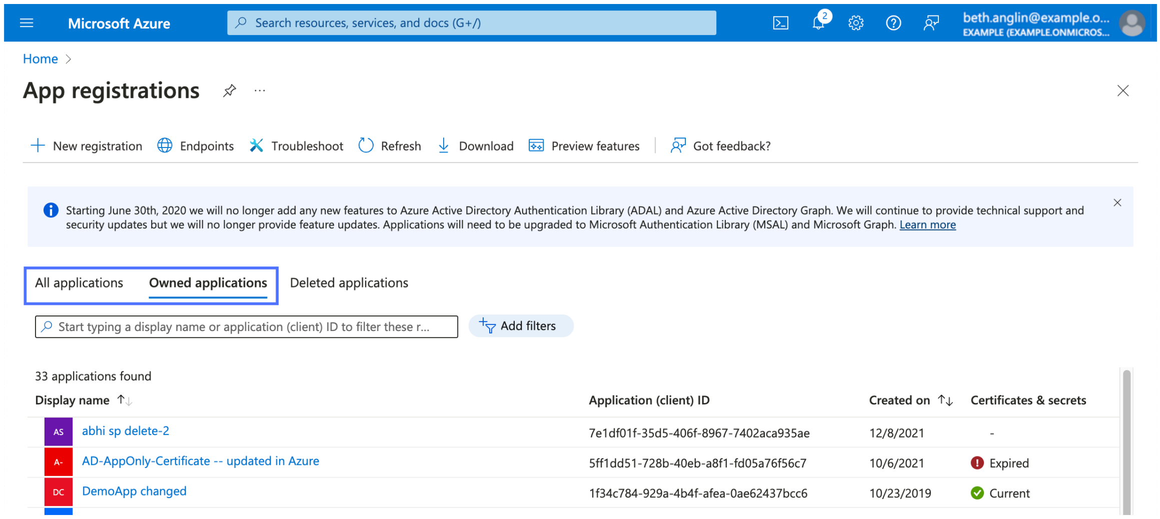

d. In the search field, enter the name or application (client) ID of the OAuth 2.0 application you created for External Content Q&A Genius Results, then select that application's display name from the list of results.

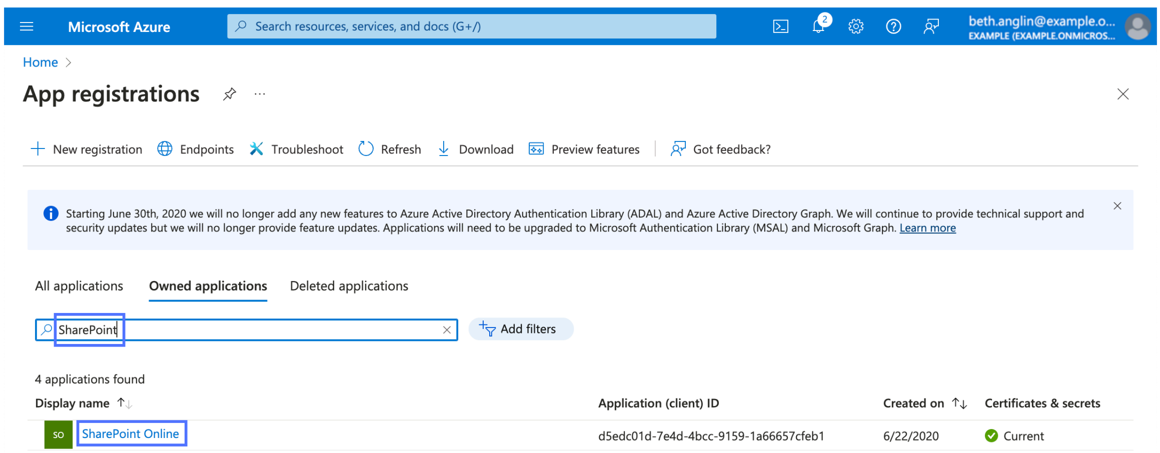

To learn how to configure an OAuth 2.0 application for External Content Q&A Genius Results, see .

2. Add an AllSites.FullControl SharePoint API delegated permission to your OAuth 2.0 application.

a. In the left panel, select **Manage** > **API permissions**.

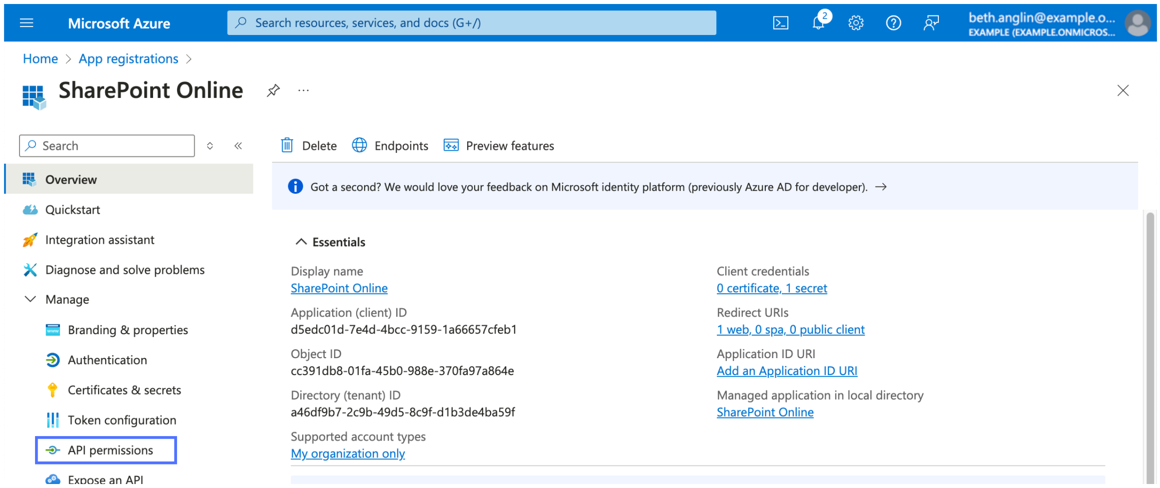

b. Under the **Configured permissions** heading, select **Add a permission**.

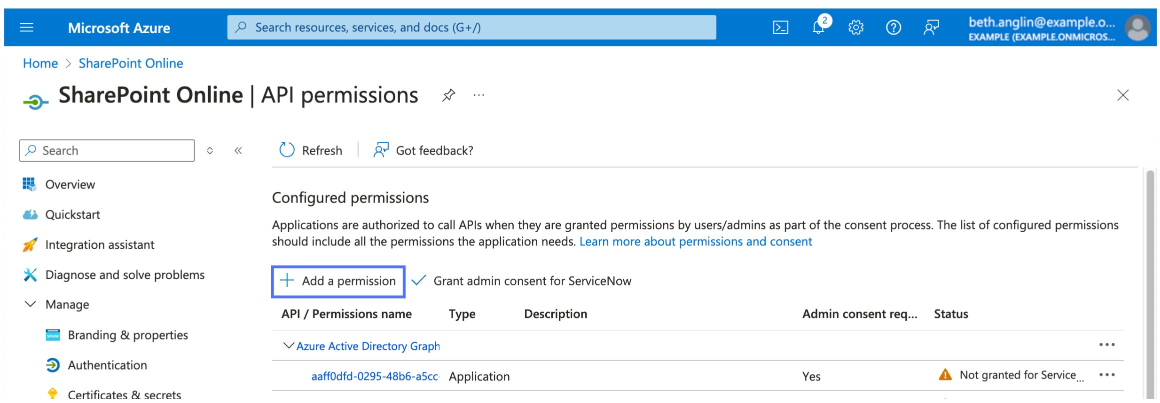

c. In the **Request API permissions** window's **Select an API** section, select **Microsoft APIs**, then select the **SharePoint** API.

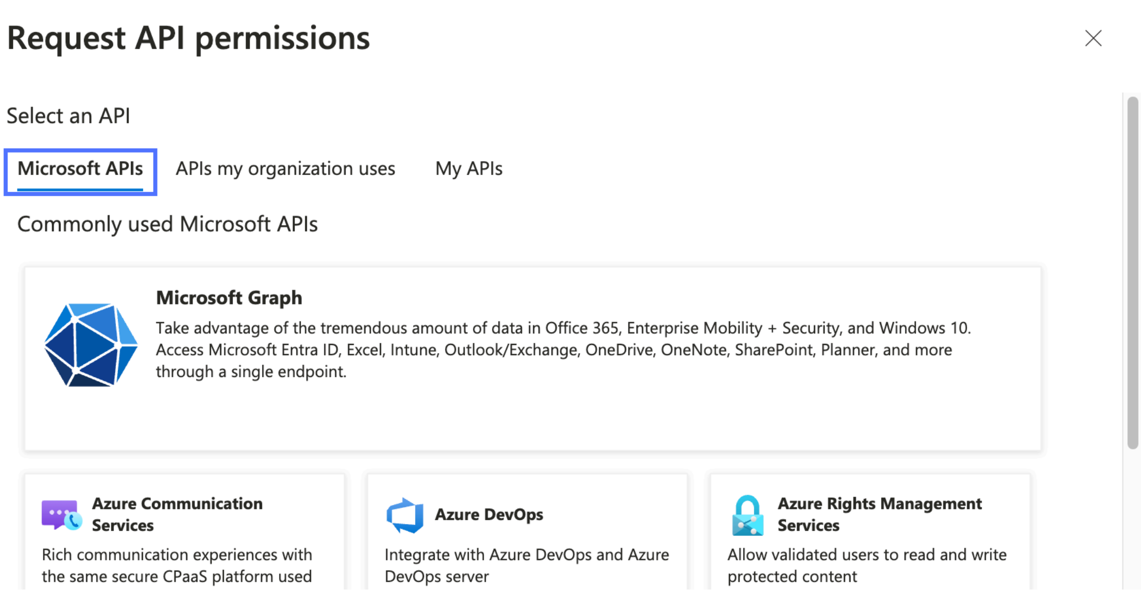

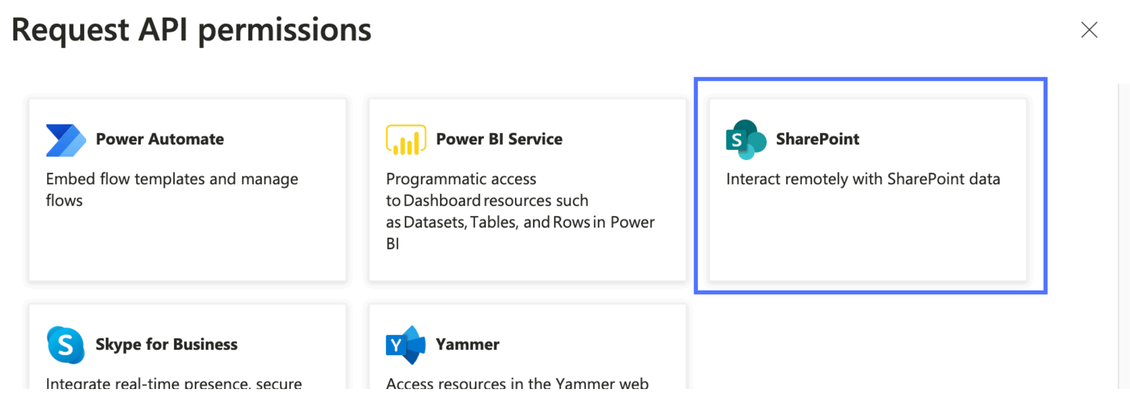

d. When prompted, select **Delegated permissions** as the type of permission to add.

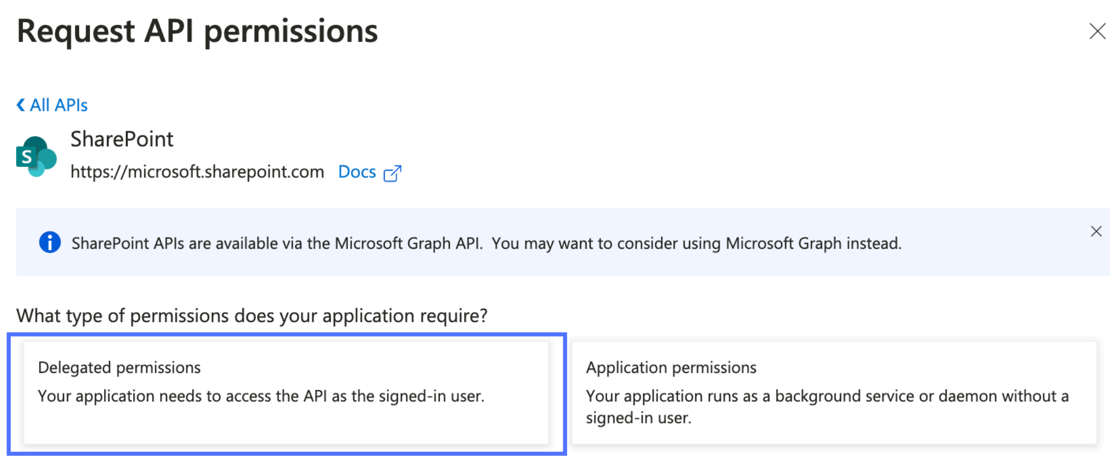

e. In the **Select permissions** section, select **AllSites** > **AllSites.FullControl**, then select **Add permissions**.

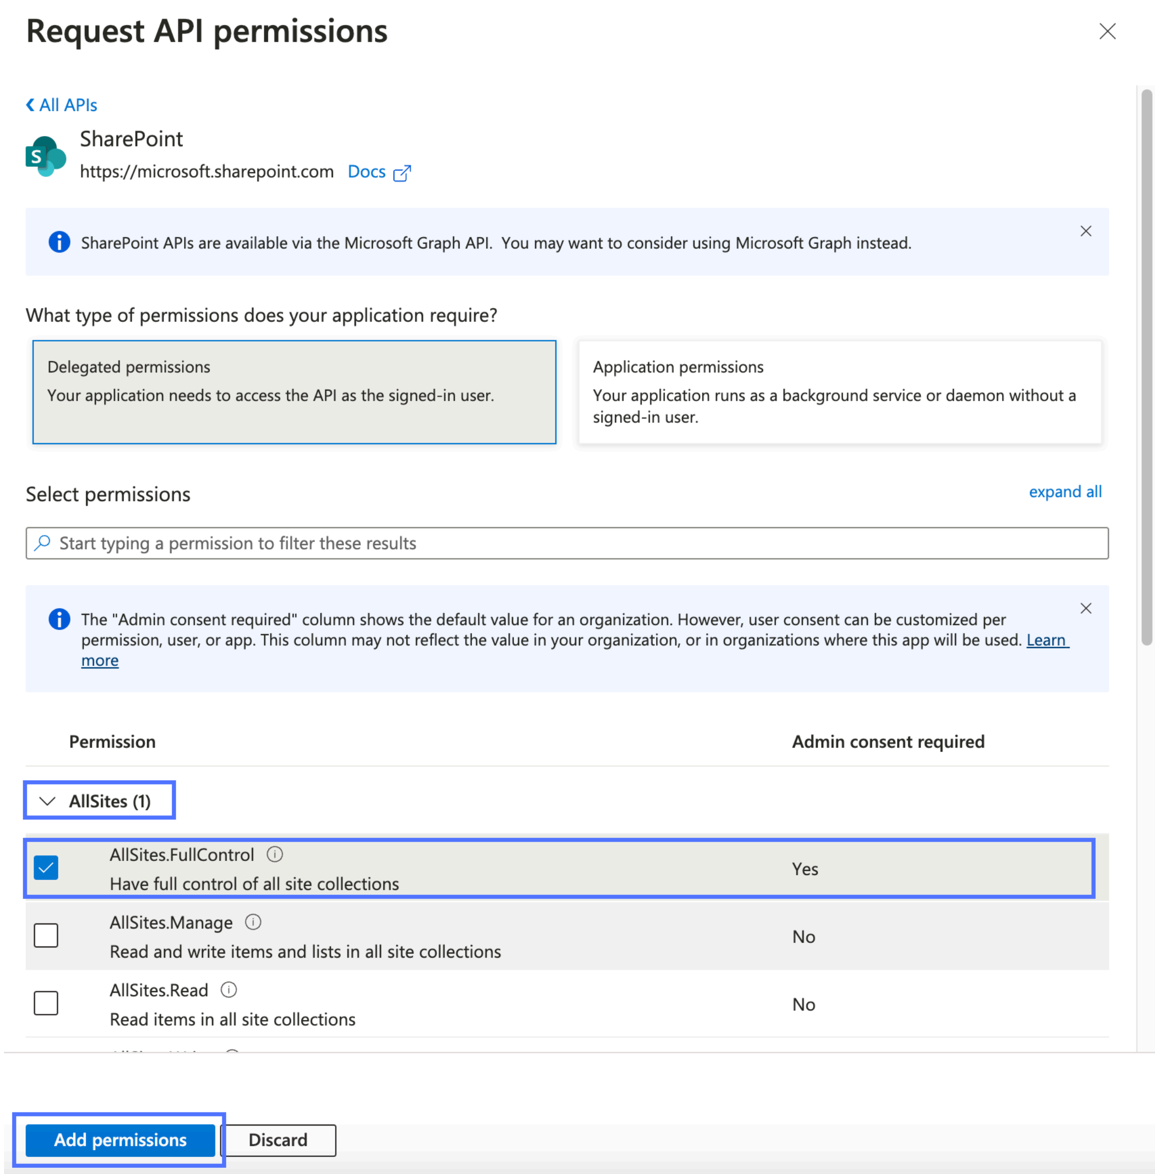

The new AllSites.FullControl SharePoint API permission appears in the list of configured permissions with status `Not granted for` *`ServiceNow`*.

f. Under the **Configured permissions** heading, select **Grant admin consent for** *`ServiceNow`*.

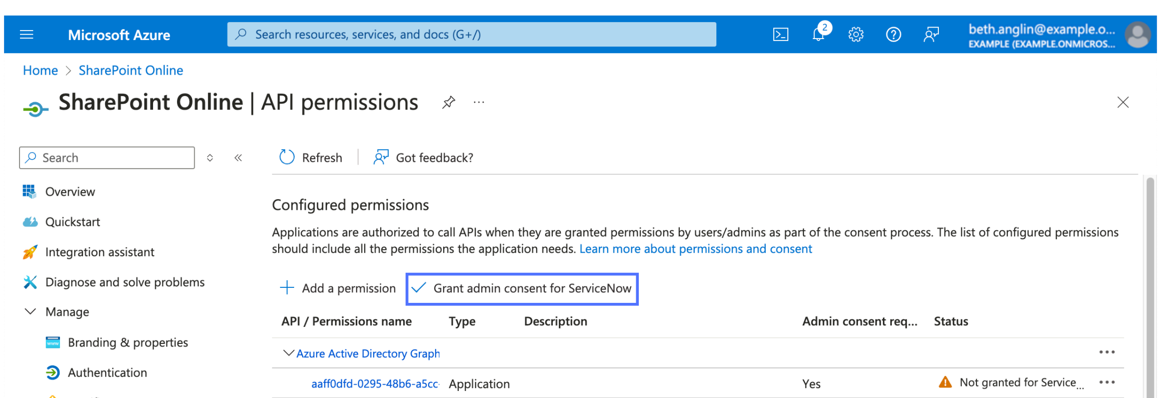

g. In the **Grant admin consent confirmation** dialog box, select **Yes**.

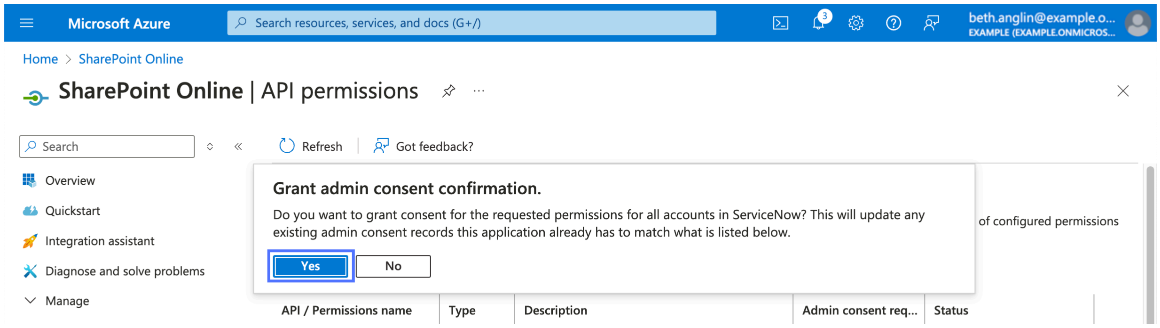

The AllSites.FullControl SharePoint API permission's status changes to `Granted for ServiceNow`.

What to do next

Configure the OAuth 2.0 settings in your ServiceNow AI Platform® instance that are needed for External Content Q&A Genius Results. For details on this procedure, see Configure OAuth settings for External Content Q&A Genius Results.

Configure OAuth settings for External Content Q&A Genius Results

Configure the settings that External Content Q&A Genius Results need to interact with your Microsoft SharePoint Online instance.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

You must have already completed the following tasks:

- Request installation of the ServiceNow IntegrationHub Enterprise Pack Installer [com.glide.hub.integrations.enterprise] plugin
-

You will need read access to the following Microsoft Azure Portal configuration settings for your Microsoft SharePoint Online instance:

- Root site URL
- Client ID
- Client secret
- Tenant ID
- Authorization URL
- Token URL
- Token revocation URL
- Redirect URL

Role required: admin

About this task

Administrators need to perform this one-time procedure to enable External Content Q&A Genius Results.

Procedure

1. Create an Application Registry record:

   a. Navigate to **All** > **System OAuth** > **Application Registry**.

   b. Select **New**.

   c. Select **Connect to a third party OAuth Provider**.

   d. On the Application Registries form, fill in the fields:

Field	Value
Name	Enter a name of your choice for the Application Registry record.
Client ID	Enter the client ID value from your Microsoft SharePoint Online configuration on Microsoft Azure Portal.
Client Secret	Enter the client secret value from your Microsoft SharePoint Online configuration on Microsoft Azure Portal.
Default Grant type	Select *`Authorization Code`*.
Authorization URL	Enter the fully-qualified authorization URL for your Microsoft SharePoint Online tenant instance in this format: `https://login.microsoftonline.com/<tenant-id>/oauth2/v2.0/authorize`. Replace `<tenant-id>` with your Microsoft SharePoint Online tenant ID, as found in your Microsoft Azure Portal settings.
Token URL	Enter the fully-qualified token URL for your Microsoft SharePoint Online tenant instance in this format: `https://login.microsoftonline.com/<tenant-id>/oauth2/v2.0/token`. Replace `<tenant-id>` with your Microsoft SharePoint Online tenant ID, as found in your Microsoft Azure Portal settings.
Token Revocation URL	Enter the fully-qualified token revocation URL for your Microsoft SharePoint Online tenant instance in this format: `https://login.microsoftonline.com/<tenant-id>/oauth2/v2.0/token`. Replace `<tenant-id>` with your Microsoft SharePoint Online tenant ID, as found in your Microsoft Azure Portal settings.
Redirect URL	Enter the fully-qualified redirect URL for the Integration Hub subflow in this format: `https://<instance-`

Field	Value
	`name>.service-now.com/oauth_redirect.do`. Replace `<instance-name>` with your ServiceNow AI Platform® instance name, as specified in your Microsoft Azure Portal settings.

e. Select **Submit**.
The new Application Registry record appears in the table list view.

2. Define an OAuth entity scope for your new Application Registry record:

a. Open your new Application Registry record in form view by selecting it.

b. In the OAuth Entity Profiles related list, verify that an OAuth entity profile record appears with **Is default** set to *`true`* and **Grant type** set to *`Authorization Code`*.

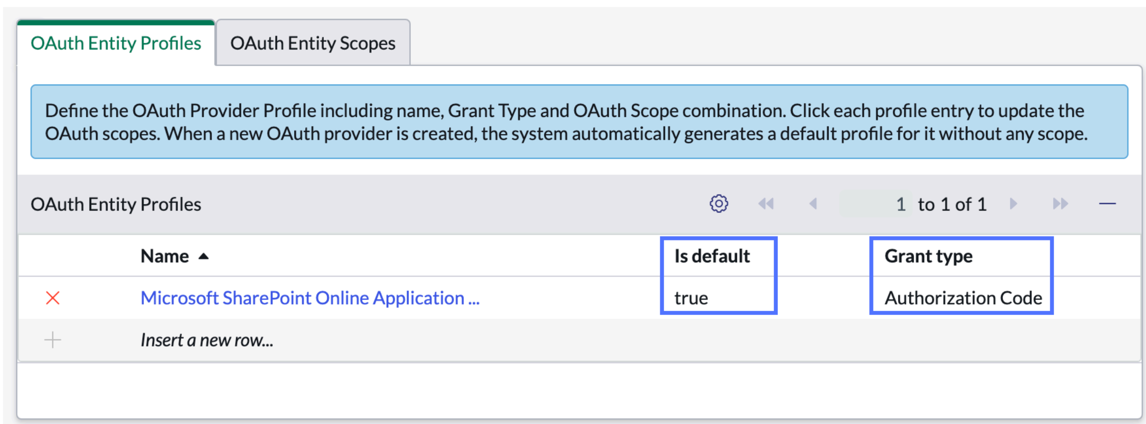

c. Open the new OAuth entity profile record in form view by selecting its name.

d. On the OAuth Entity Profile form, in the OAuth Entity Profile Scopes related list, double-click **Insert a new row**, then select the Lookup using list ( ) icon.

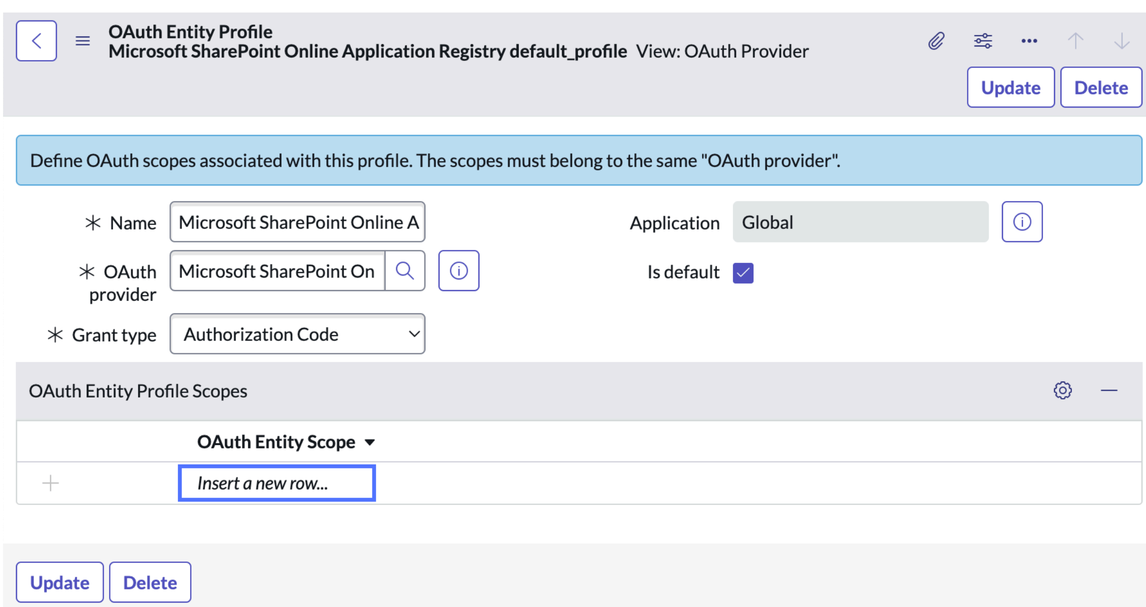

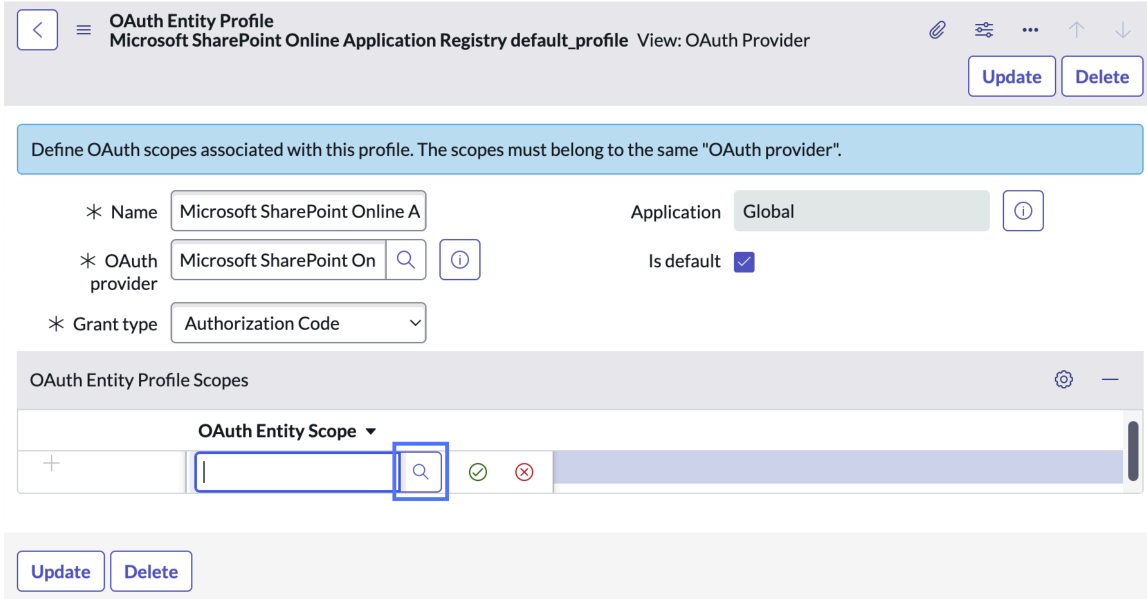

e. On the OAuth Entity Scopes form, select **New**.

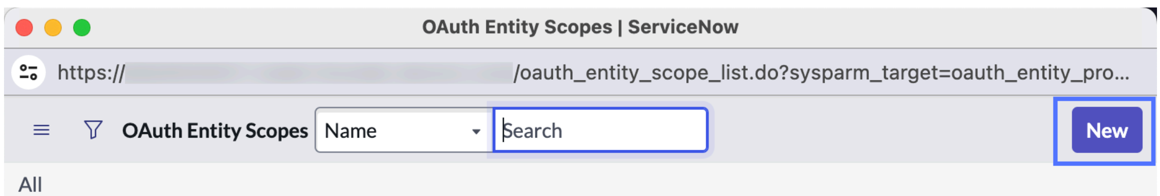

f. On the OAuth Entity Scope form, fill in the fields:

Field	Value
Name	Enter a name of your choice for the connection and credential alias.
OAuth provider	Select or enter the name of the Application Registry record that you created in step 1.e.
OAuth scope	Enter `offline_access https://<root-site>.sharepoint.com/.default`, replacing `<root-site>` with the name of your Microsoft SharePoint Online instance's root site. **Important:** Make sure to include the dot before `default`.

g. Select **Submit**.

h. In the OAuth Entity Profile Scopes related list, select the Save (Enter) icon ( ).

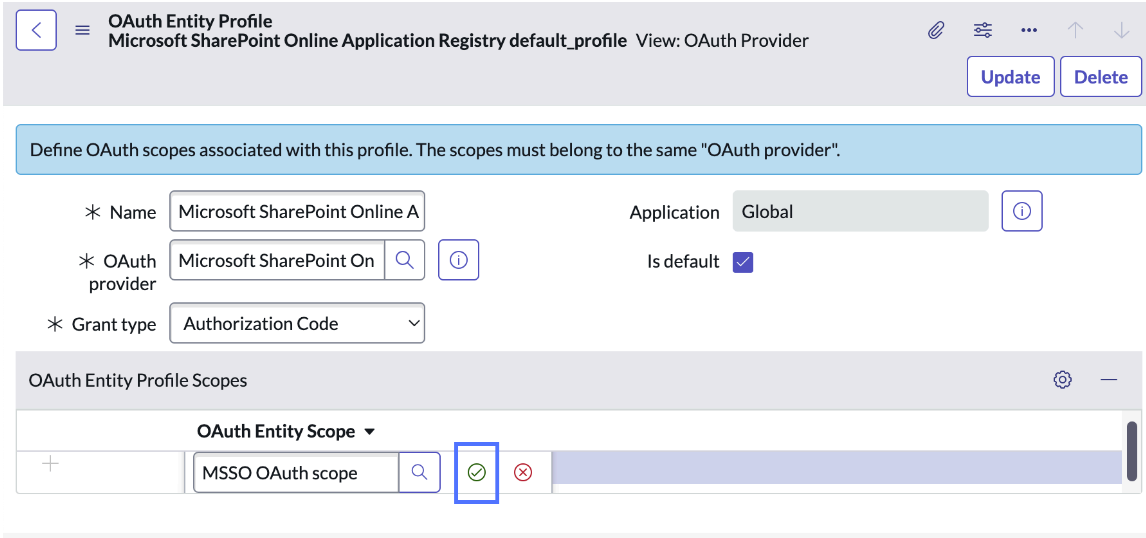

The new OAuth entity scope record appears in the OAuth Entity Profile Scopes related list.

i. Select **Update**.

3. Define a connection and credential alias record for Microsoft SharePoint Online:

   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.

   b. Open the **MicrosoftSharePointOnlineforAISearch** alias record in form view by selecting it.

   c. In the Connections related list, select **New**.

   d. On the HTTP(s) Connection form, fill in the fields:

Field	Value
Name	Enter a name of your choice for the connection and credential alias.
Connection URL	Enter the fully-qualified URL for your Microsoft SharePoint Online instance: https://<root-site>.sharepoint.com.

   e. To create your new connection record, select **Submit**.
   The new connection record appears in the Connections related list.

4. Define a new OAuth 2.0 credential associated with your Application Registry record:

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

   b. Select **New**.

   c. Select **OAuth 2.0 Credentials**.

d. If the **Integration Type** field isn't visible on the OAuth 2.0 Credentials form, configure the form layout to make the field visible.
For details on configuring a form layout, see Configuring the form layout.

e. On the OAuth 2.0 Credentials form, fill in the fields:

Field	Value
Name	Enter a name of your choice for the OAuth 2.0 credential record.
OAuth Entity Profile	Select the new OAuth entity profile record associated with the Application Registry record that you created in step 1.e.
Integration Type	Select `Personal`.

f. Select **Submit**.
The new OAuth 2.0 credentials record appears in the Credentials [discovery_credentials] table list view.

5. Specify your new OAuth 2.0 credential as the one to use when connecting to Microsoft SharePoint Online:

a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.

b. Open the **MicrosoftSharePointOnlineforAISearch** alias record in form view by selecting it.

c. In the Connections related list, select the connection record that you created in step 3.e.

d. On the HTTP(s) Connection form, select the **Credential** field's Lookup using list ( ) icon, then select the credential that you defined in step 4.

e. Select **Update**.

What to do next

You can optionally configure a Microsoft Keyword Query Language (KQL) query to restrict the set of Microsoft SharePoint Online sites that External Content Q&A Genius Results search for relevant documents. For details on defining this KQL query, see Restrict the Microsoft SharePoint Online sites that External Content Q&A Genius Results can search.

Restrict the Microsoft SharePoint Online sites that External Content Q&A Genius Results can search

Define a Microsoft Keyword Query Language (KQL) query to restrict the set of Microsoft SharePoint Online sites (paths) that External Content Q&A Genius Results can search for documents. Only documents that satisfy your query are eligible for External Content Q&A Genius Result answer generation.

Before you begin

The Now Assist in AI Search ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install Now Assist in AI Search.

Role required: admin

About this task

By default, External Content Q&A Genius Results search for relevant documents in all sites (paths) in your Microsoft SharePoint Online source system. You can restrict the set of sites

searched by specifying a Microsoft Keyword Query Language (KQL) property restriction query expression as the value for the *sn_ais_assist.sharepoint_search.query_filter* system property. Only documents from sites that satisfy your property restriction query are eligible for External Content Q&A Genius Result retrieval and answer generation.

To learn more about KQL property restriction queries and the operators they support, see https://learn.microsoft.com/en-us/sharepoint/dev/general-development/keyword-query-language-kql-syntax-reference.

Procedure

1. Navigate to the System Property [sys_properties] table's list view.
   a. Select **All**.
   b. In the **Filter** field, enter `sys_properties.list`.
   c. Press Enter.
2. Search the list for the *sn_ais_assist.sharepoint_search.query_filter* system property record.
   - If the system property record appears in the list, open it in form view by selecting it.
   - If the system property record does not appear in the list, add it, setting its **Type** to `string`. To learn how to add a new system property, see Add a system property.
3. In the system property record's **Value** field, enter a valid KQL property restriction query expression.
   As an example, suppose your Microsoft SharePoint Online source system is `https://example.sharepoint.com`.
   - To limit External Content Q&A Genius Result searches to documents in your Marketing and Sales sites, you might enter `path:"https://example.sharepoint.com/teams/Marketing/" path:"https://example.sharepoint.com/teams/Sales/"` as your property restriction query expression.
   - To prevent External Content Q&A Genius Results from searching documents in your Engineering and QA sites, you might enter `NOT (path:"https://example.sharepoint.com/teams/Engineering/" path:"https://example.sharepoint.com/teams/QA/")` as your query expression.

   Note:
   You can use KQL Boolean and inclusion/exclusion operators to combine multiple property restriction expressions into a complex query expression.

4. Select **Update**.

Result

When searching for documents in your Microsoft SharePoint Online source system, External Content Q&A Genius Results apply your property restriction KQL query as a filter. Only documents in sites that satisfy your query are eligible for retrieval and answer generation.

Sign in to Microsoft SharePoint Online to authorize External Content Q&A Genius Results

Obtain an OAuth 2.0 token from Microsoft SharePoint Online so that External Content Q&A Genius Results can find documents and generate answers when you search.

Before you begin

You only need to perform this task if your search results page displays this informational message: `To view results generated from SharePoint, you need to be signed in.`

Role required: none

About this task

External Content Q&A Genius Results use an OAuth 2.0 token to access your Microsoft SharePoint Online instance and search for relevant documents. The OAuth 2.0 token is specific to your Microsoft SharePoint Online user account, so your search only returns documents that you can access.

Perform this task to generate a valid OAuth 2.0 token in Microsoft SharePoint Online and authorize External Content Q&A Genius Results for your searches.

Note:
How often you need to sign in depends on your Microsoft SharePoint Online instance's settings. For questions about the OAuth 2.0 token's lifespan, contact your Microsoft SharePoint Online administrator.

Procedure

1. Select the **Sign in** link in the informational message on your search results page.
2. When prompted, enter your Microsoft SharePoint Online credentials to authenticate with Microsoft.

   Note:
   This login step submits your credentials directly to Microsoft. If you're unable to authenticate, contact your Microsoft SharePoint Online administrator.

3. Once you're authenticated, repeat your previous search.

   Note:
   External Content Q&A Genius Results respect your Microsoft SharePoint Online user profile security settings. If you're not seeing an expected Genius Result answer, it's a good idea to check with your Microsoft SharePoint Online administrator to make sure you have access to the appropriate source documents.

Now Assist Actions Genius Results

Now Assist Actions Genius Results display actionable answers showing Catalog Items and Virtual Agent topics that match your search query. AI Search matches records based on their similarity to your search's intent and meaning instead of looking for exact term matches.

Now Assist Actions Genius Results send content from your search's most relevant Catalog Item records and Virtual Agent topic records to the Now LLM Service. The Now LLM Service ranks and filters the selected records' content for display.

A Now Assist Actions Genius Result answer card can display content from up to three records. AI Search uses semantic vector search and legacy keyword search to select records that best match the meaning and intent of your search query. For more details on semantic vector search, see Semantic vector search in AI Search.

Supported record types

In Service Portal, Now Assist Actions Genius Results only operate on Catalog Items. In Virtual Agent, Now Assist Actions Genius Results operate on Virtual Agent topics as well as on Catalog Items. To learn more about Virtual Agent topics, see Virtual Agent.

In Virtual Agent, Now Assist Actions Genius Results rank Catalog Item and Virtual Agent topic records using a unified relevancy scoring scheme. As a result, Now Assist Actions Genius Result answers may display records of either type in any order.

Activating Now Assist Actions Genius Results

To enable the Now Assist Actions skill in an AI Search application, you can link the Now Assist Actions Genius Result configuration to your search profile for the search application. For details on this procedure, see Link a Genius Result configuration to a search profile.

For AI Search portal applications, you can enable Now Assist Genius Results using the Now Assist in AI Search Setup page. For more information on this approach, see Enable Now Assist Genius Results in AI Search portals.

Note:
Now Assist Actions Genius Results are only supported in Service Portal and Virtual Agent.

Limitations

Now Assist Actions Genius Results only support English-language searches.

Interaction with other search features

The following table describes the interactions between Now Assist Actions Genius Results and other search features.

Feature	Interaction with Now Assist Actions Genius Results
Result improvement rules	When computing Now Assist Actions Genius Result answers for a search query, AI Search applies result improvement rules normally. The effects depend on the result improvement rule's action, as follows: • *block*: Now Assist Actions Genius Results don't return blocked records as answers. • *boost* or *promote*: Boosted and promoted records are more likely to be returned as Now Assist Actions Genius Result answers.
Stop words	Now Assist Actions Genius Results use a blend of semantic vector search, which doesn't support stop words, and keyword-based search. AI Search only removes stop words from keyword-based searches, so Now Assist Actions Genius Result answers may not reflect your stop words settings.
Synonyms	Now Assist Actions Genius Results use a blend of semantic vector search, which doesn't support synonyms, and keyword-based search. AI Search expands synonyms in keyword-based searches, so your synonyms are likely to improve the relevancy of Now Assist Actions Genius Result answers.

Feature	Interaction with Now Assist Actions Genius Results
Typo handling	When computing Now Assist Actions Genius Result answers for a search query, AI Search corrects misspelled terms in the query.

Dynamic Translation for Now Assist Actions Genius Results

Dynamic Translation improves the international search experience for Catalog Items and Virtual Agent topics. When Dynamic Translation is activated, AI Search can translate searches into English before matching Catalog Items and Virtual Agent topics for inclusion in Now Assist Actions Genius Result answers, and can translate generated answers into the search user's language.

When computing Now Assist Actions Genius Result answers, the Now LLM Service only supports English-language searches for Catalog Items and Virtual Agent topics.

Dynamic Translation expands the range of languages supported by Now Assist Actions Genius Results by adding the following features:

Feature	Description
Search translation	The Now LLM Service automatically translates non-English searches into English and uses the translated search to find Catalog Items and Virtual Agent topics for inclusion in Now Assist Actions Genius Result answers.
Answer translation	When you search in a non-English language, the Now LLM Service automatically translates Now Assist Actions Genius Result answers into your search language. The Genius Result answer card displays the translated text.

To read more about the capabilities of Dynamic Translation and learn how to activate it, see Dynamic Translation.

Languages supported

Dynamic Translation works with all languages supported by the ServiceNow AI Platform®.

Now Assist in AI Search reference

Reference information for the Now Assist in AI Search application.

Components installed with Now Assist in AI Search

The Now Assist in AI Search application installs new system components including scheduled jobs and Entity View Action Mapping (EVAM) configurations and templates.

Scheduled jobs installed with Now Assist in AI Search

Scheduled job	Description
Update Semantic Cache	• Populate the second-level Now Assist Q&A Genius Results cache with results for the most frequently submitted queries found in the Search Event [sys_search_event] search signal table. For more information on this table, see Search signal tables. • Purge all unpinned entries in the Now Assist Q&A Genius Results second-level cache that have not been used in the past seven days. Search administrators can pin results in the second-level cache table to prevent them from being purged. For more details on this procedure, see Pin cached answers for Now Assist Q&A Genius Results.

Entity View Action Mapper (EVAM) configurations installed with Now Assist in AI Search

Configuration	Description
AI Search Global - Now Assist Q&A Genius Result	Contains the view configuration settings for the Now Assist Q&A Genius Result configuration in AI Search for Next Experience global search. • Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**
Now Assist Q&A Genius Result	Contains the view configuration settings for the Now Assist Q&A Genius Result configuration in a search application. • Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**
Taxonomy - Now Assist Q&A Genius Result	Contains the view configuration settings for the Now Assist Q&A Genius Result configuration in Content Taxonomy. • Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**
Virtual Agent - Now Assist Q&A Genius Result	Contains the view configuration settings for the Now Assist Q&A Genius Result configuration in Virtual Agent.

Configuration	Description
	• Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Configurations**

Entity View Action Mapper (EVAM) templates installed with Now Assist in AI Search

Template	Description
AI Search Global - Now Assist Q&A Genius Result Template	Contains the component, static value, field-mapping, and action-mapping settings for the Now Assist Q&A Genius Result configuration in AI Search for Next Experience global search. • Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Templates**
Now Assist Q&A Genius Result Template	Contains the component, static value, field-mapping, and action-mapping settings for the Now Assist Q&A Genius Result configuration in Service Portal. • Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Templates**
Virtual Agent - Now Assist Q&A Genius Result Template	Contains the component, static value, field-mapping, and action-mapping settings for the Now Assist Q&A Genius Result configuration in Virtual Agent. • Application scope: Now Assist in AI Search • Location: **All** > **Entity View Action Mapper (*EVAM*)** > **View Definitions** > **View Templates**

List of all components installed with Now Assist in AI Search

To view the complete list of components installed with Now Assist in AI Search, follow the steps in Find components installed with an application. The application's package name is *`Now Assist`* **in** *`AI Search`*.

Now Assist in AI Search properties

System properties that you can use to modify behavior for Now Assist in AI Search Genius Results.

These properties are available for Now Assist in AI Search.

Note:
To open the System Property [sys_properties] table's list view, select **All**, enter `sys_properties.list` in the **Filter** field, and press **Enter**.

Properties for Now Assist in AI Search

Property	Description
sn_ais_assist.semantic_cache_mode	Operational mode for the Now Assist Q&A Genius Results caches. Use this mode to control how AI Search uses the caches and the Now LLM Service when looking for Now Assist Q&A Genius Result answers. For details about how AI Search caches Now Assist Q&A Genius Result answers, see Caching for Now Assist Q&A Genius Results. • Type: string • Default value: off • Supported values: ◦ *off*: Use the first-level cache and the Now LLM Service to find Now Assist Q&A Genius Result answers. AI Search looks in the first-level cache for Now Assist Q&A Genius Result answers that exactly match your search query and Knowledge article result. If it doesn't find a matching answer, it sends your query and Knowledge article sys_id to the Now LLM Service for answer generation. Now Assist Q&A Genius Result answers generated by the Now LLM Service populate the first-level cache. When using Dynamic Translation, AI Search bypasses the caches and queries the Now LLM Service to generate an answer for the Now Assist Q&A Genius Result. ◦ *offline*: Use the first-level and second-level caches to find Now Assist Q&A Genius Result answers. Don't submit queries to the Now LLM Service. AI Search looks in the first-level cache for Now Assist Q&A Genius Result answers that exactly match your search query and Knowledge article result. If it doesn't find a matching answer, it uses semantic vector search to look for answers that match the meaning of your query in the second-level cache. If no cached answers match your query and Knowledge article result, AI Search returns no answer for the Now Assist Q&A Genius Result. Now Assist Q&A Genius Result answers found in the second-level cache populate the first-level cache. When using Dynamic Translation, AI Search bypasses the caches and returns no answer for the Now Assist Q&A Genius Result.

Properties for Now Assist in AI Search (continued)

Property	Description
	◦ *online*: Use the first-level and second-level caches and the Now LLM Service to find Now Assist Q&A Genius Result answers. AI Search looks in the first-level cache for Now Assist Q&A Genius Result answers that exactly match your search query and Knowledge article result. If it doesn't find a matching answer, it uses semantic vector search to look for answers that match the meaning of your query in the second-level cache. If no cached answers match your query and Knowledge article result, AI Search submits the query and article sys_id to the Now LLM Service for answer generation. Now Assist Q&A Genius Result answers generated by the Now LLM Service populate the first-level and second-level caches. When using Dynamic Translation, AI Search bypasses the caches and queries the Now LLM Service to generate an answer for the Now Assist Q&A Genius Result. • Location: Not present in the System Property [sys_properties] table by default. To override the default value, add this property by following the steps in Add a system property.
sn_ais_assist.u_llm_api_timeout	Maximum time in seconds to wait for Now Assist Q&A answer generation results from the Now LLM Service. • Type: integer • Default value: 120 • Location: Not present in the System Property [sys_properties] table by default. To override the default value, add this property by following the steps in Add a system property.
sn_ais_assist.u_question_regex	Case-insensitive Java regular expression pattern that defines searches eligible for Now Assist Q&A answer generation. As an example, the pattern `^(what|where|when|why|who|whom|which|how)\b` restricts eligibility to search queries that begin with one of the specified keywords. When this system property has a value defined, only searches in supported applications that match the specified pattern are eligible for Now Assist Q&A answer generation. When no value is defined, all searches in supported applications are eligible for Now Assist Q&A answer generation. • Type: string • Default value: none • Location: System Property [sys_properties] table

Properties for Now Assist in AI Search (continued)

Property	Description
	To learn about Java regular expression pattern syntax, see the Javadoc for the java.regex.util.Pattern class.
sn_ais_assist.u_sanitize_html	Option to remove HTML tags from Knowledge article text before sending to the Now LLM Service for Now Assist Q&A answer generation. • Type: true | false • Default value: false • Location: Not present in the System Property [sys_properties] table by default. To override the default value, add this property by following the steps in Add a system property.

[store-future: BEGIN review]

AI Search Tuning

The AI Search Tuning ServiceNow® Store application enables search administrators and selected business users to streamline evaluation and tuning of AI Search relevancy.

[End]

Install AI Search Tuning

Install the AI Search Tuning application from the ServiceNow Store.

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.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the AI Search Tuning application (sn_aisearch_global) 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.

   Visit the ServiceNow Store website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.
3. Select **Install**.

Using AI Search Tuning

Search administrators can use AI Search Tuning to create AI Search evaluations. For each evaluation, search admins can assign business subject-matter experts as search result reviewers, and view search analytics based on result ratings.

Assign the AI Search Tuning reviewer role to business users

Grant the sn_ais_data.tuning_reviewer role to business users so they can rate and review search results in tuning evaluations.

Before you begin

The AI Search Tuning ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search Tuning.

Role required: admin

About this task

To rate and review search results in AI Search Tuning evaluations, users need either the ais_admin role or the sn_ais_data.tuning_reviewer role.

The ais_admin role is restricted to AI Search administrators. Grant the sn_ais_data.tuning_reviewer role to business users that you want to specify as reviewers in AI Search Tuning evaluations.

Procedure

- Assign a role to a user
- Assign a role to a group

Create an AI Search Tuning evaluation

Create an evaluation to measure two search applications' result relevancy for a specified set of search queries. Compare the performance of AI Search and Zing text search, or compare the performance of two different AI Search applications.

Before you begin

The AI Search Tuning ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search Tuning.

Role required: ais_admin

About this task

An AI Search Tuning evaluation lets you compare search result relevancy for two search applications. You can create an evaluation for either of the following use cases:

- Compare search result relevancy for Zing and AI Search applications that have the same search sources
- Compare search result relevancy for two AI Search applications that share the same search sources

The evaluation includes the following settings:

1. The two search applications to compare — a baseline application and a comparison application
2. The set of search terms and the search language to use when evaluating search result relevancy for the two selected applications
3. The search users assigned to rate and review search results for the selected queries in the two selected applications
4. A deadline for completion of the rating and review stages

Procedure

1. Navigate to **All** > **AI Search Tuning** > **Admin Homepage**.
2. On the **Start new tuning goal** card, select **View setup**.
   The AI Search Tuning wizard appears.
3. On the wizard's **1. Set up new evaluation** page, specify basic parameters for your evaluation.
   a. In the **Evaluation details** section, enter a name and description for your evaluation.
   b. Enter a reviewer deadline date or select one using the calendar picker tool ( ).
   c. In the **Which search applications do you want to compare?** section, select the appropriate search engine for your baseline search application.
   d. Select the baseline Zing or AI Search application for your relevancy comparison.
   e. Select the AI Search application that you want to compare with your selected baseline search application.

   > **Note:**
   > You can only select an AI Search application that uses the same search sources as your baseline search application. If the selection list is empty, no evaluation is possible for the selected base system application.

   f. **Optional:** To review a selected search application's configuration settings, select its **View search application details** link.
   g. Advance to the next page by selecting **Continue to 2. Choose search terms**.
4. On the wizard's **2. Choose search terms** page, specify the source for search terms you want users to rate and review in the two selected search applications.
   - To retrieve search terms in a specific language from the Search Event [sys_search_event] table, choose **Select from logs** and choose your desired language.

   > **Note:**
   > The system filters out search terms that only include record numbers or that are 3 characters or less in length.

   - If you have the admin role and want to retrieve search terms from another table, choose **Advanced Configuration**, specify the table and column containing search terms, and optionally select **Filter search terms with condition builder** to define filter conditions that terms must satisfy.

   a. In the **Maximum number of search terms** field, enter a integer limit.
   b. In the **Maximum suggested search results per application** field, enter an integer limit.
   c. Advance to the next page by selecting **Continue to 3. Refine search terms**.
5. On the wizard's **3. Refine search terms** page, modify the set of search terms.
   a. Select **Load search terms**.
   b. Review the loaded search terms.

c. **Optional:** Use the search field and Search icon ( ) to refine the set of search terms displayed, or select the Clear input icon ( ) to view the full list of loaded search terms.
d. **Optional:** To remove unwanted search terms, select them, then select **Delete** and confirm by selecting **Delete All**.
e. **Optional:** To add new search terms, select **Add new search terms**, then enter your new terms, selecting **Add** after each, and finally select the Close dialog icon ( ).
f. Advance to the next page by selecting **Continue to 4. Assign reviewers**.

6. On the wizard's **4. Assign reviewers** page, select reviewers for your evaluation.

   a. In the **Reviewers** field, select one or more search users to rate and review search results for the evaluation.

   Note:
   You can only select users with the sn_ais_data.tuning_reviewer or ais_admin roles as reviewers.

   b. Advance to the next page by selecting **Continue to Review and Submit**.

7. On the wizard's **Review and Submit** page, save your evaluation.

   a. Confirm that your evaluation settings are correct.

   Note:
   If a setting is incorrect, select the link for the appropriate wizard page, correct the setting, and then return to the **Review and Submit** page.

   b. Select **Save**, then select **Got it**.

What to do next

Notify your selected reviewers that they can begin rating and reviewing search results for the evaluation. For more details on this procedure, see Rate and review search results in an AI Search Tuning evaluation.

Note:
Don't modify the configuration of either of your evaluation's search applications until review is complete. Mid-review changes to search application configurations can cause unexpected behaviors in the reviewing UI.

Rate and review search results in an AI Search Tuning evaluation

Select, review, and rate results for search terms included in an AI Search Tuning evaluation.

Before you begin

The AI Search Tuning ServiceNow® Store application must be installed on your instance. For details on installing this application, see Install AI Search Tuning.

Role required: sn_ais_data.tuning_reviewer

About this task

Reviewers rate the relevancy of search results for AI Search Tuning evaluations. Search result relevance is rated separately for each search term defined in the evaluation. The accumulated search result ratings determine the overall search relevancy quality for the search applications being evaluated.

Each reviewer selects search terms to review from the list of terms defined in the evaluation. Once a reviewer selects a search term, that search term becomes unavailable for selection by other reviewers for the same evaluation.

When reviewing results for a search term, the review page displays up to 10 de-duplicated Genius Results and search results from either or both of the two search applications being evaluated. Results appear in an arbitrary order, and aren't sorted by relevancy score. Reviewers specify a numeric rating of 2, 1, or 0 for each result displayed, indicating its relevance to the selected search term.

Procedure

1. Navigate to **All** > **AI Search Tuning** > **Reviewer Homepage**.
2. In the card for the AI Search Tuning evaluation you want to work on, select **Get started** or **Resume reviewing**.
3. In the **Select search terms** dialog, select search terms to rate and review.
   a. **Optional:** Use the search field and Search icon ( ) to refine the set of search terms displayed, or select the Clear input icon ( ) to view the full list of search terms.
   b. Select all search terms that you want to rate and review.
   c. Select **Start reviewing**.
4. On the **Review search results** page, rate and review search terms.
   a. Select a search term to rate and review from the Search terms list.

   Note:
   The Search terms list shows search terms in alphabetical order and highlights the selected search term. A **Reviewed** count shows the number of search results rated for the selected search term and the total number of search results it returned.

   b. For each search result displayed, select a rating value:

Rating	Description
2 - Exact	The result is an exact match for the search. It should appear with high relevancy in results for the search term.
1 - Useful	The result isn't directly relevant for the search, but might be useful to some search users. It should appear with low relevancy in results for the search term.
0 - Irrelevant	The result isn't relevant for the search. It shouldn't appear in results for the search term.

   c. **Optional:** If search results for the selected search term are missing, add them by selecting **Add results**, selecting those results, then selecting **Add**.

   Note:
   If you rate all results for a search term as irrelevant, a message appears prompting you to add missing search results.

d. **Optional:** To return to the search term selection page, select **Go back to list** in the Search terms list.

Note:
If no more search terms are available for selection, selecting **Go back to list** returns you to your list of evaluations instead.

e. When your review is complete, select **Got it**.

Search metrics and analytics for AI Search Tuning evaluations

Search metrics and analytics provide an overview of search relevance quality for evaluated search applications. An evaluation's metrics and analytics are derived from search result ratings assigned by evaluation reviewers.

Search administrators can access an evaluation's search quality graph and overall progress meter by navigating to **All** > **AI Search Tuning** > **Admin Homepage**. To view the evaluation's full analytics page with all metrics, select **View full analytics**.

Note:
AI Search Tuning only populates the search quality graph and analytics page metrics for evaluations with 20 or more reviewed search terms. The overall progress meter is always populated.

Search quality graph

This metric appears on the AI Search Tuning administrator homepage and on the full analytics page. The **Score comparison by search terms** link only appears in the card on the full analytics page.

Search quality card

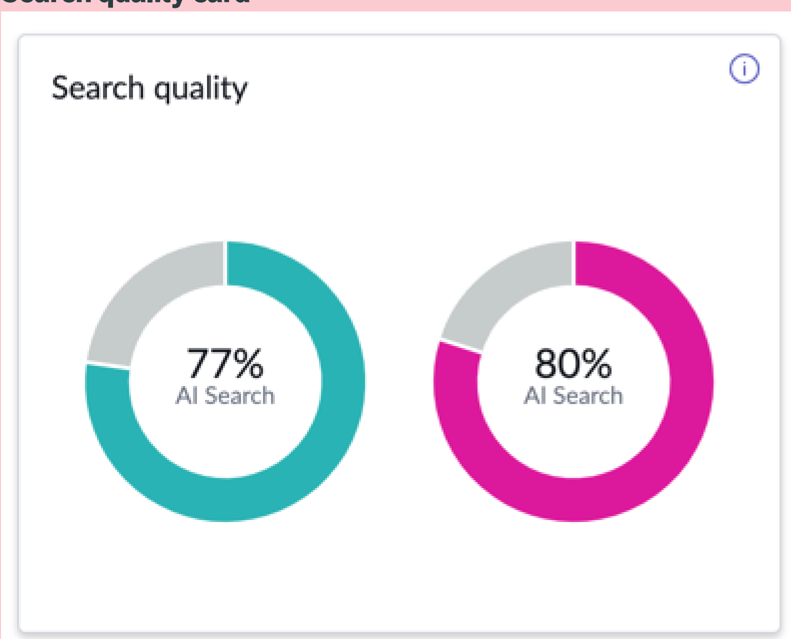

The Search quality card shows overall search relevance quality for the baseline and comparator search applications specified in an evaluation. An application's search quality score indicates how well its top five search results for reviewed search terms align with evaluation reviewers' search result ratings for those terms. Higher quality scores indicate more relevant search results and better overall search relevancy.

- In San Diego and previous releases, search quality scores are displayed as numbers between 0 and 100 inclusive.
- Starting in the Tokyo release, search quality scores are displayed as percentages.

To view score comparisons for individual search terms from the evaluation, select **Score comparison by search terms**. This link only appears in the search quality card on the full analytics page. Select **Back** to return to the full analytics page.

Overall progress graph

This metric appears on the AI Search Tuning administrator homepage and on the full analytics page.

Overall progress card

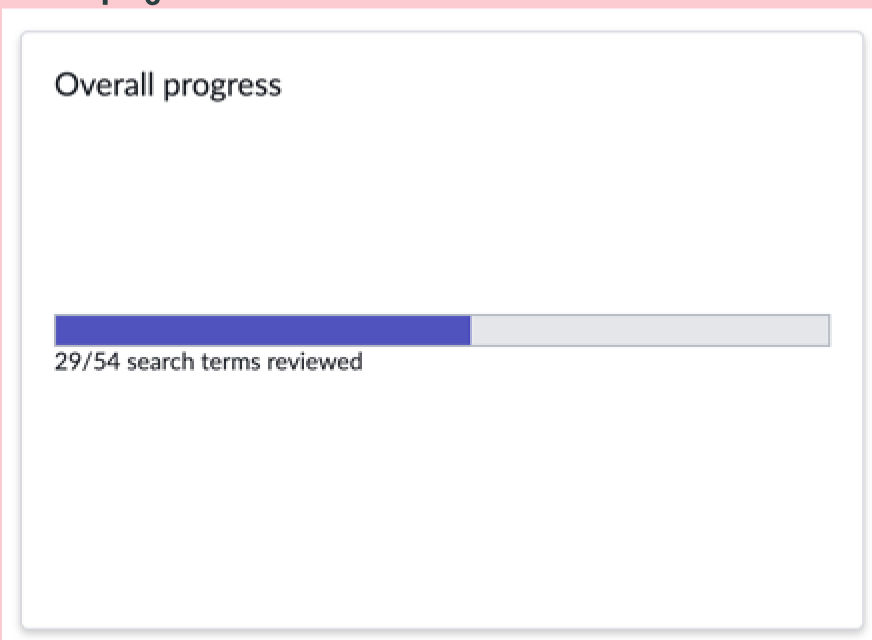

The Overall progress graph shows the count of search terms reviewed and the total number of search terms available for review in an evaluation.

Exact result in top 5 graph

This metric only appears on the full analytics page.

Exact result in top 5 card

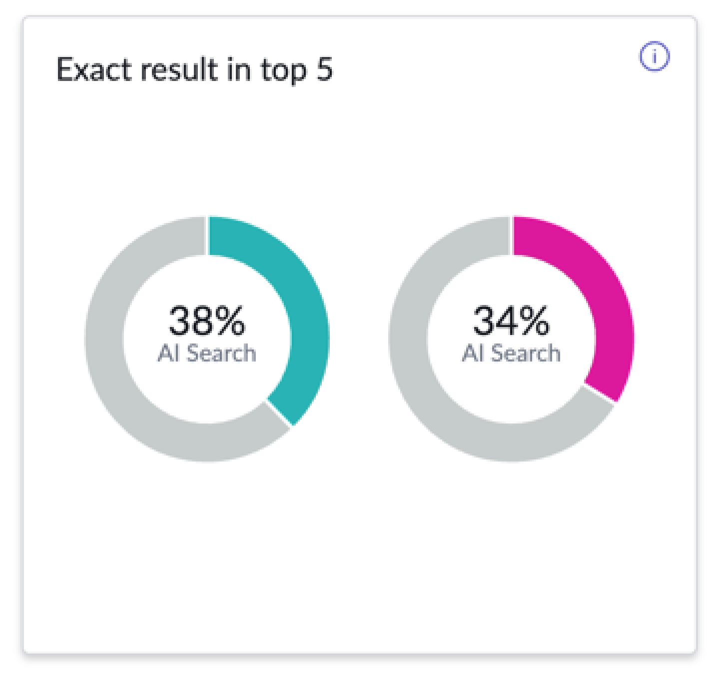

The Exact result in top 5 card shows the percentage of reviewed search terms for which a reviewer rated at least one of the top five search results as an exact match. Higher scores indicate more relevant results appearing first for evaluated searches.

AI Search Catalog Item Genius Results

This metric only appears on the full analytics page.

AI Search Catalog Item Genius Results card

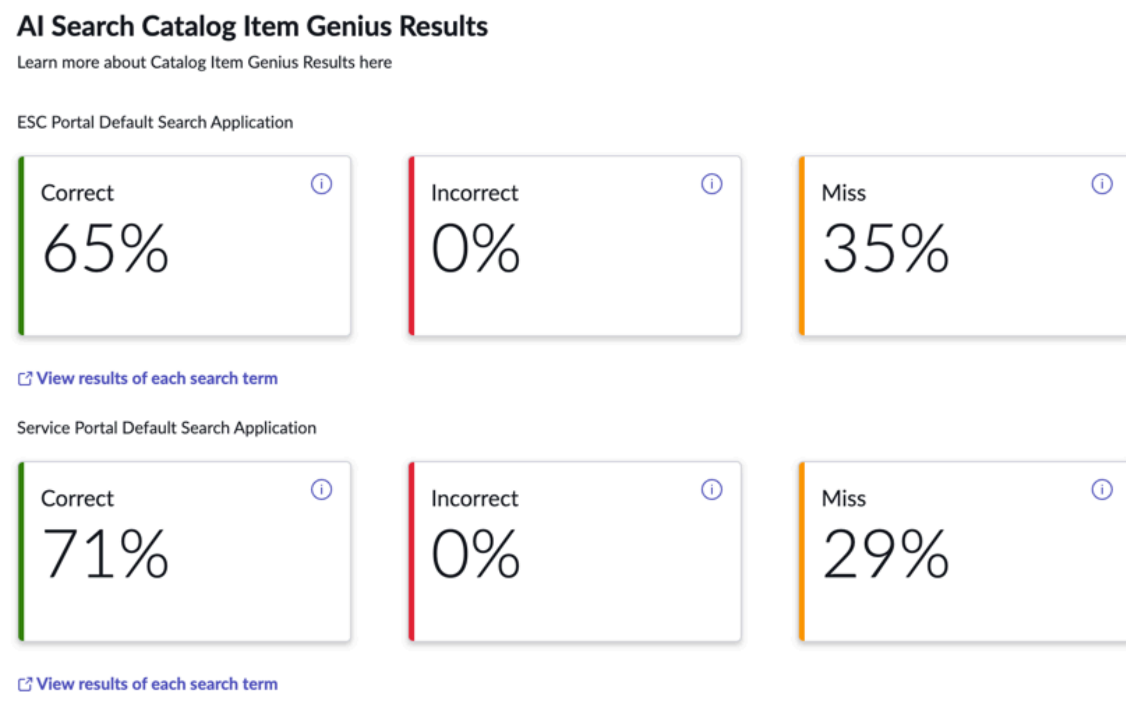

The AI Search Catalog Item Genius Results card shows correct, incorrect, and miss percentages for Catalog Item Genius Results in the two evaluated search applications. These category percentages are defined as follows.

Category	Description
Correct	The sum of the following percentages: • The percentage of reviewed search terms that returned a Catalog Item Genius Result rated as an exact or useful match. • The percentage of reviewed search terms that returned no Catalog Item Genius Result and that had no Catalog Item search results rated as exact or useful matches.
Incorrect	The percentage of reviewed search terms that returned a Catalog Item Genius Result rated as irrelevant.
Miss	The percentage of reviewed search terms that returned no Catalog Item Genius Result and that had no Catalog Item search results rated as exact or useful matches.

To view a search application's Catalog Item Genius Results details for individual search terms, select **View results of each search term** for the desired application. Select **Back** to return to the full analytics page.

AI Search Tuning reference

Reference information for the AI Search Tuning application.

Components installed with AI Search Tuning

The AI Search Tuning application installs new system components including a user role and multiple tables.

Roles installed with AI Search Tuning

Role	Description	Contains roles
AI Search Tuning reviewer [sn_ais_data.tuning_reviewer]	Rates and reviews search terms in tuning evaluations.	None

Tables installed with AI Search Tuning

Table	Description
AI Search Data Labeling Assignment [ais_labeling_assignment]	List of mappings between reviewers and search terms from tuning evaluations.
AI Search Data Labeling Label [ais_labeling_label]	List of search result ratings for search terms specified in tuning evaluations.
AI Search Data Labeling Project [ais_labeling_project]	List of tuning evaluations.
AI Search Data Labeling Term [ais_labeling_term]	List of mappings between tuning evaluations and search terms.

[store-future: BEGIN review]

Genius Results for AI Search

This ServiceNow Store application provides NLU models and Genius Result configurations that support French, German, Japanese, and Spanish search queries in AI Search.

Required roles

The following roles are associated with this application:

- ServiceNow AI Platform® administrator (admin): Install and activate this application.
- AI Search administrator (ais_admin): Review the publication status of the application NLU models. Link the application Genius Result configurations to search profiles. Clone Genius Result configurations for modification. Modify search application Entity View Action Mapping (EVAM) definitions to work with new Genius Result configurations.
- Entity View Action Mapping (EVAM) administrator (evam_admin): Modify search application EVAM definitions to work with new Genius Result configurations.

[End]

Install Genius Results for AI Search

You can install the Genius Results for AI Search application (sn_genius_results) 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.

Role required: admin

About this task

Genius Results for AI Search includes NLU models and Genius Result configurations for AI Search queries in French, German, Japanese, and Spanish.

For more information on these installed items, see Genius Result NLU models for French, German, Japanese, and Spanish and Default Genius Result configurations for French, German, Japanese, and Spanish.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Genius Results for AI Search application (sn_genius_results) 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.

   Visit the ServiceNow Store website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.
3. Select **Install**.

What to do next

The Genius Results for AI Search application includes EVAM configuration for Catalog Item Genius Result answer cards in Service Portal. For details on configuring People and Catalog Item answer cards in other search applications, follow the steps in EVAM settings for French, German, Japanese, and Spanish Genius Results in search applications.

EVAM settings for French, German, Japanese, and Spanish Genius Results in search applications

To use French, German, Japanese, and Spanish People and Catalog Item Genius Results in your search applications, you must configure Entity View Action Mapping (EVAM) settings for those applications.

Search application	EVAM configuration steps
Employee Service Center	The default Employee Service Center EVAM settings only support English-language People and Catalog Item Genius Result configurations. For instructions on updating Employee Service Center EVAM settings to accommodate the new French, German, Japanese, and Spanish People and Catalog Item Genius Result configurations, see .
Now Mobile	The default Now Mobile EVAM settings only support English-language People and Catalog Item Genius Result configurations.
Service Portal	The Genius Results for AI Search application includes Service Portal EVAM configurations for the new French, German, Japanese, and Spanish Catalog Item Genius Result configurations. You don't need to configure EVAM settings to use these Genius Result configurations in Service Portal. **Note:** Service Portal doesn't use People Genius Results.
Virtual Agent	The default Virtual Agent EVAM settings only support English-language People and Catalog Item Genius Result configurations. For instructions on updating Virtual Agent EVAM settings to accommodate the new French, German, Japanese, and Spanish People and Catalog Item Genius Result configurations, see .

Configure EVAM settings for French, German, Japanese, and Spanish Genius Results in your custom portal

Configure Entity View Action Mapping (EVAM) settings to enable French, German, Japanese, and Spanish People and Catalog Item Genius Results in your custom portal.

Before you begin

You must have already created an EVAM definition for your custom portal, and that this definition includes a view configuration bundle with view configurations for the default Catalog Item - EN and People - EN Genius Result configurations.

For details on creating a custom portal in Service Portal, see Create a portal.

Role required: evam_admin or admin

About this task

You can clone and modify the EVAM view configuration files for your custom portal to support Catalog Item and People Genius Results for French, German, Japanese, or Spanish. To create a new view configuration for one of these Genius Result configurations, duplicate your existing English-language Catalog Item or People view configuration record using the **Insert and Stay** method, then modify it with settings for the chosen language.

For more information on the **Insert and Stay** method, see Insert a record.

Procedure

1. Navigate to **All** > *Entity View Action Mapping* **(***EVAM***)** > *EVAM* **Definitions** > **View Configurations**.
2. Open the existing English-language Catalog Item or People EVAM view configuration record for your custom portal.
3. Make a note of the actions assigned to this view configuration in the EVAM View Config Action Assignment M2Ms related list.
4. To create a new duplicate record, select the additional actions icon (≡), then select **Insert and Stay**.
5. In the **Name** field, enter a name that describes the Genius Result configuration for this new record.
   For example, if you're creating a view configuration record for the Catalog Item - FR Genius Result configuration, you could enter `Catalog Item - FR Genius Results`.
6. Change the argument for the **Condition** Type filter to specify the Genius Result configuration to use.
   For example, if you're configuring Genius Results in Japanese:
   - For a Catalog Item record, specify: **[Type] [is] [Catalog Item - JA]**
   - For a People record, specify: **[Type] [is] [People - JA]**
7. In the **Table Fields** field, select the field items.
   a. In the **Selected** section, remove the existing fields.
   b. In the **Available** section, select Columns(+), select **+**, and choose the appropriate Genius Result configuration, such as Catalog Item - ES.
   c. In the **Available** section, move the appropriate fields for the record to **Selected**.

For example, for Catalog Item configurations, select: table, category, name, sys_id, picture, price, and short_description.

d. In the **Available** section, move the Type field to **Selected**.

8. To save the view configuration record, select the additional actions icon ( ), then select **Save**.
9. In the EVAM View Config Action Assignment M2Ms related list, select **Link Existing** and link each of the actions you previously noted to the view configuration record.
For example, in the Action Assignments form, you could select `request_item` as the declarative action for a Catalog Item view configuration, or `viewProfile` and `viewOrgChart` as the declarative actions for a People view configuration.

Genius Result NLU models for French, German, Japanese, and Spanish

The Genius Results for AI Search application includes NLU models that enable Genius Result configurations for French, German, Japanese, and Spanish search queries.

To view these NLU models, navigate to **All** > **NLU Workbench** > **Models**.

NLU models included in the Genius Results for AI Search application

NLU model	Description
Default AI Search Genius Results Model - DE	Provides query intents for search queries in German.
Default AI Search Genius Results Model - ES	Provides query intents for search queries in Spanish.
Default AI Search Genius Results Model - FR	Provides query intents for search queries in French.
Default AI Search Genius Results Model - JA	Provides query intents for search queries in Japanese.

Training and publishing the NLU models for default Genius Result configurations

The default Genius Result configurations with *NLU* triggers use intents from the NLU model for their language. These models are automatically trained and published after you install the Genius Results for AI Search application.

Depending on the amount of data in your instance, the system may take some time to train and publish these models. You can check on a model's current status by navigating to **All** > *AI Search* > **Search Query Settings** > **Search Intents**. When the model is published, it's ready for use.

Default Genius Result configurations for French, German, Japanese, and Spanish

The Genius Results for AI Search application includes pre-configured Genius Result configurations for People and Catalog Item search queries in French, German, Japanese, and Spanish.

To access these Genius Result configurations, navigate to **All** > *AI Search* > **Search Query Settings** > **Genius Results**.

You can link these default Genius Result configurations to your own search profiles. You can't modify the default Genius Result configurations, but you can clone them and configure the cloned configurations.

Catalog Item - DE

The Catalog Item - DE Genius Result configuration displays top search results from the Catalog Item [sc_cat_item] table for search queries in German. Each answer card shows a single CI record.

Sample Catalog Item - DE answer card

Select the answer card's **Request** action link to create a new request for the specified CI.

Catalog Item - ES

The Catalog Item - ES Genius Result configuration displays top search results from the Catalog Item [sc_cat_item] table for search queries in Spanish. Each answer card shows a single CI record.

Sample Catalog Item - ES answer card

Select the answer card's **Solicitud** action link to create a new request for the specified CI.

Catalog Item - FR

The Catalog Item - FR Genius Result configuration displays top search results from the Catalog Item [sc_cat_item] table for search queries in French. Each answer card shows a single CI record.

Sample Catalog Item - FR answer card

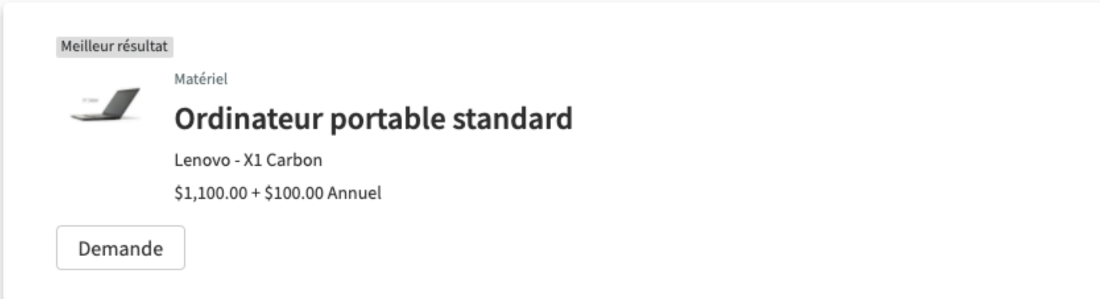

Select the answer card's **Demande** action link to create a new request for the specified CI.

Catalog Item - JA

The Catalog Item - JA Genius Result configuration displays top search results from the Catalog Item [sc_cat_item] table for search queries in Japanese. Each answer card shows a single CI record.

Sample Catalog Item - JA answer card

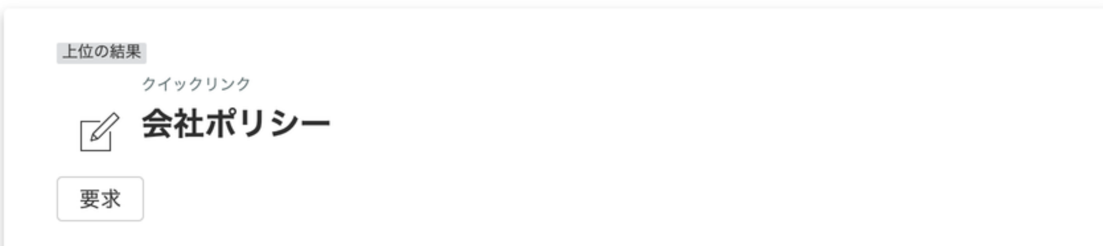

Select the answer card's 要求 action link to create a new request for the specified CI.

People - DE

The People - DE Genius Result configuration displays top search results from the User [sys_user] table for search queries in German. Each answer card shows a single user record.

Sample People - DE answer card

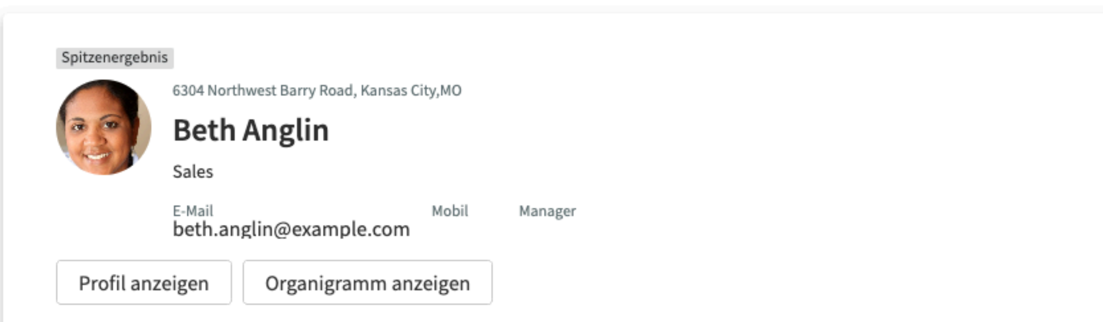

Select the answer card's **Profil anzeigen** action link to see the user's profile, or select the **Organigramm anzeigen** action link to view the user's organization chart.

People - ES

The People - ES Genius Result configuration displays top search results from the User [sys_user] table for search queries in Spanish. Each answer card shows a single user record.

Sample People - ES answer card

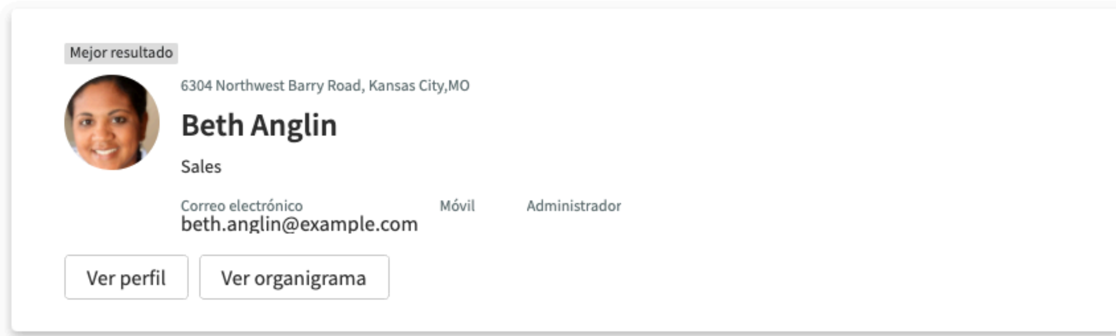

Select the answer card's **Ver perfil** action link to see the user's profile, or select the **Ver organigrama** action link to view the user's organization chart.

People - FR

The People - FR Genius Result configuration displays top search results from the User [sys_user] table for search queries in French. Each answer card shows a single user record.

Sample People - FR answer card

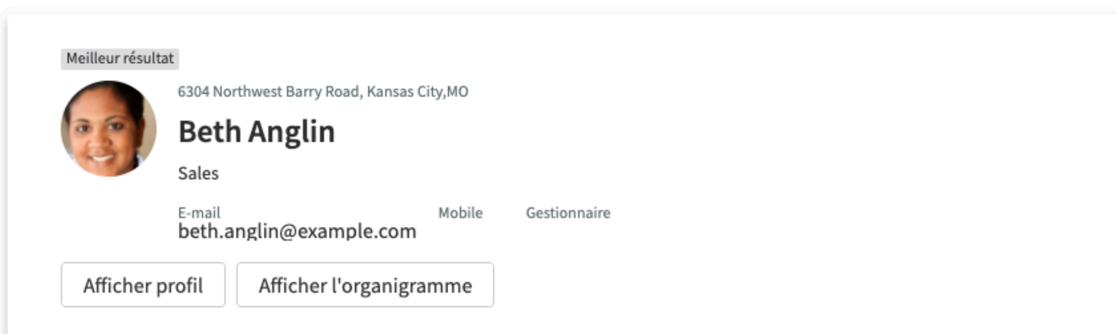

Select the answer card's **Afficher profil** action link to see the user's profile, or select the **Afficher l'organigramme** action link to view the user's organization chart.

People - JA

The People - JA Genius Result configuration displays top search results from the User [sys_user] table for search queries in Japanese. Each answer card shows a single user record.

Sample People - JA answer card

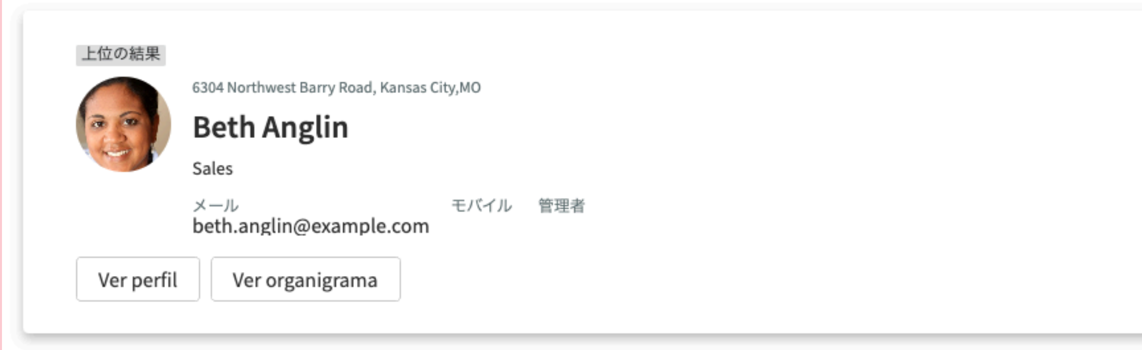

Select the answer card's プロフィールの表示 action link to see the user's profile, or select the 組織図の表示 action link to view the user's organization chart.

[store-future: BEGIN review]

AI Search Admin Console

The AI Search Admin Console application provides AI Search administrators with a no-code management tool. Use the AI Search Admin Console to configure, test, and enable AI Search solutions across different types of search applications.

Required roles

The following roles are associated with this application:

- ServiceNow AI Platform® administrator (admin):
  - Install and activate AI Search Admin Console.
  - Install additional applications to add AI Search features.
- AI Search administrator (ais_admin):
  - Manage and configure AI Search settings for search applications.
  - Review AI Search activity.
  - View AI Search documentation.
  - Access AI Search forum on ServiceNow Community.
  - See answers to common questions about AI Search.

[End]

Install AI Search Admin Console

The AI Search Admin Console application (com.sn_ai_search_admin) is available by invitation from the ServiceNow® Store. You can install the application if you have the admin role.

Before you begin

The following ServiceNow® Store applications must be installed on your instance:

- Admin Experience Framework (sn_ace) version 1.4.1 or later
- Advanced AI Search Management Tools

Note:
The Admin Experience Framework application is available by invitation only.

Role required: admin

About this task

Installing the AI Search Admin Console makes the application available to users with the ais_admin or admin roles. For information on the tasks that users can perform in the application, see Using AI Search Admin Console.

Procedure

1. In a web browser, navigate to the AI Search Admin Console page in the ServiceNow® Store by following the link in your invitation.
2. Select **View Requests**.
3. In the ServiceNow Request for App Installation form, enter your instance name and the reason for your request, then select **Request**.

Result

Your installation request is submitted. You will receive an email when this request is approved or declined. If it's approved, the application is installed on your instance.

What to do next

Access the AI Search Admin Console application by navigating to **All** > `AI Search` **Admin** > `AI Search` **Admin Home**.

Using AI Search Admin Console

Search administrators can use the AI Search Admin Console to set up and manage AI Search in search applications, track AI Search activity, add features to AI Search, and access AI Search resources.

To access AI Search Admin Console, navigate to **All** > `AI Search` **Admin** > `AI Search` **Admin Home**.

Set up and manage AI Search in search applications

Manage and configure AI Search settings for search applications that support the AI Search engine.

Before you begin

Role required: ais_admin

Procedure

Navigate to **All** > `AI Search` **Admin** > `AI Search` **Admin Home**.

Track AI Search activity

View key analytics for AI Search usage, or follow a link to the full AI Search dashboard.

Before you begin

Role required: ais_admin

About this task

The **Track** `AI Search` **activity pane** pane displays reports summarizing your AI Search configurations and trends.

Procedure

1. Navigate to **All** > `AI Search` **Admin** > `AI Search` **Admin Home**.
2. In the **Home** page's **Track** `AI Search` **activity** pane, review the dashboard reports.

3. **Optional:** To view the full AI Search dashboard, select the **Go to full dashboard** link.
For details on this dashboard's reports, see AI Search dashboard.

AI Search activity reports

Reports summarize your AI Search configurations and trends.

Title	Type	Source table	Description
Applications that need configuration	Single Score ( )	ais_publish_history, sp_portal, sys_cs_context_profile_search, sys_properties, sys_search_context_config, and sys_sg_global_search (all accessed via AISGetAppStatus script include)	Shows the number of search applications that need configuration before AI Search can be enabled.
Applications ready to turn on	Single Score ( )	ais_publish_history, sp_portal, sys_cs_context_profile_search, sys_properties, sys_search_context_config, and sys_sg_global_search (all accessed via AISGetAppStatus script include)	Shows the number of search applications that have AI Search configurations and are ready to have AI Search enabled.
Applications with AI Search on	Single Score ( )	ais_publish_history, sp_portal, sys_cs_context_profile_search, sys_properties, sys_search_context_config, and sys_sg_global_search (all accessed via AISGetAppStatus script include)	Shows the number of search applications that have AI Search enabled.
Indexed documents by month	Bar ( )	sn_ais_admin_tools_ai_search_dashboard_documents_by_mo	Shows the number of searchable records newly indexed by AI Search, grouped by month.
Documents by indexed sources	Donut ( )	sn_ais_admin_tools_ai_search_dashboard_total_indexed_docu	Shows the number of searchable records indexed by AI Search, grouped by indexed source.
Queries by search profile	Donut ( )	sys_search_event	Shows the number of search queries, grouped by search profile used.

Add features to AI Search

Expand AI Search functionality by installing the AI Search Integration Hub spoke and the AI Search Data Labeling and AI Search for Next Experience applications.

Before you begin
Role required: admin

Procedure

Navigate to **All** > `AI Search` **Admin** > `AI Search` **Admin Home**.

Access AI Search resources

Review the AI Search introductory video, follow links to the AI Search documentation and ServiceNow Community forum, or view answers to commonly asked questions about AI Search.

Before you begin
Role required: ais_admin

Procedure

Navigate to **All** > `AI Search` **Admin** > `AI Search` **Admin Home**.

AI Search Admin Console reference

Reference information for the AI Search Admin Console application.

Components installed with AI Search Admin Console

The AI Search Admin Console application installs new system components.

List of all components installed with AI Search Admin Console

To view the complete list of system components installed with AI Search Admin Console, follow the steps described in Find components installed with an application. The application's package name is `AI Search` **Admin Console**.

Contextual search

Contextual search is a search technology that focuses on the context of the query as well as the intent of the user in order to fetch the most relevant set of results. Contextual search displays related results within a form or record producer based on the text you enter in a field.

Example use cases for Incident:

- An agent, working on an Incident record, can use contextual search to find and attach knowledge articles, order a catalog item on behalf of the caller, or search for an article and link it to a related incident or problem.
- A user, creating an Incident via the Service Portal, can see related catalog items that the user can order. The user can also view related knowledge articles to help them resolve the issue on their own without involving the service desk operators.

Contextual search helps with incident deflection for the following areas:

- Incident record producers: deflecting incidents by helping end users resolve issues before they raise an incident.
- Incident forms: helping service desk staff resolve incidents quickly by providing relevant knowledge.
- Incident email notifications: helping end users resolve their incidents themselves without requiring manual intervention from service desk staff.

Contextual search plugin

Contextual search plugin is activated in the base system.

Contextual search plugin

The Contextual Search plugin (com.snc.contextual_search) is activated by default in the base system. This plugin activates related plugins if they aren't already active.

Before you begin

Role required: admin

About this task

If the related plugins aren't activated, Contextual Search activates the plugins.

Plugins for Contextual Search

Plugin	Description
JSON Service request/response model [com.snc.application.json_service]	Core components and helpers for a JSON request/ response model. Includes JSON and XML transports for NG and GlideAjax support. An extension of the processor framework.
Contextual Search — Internal [com.snc.contextual_search.internal]	An internal plugin component for Contextual Search.
GraphQL Schema for Contextual Search [com.snc.contextual_search.graphql]	Provides GraphQL schema for contextual search services.

Procedure

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

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

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

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

Related topics

List of plugins (Washington DC)

Components installed with Contextual Search

Several types of components are installed with the activation of the Contextual Search plugin (com.snc.contextual_search), including tables and user roles.

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
Available Search Resource [cxs_search_resource]	List of available search resources used in configuration of searches.
Email Configuration [cxs_table_email_config]	List of email notifications configured to have contextual search results attached.
Record Producer Configuration [cxs_rp_config]	List of record producers configured to have contextual search.
Relevant Document [cxs_relevant_doc]	Groups the Relevant Document Detail individual recorded actions by search session. **Note:** A new search session group is created every time you create or open a record or reload the page. For more information on the table, refer Feedback information on contextual search.
Relevant Document Detail [cxs_rel_doc_detail]	List of individual recorded actions. For example, **Preview**, **This helped**, **Attach**, or **Order**. For more information on the table, refer Feedback information on contextual search.
Search Context	Search contexts for contextual search.

Table	Description
[cxs_context_config]	
Search Field [cxs_table_field_config]	All form fields that have contextual search configured to it.
Search Resource [cxs_search_res_config]	Search resource configurations for the related searcher.
Search Resource Context Configuration [cxs_res_context_config]	Conditions for this search context that restrict the results returned to a user.
Search Resource Context Property [cxs_res_context_config_prop]	Property values for search resource context configuration.
Search Resource Property [cxs_search_res_config_prop]	Property values for search resource searcher configuration.
Searcher Configuration [cxs_searcher_config]	Details of available search configurations.
Table Configuration [cxs_table_config]	Lists the tables configured for table configurations.
Wizard Configuration [cxs_wizard_config]	Lists the tables configured for wizard configurations.
Base UI Configuration [cxs_ui_config_base]	Common elements for the record producer, table, and wizard configurations.

Contextual search concepts

Contextual search concepts help you to understand different contextual search components that are used in the documentation, configuration, and user interface. Understanding of the concepts makes it easier for you to comprehend the interconnection between the components.

Contextual search uses terms and concepts that can differ from industry norms due to the unique nature of the ServiceNow platform. Contextual search includes the following concepts and components:

Search resource

Resource that can be searched such as Catalog or Knowledge Articles. As the resources are scripted and contain JavaScript, they can combine results from more than one table.

Note:

- ServiceNow provides pre-defined scripted search resources. You can't create, update, or delete search resources.
- A search resource uses the Zing search engine to search a table.

Searcher

A single search resource or a group of search resources that can be searched. For example, the Knowledge and catalog searcher contains Knowledge Articles, Questions and Catalog Items as its search resources.

Note:
ServiceNow provides pre-defined searchers. You can't create, update, or delete searchers.

Additional resources

An additional resource is a single table source with query conditions that you can search. Additional resources aren't scripted as you search from a single table. For example, Resolved Incidents (within the last six months).

Note:

- ServiceNow provides pre-defined additional resources. You can't create or delete additional resources, but can update additional resources to better match your organization requirements. For example, you can update the condition and description for Resolved Incidents (Last 6 months) to change it to Resolved Incidents (Last 3 months).
- An additional resource performs the search based on the type of the source:
  - Table: The Zing search engine to used to search a table.
  - Predictive Intelligence similarity: Search is based on similarity and not on the Zing search engine.

Search result display fields

Defines the fields whose values are displayed with the search result of an additional resource. The fields can include **Assigned to**, **Opened by**, and **State**. The search result display fields appear at the bottom of a search result record.

Note:
ServiceNow provides pre-defined search result display fields. You can't create or delete additional resources, but you can update search result display fields to better match your organization requirements. For example, change which field to use as a title or the description or which fields to show at the bottom of a search result record.

Search context

Brings together a searcher and additional resources (if any) that can be searched.

Note:
ServiceNow provides pre-defined search contexts. You can also create search contexts.

Searcher text

Defined in a search context, a searcher text specifies a user-friendly group name for the searcher. When a searcher contains more than one search resource, a

group name is used to represent that group of search resources. For example, the Incident Deflection search context uses a searcher that searches Knowledge Articles, Questions and Catalog Items. The user-friendly searcher group name that represents all of them is Knowledge & Catalog.

Search on tab

Defined in the search context, **Search on tab** specifies when a contextual search should start the search. The different scenarios when search is performed are:

- If you select the **Search on tab** check box, the search only runs when you update and leave the search field (recommended).
- If you clear the **Search on tab** check box, the search runs when you pause typing.

Search actions

Specifies the actions for each such results. For example, for catalog items, the search action is **Order** and for knowledge articles it's **Attach**. All sources have the **This helped** action.

Contextual search components diagram

The contextual search components diagram helps you to understand how contextual search concepts work together to perform a search and from which module you can configure the components.

Contextual search components diagram

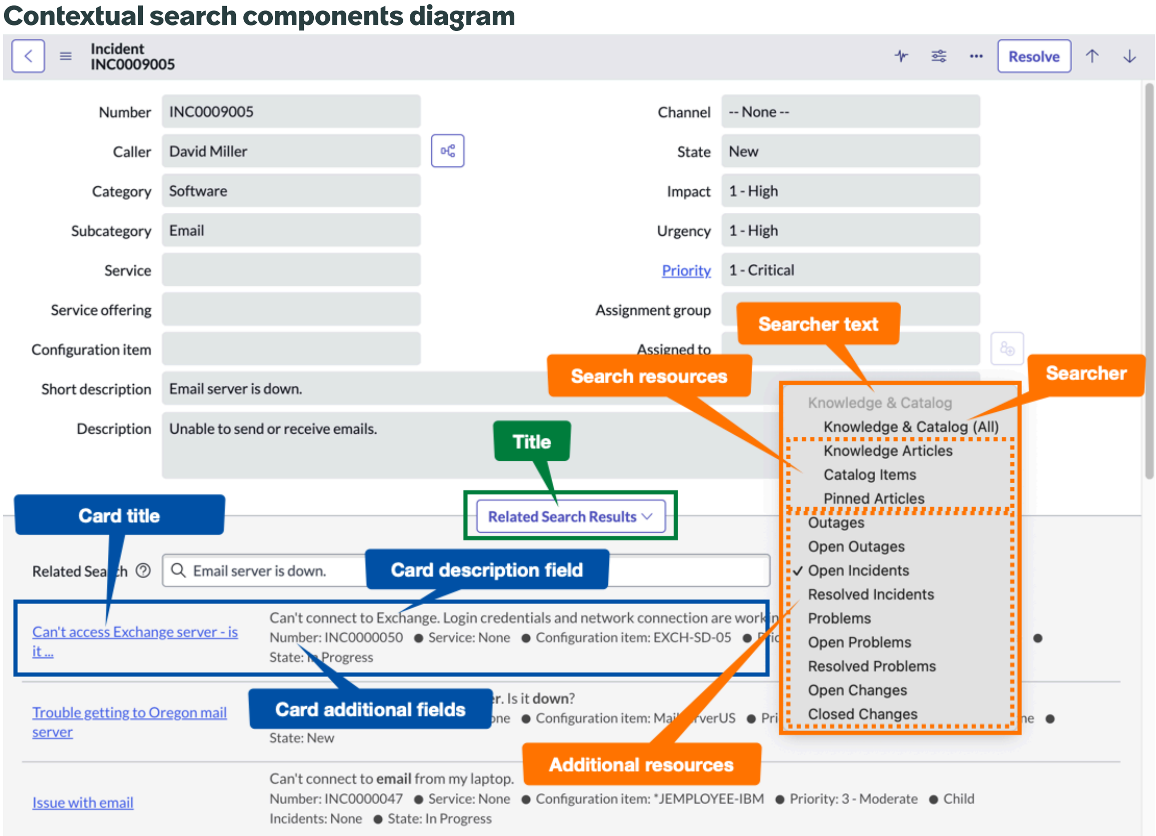

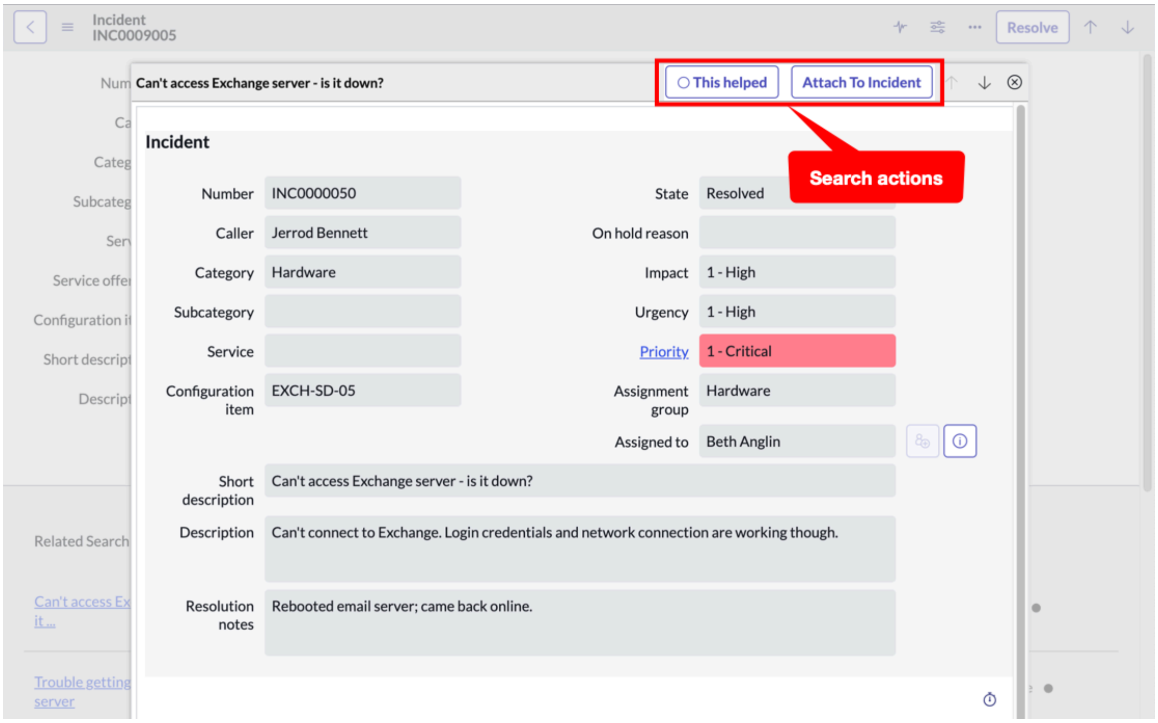

Components	Navigation
Table configuration	You can configure the search title and the form using the **All** > **Contextual search** > **Table configuration** module. For more information, refer Configure table for a contextual search.
Search Contexts	You can configure the searcher text and the searcher using the **All** > **Contextual search** > **Search Contexts** module. For more information, see Define a search context.
Search Result Display Configurations	You can configure the card title, card description field, and card additional fields using the **All** > **Contextual search** > **Search Result Display Configurations** module. For more information, see Edit search resource display field record.
Search Action Configurations	You can configure the search actions using the **All** > **Contextual search** > **Table Configuration** > **Search Action Configurations** module. For more information, see Modify or disable search actions available for contextual search.

Configuring contextual search

You can configure different components of contextual search.

Contextual search properties

Contextual search properties are used to control the fields that you can search from and the number of results displayed per page.

The contextual search properties are available at **Contextual Search** > **Properties**.

Properties for contextual search

Property label	Description
Default value for maximum number of search results returned for table and record producer configurations [com.snc.contextual_search.result.default.limit]	Default value for the maximum number of search results returned when you create a new table or record producer configuration. By default, the value is 10.
Number of characters that a configured field needs before triggering a search [com.snc.contextual_search.min_length]	Number of characters that contextual search requires before triggering a search based on the typed text in the search field. By default, the value is set to 3 characters. This property also applies to the Agent Assist search.
Number of milliseconds that a configured field will wait before triggering a search [com.snc.contextual_search.wait_time]	Number of milliseconds that contextual search waits before triggering a search based on the typed text in the search field. By default, the value is set to 1,000 milliseconds (one second).
Records returned above this threshold will log warnings [com.snc.contextual_search.result.threshold]	Triggers a warning message within the system log whenever more than this number of results are returned in a search. This logging helps if you are experiencing performance issues because of large searches. By default, the value is set to 10,000 results.
When a Form configuration is created a search field with a name matching this property will be automatically created [com.snc.contextual_search.widget.form.default_field]	The default field that is used for searching when you create a new table configuration. By default, the value is set to short_description.
Collapse the search results widget when creating a new record in a form [com.snc.contextual_search.widget.form.open_collapsed_new_records]	When the value is set to true and you create a new record, the search results widget is collapsed to save space. By default, the value is set to false.
Collapse the search results widget when opening an existing record in a form	When the value is set to true and you open an existing

Properties for contextual search (continued)

Property label	Description
[com.snc.contextual_search.widget.form.open_collapsed_existing_records]	record, the search results widget is collapsed to save space. By default, the value is set to true. **Note:** The property applies only to existing records and not to new records.
Auto expand the search results widget when search field updates in a form [com.snc.contextual_search.widget.form.expand_on_search_field_update]	When the value is set to true and you update the search field in a form, the search results widget automatically expands and runs the search. By default, the value is set to true.
Show meta data on form contextual search results [com.snc.contextual_search.widget.form.show_meta_data]	When the value is set to true, each search result includes additional metadata below each result record. By default, the property is set to true. • For Knowledge articles, you configure the metadata fields to display by modifying Knowledge search properties. As an example, you can display the article's author, article number, and knowledge base name. For details on Knowledge search properties, see Knowledge Management properties. • Configure the search result display fields for other resources by editing search result display configurations. For details on search result display configuration settings, see Edit search resource display field record.

Define a search context

Define where agents can search for a particular resource. You can specify that the search resources are predefined or additional resources outside the system.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Contextual Search** > **Search Contexts**.
2. Select **New**.
3. On the form, fill in the fields.

Search Context form

Field	Description
Name	Unique name of the search context.
Short description	Brief summary of the context.
Searcher	Lookup icon (🔍) to select a searcher for this context. The searcher determines the information source to search. ◦ Catalog only: Provides search results from service catalog items. ◦ Knowledge and catalog: Provides search results from knowledge articles and service catalog items. ◦ Knowledge only: Provides search results from knowledge articles. ◦ Knowledge, Discussions and pinned: Provides search results from knowledge articles, pinned knowledge articles, and ServiceNow Community. ◦ Knowledge, pinned knowledge and catalog: Provides search results from knowledge articles, pinned knowledge articles, and service catalog items. ◦ Knowledge and pinned knowledge: Searches knowledge articles and pinned knowledge articles. ◦ Knowledge, Pinned, Catalog: Provides search results from Knowledge, Pinned Knowledge, and Catalog Items. ◦ Knowledge and pinned: Provides search results from knowledge articles and pinned knowledge articles. ◦ Pinned Knowledge only: Provides search results from pinned knowledge articles. **Note:** The **Searcher** field is optional. You can decide whether to show or hide sources from an agent.
Searcher default search	List used to select a specific resource from the **Searcher** field to search as default.
Searcher text	Unique label for the searcher group that is used when there is more than one resource in the searcher.
Active	Option to activate the search context form.
Enable wildcard searches	Option to enable wildcard characters, such as * and %, for searches.

Field	Description
Search on tab	Option to determine when to perform the search. If you select this option, the search runs only when you enter text in a field and then leave the field. If you clear the check box, the search runs after you stop typing even though the cursor is still in the same field. **Note:** To display the search context in the Source selector on Platform, you must select this option. Search contexts with this option deselected don't appear in the Source selector on Platform.

The **Searcher default search** list is available if the **Searcher** is the default search. If an additional resource is the default search, you see the **Set Searcher as the Default Search** related link. For information on how to set the default source for search context, see Set the default source for search context.

4. Select and hold (or right-click) the form header and select **Save**.

Configure table for a contextual search

Configure the options for the contextual search that a table performs.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration** and click **New**.
2. On the form, fill in the fields.

Table Configuration form

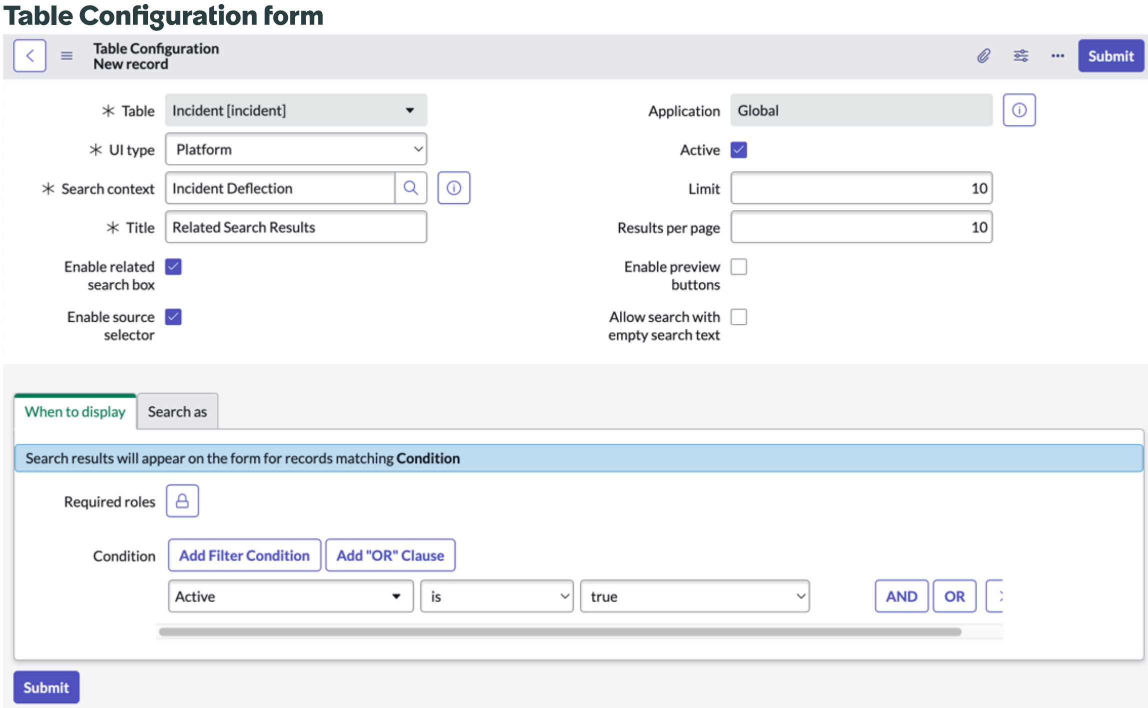

Table Configuration fields

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td>Table</td><td>Name and label of the table associated with the form. For example, Incident [incident].</td></tr>
<tr><td>UI type</td><td>User interface differs depending on whether you use the table configuration for Platform or Workspace.

Note:
The **When to display** section appears only when you select *`Platform`* from **UI type**.

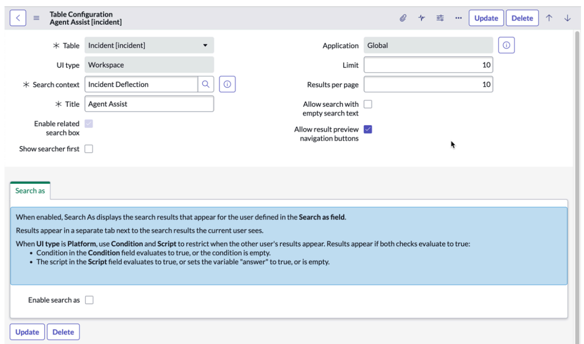
</td></tr>
<tr><td>Search context</td><td>Lookup icon to select a search context for which you are configuring the table.</td></tr>
<tr><td>Title</td><td>Unique title for the search results area that appears in the form. For example, if you enter `Knowledge results` as the value in this field, the related search area displays **Knowledge results** as the title.
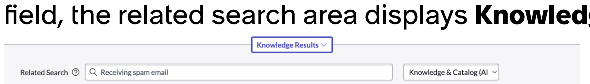
</td></tr>
<tr><td>Enable related search box</td><td>Option to show the related search box in the form.</td></tr>
<tr><td>Show searcher first</td><td>Option to display searcher resources before additional resources in the search context. This field appears only when you select *`Workspace`* from **UI type**.</td></tr>
<tr><td>Enable source selector</td><td>Option to show the list of search sources that you can search from.

Note:
This field appears only when you select **Enable related search box**.</td></tr>
<tr><td>Application</td><td>Application scope of the table configuration. The configuration is available for all applications or for scoped applications.</td></tr>
<tr><td>Active</td><td>Option to activate the table configuration.</td></tr>
<tr><td>Limit</td><td>Maximum number of search results returned.</td></tr>
<tr><td>Results per page</td><td>Number of results to display per page.</td></tr>
</table>

Field	Description
Allow expand/ collapse	Option to enable a user to collapse the search results in the form.
Enable preview buttons	Option to show a preview button (platform only) for each search result. The user can then use the preview button or the search result title to preview the result record. **Note:** The option is available from the London release. Clear the option to provide a more compact display for each search result.
Allow search with empty search text	Option to specify whether the search runs when there is no search text. **Note:** There is no search text when there is no **Search Field** specified or when there is a **Search Field** whose value is empty.
When to display	
Required roles	List of roles that grant access to the related search results widget from the table. Users must have one of these roles to see the widget when viewing the table.
Condition	Conditions that records must satisfy to appear as related search results. For example, set **[Active] [is] [true]** to only enable the active records as related search results.
Search as	
Enable search as	Option to enable searching as a different user. The message box describes the settings available when you select this 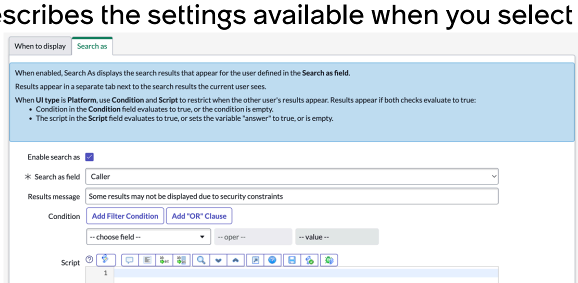 option. For more information on using **Search as**, see Enable viewing of search results for the current and another selected user.

3. Click **Submit**.
   The Search Fields, Filter Configurations, Email Configurations, and Search Action Configurations related lists appear. By default, the **Short description** field is added to the Search Fields related list.

   Note:
   Search results only appear for Search Fields that are visible on the form.

Filter configuration for contextual search

Define filter configurations to map fields from the form to provide a more targeted search result set.

Create a filter configuration using mapping

Create filter conditions using mapping that you can associate with a table configuration. The search result displays result that meets the filter conditions.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. Open a table configuration record on which to create the filter.
3. In the Filter Configurations related list, click **New**.
4. On the form, fill in the fields.

Filter Configuration form

Field	Description
Resource configuration	Lookup icon to select a resource such as Knowledge Articles.
Table configuration	[Read-only field]. Table for which you are applying the filter configuration.
Scripted filter	Option to define the filter using a script. **Note:** Clear this option since you want to create a filter configuration using mapping.
Active	Option to activate the filter configuration.

5. Click **Submit**.
   The Table Configurations form appears.
6. In the Filter Configuration related list, click the filter that you have created.
7. In the Filter Conditions related list, click **New**.

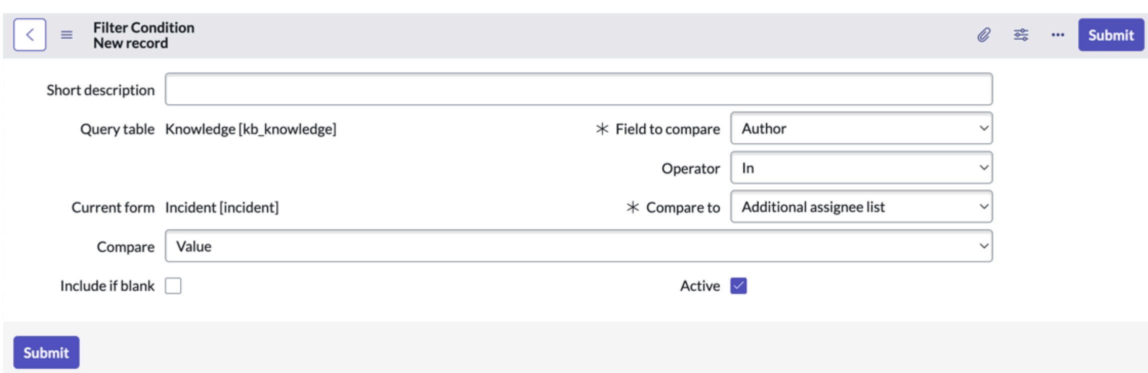

8. In the form, fill in the fields.

Field	Description
Short description	Brief summary of the filter condition.

Field	Description
Query table	[Read-only] Table from which the filter condition is queried.
Current form	[Read-only] Form for which you created the filter conditions.
Field to compare	Value of the field is compared with the value of the **Compare to** field.
Operator	Operation that should be performed on the **Field to compare** and **Compare to** field.
Compare to	Value of the field is compared against the value of **Field to compare**.
Compare	Comparison of the field can be based on the actual value stored on the database or the value that is displayed on the interface. The available values are: **Value** and **Display Value**.
Include if blank	Option to include a condition when the value of the **Compare to** field is blank.
Active	Option to activate the filter configuration.

9. Click **Submit**.

Create a filter configuration using scripts

Create filter conditions using script that you can associate with a table configuration. The search result displays results that meet the filter conditions.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. Open a table configuration record on which to create the filter.
3. In the Filter Configurations related list, click **New**.
4. On the form, fill in the fields.

Filter Configuration form

Field	Description
Resource configuration	Lookup icon 🔍 to select a resource such as Knowledge Articles.
Table configuration	[Read-only field]. Table for which you want to apply the filter configuration.
Scripted filter	Check box to select if you want to enter scripts in the **Script** field. **Note:** The **Script** field appears only when you select **Scripted filter**.
Application	Application scope for your filter configuration.

Field	Description
Active	Check box to activate the filter configuration.

5. Enter a script corresponding to the resource configuration.

Note:

Use one of these dynamic filter scripts to set up your dynamic filter.

6. Click **Submit**.

Configure search resource context properties

Configure the properties to limit the search to a specific catalog or knowledge base. For example, if you want to limit the knowledge search to an IT knowledge base or a service catalog search to software service catalog, you can configure the properties.

Before you begin

Role required: admin

About this task

The following procedure uses Incident Deflection as an example.

Procedure

1. Navigate to **All** > **Contextual Search** > **Search Contexts**.
2. Open the Incident Deflection record.
3. View the Resource Configurations related list.

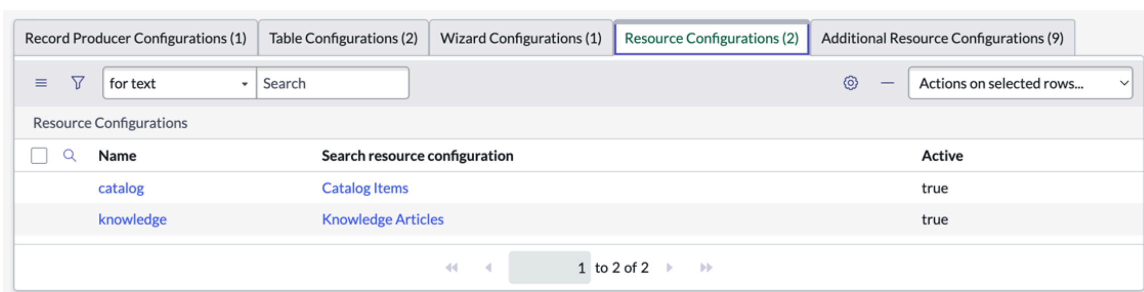

In the Resource Configurations related list, edit a resource configuration record to view the properties of the search resource.

The **knowledge** record displays the following knowledge resource configuration properties:

Knowledge resource configuration properties

Property	Description
Condition	Conditions to specify one or more knowledge bases to search.
Exclude pinned articles	Pinned article is a knowledge article. This property helps to decide whether to exclude pinned articles from knowledge base search or not.

The **catalog** record displays the following catalog resource configuration properties:

Catalog resource configuration properties

Property	Description
Catalog name	A catalog name to search the catalog items.
Search operator	Value of the operator is used in the Zing search engine to instruct how to combine the AND and OR conditions when parsing the search string. The integration of the search strings helps to score the displayed results. The value can only be one of the following values: ◦ **IR_AND_OR_QUERY** (the default): Display results with exact matches of all terms if the number of results is greater than the value of the `glide.ts.query.and_or_limit` property (default 0) or the (overriding) value of the table custom attribute `text_search_and_or_limit` (when added). Otherwise display results with any matches of any terms. ◦ **IR_AND_QUERY**: Display results with exact matches of all terms only. ◦ **IR_OR_QUERY**: Display results with any matches of any terms.

4. To edit the value of any property, click the property name and in the Search Resource Context Property form, update the **Value** field to `true` or `false`.

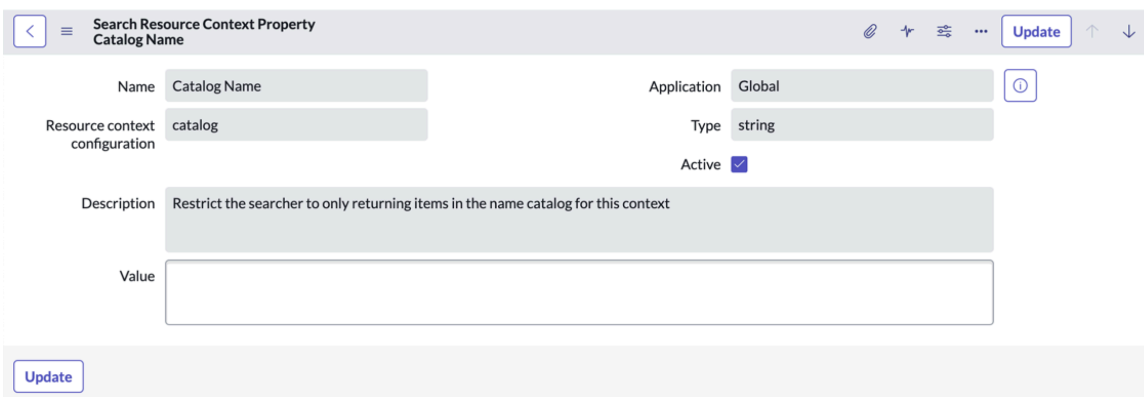

5. Click **Update**.

Managing contextual search

Contextual search provides a way to display results that can be relevant to the issue for which you are creating an incident or requesting a service.

The section covers topics that help you to:

- Understand the actions that are taken on the search results.
- Use contextual search results in a record producer within the Service Portal or as an agent in a form.

Set a search context as default

If you use the contextual search REST API and don't specify which search context to use, the API uses the search context that you have set as default. The search context that you use in the API takes preference over the one marked as default.

Before you begin

Role required: admin

About this task

Procedure

1. Navigate to **All** > **Contextual Search** > **Search Contexts**.
2. Open the search context record that you want to set as default.
3. Select **Set As Default**.

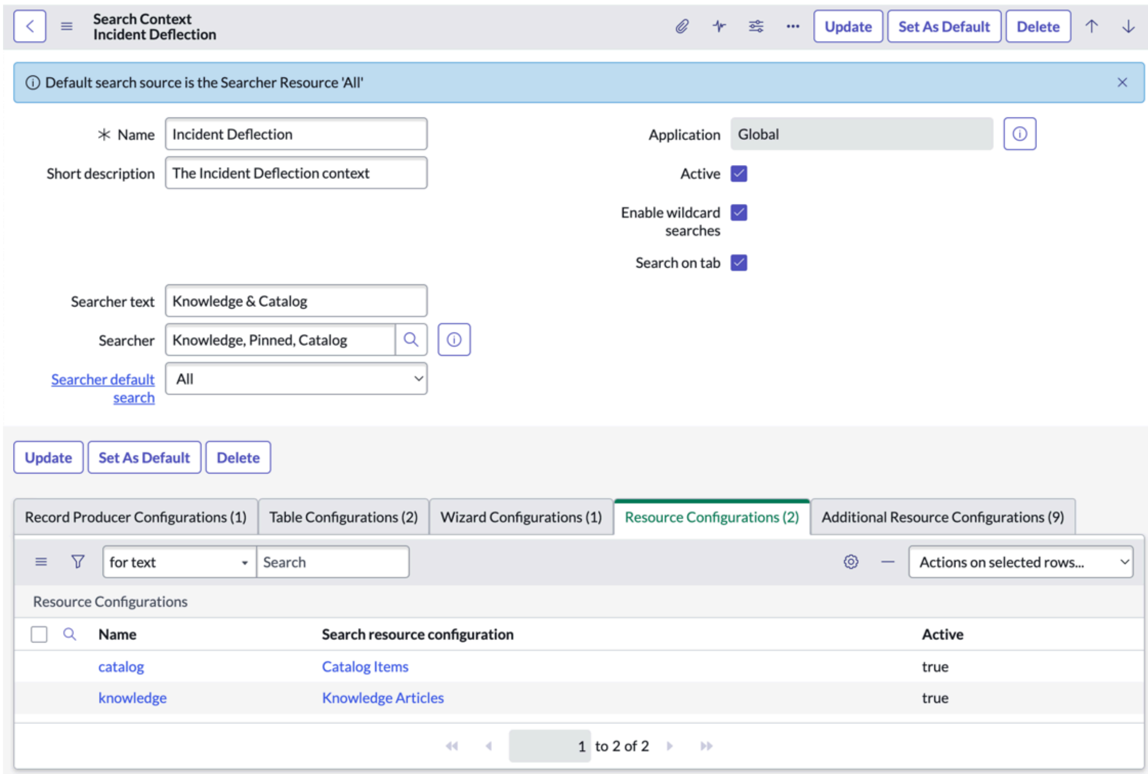

Set the default source for search context

You can decide which source to set as the default search source.

Before you begin

Role required: admin

About this task

- If you set an additional resource as the default source that source is considered the default search source.
- If you don't mark any additional resource as the default source, then the searcher source is considered the default search source.
- If you don't have a searcher and you have not marked any additional resource as the default source, the default is the additional resource with the lowest value in the **Order** field.

Procedure

1. Navigate to **All** > **Contextual Search** > **Search Contexts**.
2. Open the search context record to change its default source.
3. To set an additional resource as the default search source:

a. Click the Additional Resource Configuration related list.

b. Select and hold (or right-click) the search resource configuration record and click **Set as Default Search**.

4. To set a searcher source as the default search source:

a. Click the **Set Searcher as the Default Search** related link to set **Searcher** as the default search source.

Note:
The **Set Searcher as the Default Search** related link appears only if an additional resource is the default search source.

b. To search all searcher sources or a specific searcher source from the list, select **All**.

Note:
The **Searcher default search** field appears only when the **Searcher** field contains more than one source.

View a searcher

View the details of a searcher to know which search resources it searches and which search contexts use it.

Before you begin

Role required: admin

About this task

ServiceNow provides pre-defined searchers. You can't create, update, or delete searchers.

Procedure

1. Navigate to **All** > **Contextual Search** > **Searchers**.
2. Edit a searcher record.

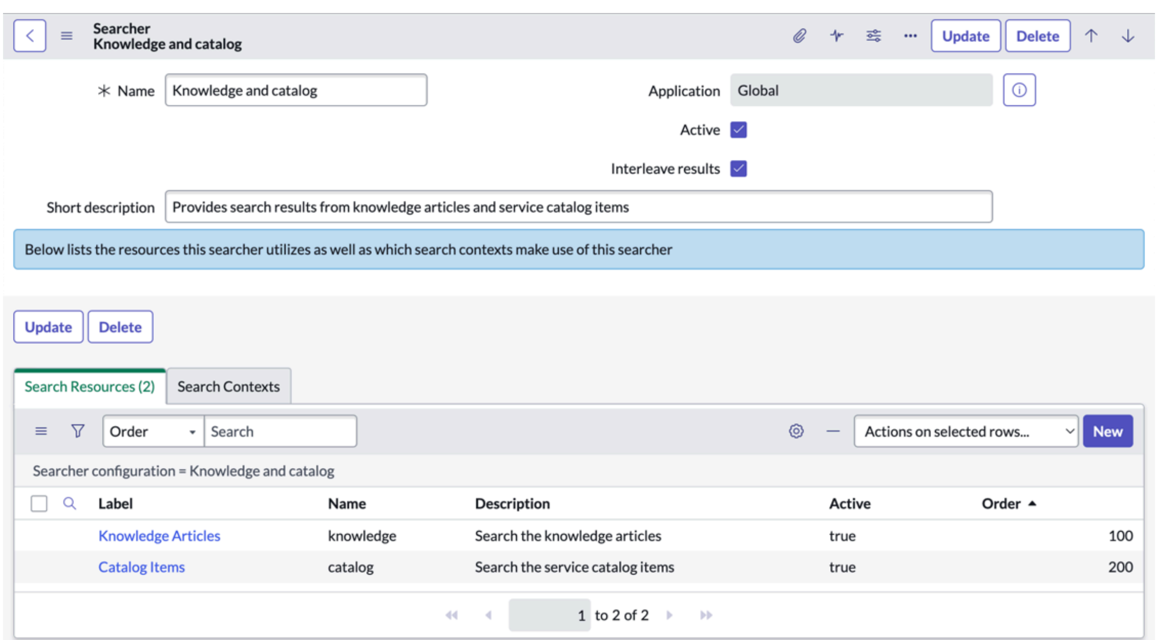

3. **Optional:** In the Search Resources related list, view the resources searched by this searcher.
4. **Optional:** In the Search Contexts related list, view all the search contexts that use this specific searcher.

Show the related search box in a form

Enable agents to search for related content without having to modify the text in the **Short description** field.

Before you begin
Role required: admin

About this task
Starting in the London release, the related search box is available in forms by default for new instances. If your form doesn't display the related search box, you can follow this procedure to show it.

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. Open the record for the table whose form you want to show the related search box on.
3. Select the **Enable related search box** option.
   The **Enable source selector** option appears.
4. Select and hold (or right-click) the header form and select **Save**.

Add multiple search fields for contextual search

Add multiple search fields if you want to search from a field other than the short description without using the related search box.

Before you begin
Role required: admin

About this task
Search results for the default search field appear automatically when you open a record. You can configure contextual search to search for results even when there is no text in the default search field. Search results in other associated fields only appear when you have updated content in those fields.

As an example, in the base system the **Short description** field is the default search field for the Incident form, so search results based on the field are displayed automatically. If you configure **Description** as an additional search field, contextual search displays search results for **Description** when you open a record with an empty **Short description**.

> **Warning:**
> Adding large fields can significantly affect contextual search performance. For best results, only add search fields that are restricted to 100 or fewer characters.

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. Open the table configuration record for which you want to add multiple search fields.
3. In the Search fields related list, click **New**.
4. On the form, fill in the fields.

Search Field form

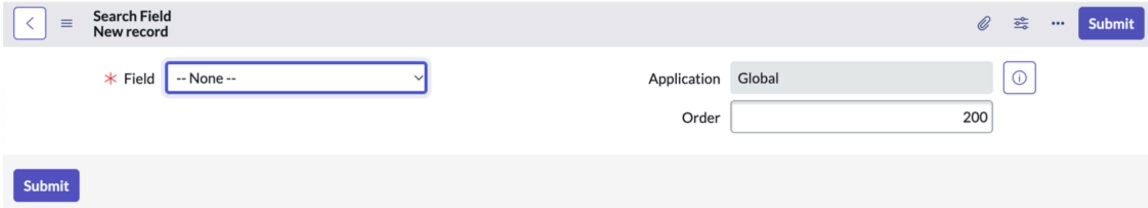

Field	Description
Field	The search field to be added in the search fields list. **Note:** You can only select a text field.
Application	Application scope of the table configuration. The configuration is available for all applications or for scoped applications.
Order	Order in which the field appears in the search fields list. If there are multiple search fields, this field indicates which searched field is executed first.

5. Click **Submit**.
   The new search field appears in the table configuration record's Search Fields related list.

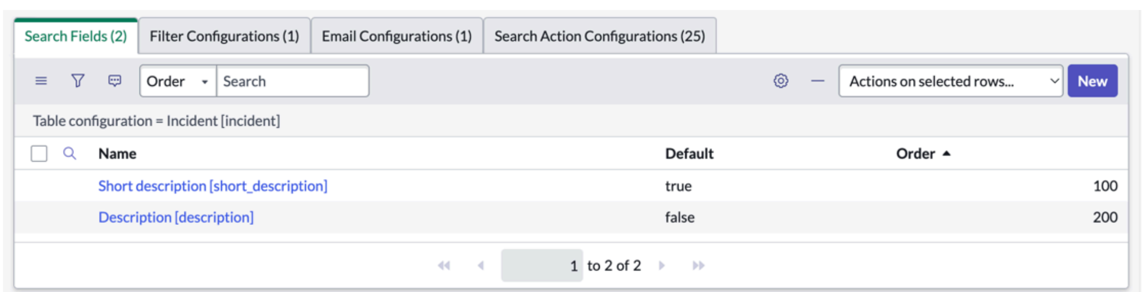

What to do next

Use the **Default** column to set one search field as the default search field.

Add additional resources to the source selector of a form

Add additional resources when you want the user to search beyond Knowledge and Service Catalog search resources. For example, add additional resources for **Open Incidents**, **Resolved Incidents**, or **Similar Open Incidents**.

Before you begin

Role required: admin

About this task

Additional resources are pre-defined. An additional resource is either a searchable source table with query conditions or a reference to a Predictive Intelligence Similarity solution definition.

Note:
You can only view the Predictive Intelligence resources if you are licensed for Predictive Intelligence and have activated Predictive Intelligence for Contextual Search.

Procedure

1. Navigate to **All** > **Contextual Search** > **Search Contexts**.
2. Open the search context record for which you want to see the additional resources.
3. In the Additional Resource Configurations related list, select **Edit**.
4. **Optional:** Filter the additional resources displayed in the **Collection** field by defining filter conditions and selecting **Run filter**.
   As an example, you can display only additional resources with the Predictive Intelligence Similarity type.
5. From the **Collection** field, move the required additional resources to the **Additional Resource Configurations List** field.
   For example, if you want to move additional resources whose type is **Predictive Intelligence Similarity**, then move **Similar Incidents**, **Similar Knowledge Articles**, **Similar Open Change Requests**, **Similar Open Incidents**, **Similar Open Problems**, and **Similar Resolved Incidents** from the **Collection** field to the **Additional Resource Configurations List** field.

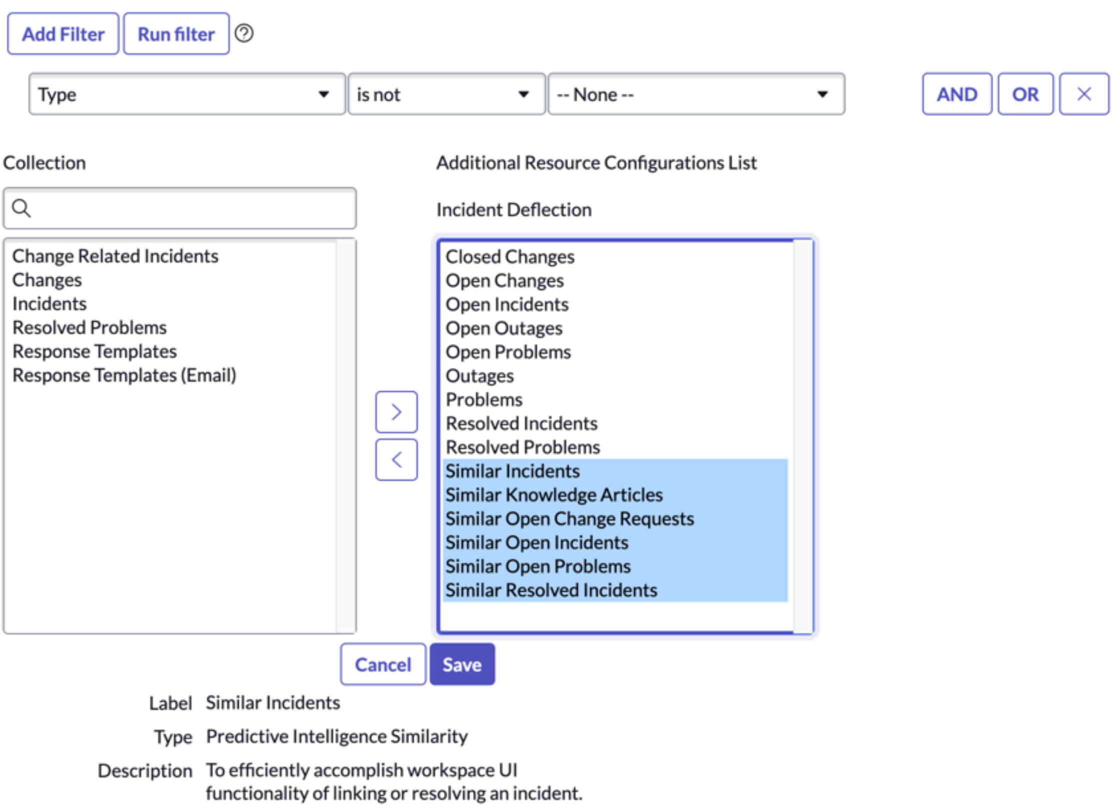

Note:
You can view the Predictive Intelligence sources only if you have license for Predictive Intelligence.

6. Select **Save**.

The additional resources appear under **Additional Resource Configurations**.

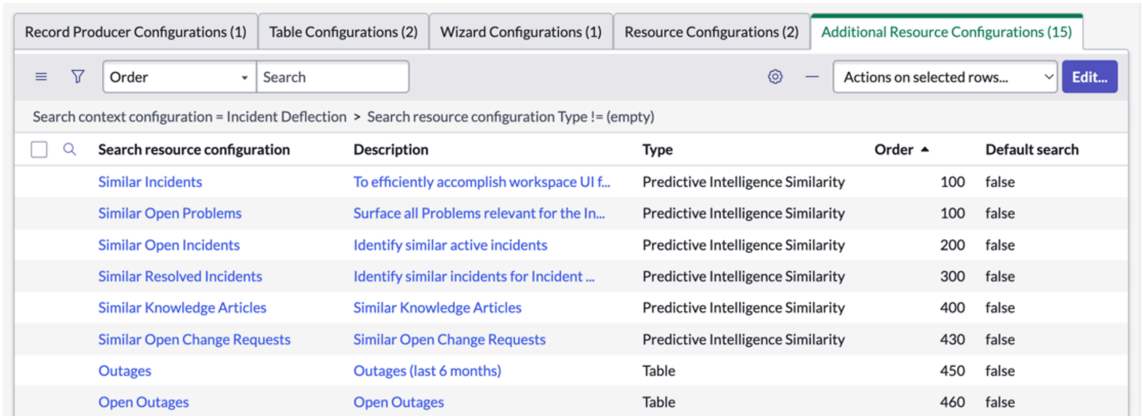

Search resource configuration	Description	Type	Order	Default search
Similar Incidents	To efficiently accomplish workspace UI f...	Predictive Intelligence Similarity	100	false
Similar Open Problems	Surface all Problems relevant for the In...	Predictive Intelligence Similarity	100	false
Similar Open Incidents	Identify similar active incidents	Predictive Intelligence Similarity	200	false
Similar Resolved Incidents	Identify similar incidents for Incident ...	Predictive Intelligence Similarity	300	false
Similar Knowledge Articles	Similar Knowledge Articles	Predictive Intelligence Similarity	400	false
Similar Open Change Requests	Similar Open Change Requests	Predictive Intelligence Similarity	430	false
Outages	Outages (last 6 months)	Table	450	false
Open Outages	Open Outages	Table	460	false

Related topics

Additional search resources

Similarity solution definitions

Edit search resource display field record

Use the search result display fields to provide titles, descriptions, and additional details for contextual search results.

Before you begin

Role required: admin

About this task

ServiceNow provides pre-defined additional resource or search sources to the user. You can't create or delete search result display configurations. Each table referred by an additional resource has a search result display configuration record. You can edit the values in the **Card title field**, **Card description field**, and the **Card additional fields**.

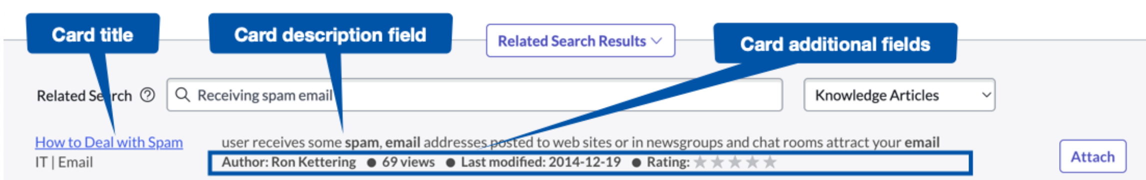

Note:
Ensure that you select the contextual search property **Shows meta data on form contextual search results** to view **Card additional fields** in the search result.

Procedure

1. Navigate to **All** > **Contextual Search** > **Search Result Display Configurations**.
2. Open the search result display configuration record that you want to modify.

   Note:
   Search result display configuration records with the Platform UI type do not affect search for Knowledge articles, Catalog Items, or pinned articles. You can configure settings for platform Knowledge article searches using Knowledge search properties.
3. In the form, edit the fields.

Search Result Fields form

Field	Description
UI type	[Read-only] Interface where the values of the search result display fields should appear. It can appear in platform or workspace.
Table	[Read-only] Name of the table in which the values of the search result display fields should appear. For example, if you mention Incident [incident] as the table, then the additional resource field values appear in the search result for the Incident [incident] table.
Application	[Read-only] Application scope of the record. The record is available for all applications or for scoped applications.
Card View	
Card icon	[Workspace] Icon name displayed on the search result card view.
Card label	[Workspace] Label displayed on the search result card view.
Card title field	[Platform or Workspace] Field whose value is shown as title in the search result title.
Card additional fields	[Platform or Workspace] List of field names displayed on the search result to provide more context to the search result. The card additional fields appear at the bottom of a search result record.
Card description field	[Platform or Workspace] Name of the field whose value is shown as the search result description.
Card display updated date	[Workspace] Check box to display the latest date on the search result card view.
Details View (only when the UI type is Workspace)	
Detail title field	Field whose value is shown as title in the search result detail view.
Detail link field	Field used as clickable link on the search result detail view.
Detail additional fields	List of field names displayed on the search result detail view.
Detail show last work note	Check box to display the latest work note on the search result detail view.

4. Select **Update**.

Specify fields for the Cxs_popup view

Edit the Cxs_popup view to specify the fields to be shown for addition resources in the preview window of a task record. The preview helps you with the summary of the record.

Before you begin

Role required: admin

About this task

If a Cxs_popup view isn't defined for the additional resource table, the sys_popup view is used. If the sys_popup view is also not defined, the default view is considered.

Procedure

1. Navigate to **Incident** > **All**.
2. Open the incident record whose view you want to change to Cxs_popup.
3. Select the additional actions icon ( ) and in the **View** option, select **Cxs_popup**.

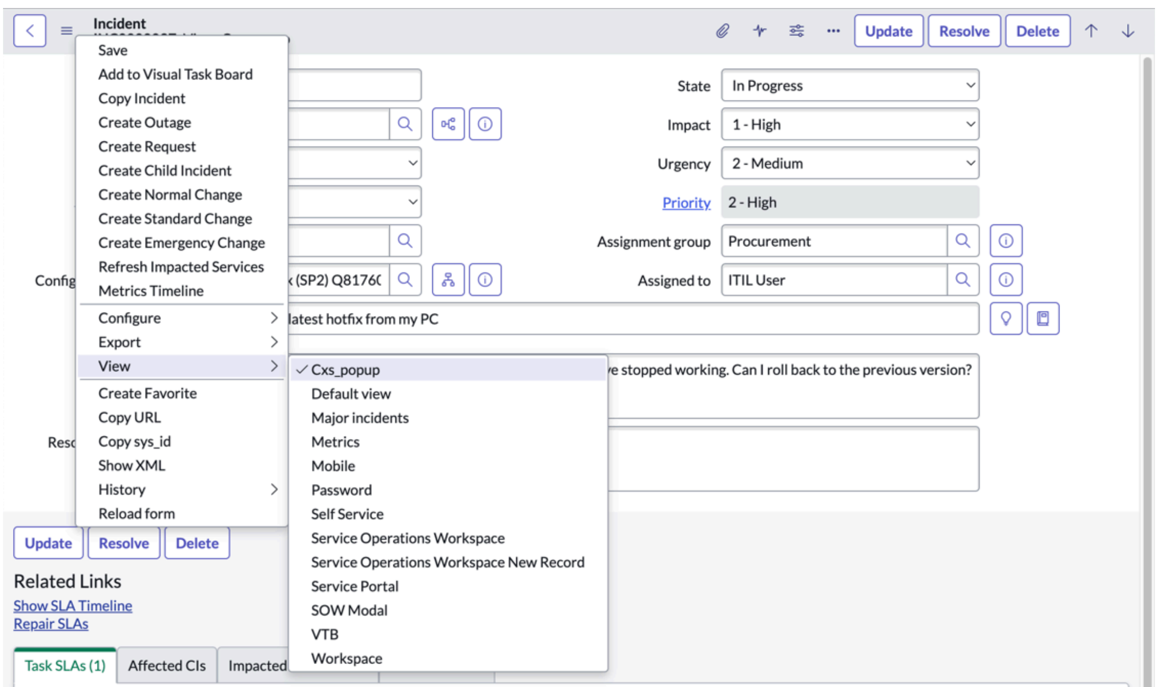

The view of the incident record changes to Cxs_popup view.

4. Select the additional actions icon ( ), and in the **Configure** option, select **Form Layout**.
5. From the **Available** list of fields, select the fields that you want to be available in the Cxs_popup view and move them under **Selected**.
6. Select **Save**.

Modify or disable search actions available for contextual search

Modify or disable search actions to restrict the actions that a user can perform on a search result.

Before you begin

Role required: admin

About this task

The Search Action Configurations related list contains all the available Search actions available for the selected Search context. Search actions can be modified and disabled but not deleted and new actions can't be created.

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. Edit a table configuration record.
3. In the Search Action Configurations related list, view the list of search actions available for the contextual search.

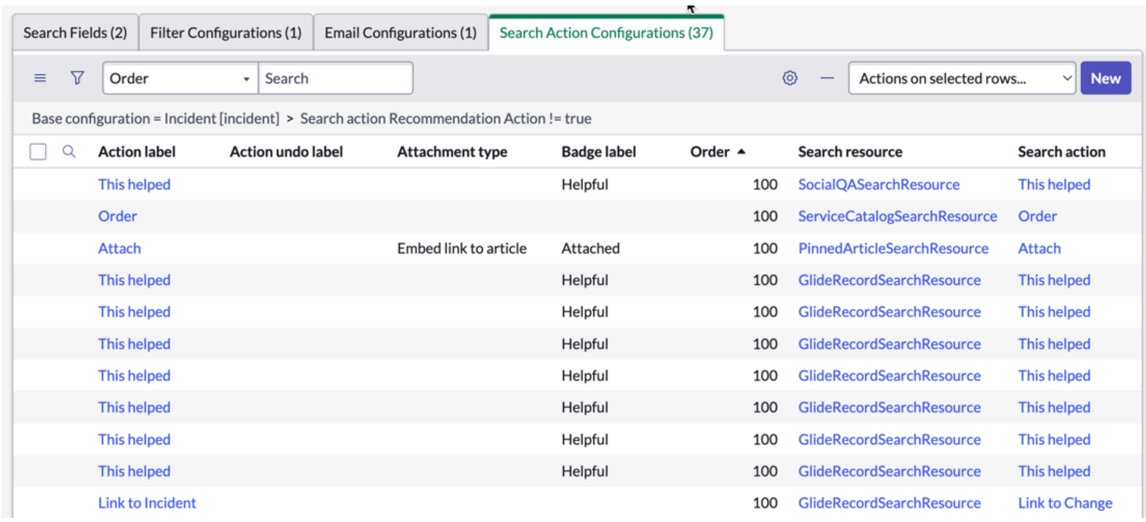

Action label	Action undo label	Attachment type	Badge label	Order	Search resource	Search action
This helped			Helpful	100	SocialQASearchResource	This helped
Order				100	ServiceCatalogSearchResource	Order
Attach		Embed link to article	Attached	100	PinnedArticleSearchResource	Attach
This helped			Helpful	100	GlideRecordSearchResource	This helped
This helped			Helpful	100	GlideRecordSearchResource	This helped
This helped			Helpful	100	GlideRecordSearchResource	This helped
This helped			Helpful	100	GlideRecordSearchResource	This helped
This helped			Helpful	100	GlideRecordSearchResource	This helped
This helped			Helpful	100	GlideRecordSearchResource	This helped
This helped			Helpful	100	GlideRecordSearchResource	This helped
Link to Incident				100	GlideRecordSearchResource	Link to Change

To disable a search action, set the value of **Active** to **false**.

Note:
You can't disable the **This helped** action.

The search actions that you can modify are as follows:

- **Action label**: Label on the action button on the preview window. For example, **This helped** or **Attached To Incident**.
- **Attachment type**: Method to attach a knowledge article with the record. You can embed a link (recommended) to the article or embed a copy of the article.

- ○ **Show on new record**: Enables the **Attach** button to appear when you are creating a record. Prior to the London release, the user had to fill out and save the form, then re-expand the related search results to access the **Attach** button.

  Note:
  This option only applies to tables with the Platform UI type. It defaults to **false** to provide a consistent user experience between Platform and Workspace tables.

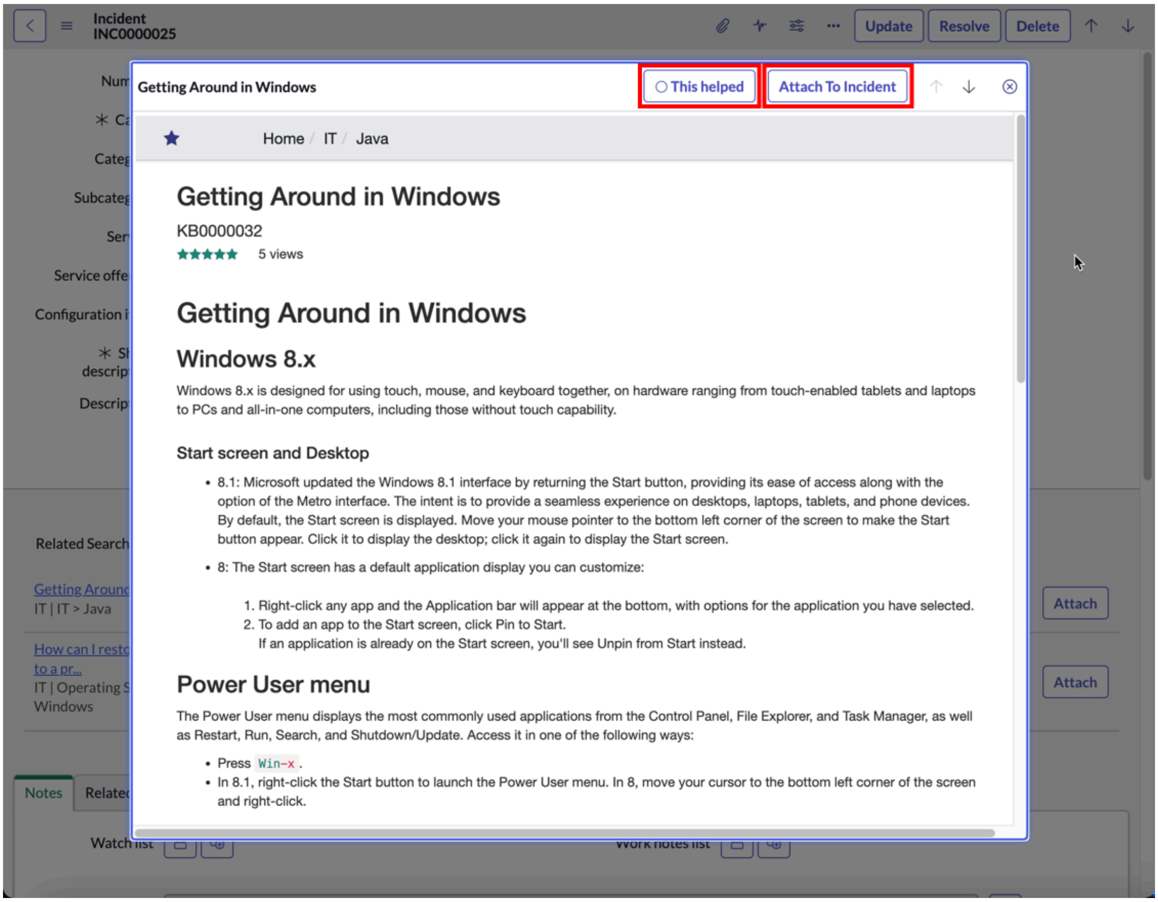

Enable viewing of search results for the current and another selected user

Enable the search as functionality when you want the ability to view search results common to the current user and another user that you specify.

Before you begin

Role required: admin

About this task

For example, an HR administrator may have some confidential information. A user might have partial access to the information. The Search as functionality applies security access to display results that are common to both the logged in or current user and the user referenced in the **Search as field**. The search result for the other user may have fewer entries than what that current user can actually view.

Note:
No search result is displayed for **Predictive Intelligence Similarity** additional resources. Record producers only support searching for the current user.

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. Edit a table configuration record.
3. In the Search as section, fill in the fields.

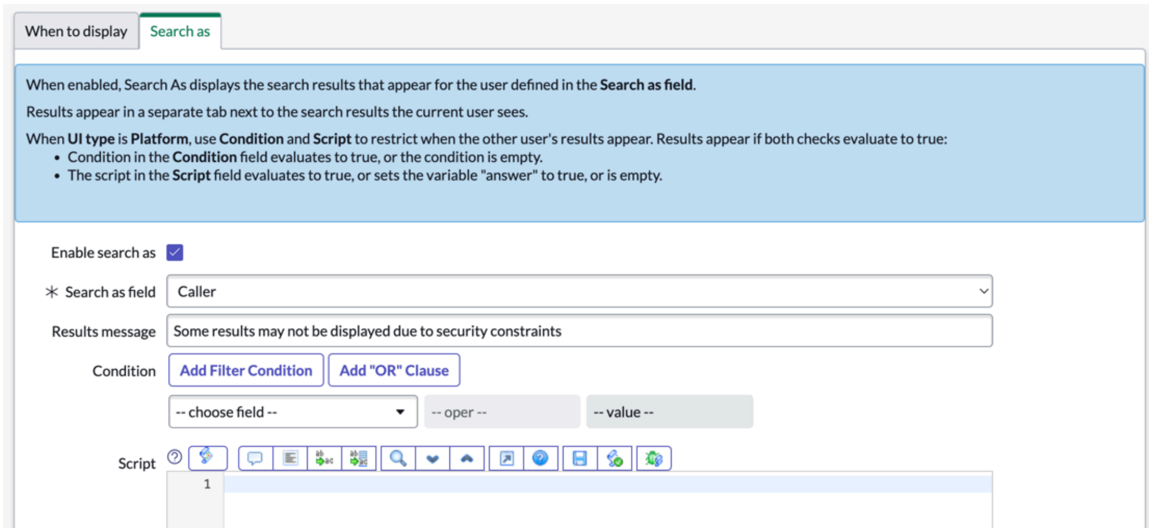

Search as configuration fields

Field	Description
Enable search as	Option to enable the Search as functionality.
Search as field	Field defining the other user to search as. This field appears only when you select **Enable search as**.
Results message	Message to display with search results when searching as a different user. This field appears only when you select **Enable search as**.
Condition	Condition restricting when the search result of the user mentioned in the **Search as field** should appear. This field appears only when you select **Enable search as**. **Note:** Both **Condition** and **Script** must evaluate to *true* for search results to appear. An empty condition or script also evaluates to *true*.
Script	A condition script, enabling administrators to implement more powerful conditions. For example, to restrict access to users that are members of a group, use `gs.getUser().isMemberOf("<group name>")`. This field appears only when you select **Enable search as**. **Note:** Both **Condition** and **Script** must evaluate to *true* for search results to appear. An empty condition or script also evaluates to *true*.

Result

The **My Results** tab displays search results for the currently logged-in user. The **<user-name> Results** tab displays search results for the user referenced in the **Search as field**.

Define contextual search for record producer

Define a record producer to trigger a search based on the text that you enter in a variable field. Search results appear at the bottom of the record producer form.

Before you begin

Role required: admin

About this task

The base system **Create New Incident** record producer is configured to use contextual search. If you have customized this record producer, you can configure contextual search to link to a specific field on your record producer.

Note:
Only one variable within a record producer can use contextual search.

Procedure

1. Navigate to **All** > **Contextual Search** > **Record Producer Configuration**.
2. Select **New**.
3. On the form, fill in the fields.

Record producer configuration form

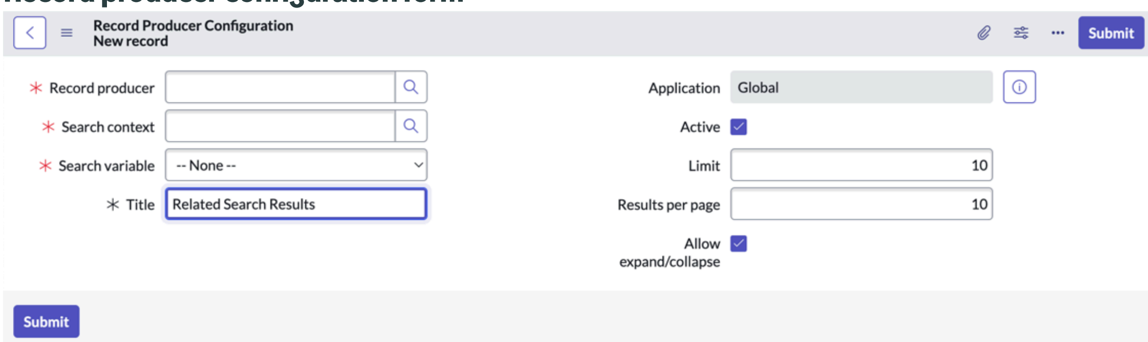

Record producer configuration form

Field	Description
Record producer	Unique name of the record producer where you want to add the contextual search.
Search context	Name of the search context to use.
Search variable	Variable within the selected record producer that uses the contextual search.
Title	Title for the search results area that appears in the form.
Active	Option to activate the record producer configuration.
Limit	Maximum number of search results returned.
Results per page	Number of search results to display per page.
Allow expand/collapse	Option to allow expanding or collapsing the search results in the form.

4. Select **Submit**.

Define email configuration for contextual search

Define email configurations to allow the system to include search results with an email notifications.

Before you begin

Role required: admin

About this task

When Knowledge articles are associated with a record, those search results can be included in email notifications. For example, email notifications sent when an incident is created can include Knowledge search results. An article in the search results may help the user who raised the incident to resolve the issue independently.

Procedure

1. Navigate to **All** > **Contextual Search** > **Table Configuration**.
2. In the Email Configurations related list, click **New**.
3. In the form, fill in the fields.

Email Configuration form

Field	Description
Email notification	Lookup icon to select an email notification to attach the search results. You can select only from the notifications that are on the same table as the table configuration record.
User field	A reference field that helps to narrow down on the search results. For example, if you select **Opened by**, the search results include only the articles that the **Opened by** user can access based on user criteria.
Application	Application scope of the email configuration. The configuration is available for all applications or for scoped applications.
Limit	Maximum number of search results returned.

4. Click **Submit**.

Provide knowledge in incident email notification

Contextual search results are included in email notifications that are sent to users who create a new incident.

Before you begin

Role required: admin

About this task

The email notification provides links to knowledge articles that may help users resolve their issues faster. For example, if a user raises an incident when the service desk staff isn't available, the email notification provides knowledge links that may help the user.

By default, contextual search results are based on the short description in the incident. Within the automated email response, contextual search adds links to relevant knowledge articles.

Procedure

1. A customer sends an email to IT support with the subject `My laptop keeps crashing`.
2. An incident is created based on this email.
3. The email subject is inserted into the **Short description** field of the new incident.
4. The automated email notification sent to the user includes search results based on `My laptop keeps crashing`.
   By default, notifications provide three article links.
5. The customer may then be able to use the search results to resolve the incident.

   Note:
   The knowledge links provided are filtered to ensure that the articles are accessible to the user who submitted the email.

What to do next

You can configure notification options, such as changing the number of links provided with notifications. You can also configure contextual search functionality to match the email notifications of your organization or you can use contextual search with notifications for other records.

Related topics

Edit an email notification for the search results

Edit an email notification for the search results

Include search results in an email notification sent to users.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Open the email template that you want to modify to include search results.
3. Select **What it will contain**.
4. In the **Message HTML** field, add the following script:
   ```
   ${mail_script:cxs_EmailSearchResults}
   ```

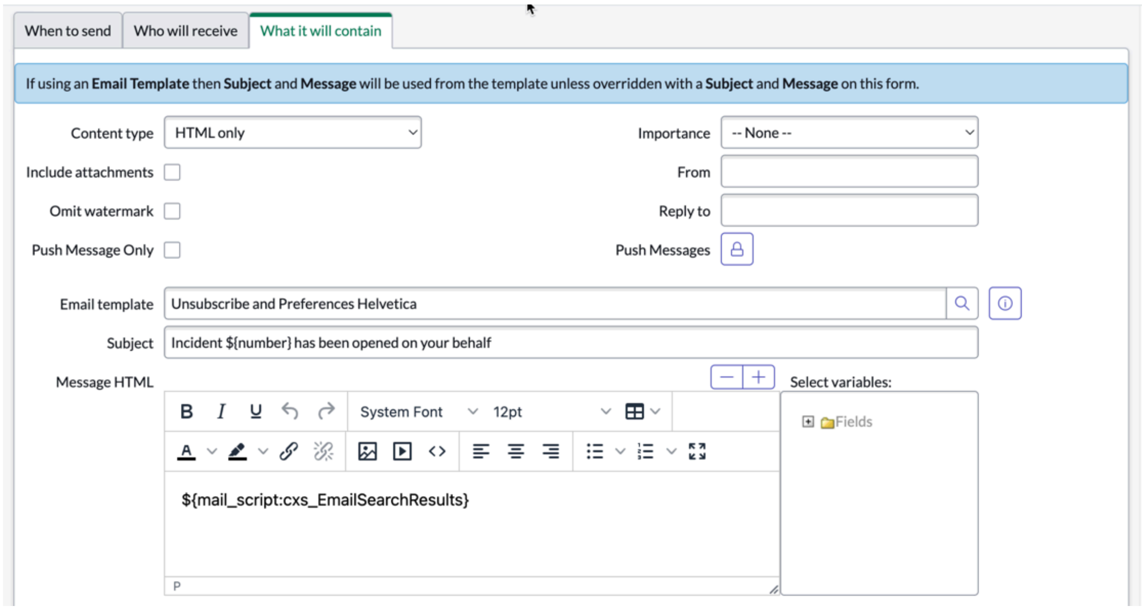

5. Select and hold (or right-click) on the header form and click **Save**.

Attach a Knowledge article

Attach a Knowledge article to an incident to solve issues quickly by going through the in-depth troubleshooting steps and other detailed information in the article.

Before you begin

Role required: none

About this task

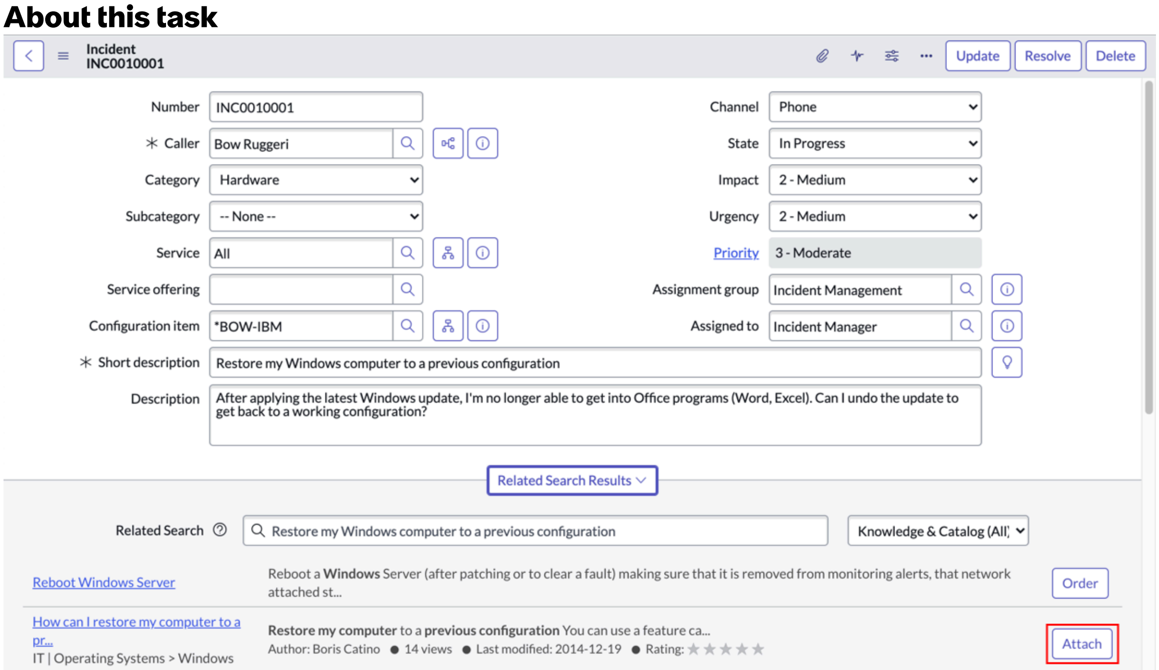

Procedure

1. Navigate to **All** > **Incident** > **All**.
2. Edit an active incident record.
3. Select **Related Search Results**.

Search results display in the Related Search Results pane.

4. **Optional:** To display only Knowledge article search results, open the source selector and select **Knowledge Articles**.
5. Select **Attach** in the search result for the Knowledge article that you want to attach to the incident record.
   A note about the Knowledge article is copied or appended to the **Additional comments (Customer visible)** field on the incident record. To set a different destination field on the incident record, see Specify field for attached Knowledge article links.

   Note:
   Attachments from the selected Knowledge article aren't attached to the incident record.

6. To add the note to the activity stream, select **Post**.
   Alternately, you can select and hold (or right-click) the form header, then select **Save**.

Result

Details of the selected Knowledge article appear in the designated field.

Specify field for attached Knowledge article links

Specify which field to add a note to when you attach a Knowledge article to a record.

Before you begin

Role required: admin

About this task

When you attach a Knowledge article to a record, by default a link to the Knowledge article is added to the **Additional comments (Customer visible)** field. You can specify the field for this link based on the table. For example, a Problem record doesn't have the **Additional comments (Customer visible)** field on the form, so the **Attach** action is configured to add the note to the **Work notes** field.

You can specify the field either globally or locally. The local field for a table overrides the global default field value.

To specify the global field, navigate to `Knowledge` > **Administration** > **Properties** > **Other Knowledge Properties** and specify the name of the field in the `glide.knowman.attach.fields` property.

You can specify the local field using the following procedure.

Procedure

1. Navigate to **All** > `Contextual Search` > **Table Configuration**.
2. Open the table configuration that you want to change attachment fields for.
3. Select the Search Action Configurations related list.

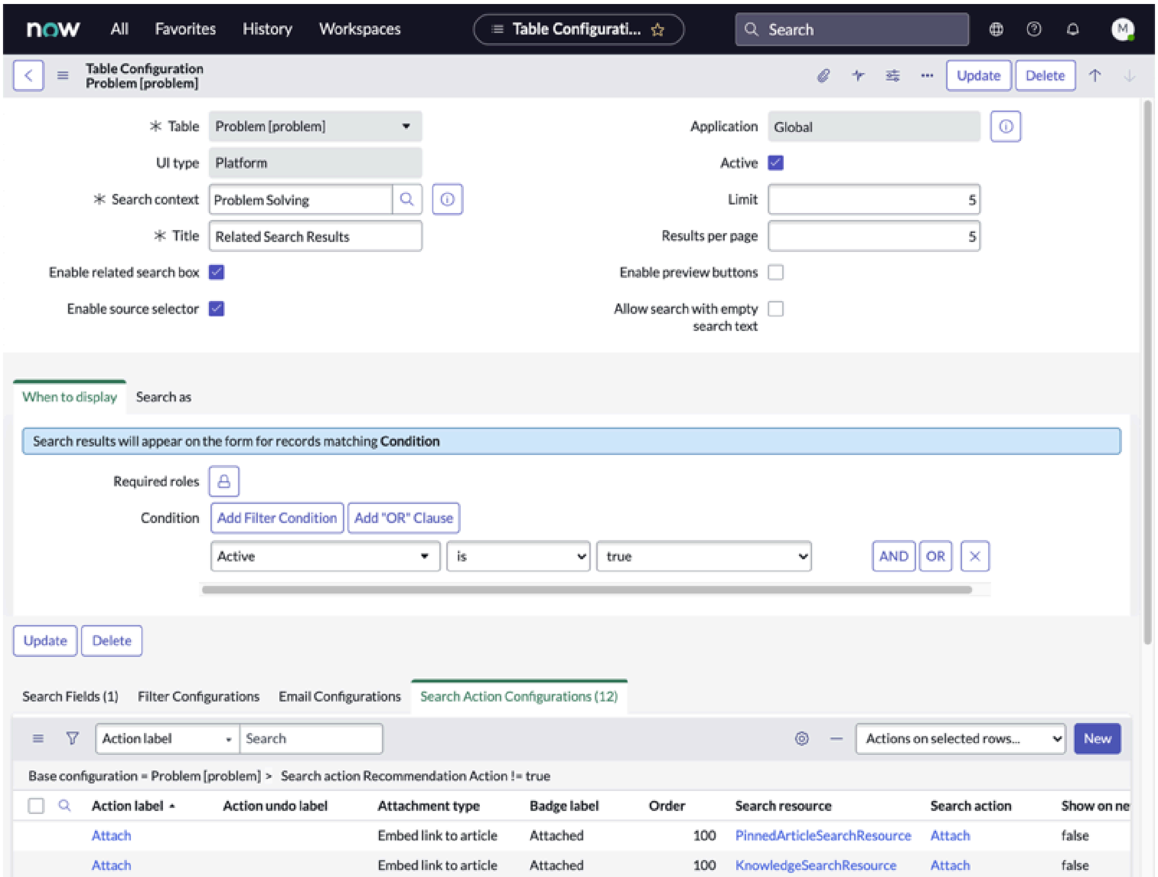

4. For each search action that you want to change the default field for, perform the following steps:
   a. Open the search action.
   b. Select the **Use custom field for attach note** option, if not already selected.
   c. In the **Attach note field**, enter the name of the field that you want to copy the article information into, such as `comments` or `work_notes`.
   d. Select **Update**.

   Note:
   - If you don't want to add a note to any field, leave the **Attach note field** blank.
   - If the field you have specified in the **Attach note field** isn't found on the form, then the note isn't added to any field.

What to do next

Clear the system cache to make sure the changed field specification applies. You can clear the system cache by appending `cache.do` to the instance URL. For example, `https://<instance name>.service-now.com/cache.do`.

Warning:

Clearing the system cache can affect overall performance and degrade system response times. Don't run cache flushes during business hours, and don't trigger cache flushes automatically.

Feedback information on contextual search

When a user clicks any of the action buttons such as **This helped**, **Order**, or **Attached**, the feedback information is stored in specific tables. The information helps you to understand the effectiveness of the searches. You can also create custom reports by querying these tables.

Table [name]	Description
Relevant Document Detail [cxs_rel_doc_detail]	The list of individual recorded actions. For example, Preview, This helped, Attach, or Order. The key fields in the Relevant Document Detail table are as follows: • **Original search terms**: The original search terms. This field is used only when the fulfiller has overridden the original search terms using the related search box. • **Relevance type**: The recorded action. • **Relevant document**: The reference to the search session group in the Relevant Document table. • **Created by**: The user who performed the action. • **Position**: The position of the recorded action in the result list. • **Search Resource**: The reference to the search resource for that particular result. Used only for a Predictive Intelligence Similarity result.
Relevant Document [cxs_relevant_doc]	Groups the Relevant Document Detail individual recorded actions by search session. **Note:** A new search session group is created every time you create or open a record or reload the page. The key fields in the Relevant Document table are as follows: • **Search session**: The search session group. • **Displayed on**: The table or record producer where the results were displayed. • **Relevant to**: The reference to the table record when **Displayed on** is **Table**. • **User/Created by**: The user who performed the action.

Run a report on contextual search usage

Run a report on usage such as how many times a knowledge article is attached to a record to know the effectiveness of the contextual search results.

Before you begin

Role required: admin

About this task

A service desk manager can generate a report on user issues that were resolved without an incident being raised. The information can help to create targeted knowledge based on trends. You can analyze commonly raised incidents and then create knowledge articles around those incidents which can help resolve or deflect similar incidents in future.

Procedure

1. Navigate to **All** > **Reports** > **View/Run**.
2. Select **Create a report**.
3. Under **Data**, fill in the fields.

Field	Description
Report name	Unique name for the report.
Source type	Source from which the data is populated. The default value is **Data source**. Select **Table**.
Table	Select the Relevant Document Detail table [cxs_rel_doc_detail].

4. Select **Next**.
5. Under **Type**, select **Next**.
6. Under **Configure**, select **Choose columns** and move the required fields from the **Available** box to the **Selected** box.

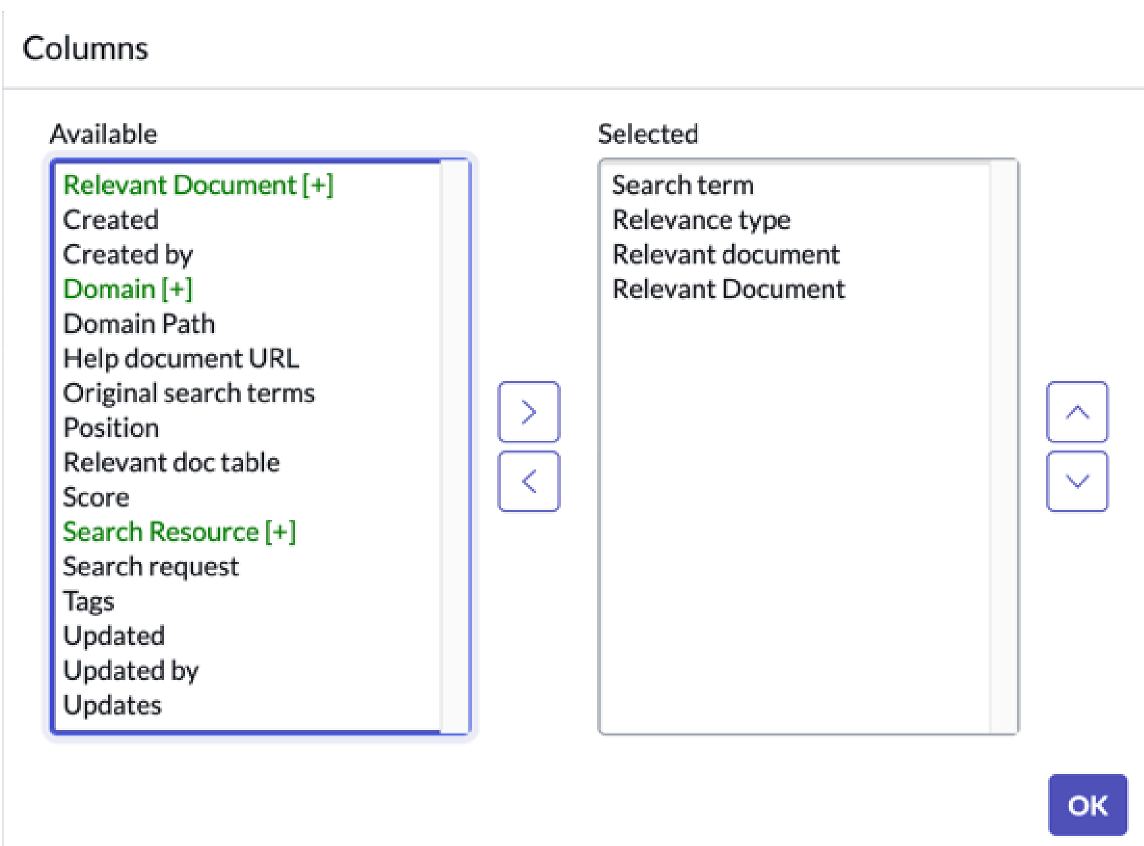

7. Select **OK** and then select **Run**.

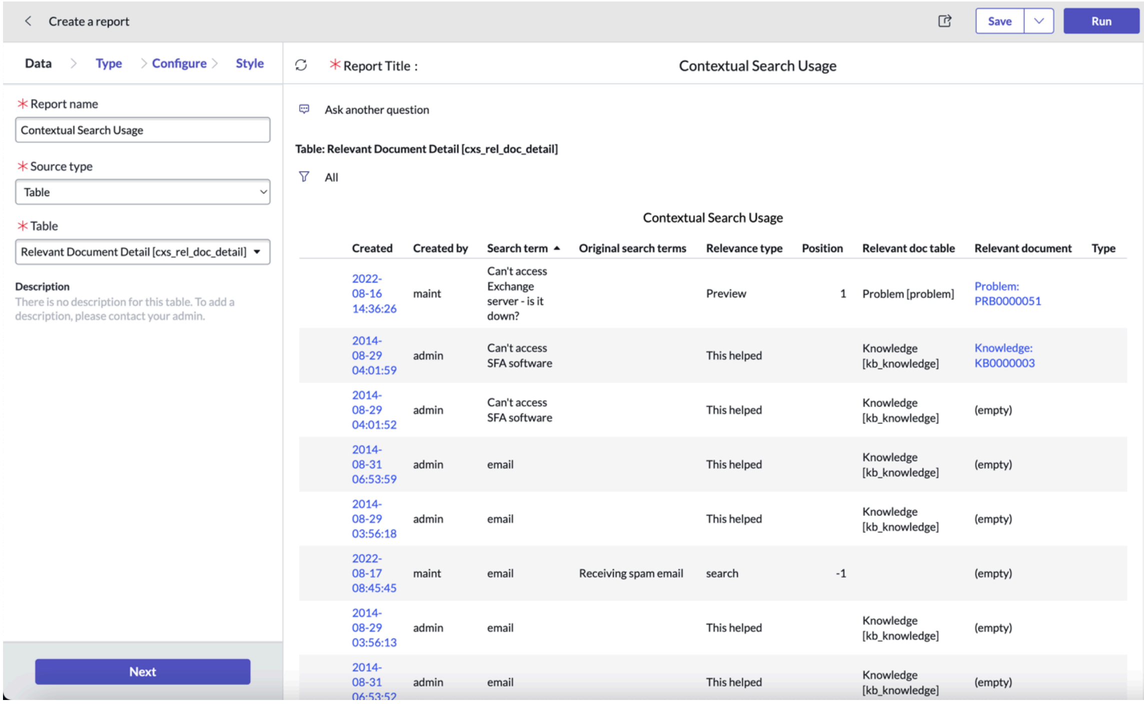

Note:
When users select the **Search on tab** check box in the Search Context [csc_context_config] table and then modify and leave the search field, an action gets logged in the Relevant Document Detail table [cxs_rel_doc_detail]. The information in the table can help you find the instances where users ran a search but did not take any action. In those cases, the value of the **Position** column is **-1**.

Domain separation and Contextual Search

Domain separation is supported in Contextual Search. 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.

How domain separation works in Contextual Search

By default, Contextual Search Suggestions results that appear use the domain of the currently logged-in user.

Search as a different user

In addition to being able to search for results using the domain of the logged-in user, the administrator can also configure Contextual Search to Search as a different user .

When **Search as** is configured, two tabs appear in the search results:

1. **My Results** tab searches using the currently logged-in user domain.
2. The **Search as** user tab searches using the domain of the User Reference field selected by the administrator. For example, selecting a request's **Opened by** or **Caller** field searches using the domain of the user who made the request.

Note:

- The **Search as** user results apply security access to display results that are common to both the logged-in user and the **Search as** user. Hence, the search result view for that user may have fewer entries than what that user can actually view.
- Record producers only support searching as the currently logged-in user.

Example: An Incident displaying search results for the currently logged-in user on the **My Results** tab.

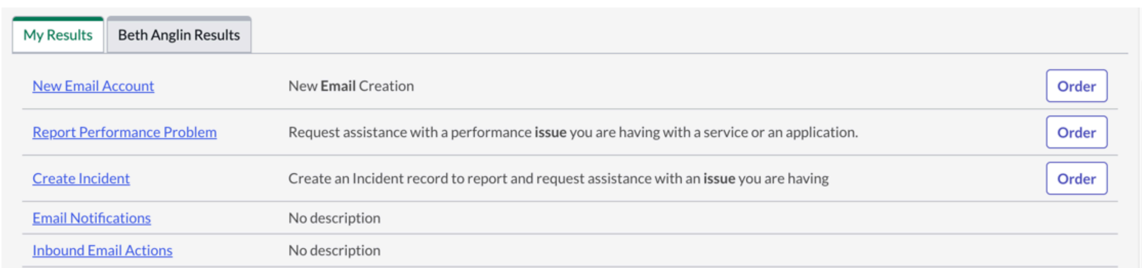

Example: The same Incident displaying search results for the user from the **Caller** reference field on a second tab.

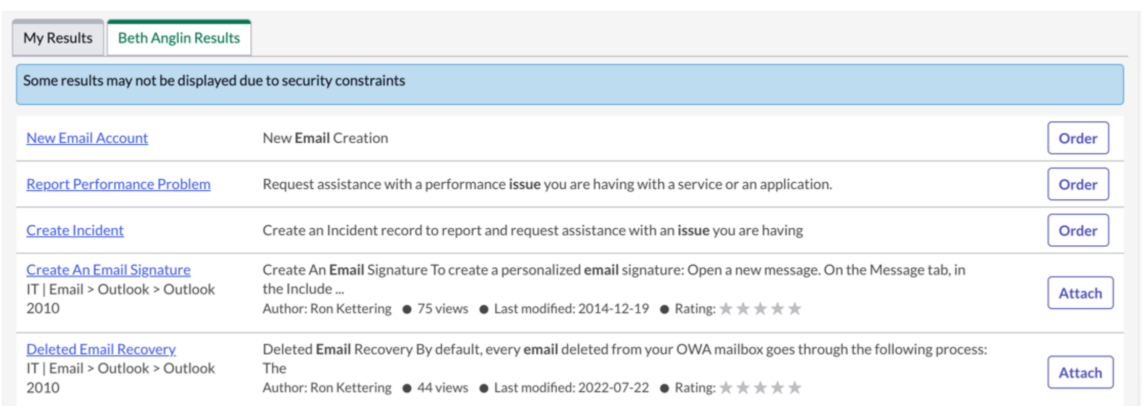

Related topics

Enable viewing of search results for the current and another selected user

Contextual Search roles

Contextual Search is installed with these roles.

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

Contextual Search feedback reader [cxs_feedback_reader]

Read records on the Relevant Document [cxs_relevant_doc] and Relevant Document Detail [cxs_rel_doc_detail] feedback tables.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

For details on these tables, see Feedback information on contextual search.

Intelligent Search for CMDB

Use everyday natural language query (NLQ) in a search string to query for a set of CIs in the CMDB. Intelligent Search for CMDB, parses, resolves ambiguities, and converts your search string into a valid CMDB query. Complex search strings open fully constructed on a canvas of CMDB Query Builder where you can continue and refine, or run.

Intelligent Search for CMDB is supported only on English language instances and isn't supported when the instance preferred language is set to a language other than English.

Integration with CMDB Workspace

Intelligent Search for CMDB is integrated into Home view in the CMDB Workspace store app. For information about how to use Intelligent Search, sample search strings, and helpful tips, see (Intelligent Search section).

Integration with CMDB Query Builder

Intelligent Search for CMDB is integrated with the CMDB Query Builder in the ServiceNow AI Platform. This integration is controlled by the system property *`glide.cmdb.query.nlq.activated`*, which is set to **true** by default. Intelligent Search for CMDB lets you use natural language processing in the CMDB Query Builder to find CIs and their relationships using Intelligent Search for CMDB functions.

Using Intelligent Search for CMDB

Intelligent Search is tailored to the CMDB, searching only through the CMDB class hierarchy for tables, and for CIs and their relationships.

Use the Intelligent Search search field to construct a search string using everyday natural language. Your queries can span multiple CMDB classes and involve many CIs that are connected by different relationships. After resolving any ambiguities with table names or relationship types, Intelligent Search converts your search string into a query that the CMDB can run. The CMDB query is constructed dynamically as you type into the search box and spell checker is applied if needed. A dynamic list of relevant suggestions appears as you type, with items such as table names, matching single words or part phrases in the typed-in text.

Use Intelligent Search:

- **Search tips**: Shows details and tips about the usage, and examples for single and multi-table search, advanced filtering, and relationships in Intelligent Search. The Relationships tab contains a link to the CMDB Implicit Relationships table.
- **Search**: Depending on whether the search string is already fully converted into a valid CMDB query and whether the search is for a single or multiple tables.
  - If the search string has no ambiguities with the table name or relationships, then the query runs and the results appear in a list view format.

    If the constructed CMDB query contains more than a single table, then the **View in Query Builder** button appears. Click the button to open the CMDB Query Builder with your query fully constructed on the Query Builder canvas. You can use the Query Builder to continue editing the query.
  - If there are any ambiguities with table names or relationship types in the search string, then the search string can't be converted into a valid CMDB query. In this case, the Refine your query dialog box appears to continue and parse your search string into a valid CMDB query. The dialog box contains suggested synonyms and labels for phrases in your search string. Use the drop-down lists to select the synonyms that match your intended search and then click **Go** to run the query.
  - If Intelligent Search is unable to convert your search string into a valid CMDB query, then clicking **Search** does not generate any query results. Instead, a feedback form appears. Fill out the form and click **Submit Feedback** to send your feedback to ServiceNow analysis.
- **Results Feedback**: Submit feedback to ServiceNow analysts, to express your assessment of the results. Choose descriptions that capture any gap between the results and your expected results, and add any helpful details.

Sample searches

When you click the search box, the drop-down list of pre-defined sample searches appears. The list consists of more common searches, or searches that are more difficult to construct such as searches that involve application services. Run any of those searches to get started.

- Sample searches are stored in the NLQ Sample Search [sn_cmdb_ws_nlq_sample_search] table
- Referenced tables are stored in the NLQ Sample Search Table [sn_cmdb_ws_nlq_sample_search_table] table

CMDB Admins (sn_cmdb_admin user role) can modify a sample search by directly editing its record in the NLQ Sample Search table. Click **All** and then in the Filter navigator, enter `sn_cmdb_ws_nlq_sample_search.list`. In the NLQ Sample Searches list view edit the record for a search that you want to modify.

Any modification to sample searches is reflected in both, the CMDB Workspace and the CMDB Query Builder.

Synonyms

The NLQ Synonym [nlq_synonym] table is pre-populated with synonyms for natural language strings for CMDB table and column names, and relationships. This table is used to match natural language search words to the CMDB query language. For example, the phrase ‘linux server’ has synonyms such as ‘Linux Server’, ‘Server’, and ‘Virtual Machine Instance’.

For details about viewing and adding synonyms customized to your business needs, see NLQ synonyms.

CMDB Implicit Relationships

You can help Intelligent Search find more results by defining some of the relationships between classes as implicit relationships. Implicit relationships can be useful in queries that involve service offering and application services.

NLQ admins can create new implicit relationships by navigating to **All** > **NLQ** > **CMDB Implicit Relationships**.

An implicit relationship defines the relationship between two tables and includes any filters you want to apply. When creating an implicit relationship, you set the following items:

- From table (from_table): The class that acts as the parent
- Filters: Conditions that are applied to the columns of the from_table
- To table: The class that acts as the child
- Relationship: How the from_table interacts with the to_table. For example, Contains: Contained by means the from_table contains the to_table
- Skipped table: The class that is implied and not captured by the CMDB Query Builder

For example, in CMDB Query Builder, you want to see your service offerings that have had a P1 incident in the last 10 days. However, if you were to type `show me all business service offerings with p1 incidents in the last 10 days`, NLQ wouldn't understand the relationship.

Implicit relationships are stored in the NLQ CMDB Implicit Relationship [nlq_cmdb_implicit_relationship] table and are used in the CMDB Workspace and if integrated, also in CMDB Query Builder.

For more information about NLQ in the ServiceNow AI Platform, see Natural Language Query (NLQ).

Time configuration

A standard feature of the ServiceNow AI Platform, ServiceNow's time configuration system handles scheduling across your entire enterprise, enabling you to manage projects efficiently, track and schedule events, and more.

Get started

Time configuration includes the scheduling of events and tracking of time. There are several functions available which enable tracking of time and using that information across applications.

Explore	Use	References
Learn about time configuration concepts and features.	Discover how to use various time configuration options.	Find useful time configuration reference materials and information.

Exploring time configuration

Time configuration helps you easily handle any necessary timing or scheduling aspects of your online enterprise.

Records stored on the ServiceNow AI Platform are kept with times measured down to the millisecond. Time configuration includes multiple applications to help you use those records correctly, with scheduling and monitoring activities in your business' day-to-day operations.

Date and Date/Time fields

Records in the ServiceNow AI Platform store date and time values in several different types of fields. These values are stored as integer numbers, in milliseconds, and appear in the appropriate date or date/time format.

Date and Date/Time field types

The ServiceNow AI Platform stores date and date/time information in records using the following field types:

Date and time field types

Field type	Dictionary XML type	MySQL DB type
Date	*glide_date*	DATE
Date/Time	*glide_date_time*	DATETIME
Time	*glide_time*	DATETIME
Duration	*glide_duration*	DATETIME
Due date	*due_date*	DATETIME

For the full listing of field types, see Introduction to Fields.

Global date and time field format

You define default date and time formats globally using system properties.

Date format

The *glide.sys.date_format* property defines the date format. An administrator can modify the property by navigating to **System Properties** > **System**. The default format is: yyyy-MM-dd.

Note:
A user can override the global date or time format with a personal preference.

Modifying the `glide.sys.date_format` property changes the date or time format globally. When modifying the standard date format, also verify the format using a Validate date and time script. You can use the same pattern strings as the `java.text.SimpleDateFormat` class, with the following exception:

- You cannot append the 'z' character to include the time zone when adding it to a script.
- If the resulting time zone is not three characters, an Invalid Date appears error appears when you validate the script, because the Date/Time Validation Script (sys_script_validator) fails. This failure occurs even though the script debugger shows that the content of the value meets the requirement of 'g_user_date_time_format'.

MM is months, while *mm* indicates minutes. The format string consists of the following abbreviations.

Field	Full form	Short form
Year	yyyy (4 digits)	yy (2 digits), y (2 or 4 digits)
Month	MMM (name or abbr.)	MM (2 digits), M (1 or 2 digits)
Day of Month	dd (2 digits)	d (1 or 2 digits)

Note:
It is recommended that you use a `yyyy` format-based date. For example, use `MM-dd-yyyy`, `yyyy-MM-dd`, or `dd-MM-yyyy`.

- If your system date format is set to use the `yy` year format, the system considers dates 20 years later or 80 years earlier from the current date as a date in the past or future leading to an unexpected behavior.
- For example, if you set the year as `51` for 2051, the system considers the year as 1951. Similarly, if you set the year as `37` for 1937, the system considers the year as 2037.

Time format

An administrator can modify the `glide.sys.time_format` property by navigating to **System Properties** > **System**. Modifying the property changes the date or time format globally. When modifying the standard time format, also verify the format using a Validate date and time script. You can use the same pattern strings as the `java.text.SimpleDateFormat` class, with the following exception: Appending 'z' to include the time zone is not supported.

The format string consists of the following abbreviations.

Field	Full form	Short form
Hour (1-12)	hh (2 digits)	h (1 or 2 digits)
Hour (0-23)	HH (2 digits)	H (1 or 2 digits)
Minute	mm (2 digits)	m (1 or 2 digits)
Second	ss (2 digits)	s (1 or 2 digits)

Add the character `a` to the end of the time format string to indicate AM or PM. This option shows AM or PM whether you are using 12-hour time (hh) or 24-hour time (HH).

The default format is: HH:mm:ss.

Note:
A user can override the global date or time format with a personal preference.

Related topics

Schedule entry fields

Personalize the system date format

You can personalize the format in which date values appear in your instance.

About this task

Personalizing the date format does not change global settings or impact the way other users see date values.

Note:
It is recommended that you use a `yyyy` format-based date. For example, use `MM-dd-yyyy`, `yyyy-MM-dd`, or `dd-MM-yyyy`.

- If your system date format is set to use the `yy` year format, the system considers dates 20 years later or 80 years earlier from the current date as a date in the past or future leading to an unexpected behavior.
- For example, if you set the year as `51` for 2051, the system considers the year as 1951. Similarly, if you set the year as `37` for 1937, the system considers the year as 2037.

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.
2. In the **Date format** field, select an option.
3. Click **Update**.

Personalize the system time format

You can personalize the format in which time values appear in your instance.

Before you begin

An administrator must add the **Time format** field to the Self-Service view of the User form. For more information, see Configuring the form layout.

About this task

Personalizing the time format does not change global settings or impact the way other users see time values.

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.
2. In the **Time format** field, select an option.
3. Click **Update**.

Configure the date picker for the list editor

A system property enables you to choose between two date picker configurations for the list editor.

Before you begin

Role required: admin

Procedure

1. Navigate to `sys_properties.list`.
2. Search for the property named *`glide.ui.list_edit.show_calendar_only`*.
3. Set the property **Value** to either of the following options.

<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>false</td><td>Calendar as well as a field appears for a date picker for manual date entry.
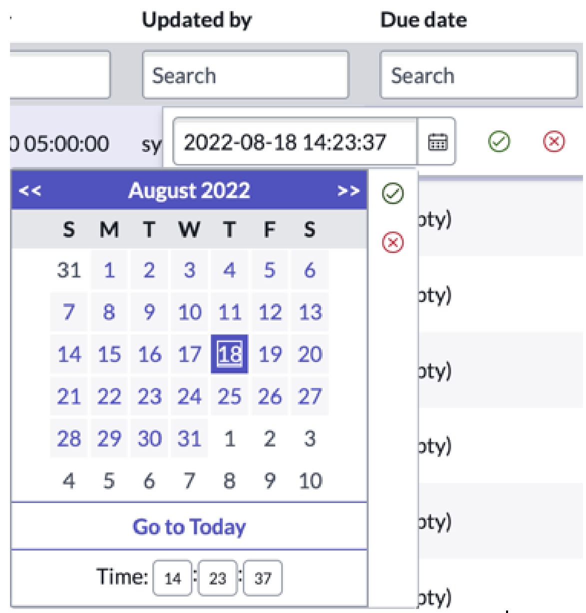
</td></tr>
<tr><td>true</td><td>Calendar only appears for a date picker, which is the default behavior in Core UI and UI15.
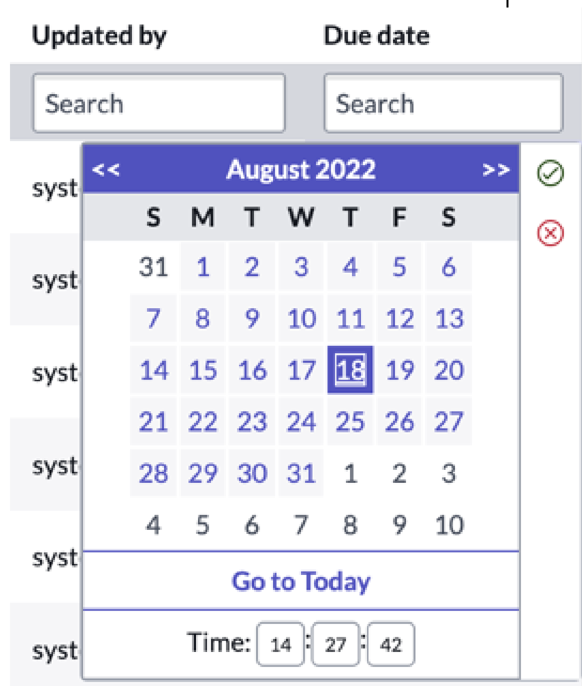
</td></tr>
</table>

Related topics

[List editor administration]

Default date and time fields

Certain time fields are provided by default to store particular date and time fields.

Global timestamp fields

All records inherit certain time stamp fields from the Global [global] table.

- **Created**
- **Updated**

These fields are automatically populated with the correct date and time, expressed in Greenwich Mean Time (GMT), rather than in the logged in user's local time zone.

Planned Task time fields

The Planned Task plugin provides a table (Planned Task [planned task]) with standard fields for measuring a planned task's time. For more information, see Planned Task.

Task fields for measuring work time

Use default task fields to measure progress and resolution for certain records.

The following base system fields are provided on certain tables to track how long it takes to close tickets:

- Time worked: A timer that runs when you view the record, and pauses while you close it, or when you manually pause it. Used to track the time spent by the help desk while working on the record.
- Resolve time: A calculated field that measures the time from the moment the record is opened, to the moment the record is closed. Used to track how long it takes to resolve the record.

These fields provide different metrics for request response.

The following additional tools are available for tracking work time:

- **Service level agreements (SLAs)**: Measures how long it takes a record fulfill certain conditions (such as an incident being marked **Resolved**).
- **Time cards**: Use the **Time worked** field to break down how much time was spent by day of the week.

Time worked

The Task [task] table provides a time-tracking field called **Time worked**.

Time worked field

This field measures how long a record has been viewed to track work time on a ticket. Any table that extends Task can use this field. To add the field, configure the form.

As the record is viewed, the timer counts upward. To pause the timer, click the stop icon ( ).

To resume the timer, click the start icon ( ).

When you save a task, the amount of new time in the timer is used to generate a record on the Time Worked [task_time_worked] table. You can view this table as a related list on the task form.

By default, the time in the **Time worked** field is a cumulative value stored in the task record. If you modify a Time Worked record, the changes are not reflected in the task timer.

You can set the property *com.snc.time_worked.update_task_timer* to enable updating of the task timer value, based on changes to the Time worked records. You do this using the *Update task timer* business rule.

You can also enable the dictionary attribute *time_worked_alert* so that updates to the time worked field make the form dirty. By default the attribute is set to false.

Resolve time

The **Resolve time** field is available on the Incident [incident] and Request [sc_request] tables. This field allows for easy reporting on how long it takes for requests to be closed, and is stored as an integer number of seconds.

Business rule calculation

Specific business rules calculate the **Resolve time** field when the record is resolved or marked closed, and measures the difference between the **Opened** and **Closed** dates. The field is stored in the system as an integer number of seconds.

- On the Incident table, the field is calculated on closure by the business rule *mark_closed*. The following lines of code calculate the resolve time:

```
if (dataChange || current.calendar_stc.nil())
      current.calendar_stc = gs.dateDiff(opened, resolved, true);
```

- When the incident is closed directly, the calculation is based on the *mark_closed* business rule. The following lines of code calculate the resolve time:

```
if (dataChange || current.calendar_stc.nil())
        current.calendar_stc = gs.dateDiff(opened, closed, true);
```

Display resolve time as a duration

You can display the resolve time as a human-readable duration rather than an integer representing a duration in seconds.

Procedure

1. From the Incident or Request form, right-click the form header and click **Configure** > **Table**.
2. From the **Columns** embedded list, click **Resolve time**.
3. From the **Attributes** related list, click **New**.
4. Enter these values.

Dictionary attribute values

Field	Value
Attribute	Format
Dictionary entry	Resolve time
Value	glide_duration

5. Click **Submit**.

   The resolve time appears in forms and lists as a duration, expressed in days, hours, and minutes.

   Note:
   This attribute does not change the field data format, which remains an integer representing a duration in seconds. In reports and data exports, it still appears as the actual number of seconds, rather than as a duration.

Export date and time formats

Because some export formats are intended for human consumption and others for database usage, various methods are used to provide date and time field information in different formats.

Excel

Date, Date-Time, and Time fields are all exported as their display values, and appear in a custom format, instead of the system date format. However, duration fields, export as the value stored in the database, which is an integer value of seconds.

Note:
If the date and time format is `hh:mm:ss` in the *`glide.sys.date_format`* System Properties setting, and you export time values to Excel, they appear in 24-hour military time.

To display the exported values in standard 12-hour am/pm time formats, select the `1:30PM time` format in **Format Cells** > **Time** in Excel.

XML

All Date and Time fields export as the value stored in the database.

PDF

All Date and Time fields (including Duration) export as their display value.

CSV

All Date and Time fields export as the value stored in the database.

Schedules

Use schedules to specify when service level agreements or inactivity monitors are active, or to specify when on-call rotations should take effect. Schedules are rules that include or exclude time for various actions or tasks.

For example, if a service level agreement is set to an 8-5 Weekdays schedule, the SLA only counts time during those hours. Generate and validate schedules in the **System Scheduler** > **Schedules** menu.

Note:
The legacy functions that you may have used, on the **System Scheduler** > **Schedules** and Fiscal Calendar menus, to define schedules and fiscal calendars, continue to be supported. However, you should use the Business Calendar menu in their places to define calendars and schedules and then associate them with specific ServiceNow AI Platform packages. To learn more about defining and using business calendars, see Creating business calendars.

Default schedules

Default schedules are available in a base system.

Name	Schedule
8-5 weekdays	Repeats every week on weekdays Monday through Friday.
8-5 weekdays excluding holidays	Repeats every week on weekdays Monday through Friday. Includes the child schedule U.S. Holidays.
Application	Repeats every week on Sunday.
Application FLX	Repeats every week on Sunday.
Blackout Wednesdays (GMT)	Repeats every week on Wednesday.
Database Server FLX	Repeats every week on Saturday.
Default MS Project	Repeats every week on weekdays Monday through Friday from 8:00am to noon and 1:00-5:00pm.
Global Infrastructure	Repeats every week on Saturday.
MySQL Database Service	Repeats every week on Saturday.
Network	Repeats every week on Saturday.
Project Management Schedule	Repeats every week on weekdays Monday through Friday from 8:00am to noon and 1:00-5:00pm. Default schedule for the Project Management application.
Resource Management Schedule	Repeats every week on weekdays Monday through Friday from 8:00am to noon and 1:00-5:00pm. Default schedule for the Resource Management application.
Server	Repeats every week on Saturday.
Servers San Diego	Repeats every week on Wednesday.
Software Blackout	Repeats every week on Wednesday.
U.S. Holidays	Lists 12 U.S. holidays.
WebServer FLX	Repeats every week on Sunday.
Weekends	Repeats every week on Saturday for two days.

Related topics

Schedules

Holidays

You can define each individual holiday as a schedule entry to create exceptions to existing schedules.

For instance, if an SLA requires an incident be resolved within three business days excluding Christmas, create a schedule entry for Christmas. Creating this entry ensures that the SLAs do not count Christmas when calculating elapsed time, even if it falls within the work week.

Because schedules can be included in other schedules through a parent-child relationship, it is also possible to create a holiday schedule and include it in other schedules to keep holidays consistent. The following example shows a holiday schedule.

US Holidays

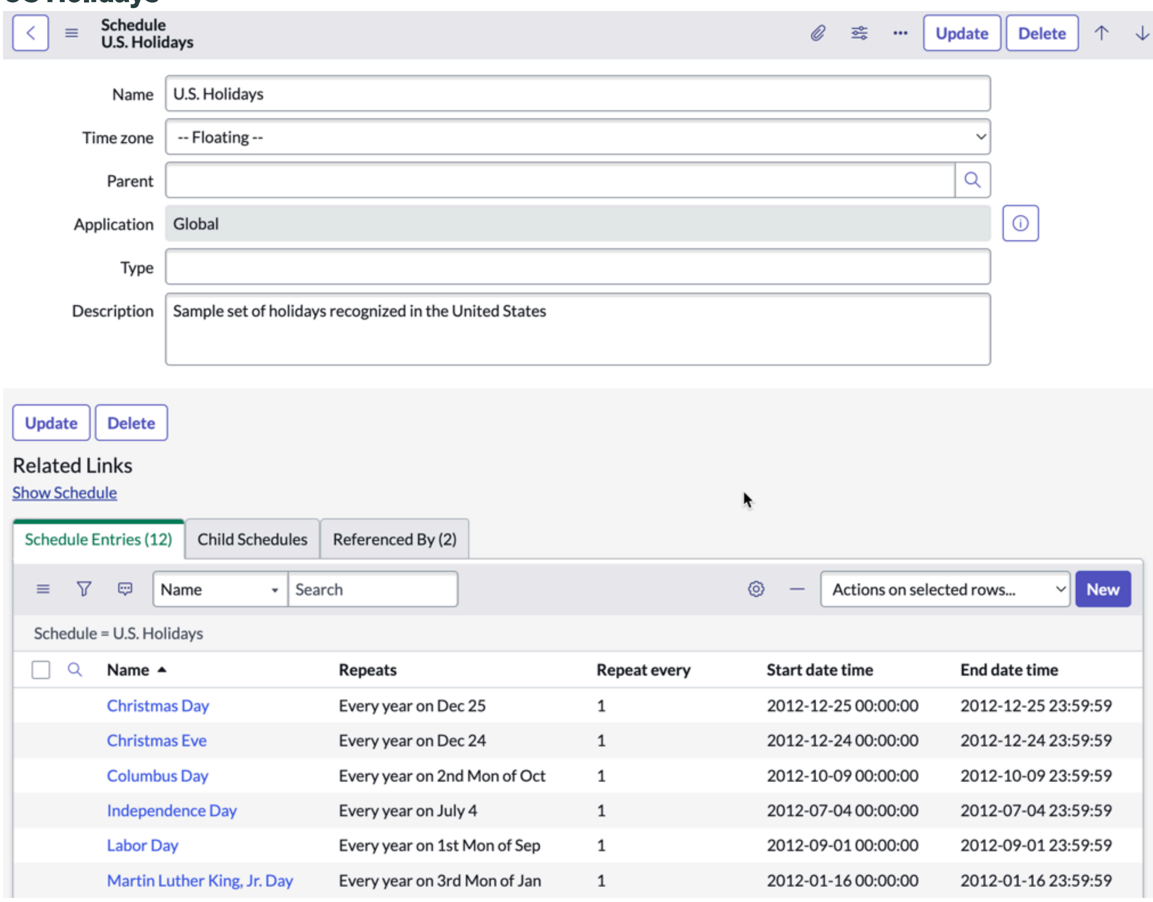

The following example shows a schedule that includes the preceding holiday schedule.

Child Schedule

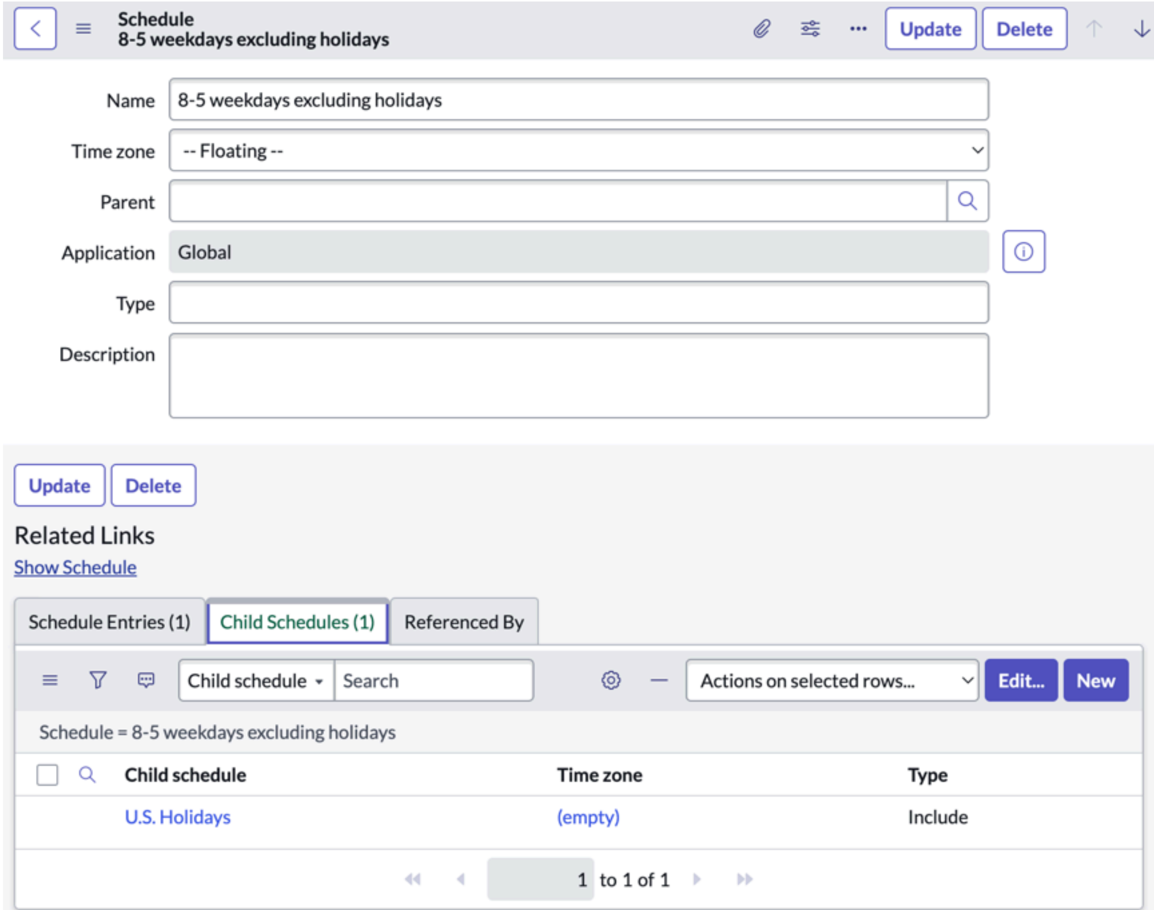

Related topics

[Schedules](#)

[Define a schedule](#)

Create a holiday schedule for multiple regions

You can create holiday schedules for multiple regions that follow the same work schedule but have different holidays.

About this task

The following method supports multiple regions with the same work schedule (for example, an 8-5 weekdays schedule) but with different holiday schedules.

Procedure

1. Create a holiday schedule for each region.
   For example, U.S. Holidays, British Holidays, and Australian Holidays.
2. Add the work schedule as a child schedule to each region's holiday schedule.

   This method requires making <number of schedules> + 1 total schedules. If you make the regional holiday schedule a child schedule of the work hours schedule, you must create a separate work hours schedule for each region. The total number of schedules in this case is <number of schedules> x two schedules.

Related topics

Schedules

Define a schedule

Parent and child schedules

Schedules can have one of two parent-child relationships with other schedules.

Parent field

When a schedule record lists a value for the **Parent** field, schedule entries from the parent schedule apply to both the parent schedule and the child schedule. By default, there are no sample schedules that use the **Parent** field.

Child schedule

When a schedule record has one or more child schedules in the **Child Schedules** related list, schedule entries from the child schedule apply to the containing schedule. By default, there are several sample schedules that use child schedules. For example, see the **8-5 weekdays excluding holidays** schedule that includes the **U.S. Holidays** schedule.

Parent and child schedules cannot contain conflicting schedule entry types. For example, a schedule containing maintenance schedule entries cannot also contain blackout schedule entries. Nor can a maintenance schedule have a child schedule containing blackout schedule entries.

Note:
To create a Holiday schedule and include it in a parent schedule, you must add it as a child schedule and with an Include type.

Parent schedules must have at least one entry that is not an **Excluded** type. If parent schedules are only exclusionary, they are invalid.

Note:
The **Show Schedule** related link shows schedule entries from the current schedule and the child schedule record. For example, when showing the **8-5 weekdays excluding holidays** schedule, holidays are also shown as excluded because the holiday schedule is a child schedule.

Related topics

Schedules

Define a schedule

Define a schedule

Configure schedules using schedule and schedule entry records.

Before you begin

Role required: schedule_admin or admin

About this task

- Schedule records specify a time zone and a type of schedule and use one or more schedule entries. Schedule records are saved in the Schedule [`cmn_schedule`] table.
- Schedule entry records specify the time periods that are included or excluded from a schedule. Schedule entries are saved in the Schedule Entry [`cmn_schedule_span`] table.

Alternately, you can also use the Business Calendar functions to create schedules. To learn more about creating business calendars and schedules, see Creating business calendars.

Warning:
If you are using Performance Analytics, use caution when changing time spans. Any change you make to the time span on a business or fiscal calendar can invalidate Performance Analytics data. Consult with a Performance Analytics administrator before making such changes.

Procedure

1. Navigate to **All** > **System Scheduler** > **Schedules** > **Schedules**.
2. Select **New**, enter a unique and meaningful **Name** and **Description** and then fill in the form.

   For detailed information on Schedule form fields, see Schedule fields.

   Note:
   If there are no active entries defined for the current schedule, a warning message appears in the Schedule form. If your schedule is a child schedule that contains only exclusions, ignore the message because exclusions are non-active entries.

3. Right-click the header bar and select **Save**.

   Note:

   - In the Related link, select **Show Schedule** to view the schedule in the calendar using the Next Experience user interface. For more information on this framework, see Next Experience UI. To create schedule entries in the Next Experience UI, refer to Create schedule entries in the Schedule calendar.

     To use the classic schedule calendar, you can add the *glide.schedule.enable_classic_calendar* property and set it to **true**. For information on how to add a system property, see Add a system property.

   - If you create a schedule of type **maintenance** and save the record, a UI policy hides the **Type** field from the form. To view or change the **Type** value, use the list of schedules instead of the schedule form. Add the **Type** column if necessary. You can click the cell for the value in the **Type** column and modify the value in the list view.

4. Configure one or more schedule entries.

   For detailed information on Schedule Entry form fields, see Schedule entry fields.

Related topics

Schedules

Schedule fields

The Schedule form contains the following fields.

Field	Description
Name	Unique name for the schedule.
Time Zone	Time zone for the schedule. If you select **Floating**, the time zone is relative to whatever process is accessing the item at the time.

Field	Description
	• For example, if a resource manager in Amsterdam sets a floating schedule for 8:00A.M. to 5:00P.M., a user in San Jose sees the schedule as 8:00 A.M. to 5:00 P.M. • When a schedule is defined in a specific time zone, users in different time zones see the schedule with their own time zone applied.
Parent	Parent schedule that constraints the new schedule.
Application	Application scope of the schedule.
Type	Text label that describes the purpose of the schedule. You can also use one of these system terms to determine how to process certain schedules: • **Excluded**: Excludes time periods from SLA counts. • **Maintenance**: Specifies time periods where change management activities are allowed. A schedule containing maintenance schedule entries cannot also contain blackout schedule entries. • **Blackout**: Excludes time periods from change management schedules. A schedule containing blackout schedule entries cannot also contain maintenance schedule entries.
Description	Description of the schedule.

Related topics

Schedules

Define a schedule

Schedule entry fields

You can only associate a schedule entry with one schedule. The **Schedule Entries** related list contains the definitions of the time periods you want to include in or exclude from the schedule.

Field	Description
Name	Unique name for the schedule entry.
Type	Label that describes the purpose of the schedule. The system also uses the schedule type to determine how to process certain schedules. • **Excluded**: Excludes time periods from SLA counts. • **Maintenance**: Specifies time periods where change management activities are allowed. A schedule containing maintenance schedule entries cannot also contain blackout schedule entries. • **Blackout**: Excludes time periods from change management schedules. A schedule containing blackout schedule entries cannot also contain maintenance schedule entries.
Show As	Option that indicates how the schedule entry should appear in calendar applications and how it should interact with other schedule entries. • **Busy**: Schedule entry is not available for scheduling a work task. • **At no extra charge**: Schedule entry is available for scheduling a work task.

Field	Description
	• **Tentative**: Schedule entry is tentatively available for scheduling a work task. • **On call**: Schedule entry is on call, which provides a means to determine which member of a user group is available to complete a task.
When	Date and time to which the schedule entry applies. If the schedule entry applies to a full 24-hour day, select the **All day** check box. **Note:** Changing the global date and time format in the `glide.sys.date_format` System Properties setting has significant impact on schedule calculations made when you select the **All day** check box. To learn more, see Global date and time field format.
Repeats	Repetition interval for the schedule entry, if any. If you select a repetition interval, other fields appear so you can further specify the repeat interval.
Repeat every	Scheduling repetition frequency - weekly, monthly, or yearly. This field is only visible when the **Repeats** field has a value of **Daily**, **Weekly**, **Monthly**, or **Yearly**.
Repeat on	Days of the week a weekly schedule repeats on. This field is only visible when the **Repeats** field has a value of **Weekly**.
Monthly type	Monthly schedule repetition frequency. This field is only visible when the **Repeats** field has a value of **Monthly**. Monthly repeat options include: • Repeat on a specific day of the month • Repeat on a specific day in a specific week of the month • Repeat on the last day of the month • Repeat on a specific weekday in the last week of the month
Yearly type	Yearly schedule repetition frequency. This field is only visible when the **Repeats** field has a value of **Yearly**. Yearly repeat options include: • Repeat on a specific day of the year • Repeat on a floating day
Float week	Week of the month a floating yearly schedule repeats on. This field is only visible when the **Yearly type** field has a value of **Floating**.
Float day	Day of the week a floating yearly schedule repeats on. This field is only visible when the **Yearly type** field has a value of **Floating**.
Month	Month of the year a floating yearly schedule repeats on. This field is only visible when the **Yearly type** field has a value of **Floating**.
Repeat until	Repetition end date. If you leave this field blank, the schedule repeats indefinitely.
Type	[Optional] Schedule entry description.

Related topics

Schedules

Define a schedule

Schedule for the fifth instance of a week date

You can select a date that computes to the fifth instance of a weekday for a date near the end of a month for a repeating monthly schedule,

About this task

Three options are available for handling months that do not have a matching Fifth instance of the selected day.

- Last: Selects the last instance of the weekday in the month.
- Next: Selects the first instance of the weekday in the next month.
- Strict: Skips any month without a matching Fifth instance and selects only months that have a matching Fifth instance.

Use the `glide.schedules.fifth` system property to control how a schedule entry that selects the fifth occurrence of a weekday behaves in months containing only four occurrences of that day.

- Use of this property is only valid when the `glide.schedules.repeat_nth` property is set to **Day**.
- The following example illustrates computing what day of the month a schedule repeats on when the schedule starts on the Fifth instance of a weekday in the month.

Procedure

1. Navigate to `sys_properties.list`.
2. Open the `glide.schedules.fifth` property.
3. Verify that the **Value** is set to **last**.
4. Navigate to **System Scheduler** > **Schedules** > **Schedules**, define a new schedule, and click **Submit**.
5. Open the new schedule and in the **Schedule Entries** related list, create a new entry with the following parameters:
   - When: November 29, 2012 at 10:00 to November 29, 2012 at 11:00
   - Repeats: Monthly
   - Monthly type: Day of the Week
   - Starting: November 29

   Note:
   In this example, November 29 is the fifth Thursday in the month.
6. Click **Submit**.
7. Open the same schedule entry.

   Note:
   The message **Every month on the Fifth Thu** appears on the form.

   The schedule for the first three months is computed as:

   - November 29, 2012 (Fifth Thursday of the month)
   - December 27, 2012 (Last Thursday of the month)
   - January 31, 2013 (Fifth Thursday of the month)

8. If you set the **Value** on the `glide.schedules.fifth` property to **next** instead of **last** in step 3, the schedule for the first three months is computed as:
   - November 29, 2012 (Fifth Thursday of the month)
   - January 3, 2012 (First Thursday of the next month since December 2012 does not have five Thursdays)
   - January 31, 2013 (Fifth Thursday of the month)
9. If the **Value** on the `glide.schedules.fifth` property is set to strict instead of **last** in step 3, the schedule for the first three months is computed as:
   - November 29, 2012 (Fifth Thursday of the month)
   - No meeting (December 2012 skipped because it does not have five Thursdays)
   - January 31, 2013 (Fifth Thursday of the month)

Related topics

Schedules

Define a schedule

Repeat a monthly schedule

Repeat a monthly schedule

Repeat a monthly schedule to reduce the amount of time required to define schedules and to provide consistency in scheduling.

About this task

You can specify the following options for monthly schedules that start on a particular day of the month (**Monthly type** is set to **Day of the month.**). The following methods are available to compute what day of the week a monthly schedule repeats on:

- **Day**: Calculates the day of the week to repeat on by determining the order of the selected starting date within the month. For example, if you specified that the starting date appears on the first Monday in the month, the schedule repeats every first Monday of every month.
- **Week**: Calculates the day of the month to repeat on by determining what week number the selected starting date appears in the month. For example, if the starting date is a Monday during the second week of the month, the schedule repeats the second Monday of every month.

Use the system property `glide.schedules.repeat_nth` to determine what method your instance uses to compute what day a repeating monthly schedule occurs on. By default, instances use the more accurate Day method.

> **Note:**
> Use the **Week** method to maintain backwards compatibility with customized schedule logic. To learn more about calculations for the starting day of each month, for monthly schedules that start on the fifth instance of a day, see Fifth Instance of a Day of the Week.

The following example illustrates computing what day of the week a monthly schedule repeats on.

Procedure

1. Navigate to `sys_properties.list`.
2. Open the `glide.schedules.repeat_nth` property.
3. Verify that the **Value** is set to **day**.

4. Navigate to **System Scheduler** > **Schedules** > **Schedules**, define a new schedule, and click **Submit**.
5. Open the new schedule and in the **Schedule Entries** related list, create a new entry with the following parameters:
   - When: November 5, 2012 at 10:00 to November 5, 2012 at 11:00
   - Repeats: Monthly
   - Monthly type: Day of the Week
   - Starting: November 5

     Note:
     In this example, November 5 is the first Monday in the month, but it is in the second week.
6. Click **Submit**.
7. Open **Schedule Entry**.

   Note:
   The message **Every month on the first Mon** appears on the form. The first few dates this schedule runs on are:
   - November 5, 2012 (First Monday of the month)
   - December 3, 2012 (First Monday of the month)
   - January 7, 2012 (First Monday of the month)
8. If you set the **Value** on the `glide.schedules.repeat_nth` property to **week** instead of **day** in step 3, the first few dates this schedule run on are:
   - November 5, 2012 (Schedule starts on Monday in the second week of the month)
   - December 10, 2012 (Second Monday in the month)
   - January 14, 2012 (Second Monday in the month)

Related topics

Schedules

Define a schedule

Using schedules and calendars

Specific applications within the ServiceNow AI Platform generate graphical calendar displays based on schedule pages.

These schedule pages appear in daily, weekly, or monthly views. Currently, the applications using schedule pages include:

- Project Management
- Maintenance Schedules
- Group On-Call Rotation
- Field Service Management

Schedule pages are records that contain the scripts that determine the functionality of the graphical display.

- Because of the heavy degree of scripting involved in a schedule page, you should, for most instances, use the default schedule pages in the base platform.
- The schedule page uses a URL with a series of parameters attached to generate the graphical display.
- A schedule page controls calendar content (`cmn_schedule_page`) record.

To access Schedule Pages, navigate to **System Scheduler > Schedules > Schedule Pages**. The Schedule Pages form provides the following fields:

Field	Field Type	Description
Name	String	General name used to identity the current schedule page.
Schedule type	String	Schedule type is a string that is used to uniquely identity the schedule page via the "`sysparm_page_schedule_type`" URI parameter. For example, a schedule page cou follows: **/show_schedule_page.do? sysparm_page_schedule_type=gantt_chart&sysparm_timeline_task_id=d530bf907f0(** Alternatively, you can access the schedule page can by setting the "**sysparm_page_sys_i** unique 32 character hexadecimal system identifier of the schedule page.
View Type	Choice	Each view type displays different field combinations. There are two options available: • **Calendars** • Schedule Pages
Description	String	General description that provides additional information about the current schedule page. necessary.
Init function name	String	**Note:** This functionality is only used by **Calendar** type schedule pages. The init function name specifies the name of the JavaScript function to call inside the **Clie** calendar type schedule pages.
HTML	String	**Note:** This functionality is only used by **Calendar** type schedule pages. Jelly parses the HTML field, which is a scriptable section. It then injects it into the display p the calendar. It can be used to pass in variables from the server and define extra fields are
Client script	String	The client script is a scriptable section that allows for configuring options of the schedule different depending on the schedule page view type.
Server AJAX processor	String	**Note:** This functionality is only used by **Calendar** type schedule pages. The Server AJAX processor is specific to calendar type schedule pages that is used to retu items and spans that display.

Invoke the Schedule page and view a calendar

You can use the URL from arguments module, a field decoration (dictionary attribute `ref_contributions`), or a UI Action linking to `show_schedule.do`to invoke the Schedule page.

For example, the On-call calendar module generates the calendar from the following URL:

```
$ocf.do?sysparm_start_date=2016-03-01&sysparm_current_view=month
ly&sysparm_include_view=monthly,daily
```

This URL takes the user to the monthly calendar view dated March 1, 2016.

- The URL component `sysparm_current_view = monthly || weekly || daily` determines the current calendar view. If `sysparm_current_view` is given an invalid or empty value, it defaults to monthly view and rewrite the URL to that view. Valid values are monthly or weekly or daily.
- The URL component `sysparm_include_view = monthly,weekly,daily` determines which calendar views are available. If `sysparm_include_view` is given invalid or empty values, only the valid views appear. If all values are invalid, the default three views appears and rewrites the URL to those views. Valid values are `monthly`, `weekly`, and `daily`.

  Note:
  The `sysparm_include_view` is only available for the new on-call calendar view.

- The URL component `sysparm_group_id=group_sys_id` determines the specified group to filter on. If `sysparm_group_id` is specified, it filters by the specified group. Valid value is a group sysId.
- The URL component `sysparm_start_date=2016-03-02` determines the date based on which every view appears. If `sysparm_start_date` is specified, it opens the calendar contesting that particular date in which every view is selected. The format is YYYY-MM-DD. If an invalid format or date is specified, it uses the current date and rewrites the URL to that date.

Note:
The `sysparm_zoom` URL component has been replaced with `sysparm_current_view` in OnCallRotation only.

Schedule calendar

Use the Schedule calendar to create, view, and edit events.

Schedule calendar

The schedule calendar is available in day, week, or month views in calendar and time line views. To create events in the calendar, see Create schedule entries in the Schedule calendar.

The image shows the calendar user interface.

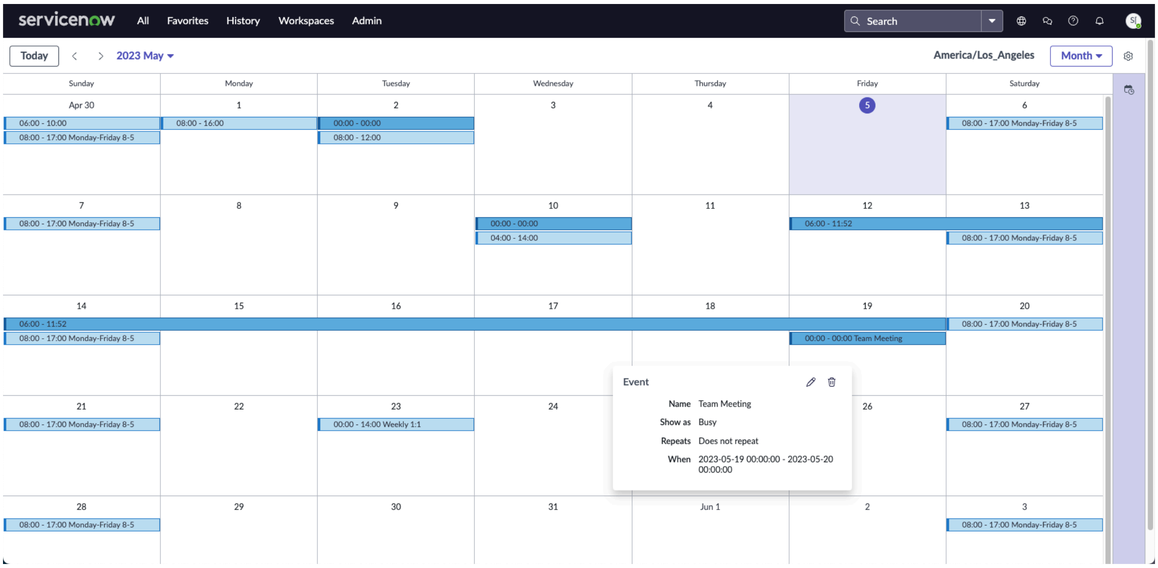

Note:
To use the classic schedule calendar, see Define a schedule.

Domain support and schedules

Domain separation is supported in schedules. 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. Activate the Domain Support [com.glide.domain] plugin to enable the domain separation functionality for schedules.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application’s service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer’s message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Overview

The records in the Schedule [cmn_schedule], Schedule Page [cmn_schedule_page], and Timeline Page [cmn_timeline_page] tables have a defined domain.

- The child tables use the *`domain_master`* attribute to derive the domain from the parent table.
- You can find the *`domain_master`* attribute on the dictionary record for the respective table.

The following diagram illustrates the scope of domain separation in different schedule tables:

Domain support for schedules

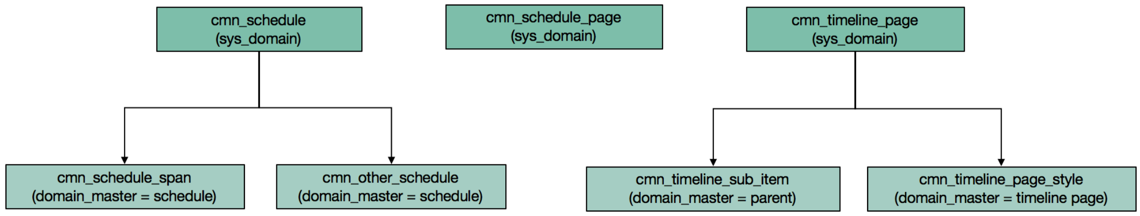

Custom domain support implementations

Domain separation support does not automatically occur when you migrate to a new release containing a custom implementation of domain support for tables such as Schedule Entry [cmn_schedule_span]. This action avoids changing any specific configurations that you may have in place.

To implement the base system domain support for schedules, a sys.script utility is provided. To run this utility, navigate to **Background** > **Scripts – Background**. The script is listed under the com.glide.schedules plugin as `fix_schedule_domain_support.js`.

- The utility attempts to add the **Domain** [sys_domain] column to the Schedule [cmn_schedule], Schedule Page [cmn_schedule_page], and Timeline page [cmn_timeline_page] tables.
- It then attempts to add the domain_master attribute to the Schedule Entry [cmn_schedule_span], Other Schedule [cmn_other_schedule], Timeline Sub Item [cmn_timeline_sub_item], and Timeline Page Span Style [cmn_timeline_page_style] tables.
- If the script finds existing records between a child and parent record that have differing domain, the script does not introduce the *`domain_master`* attribute to the child table.

For example, consider the relationship of the Schedule [cmn_schedule] (parent) and Schedule Entry [cmn_schedule_span] (child) tables. If the **Domain** [sys_domain] column is available on both tables in the upgrading instance, the utility must migrate to the base system implementation of domain support for schedules.

- If the script detects records where the child Schedule Entry [cmn_schedule_span] domain differs from its parent Schedule [cmn_schedule] domain, it stops executing and logs a warning message.
- If the script does not find differing records, it deactivates and limits read access to the **Domain** [sys_domain] and **Domain Path** [sys_domain_path] columns on the Schedule Entry [cmn_schedule_span] table.
- Finally, the script adds the domain_master=schedule attribute to the dictionary file for the Schedule Entry [cmn_schedule_span] table.

Note:
The *`domain_master`* attribute ensures that the child and parent record domains remain the same as the domain for the child that is derived from the specified reference field.

System scheduler

The System Scheduler application contains two separate engines for scheduling: the Scheduled Jobs engine and the Schedules engine.

Scheduled jobs

Scheduled Jobs executes any work that must be performed at a specific time, or on a recurring basis. The Scheduled Jobs module links to the Schedule [sys_trigger] table. Manipulating

records on the Schedule table is not recommended. Use this table to view existing base system scheduled jobs.

The Scheduled Jobs module in System Definition is an admin-friendly interface for scheduling work. Use this module to create new scheduled jobs. For more information, see Create a Scheduled Job.

Schedules

Schedules are rules that include or exclude time on a calendar. They are used by service levels, inactivity monitor, and group on-call rotation. For instance, you can define a schedule to restrict service levels to only apply to weekdays during business hours. You can also exclude holidays from an on-call rotation.

For more information, see Use Schedules.

Scheduled jobs

Scheduled Jobs are automated pieces of work that can be performed at a specific time or on a recurring schedule.

You can automate the following kinds of tasks:

- Automatically generate and distribute a report
- Automatically generate and schedule an entity of records, such as an incident, change item, configuration item, from a template
- Run scheduled jobs from scripts or business rules
- Scheduling at the end of the month
- Scheduling for weekdays

For developer training, see Scheduled Script Executions and Events Objectives on the ServiceNow Developer Site.

Scheduled job states

A scheduled job can be in any one of the following states.

Ready

Job is ready to run at the next scheduled interval.

Running

Job is in the process of carrying out a task.

Queued

Job has been added to the scheduler queue and is waiting to run.

Error

Error occurred while running the job.

Schedule job errors

If you encounter errors generating a scheduled job, visit Troubleshoot a scheduled job through Application Insights for additional information.

Create a scheduled job

Create a scheduled job on the Schedule Job [sysauto] table.

Before you begin
Role required: admin

About this task
Create all new scheduled jobs using this method.

Procedure
1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Click **New**.
3. Select the appropriate type of scheduled job.
   - Automate the generation and distribution of a report
   - Automatically generate something (a change, an incident, a ci, etc) from a template
   - Automatically run a script of your choosing

Result
The fields presented depend on the type of scheduled job required.

Related topics
Personalize the system date format

Set a system time zone

Special cases in job schedules

View a schedule item

Automate generation and distribution of a report

Generate and distribute scheduled reports via email.

Before you begin
Roles required: admin
A user with the report_scheduler role can also create a scheduled report through a different navigation path. For more information, see Schedule emails of reports.

About this task
Scheduled reports can be distributed in PDF, CSV, or XLS format. Graphical reports can be distributed in PNG or PDF format. Multilevel pivot reports can only be scheduled in PDF format.

Note:
- It is not possible to schedule Calendar, Map, or Single Score reports. You can schedule Pivot Table reports only if the plug-in `com.snc.whtp` is enabled.
- Data may not appear on reports created by an individual whose user account is deactivated. To ensure that the desired data appears, an active user must recreate the scheduled report.

Procedure
1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Click **New**.
3. Select **Automate the generation and distribution of a report**.
4. Populate the following fields:

Field	Description
Name	Name of the scheduled report.
Report	Reference to the report you are scheduling. You can use the Search icon to locate the report. **Note:** Automatic email support is not available for Calendar, Map, and Single Score reports.
Users	Individual users who receive the report at the scheduled date and time. Users must have **Notification** set to **Enabled** on their user records to receive reports.
Groups	Groups that receive the report at the scheduled date and time.
Email addresses	Comma-separated list of email addresses of report recipients who are not in the system.
Application	Name of the scoped application that contains the report. **Global** appears if the report is in the global scope.
Active	When active, scheduled reports are delivered.
Run	Time interval to use for running the scheduled job: ○ **Daily**: Runs daily, at a designated time. ○ **Weekly**: Runs on a weekly basis, at a designated time and day of the week. ○ **Monthly**: Runs on a monthly basis, at a designated time and day of the month. ○ **Periodically**: Runs on a designated repeating interval. ○ **Once**: Runs for a single occurrence only. ○ **On Demand**: Runs immediately on demand. ○ **Business Calendar: Entry Start**: Runs on the starting entry dates for the business calendar that you select in the **Business Calendar** field. A scheduled job runs for the starting date of each of the business entries that you defined for the business calendar. For example, if the business calendar represents a fiscal year, and the starting date of each entry is a fiscal month, the scheduled job runs on the first day of each month. ○ **Business Calendar: Entry End**: Runs for the ending date for the business calendar that you select in the **Business Calendar** field. This selection runs

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td></td><td>in the same manner as Business Calendar: Entry Start, but for the end dates of the associated business calendar entries.

Note:
When you select Business Calendar: Entry Start or Business Calendar: Entry End, you can apply an offset factor to schedule the job to run before or after the time span of the selected business calendar. If you select a business calendar and all its business calendar spans are in the past, the job would never run as this is an invalid configuration. To learn more, see the Offset type and Offset fields.

To learn more about creating and using business calendars and defining business calendar entries, see Creating business calendars andDefine business calendar entries.</td></tr>
<tr><td>Time zone</td><td>Time zone to use with the Time field entry when you specify the time at which the scheduled job should run. Select a time zone entry:
-None-: Use the default time zone for the logged-in user who is creating the scheduled job.

For example, the scheduled job runs at 04:45 p.m. US/Pacific time if it is the user's assigned time zone, and you enter <code>16:45</code> into the Time field.
Use System Time Zone: Use the default system time zone that is specified for the instance in which it runs.

For example, the scheduled job runs at 10:15 p.m. London time if Europe/London is the default system time zone for the instance, and you enter <code>22:15</code> in the Time field.
Actual time zone.

For example, the scheduled job runs at 1:30 p.m. in the US Eastern time zone if you select US/Eastern, and enter <code>13:30</code> in the Time field.</td></tr>
<tr><td>Time</td><td>Time of day at which the scheduled job should run, expressed in hours, minutes, and seconds on a 24-hour clock. The selection that you make in the Time zone field determines the time zone for this entry.</td></tr>
<tr><td>Day</td><td>Day on which the scheduled job should run.
If Run is set to Weekly, select the day of the week. For example, select Wednesday.
If Run is set to Monthly, select the day of the month. For example, select 25 for the 25th day of the month.
This field appears only if you select Monthly or Weekly in the Run field.</td></tr>
<tr><td>Repeat Interval</td><td>Duration of the repeat interval for each scheduled job execution. Enter the duration in the number of days, hours, or minutes. For example:
To run the scheduled job every four days, enter <code>04</code> in the Days field.
To run it every 26 hours, enter <code>26</code> in the Hours field.
If it should repeat at an interval of 13:30:25, enter <code>13</code> in the Hours field, and then enter <code>30</code> and <code>25</code> in the two unlabeled fields after it.</td></tr>
</table>

Field	Description
Starting	Date and time of the first scheduled job generation. Select the calendar date and time. This field appears only if you select **Periodically** in the **Run** field.
Business Calendar	Business calendar entry that you are using to determine the business calendar start or end date for the scheduled job. This field appears only if you select **Business Entry: Start Date** or **Business Entry: End Date** in the **Time** field.
Offset type	Type of time offset, if any, to apply to the business calendar that you selected for scheduling this job: ◦ **Past**: Apply an offset factor to schedule the job to run before the start of the time span of the selected business calendar. ◦ **Future**: Apply an offset factor to schedule the job to run after the end of the time span of the selected business calendar. ◦ **--None--**: Do not apply a time offset when scheduling this job. Adding an offset factor enables you to schedule the job to run before, or after, the formal time span that is defined in the business calendar for the following use cases:
Offset	Amount of time offset, expressed in days, hours, minutes, and seconds, to apply to the business calendar that you selected for scheduling this job. For example, if you want to schedule the job to start three days, 14 hours, 10 minutes, and 45 seconds before the business calendar start date, do the following actions: ◦ Select **Past** in the **Offset type** field. ◦ Enter **3** in the **Days** field. ◦ Enter **14**, **10**, and **45** in the **Hours** field. The **Offset Days** and **Hours** fields appear only if you select **Business Calendar: Entry Start** or **Business Calendar: Entry End** in the **Run** field, and **Past** or **Future** in the **Offset type** field.
Priority	Numerical priority for the scheduled job: ◦ Set essential jobs to a priority value below 100. ◦ Set nonessential jobs to a priority above 100. ◦ If 70 percent or more of all scheduled jobs are **Overdue**, any jobs that are marked with a value above 100 do not run.
Conditional	When turned on, specify a scripted condition for generating the report.
Omit if no records	When turned on, empty reports are not distributed.
Condition	Conditional script that determines if a scheduled job should run. The last expression of the script should evaluate to a Boolean (true/false) value. This text box appears only if you select **Use conditions**. For more information about scripts on the ServiceNow platform, see Scripts.

Field	Description
	Warning: Conditional scripts for scheduled report emails are executed in the sandbox. Therefore, function definitions are not allowed. Some API calls and keywords are also not allowed. For more information, see Script sandbox property.
Subject	Text that appears in the subject line of the distribution email.
Introductory message	(Optional) Add a message to the report. If the report output type is Embedded PNG, ◦ Use the tag `${report:png}` in the message body to position the report in the message. Otherwise, the report appears at the bottom of the message. ◦ Use the tag `${report:include_with}` to position other reports included with the email. Otherwise, these reports appear at the bottom of the message.
Type	Report output type. Graphical reports can be PNG or PDF files. List reports can be PDF, Excel, or CSV files. When scheduling a graphical report, select the output type **PDF** or **PDF-landscape** to include the chart grid data. When scheduling a List report, select output type **Excel** or **CSV**. Select **Embedded PNG** to embed the report visualization in the body of the email. Large images are scaled to fit the email. For more information, see the section Report output formats. **Note:** It is only possible to schedule multilevel pivot reports in PDF output.
Zip output	When turned on, the report is sent as a zip file attachment to the email.
Include with	One or more additional reports to include with the email. It is not possible to order the reports within the email. You can use the Search icon to locate the additional reports.
Page size (Multilevel pivot report only)	Select from A3, A4, Letter, or Legal size. To specify the dimensions for a different paper size, select **Custom** and enter the **Page height** and **Page width** in pixels.
Page height (in pixels) (Multilevel pivot report only)	Shows when **Page size** is set to **Custom**. For non-standard paper sizes, multiply the page height in inches by 72 and enter the value in this field.
Page width (in pixels) (Multilevel	Shows when **Page size** is set to **Custom**. For non-standard paper sizes, multiply the page width in inches by 72 and enter the value in this field.

Field	Description
pivot report only)	

5. Click **Submit**.

What to do next

To edit the scheduled email of a report job, open the job from the Scheduled Jobs list. To stop scheduled emails of a report, delete the relevant job's row from the Scheduled Jobs list.

Related topics

Create a scheduled job

Personalize the system date format

Set a system time zone

Report output types

You can export reports in certain output formats. You can schedule these reports for regular export.

Report output formats

Report format	Description
PDF	Generate a PDF in portrait or landscape orientation. PDF reports include the chart grid data. Map reports cannot be exported to PDF format. **Note:** To export Multilevel pivot tables to PDF, you must enable the Webkit HTML to PDF (com.snc.whtp) plugin.
Excel	Report visualization shows as a Microsoft Excel (XLS) spreadsheet.
PNG	Report visualization shows as a Portable Network Graphic (PNG) file. Also see how to embed reports as images in a scheduled report email.
CSV	Report visualization shows as a comma-separated value (CSV) plain-text file.

"Run as" user and report emails

If you are on a domain-separated instance, scheduled reports are generated using the "Run as" user's permissions.

By default, the user who sets up the scheduled email report is the "Run as" user. This user's domain access and report_view ACL permissions apply when generating the scheduled reports.

If the data in the scheduled report is different than expected, the user who set up the scheduled report is not an appropriate "Run as" user. An admin or report_admin can specify a different "Run as" user.

Warning:
The "Run as" user requires the appropriate report_view ACL permission to schedule the report. There may also be a security risk in setting this user. Consider consulting Customer Service and Support if setting the "Run as" user.

Related topics

Report_view access control

Automatically generate something from a template

Schedule the generation of entities, which include changes, incidents, and CIs.

Before you begin

Roles required: template_scheduler

Procedure

1. Navigate to Scheduled Entity Generations [sysauto_template] table.
2. Click **New**.
3. Populate the following fields:

Schedule entity generation

Field	Description
Name	Name that identifies this scheduled job.
Active	Option that indicates that scheduled job is active and should be executed at the specified date and time.
Application	Name of the application that contains the entity. **Global** appears if the entity is in the global scope.
Conditional	Option for enabling the running of the scheduled job if certain conditions are met in the associated script.
Condition	Conditional script that determines if a scheduled job should run. The last expression of the script should evaluate to a Boolean (true/false) value. This text box appears only if you select **Use conditions**.
Run	Time interval to use for running the scheduled job: ◦ **Daily**: Runs daily, at a designated time. ◦ **Weekly**: Runs on a weekly basis, at a designated time and day of the week. ◦ **Monthly**: Runs on a monthly basis, at a designated time and day of the month. ◦ **Periodically**: Runs on a designated repeating interval. ◦ **Once**: Runs for a single occurrence only. ◦ **On Demand**: Runs immediately on demand. ◦ **Business Calendar: Entry Start**: Runs on the starting entry dates for the business calendar that you select in the **Business Calendar** field. A scheduled job runs for the starting date of each of the business entries that you defined for the business calendar. For example, if the business calendar represents a fiscal year, and the starting date of each entry is a fiscal month, the scheduled job runs on the first day of each month. ◦ **Business Calendar: Entry End**: Runs for the ending date for the business calendar that you select in the **Business Calendar** field. This selection runs in

Field	Description
	the same manner as **Business Calendar: Entry Start**, but for the end dates of the associated business calendar entries. **Note:** When you select **Business Calendar: Entry Start** or **Business Calendar: Entry End**, you can apply an offset factor to schedule the job to run before or after the time span of the selected business calendar. If you select a business calendar and all its business calendar spans are in the past, the job would never run as this is an invalid configuration. To learn more, see the **Offset type** and **Offset** fields. To learn more about creating and using business calendars and defining business calendar entries, see Creating business calendars andDefine business calendar entries.
Time	Time of day at which the scheduled job should run, expressed in hours, minutes, and seconds on a 24-hour clock. The selection that you make in the **Time zone** field determines the time zone for this entry.
Time zone	Time zone to use with the **Time** field entry when you specify the time at which the scheduled job should run. Select a time zone entry: ◦ **-None-**: Use the default time zone for the logged-in user who is creating the scheduled job. For example, the scheduled job runs at 04:45 p.m. US/Pacific time if it is the user's assigned time zone, and you enter `16:45` into the **Time** field. ◦ **Use System Time Zone**: Use the default system time zone that is specified for the instance in which it runs. For example, the scheduled job runs at 10:15 p.m. London time if Europe/London is the default system time zone for the instance, and you enter `22:15` in the **Time** field. ◦ Actual time zone. For example, the scheduled job runs at 1:30 p.m. in the US Eastern time zone if you select **US/Eastern**, and enter `13:30` in the **Time** field.
Day	Day on which the scheduled job should run. ◦ If **Run** is set to **Weekly**, select the day of the week. For example, select Wednesday. ◦ If **Run** is set to **Monthly**, select the day of the month. For example, select 25 for the 25th day of the month. This field appears only if you select **Monthly** or **Weekly** in the **Run** field.
Repeat Interval	Duration of the repeat interval for each scheduled job execution. Enter the duration in the number of days, hours, or minutes. For example: ◦ To run the scheduled job every four days, enter `04` in the **Days** field. ◦ To run it every 26 hours, enter `26` in the **Hours** field. ◦ If it should repeat at an interval of 13:30:25, enter `13` in the **Hours** field, and then enter `30` and `25` in the two unlabeled fields after it.

Field	Description
Starting	Date and time of the first scheduled job generation. Select the calendar date and time. This field appears only if you select **Periodically** in the **Run** field.
Business Calendar	Business calendar entry that you are using to determine the business calendar start or end date for the scheduled job. This field appears only if you select **Business Entry: Start Date** or **Business Entry: End Date** in the **Time** field.
Offset type	Type of time offset, if any, to apply to the business calendar that you selected for scheduling this job: ◦ **Past**: Apply an offset factor to schedule the job to run before the start of the time span of the selected business calendar. ◦ **Future**: Apply an offset factor to schedule the job to run after the end of the time span of the selected business calendar. ◦ **--None--**: Do not apply a time offset when scheduling this job. Adding an offset factor enables you to schedule the job to run before, or after, the formal time span that is defined in the business calendar for the following use cases:
Offset	Amount of time offset, expressed in days, hours, minutes, and seconds, to apply to the business calendar that you selected for scheduling this job. For example, if you want to schedule the job to start three days, 14 hours, 10 minutes, and 45 seconds before the business calendar start date, do the following actions: ◦ Select **Past** in the **Offset type** field. ◦ Enter **3** in the **Days** field. ◦ Enter **14**, **10**, and **45** in the **Hours** field. The **Offset Days** and **Hours** fields appear only if you select **Business Calendar: Entry Start** or **Business Calendar: Entry End** in the **Run** field, and **Past** or **Future** in the **Offset type** field.
Generate this	Reference to the template record from which you are generating a scheduled job.

Related topics

Create a scheduled job

Personalize the system date format

Set a system time zone

Automatically run a script of your choosing

Schedule both conditional and non-conditional scripts. If Domain Separation is installed in the instance, you can also select, filter, sort, and schedule scripts based on their assigned domains.

Before you begin

Roles required: admin

About this task

The following is an example of a conditional script. It runs the scheduled job only if there are active incidents older than 30 days.

```
// Only run this Scheduled Job if there are active Incidents
 over 30 days old
var ga = new GlideAggregate('incident');
ga.addAggregate('COUNT');
ga.addQuery('active', 'true');
ga.addQuery('sys_created_on', '<', gs.daysAgo(30));
ga.query();
ga.next();
ga.getAggregate('COUNT') !== '0'
```

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Select **New**.
3. Select **Automatically run a script of your choosing**.
4. On the form, fill in the fields.

Schedule script execution

Field	Description
Name	Name that identifies this scheduled job.
Active	Option that indicates that scheduled job is active and should be executed at the specified date and time.
Application	Name of the application that contains the script. **Global** appears if the script is in the global scope.
Conditional	Option for enabling the running of the scheduled job if certain conditions are met in the associated script.
Condition	Conditional script that determines if a scheduled job should run. The last expression of the script should evaluate to a Boolean (true/false) value. This text box appears only if you select **Use conditions**. **Warning:** Conditional scripts for scheduled report emails and Performance Analytics data collection jobs are executed in the sandbox. Therefore, function definitions are not allowed. Some API calls and keywords are also not allowed. For more information, see Script sandbox property. After upgrade, jobs with conditional scripts that contain these disallowed API components finish with errors.
Run	Time interval to use for running the scheduled job: ◦ **Daily**: Runs daily, at a designated time. ◦ **Weekly**: Runs on a weekly basis, at a designated time and day of the week. ◦ **Monthly**: Runs on a monthly basis, at a designated time and day of the month. ◦ **Periodically**: Runs on a designated repeating interval. ◦ **Once**: Runs for a single occurrence only. ◦ **On Demand**: Runs immediately on demand.

Field	Description
	◦ **Business Calendar: Entry Start**: Runs on the starting entry dates for the business calendar that you select in the **Business Calendar** field. A scheduled job runs for the starting date of each of the business entries that you defined for the business calendar. For example, if the business calendar represents a fiscal year, and the starting date of each entry is a fiscal month, the scheduled job runs on the first day of each month. ◦ **Business Calendar: Entry End**: Runs for the ending date for the business calendar that you select in the **Business Calendar** field. This selection runs in the same manner as **Business Calendar: Entry Start**, but for the end dates of the associated business calendar entries. **Note:** When you select **Business Calendar: Entry Start** or **Business Calendar: Entry End**, you can apply an offset factor to schedule the job to run before or after the time span of the selected business calendar. If you select a business calendar and all its business calendar spans are in the past, the job would never run as this is an invalid configuration. To learn more, see the **Offset type** and **Offset** fields. To learn more about creating and using business calendars and defining business calendar entries, see Creating business calendars andDefine business calendar entries.
Day	Day on which the scheduled job should run. ◦ If **Run** is set to **Weekly**, select the day of the week. For example, select Wednesday. ◦ If **Run** is set to **Monthly**, select the day of the month. For example, select 25 for the 25th day of the month. This field appears only if you select **Monthly** or **Weekly** in the **Run** field.
Repeat Interval	Duration of the repeat interval for each scheduled job execution. Enter the duration in the number of days, hours, or minutes. For example: ◦ To run the scheduled job every four days, enter `04` in the **Days** field. ◦ To run it every 26 hours, enter `26` in the **Hours** field. ◦ If it should repeat at an interval of 13:30:25, enter `13` in the **Hours** field, and then enter `30` and `25` in the two unlabeled fields after it.
Time zone	Time zone to use with the **Time** field entry when you specify the time at which the scheduled job should run. Select a time zone entry: ◦ **-None-**: Use the default time zone for the logged-in user who is creating the scheduled job. For example, the scheduled job runs at 04:45 p.m. US/Pacific time if it is the user's assigned time zone, and you enter `16:45` into the **Time** field. ◦ **Use System Time Zone**: Use the default system time zone that is specified for the instance in which it runs.

Field	Description
	For example, the scheduled job runs at 10:15 p.m. London time if Europe/ London is the default system time zone for the instance, and you enter `22:15` in the **Time** field. ◦ Actual time zone. For example, the scheduled job runs at 1:30 p.m. in the US Eastern time zone if you select **US/Eastern**, and enter `13:30` in the **Time** field.
Time	Time of day at which the scheduled job should run, expressed in hours, minutes, and seconds on a 24-hour clock. The selection that you make in the **Time zone** field determines the time zone for this entry.
Day	Day on which the scheduled job should run. ◦ If **Run** is set to **Weekly**, select the day of the week. For example, select Wednesday. ◦ If **Run** is set to **Monthly**, select the day of the month. For example, select 25 for the 25th day of the month. This field appears only if you select **Monthly** or **Weekly** in the **Run** field.
Repeat Interval	Duration of the repeat interval for each scheduled job execution. Enter the duration in the number of days, hours, or minutes. For example: ◦ To run the scheduled job every four days, enter `04` in the **Days** field. ◦ To run it every 26 hours, enter `26` in the **Hours** field. ◦ If it should repeat at an interval of 13:30:25, enter `13` in the **Hours** field, and then enter `30` and `25` in the two unlabeled fields after it.
Starting	Date and time of the first scheduled job generation. Select the calendar date and time. This field appears only if you select **Periodically** in the **Run** field.
Business Calendar	Business calendar entry that you are using to determine the business calendar start or end date for the scheduled job. This field appears only if you select **Business Entry: Start Date** or **Business Entry: End Date** in the **Time** field.
Run as [Optional]	Select another user to run the script execution as. Configure the form to add this field if it is not present.
Run this script	Name of the script to run at the scheduled date and time. For example, copy script logic from a business rule, or call a script include.

5. If Domain Separation is installed in this instance, the following fields also appear for selecting, sorting, and scheduling jobs by domain.

Note:
The Domain Support - Domain Extensions Installer plugin installs these fields. To learn more, see Request domain separation.

Field	Description
Domain iterator	Check box that enables this job to run across multiple domains. If selected, the **Domain source table** and **Domain source filter** fields appear.

Field	Description
Domain source table	Name of the domain-separated table that is the source of the domains in which the scheduled job should run. For example, select **Department (cmn_department)** if you want to use it as the source of the assigned domains used to schedule scripts. The scheduled job determines these domains from the sys_domain field of the source table records. To learn more, see Domain assignment. **Note:** It's a good idea to create a domain-separated table solely dedicated to sourcing domains, unless you have a use case that warrants using an existing one. If you do create a new one, simply select it using this field.
Domain source filter	Optional condition you can specify to filter the records queried from the selected domain source table. **a.** Select the field to use for filtering records. **b.** Add filtering conditions and operators to narrow the domain selection criteria. To learn more, see Filters.

The following processing takes place when scheduling a job in a domain-separated instance:

- It first queries the selected domain source table.
- If you specified a condition in the **Domain source filter** field, it applies the filtering condition to narrow down the resulting set of records.
- Finally, it retrieves its list of unique domains from the sys_domain columns of these records.

6. Click **Update** to update the scheduled job, **Execute Now** to execute the scheduled script immediately, or **Delete** to delete the job.

Related topics

Create a scheduled job

Personalize the system date format

Set a system time zone

Domain separation for service providers

Special cases in job schedules

Some special cases require care in job scheduling.

End of the month schedules

Because months have different lengths, take care when scheduling jobs for the end of the month.

- Scheduling an event for the 29th or 30th is not recommended, because the scheduled job is executed in months (like February) which are shorter than those dates.
- If an event is scheduled for the 31st, it executes on the last day of the month, even if the month is shorter.

For example, something scheduled to run on the 31st of the month runs on February 28 or February 29 in a leap year.

Weekday schedules

For scheduled scripts, use the following script to run only on weekdays:

Warning:
Conditional scripts for scheduled report emails are executed in the sandbox. Therefore, function definitions are not allowed. Some API calls and keywords are also not allowed.

```
var isWeekday;
var today = new Date();
var dayOfWeek = today.getDay(); // Get day of the week(0 =
 Sunday, 1 = Monday, ... , 6 = Saturday)
switch (dayOfWeek) {
case 0: // Sunday
case 6: // Saturday
isWeekday = false;
break;
default:
isWeekday = true;
}
answer=isWeekday;
```

If the conditional script on a Scheduled Report is more complex and you need to make use of our Glide classes, then please use the following steps as a workaround:

1. Create a Scheduled Job and complete all the conditional logic of the Scheduled Report inside the Scheduled Job.
2. If all the conditional logic is satisfied, you can trigger the Scheduled Report with the following script inside the Scheduled Job.

```
var schRpGr = new GlideRecord("sysauto_report");
schRpGr.get("<sys_id of the scheduled report>");
gs.executeNow(schRpGr);
```

Scheduled jobs from scripts

To execute a scheduled job triggered by an event, use the following script:

```
//Execute a scheduled script job
var rec = new GlideRecord('sysauto_script');
rec.get('name', 'YOUR_JOB_NAME_HERE');
SncTriggerSynchronizer.executeNow(rec);
```

You can run the script using one of several tables:

- scheduled_import_set (Scheduled Import Sets)
- sysauto_script (Scheduled Script Execution)
- sysauto_template (Scheduled Template Generation)
- sysauto_report (Scheduled Report)

Note:
`SncTriggerSynchronizer` does not provide methods to execute scheduled jobs in the future.

Running scheduled scripts and jobs imported from another instance

Scheduled Script Execution and Scheduled Jobs are categorized as data records in the ServiceNow AI Platform, which means they are excluded from update sets. To determine if a table is included, navigate to **All** > **System Definition** > **Dictionary** and view the **Attribute** value for the table in question. Only tables with the attribute update_sync=true are included in update sets.

You can optionally export and import data records via XML. However, to prevent unexpected data changes, Schedule Item [sys_trigger] records are not created for data records imported from an XML file such as an update set. To run scheduled jobs or scripts imported from another instance, update the corresponding job or script record in the target instance.

View a schedule item

You can view scheduled items, which are individual instances of a scheduled job, from the system scheduler.

Before you begin

Role required: admin

About this task

Procedure

Navigate to **All** > **System Scheduler** > **Today's Scheduled Jobs**.

> **Note:**
> It is inadvisable to modify the schedule items themselves. It is best to modify the scheduled jobs that contain them.

Each schedule item appears on the table listing.

Related topics

Scheduled jobs

Create a scheduled job

Event scheduling

You can use various tools for scheduling actions or tasks to happen in the future.

Maintenance schedules

Changes to the CMDB can be managed through the Maintenance Schedules plugin, which enables changes to be proposed and viewed through a timeline.

On-call rotation

The Group On-Call Rotation plugin enables a schedule to be defined to determine what users are primary contacts during particular hours of the day.

Scheduled reports

Once you define reports, you can schedule them for emailing at a specific time, or at regular intervals, using the reporting interface.

Scheduled workflows

Workflows provide a robust system for automating advanced multi-step processes. Workflows are triggered by conditions, like business rules, or they can be scheduled for a particular time/recurring schedule, like scheduled jobs.

Scheduled jobs

Scheduled jobs are scripts you can specify should automatically generate at a specific date and time, or on a repeating basis.

Event registry

You can use events to schedule actions or tasks to occur when conditions are fulfilled.

For example:

kb.view

An event triggered when a user views a Knowledge Base article. You use it to trigger the script action Knowledge View that create a Knowledge Use record every time an article is viewed.

incident.commented

An event triggered when a user comments on an article, used by two incidents commented email notifications.

Inactivity Monitors

If a record has not been updated for a defined length of time, an inactivity monitor triggers an event.

Script Actions

Scripts that trigger when an event is recorded in the log. You can set scripts to execute whenever a particular activity occurs in the platform, rather than at a particular time, such as scheduled jobs, or in response to particular conditions, such as Business rules.

Notifications

You can also use events to trigger Email Notifications when an event is recorded in the log.

Register an event

You can register an event for a specific table and a business rule that fires the event.

Before you begin

Role required: admin

Procedure

1. To register an event, browse to **System Policy** > **Events** > **Registry**, and then click **New**.
2. On the form fill in the fields.

Event Registration Form Completion

Field	Description
Event name	Name of your new event.
Table	Database table for this event. **Note:** Only tables and database views that are in the same scope as the event appear on the listing.
Queue	Name of the queue that the event is placed into when triggered.

Field	Description
	Use only lowercase letters, no spaces, and no special characters except underscore (_). For example, `my_queue`. See Using custom queues to process events.
Priority	Order in which messages will be processed. Lower values have higher priority. **Note:** The com.glide.sysevent.priority.enabled system property is enabled by default. If you disable it, events processing is not done by priority.
Caller Access	Restricted caller access settings. ◦ Caller Restriction: calls to the resource must be manually approved. ◦ Caller Tracking: Calls to the resource are automatically approved.
Fired by	Name of the business rule that runs the event. This field is for reference only and is not used by any process. Make sure that there is enough information to locate your event again.
Description	Short description of the purpose of the event.

Set an inactivity monitor

You can set an inactivity monitor that triggers an event for a task record when the task has been inactive for a certain amount of time.

Before you begin

Role required: admin

About this task

If the task remains inactive, the monitor repeats at regular intervals. User updates to the task record restart the monitor. If reset conditions you defined for it have not been met, the monitor does not restart when you update the task record. Inactivity monitors only apply to records on tables that extend the `Task` table, or the `Task` table itself.

When an inactivity monitor triggers, it generates an event in the form `<tablename>.inactivity` (for example, `incident.inactivity`). Define an Email and SMS notifications or script action to drive further action for the inactivity monitor.

A record's activity is only based on user updates. System updates do not count as activity.

Inactivity monitors only support one matching inactivity monitor per task.

If multiple inactivity monitors do match a task record, one of the inactivity monitors will delete all other sys_trigger jobs, leaving only its own, preventing the others from triggering. Since inactivity monitors are driven by condition builder, the definition time can't identify inactivity monitors from intersecting with each other per task record.

Procedure

1. Navigate to **All** > **System Policy** > **SLA** > **Inactivity Monitors** and click **New**.
2. Give the inactivity monitor a name.
3. Specify the type of record to monitor in the **Table** field.
4. Specify how long the inactivity monitor should wait before sending each notification in the **Wait** field.

5. Specify any additional conditions in the **Condition** field.
   At least one condition must be specified for the inactivity monitor to work.
   - If you are using incident auto-close, add a Reset condition because, if you don't, auto-close won't work properly. To prevent this reset, set a **Reset Condition** on your inactivity monitor. To learn more about Incident auto-close, see Configure incidents to close automatically.
   - All conditions need to be True per user update for the Inactivity Monitor to start running. For example, Set this condition only works if the due date is earlier than the current date when user updates the task.
6. Specify an Order if multiple inactivity monitors might have their conditions met for a given record - the one with the lowest order is used.
7. Click **Save**.

   Note:
   If you change conditions on an inactivity monitor, the monitor stops tracking previously tracked records. It does not track records that you created before the inactivity monitor, even if the record meets all other conditions.

Escalation intervals and pause conditions

Escalation Intervals and Pause Conditions are not relevant to an inactivity monitor.

The related list and field are available because the inactivity monitor table extends the table used for SLAs, but these elements are not used in any way when an inactivity monitor attaches or is triggered.

Using time configuration

Handle scheduling, set up events, track metrics, and more with this feature.

Use time configuration apps for everything from basic time and task tracking, to calendars for business and fiscal events, all customizable to any given time zone.

Creating business calendars

Create business calendars in the ServiceNow AI Platform so that you can schedule events for your users and track their time across applications. You define schedules and calendars of all types, including fiscal calendars, in the Business Calendar menu.

Business calendar components

Use the following forms when you define calendars and schedules in the Business Calendar menu:

Business calendar definition record

Identifies an individual business calendar where you can specify a reference name, labeling, parent business calendar (if any), and a time zone. A business calendar can represent any time period, and its definition record is saved in the Business Calendar [business_calendar] table.

To learn more about creating a business calendar definition, see Create a business calendar.

Business calendar entries

Defines the individual entries that make up a business calendar, and saves them in the Business Calendar Entry [business_calendar_span] and Business Calendar Entry Name (business_calendar_span_name) tables. In contrast to legacy schedule

entries, which have repeating and excluding time spans, you must create individual business calendar entries for all time spans in the calendar.

To learn more about creating entries for a business calendar, see Define business calendar entries and Create business calendar entry names.

Business calendar filter options

Defines the date and time filtering options for the business calendar that you are creating. These filtering options are in addition to the standard date and time filtering in the ServiceNow AI Platform.

To learn more about creating filtering options for a business calendar, see Create a business calendar.

Calendars for package

Pairs a business calendar with a ServiceNow AI Platform package so that you can enable the calendar for use in that functional area. A package is a collection of scopes and applications that make up a functional area, such as HR Service Delivery.

To learn more about associating a business calendar with a package, see Pair business calendars with packages.

> **Warning:**
> If you are using Performance Analytics, use caution when changing time spans. Any change you make to the time span on a business or fiscal calendar can invalidate Performance Analytics data. Consult with a Performance Analytics administrator before making such changes.

Fiscal calendar example

In this example, you create a fiscal calendar with 4 quarters and 12 fiscal periods.

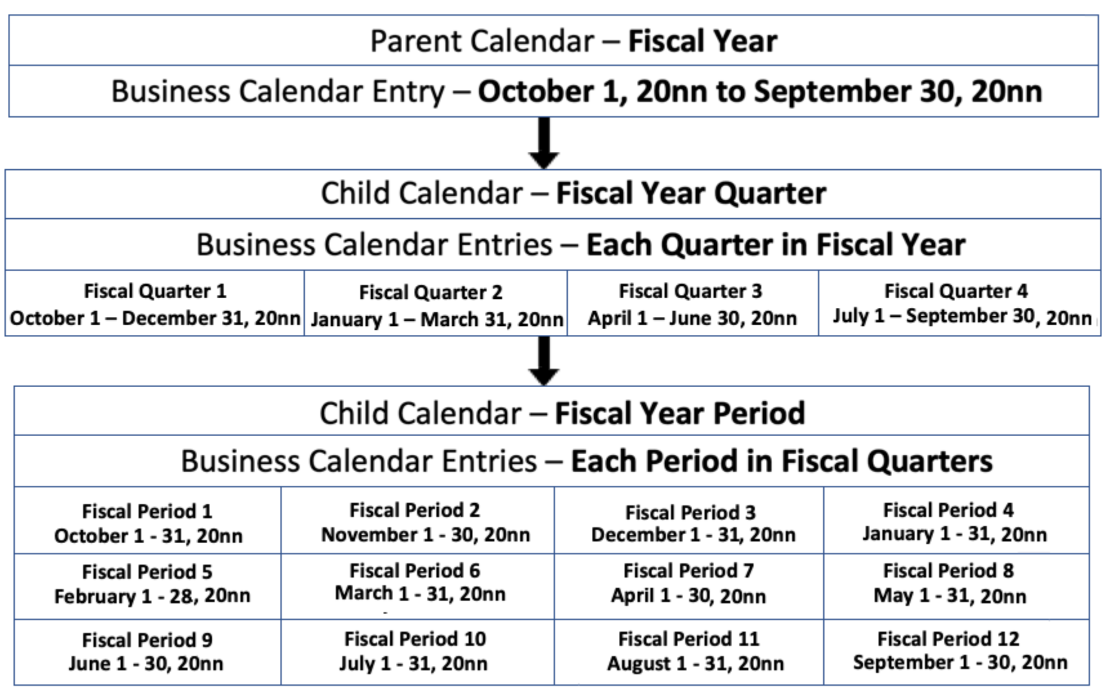

- When you navigate to **Business Calendar** > **Business Calendars**, you create:
  - A business calendar with a title of Fiscal Year.
  - A calendar for quarters that belong to the fiscal year.
  - A calendar for periods with a parent of Fiscal Year or Fiscal Quarter.

  When you create the Fiscal Year Quarter child calendar, you designate that the Fiscal Year is its parent calendar. For the child Fiscal Year Period calendar, you designate Fiscal Year Quarter as its immediate parent.
- When you create each calendar, use the Business Calendar Entries related list, or navigate to **Business Calendar** > **Business Calendar Entries** to create business calendar entries for it. For each business calendar entry, create a reference name, for example, Fiscal Period 1, and designate start and end date times for each time span within the calendar.
  - For example, the parent Fiscal Year calendar has a business calendar entry for the time span of October 1, 20nn 12:00 a.m. through September 30, 20nn 23:59:59 p.m.
  - The child Fiscal Year Quarter calendar consists of four business calendar entries, each with the start and end dates of each quarter that makes up the fiscal year.

    Note:
    You can also create 13 period, 445, 454 and 544 types of fiscal calendars. To learn more about these types of fiscal calendars, see Defining fiscal calendars.
  - The child Fiscal Year Quarters calendar consists of multiple business calendar entries, with the start and end date times for each period in the fiscal year.
- Use **Business Calendar** > **Calendars for Package** to pair the parent Fiscal Year 20nn calendar with ServiceNow AI Platform packages, such as HR Service Delivery or Finance Close Automation.

Business calendar schedule example

In this example, you create a work schedule for the month of December.

- The work schedule consists of a parent calendar, a child calendar for the working weeks in December, and another child calendar for the working days in December.
- The December Workdays calendar consists of business calendar entries for each working day. You do not include non-work days such as Saturdays, Sundays, and company holidays.

Parent Calendar – **December Work Schedule**
Business Calendar Entry – **December 1 to December 31**

Child Calendar – **December Work Weeks**				
Business Calendar Entries – **Each Working Week in December**				
Week 1 December 1-7	**Week 2 December 8-14**	**Week 3 December 15-21**	**Week 4 December 22-28**	**Week 5 December 29-31**

Child Calendar – **December Workdays**				
Business Calendar Entries – **Each Working Day in December**				
Mon December 2 8 am – 5 pm	**Mon December 9 8 am – 5 pm**	**Mon December 16 8 am – 5 pm**	**Mon December 23 8 am – 5 pm**	**Mon December 30 8 am – 5 pm**
Tues December 3 8 am – 5 pm	**Tues December 10 8 am – 5 pm**	**Tues December 17 8 am – 5 pm**	**Tues December 24 8 am – 5 pm**	**No Entry – New Year's Eve**
Wed December 4 8 am – 5 pm	**Wed December 11 8 am – 5 pm**	**Wed December 18 8 am – 5 pm**	**No Entry - Christmas**	
Thur December 5 8 am – 5 pm	**Thur December 12 8 am – 5 pm**	**Thur December 19 8 am – 5 pm**	**Thur December 26 8 am – 5 pm**	
Fri December 6 8 am – 5 pm	**Fri December 13 8 am – 5 pm**	**Fri December 20 8 am – 5 pm**	**Fri December 27 8 am – 5 pm**	

- Use **Business Calendar** > **Business Calendars** to create a parent business calendar with a title of December Work Schedule, and a child business calendar for December work weeks. When you create the child calendar, you designate that the December Work Schedule is its parent calendar.

  Note:
  You can make the December Work Schedule even more granular by creating a subordinate child calendar for December workdays. In this case, its immediate parent is the December Work Weeks calendar.

- When you create each calendar, use the Business Calendar Entries related list, or navigate to **Business Calendar** > **Business Calendar Entries** to create at least one business calendar entry for it. For each business calendar entry, you give it a reference name, and designate the start and end date times for each time span within the calendar.
  - For example, the parent December Work Schedule calendar is for a time span of December 1, 20nn 12:00 a.m. through December 31, 20nn 23:59:59 p.m.
  - The child December Work Week calendar consists of five business calendar entries. Each entry is for the start and end dates of each calendar week in December.
  - The child December Workdays calendar consists of multiple business calendar entries. Each entry is for the start and end date times for each working day in December. You do not create business calendar entries for company holidays such as New Year's Eve or New Year's Day.
- Use **Business Calendar** > **Calendars for Package** to pair the parent business calendar December Work Schedule with the appropriate ServiceNow AI Platform packages, such as Customer Service Management.

Using business calendars in place of legacy calendar and schedule definitions

The legacy functions that you may have used, on the **System Scheduler** > **Schedules** and Fiscal Calendar menus, to define schedules and fiscal calendars, continue to be supported. However,

you should use the Business Calendar menu in their places to define calendars and schedules and then associate them with specific ServiceNow AI Platform packages.

Note:
For Strategic Portfolio Management (SPM) applications, you must define Fiscal Calendars using the Generate a fiscal calendar feature.

- If you use the legacy functions to create fiscal calendars and schedules, the resulting definition records are internally flagged as belonging to the legacy [cmn_schedule] table hierarchy. These records appear on the lists on the Business Calendar menu, as well as on the lists on the respective menus for the legacy functions used. For example, if you defined fiscal calendars using legacy fiscal calendar functions, the resulting records appear on lists on the Business Calendar Fiscal Calendar menus.
- If you select a legacy fiscal calendar or schedule from these lists, you can only modify it using the legacy function that you originally used to create it. For example, if you select a legacy fiscal calendar, you can only modify it by using the functions on the Fiscal Calendar menu. You can't modify it with the Business Calendar menu functions.

Related topics

Using schedules and calendars

Create a business calendar

Create a definition record for an individual parent or child business calendar. You create this record so that you can easily identify a specific business calendar when you use it in a ServiceNow AI Platform package or functional area. A business calendar can represent any length of time, such as a week, a month, a year, or a fiscal period.

Before you begin

Role required: admin

About this task

When you create a business calendar, you specify a reference name, labeling, parent business calendar, if any, and a time zone for it.

1. You first create a business calendar record and then create business calendar entries for it in the Business Calendar Entries related list. To learn more about how business calendar entries work, see Define business calendar entries and Create business calendar entry names.
2. Using the Business Calendar Filter Options related list, you can also define the date and time filtering options for the business calendar that you are creating.

Note:
Fiscal calendars or schedules you may have created in the **Fiscal Calendar** or **System Scheduler** > **Schedules** menu also appear on the lists on the **Business Calendar** menu. If you select a fiscal calendar or schedule, it appears in the legacy function that you originally used to create it. You can't view or update legacy fiscal calendars or schedules using this form.

Procedure

1. Navigate to **All** > **Business Calendar** > **Business Calendars**.
2. Click **New**.
3. On the form, fill in the fields.

Business calendar form

Field	Description
Name	Unique reference name and identifier for the business calendar record. The name should be unique per domain.
Label	Unique singular label for identifying the business calendar when you perform filtering. For example, you can assign a week-long work calendar with a label of `Workweek`.
Plural Labels	Unique plural label for identifying the business calendar. For example, you can assign a week-long work calendar with a label of `Workweeks`. Use plural labels when you do filtering that is based on more than one unit in the past or future. For example, Last 2 Fiscal Quarters.
Parent	Parent business calendar to the current business calendar that you are creating. To search for an existing business calendar, click ◦ For example, you created a business calendar for 2020, and an individual child calendar for each month within the year, each with its own business calendar entries. As you create a business calendar for each month, you would select 2020 as the parent calendar for each. ◦ These entries create the relationship between a parent and its child business calendars. Leave this field blank if you are creating a parent business calendar itself.
Application	Application that is associated with this business calendar. If you are working on an application or are creating a business calendar from an application record, the field defaults to that application. Otherwise, the field defaults to **Global**. Any records that are created from the table record, such as modules and security rules, are assigned to this application by default. To learn more about application scoping, see Application scope
Time zone	Time zone to use for the business calendar when you specify the starting and ending times for each associated business calendar entry. Select a time zone entry: ◦ **-None-**: Use the time zone specified in the logged in user's profile. ◦ **Use System Time Zone**: Use the default system time zone specified for the instance. ◦ **– Floating --**: Time zone is relative to whatever is accessing the business calendar entry at any given time. For example, if an administrator in Amsterdam creates a floating business calendar for 08:00 a.m. to 5:00 p.m., a person in San Jose, California, sees it as 08:00 a.m. to 5:00 p.m. When you define a business calendar in a specific time zone, users in different time zones see the business calendar in their own time zone. ◦ Actual time zone. For example, Europe/London, US/Pacific, US/Central.
Description	Description of the business calendar.

4. Click **Submit**.
5. In the Business Calendars list, select the business calendar that you created.
6. In the Business Calendar Entries related list, create entries that define each span of time within the business calendar.

Create at least one business calendar entry for the selected calendar.

Note:
To learn more about creating business calendar entries, see Define business calendar entries and Creating business calendars.

7. In the Business Calendar Filter Options related list, click **New** to create filtering options for the business calendar.
   You can define the date and time filtering options for the business calendar that you are creating. These filtering options are in addition to the standard date and time filtering in the ServiceNow AI Platform.

   Note:
   To learn more about creating business calendar filtering options, see Define business calendar filtering options.

8. Click **Submit**.

Result
A business calendar record is saved in the Business Calendar [business_calendar] table.

- If you create business calendar entries in the Business Calendar Entries related list, records are created in the Business Calendar Entry (business_calendar_span) table. The records are associated with the business calendar.
- If you created business calendar filter options in the Business Calendar Filter Options related list, a record is saved in the Business Calendar Filter Option (business_calendar_filter_option) table. The record is associated with the business calendar.

Create a business calendar group

For Performance Analytics and Reporting to use business calendars, organize several calendars into a business calendar group. An application might also contain a business calendar group.

Before you begin
Role required: admin

About this task
You can base a Performance Analytics indicator or a report on business calendars instead of the standard calendar. The business calendars themselves are too granular in scale, and have to be grouped. The report or indicator then uses that group. Work with your Performance Analytics team to design the business calendar groups.

Note:
Fiscal calendars are also supported. When you generate a fiscal calendar, a corresponding business calendar group is also generated automatically.

Procedure

1. Navigate to **All** > **Business calendars** > **Business calendar group**.
2. Select **New**.
3. Name the group and save it.
   The **Business calendars** related list appears.
4. In the **Business calendars** related list, select **Edit**.
5. Choose from existing business calendars to add to the group.
   You can also create a new business calendar from this list.

> **Important:**
> Do not select business calendars that cover a period shorter than a 24-hour day for use with Performance Analytics indicators. Performance Analytics does not store times, only dates. If only reports will use the group, you can select a shorter period.

6. Select **Update** when finished.

Example: Gregorian calendar group

The Gregorian Calendar group comes with the base system. It includes the Week, Month, Quarter, and Year business calendars, which also come with the base system.

Gregorian Calendar group

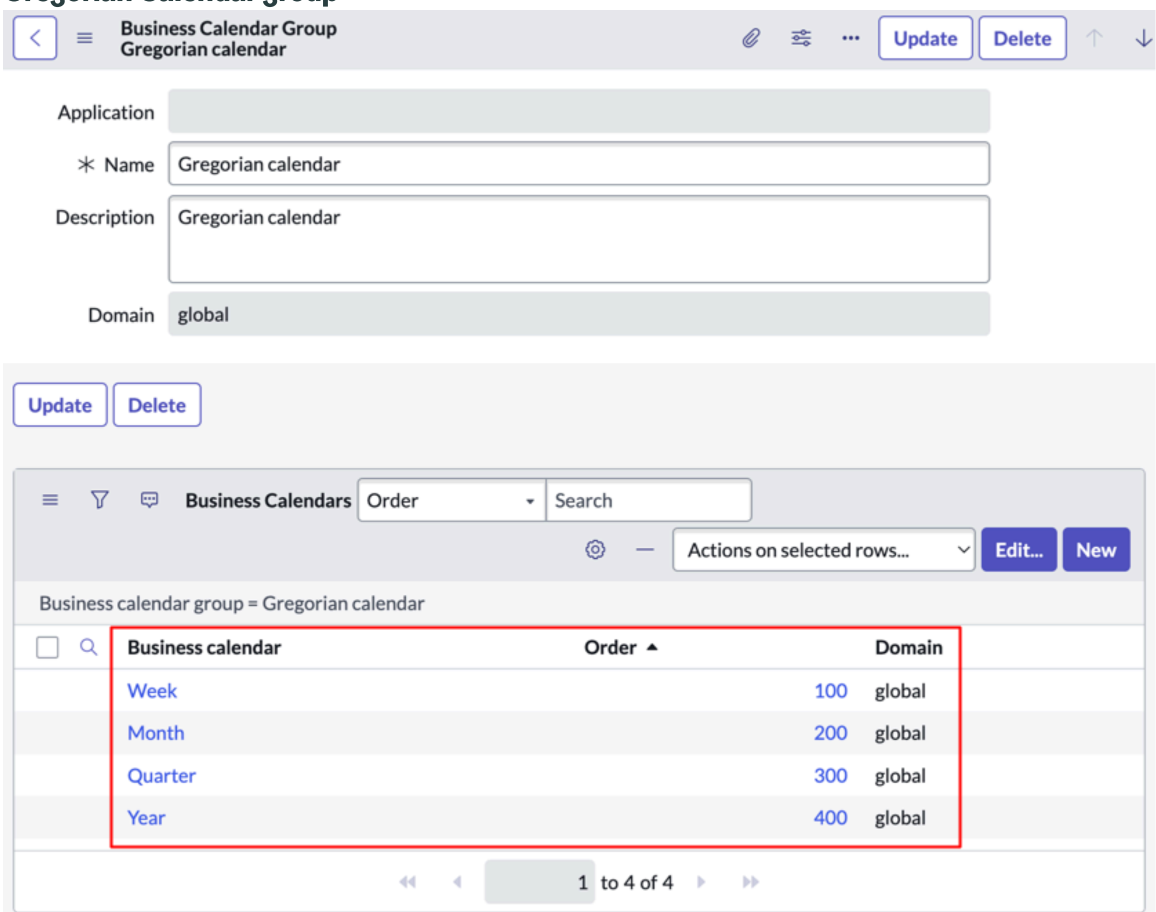

Each business calendar record also shows which groups it belongs to.

servicenow

The 'Week' business calendar

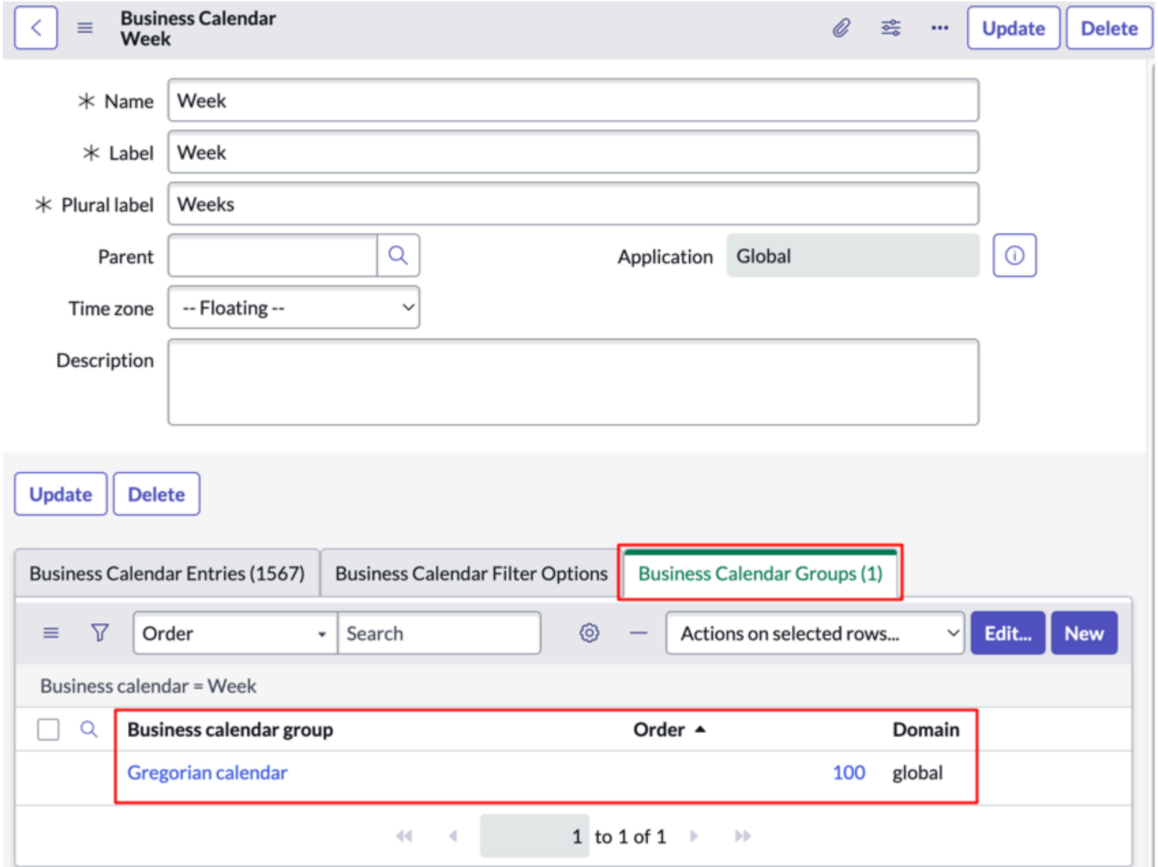

Related topics

"Indicators with business calendars" in Create an automated indicator

Define business calendar entries

Create business calendar entry records that define the individual time periods that make up a business calendar. You create a business calendar entry for every individual time span in the calendar.

Before you begin

Role required: admin

About this task

In contrast to legacy schedule entries that have repeating and excluding time spans, you have to create individual business calendar entries for every business calendar time span. For example, if you are creating a Fiscal Year 2020 business calendar with 13 fiscal periods, you have to create individual business calendar entries for each fiscal period. For each fiscal period, you specify start and end date times.

Note:
Fiscal calendars or schedules you may have created in the **Fiscal Calendar** or **System Scheduler** > **Schedules** menu also appear on the lists on the **Business Calendar** menu. If you select a fiscal calendar or schedule, it appears in the legacy function that you originally used to create it. You can't view or update legacy fiscal calendars or schedules using this form.

Note:
The easiest way to create business calendar entries and associate them with a business calendar is to use the Business Calendar Entries related list in the Business Calendar form. To learn more about creating a business calendar, see Create a business calendar.

Procedure

1. Navigate to **All** > **Business Calendar** > **Business Calendar Entries**.
2. Click **New**.
3. On the form, fill in the fields.

Business calendar entry form

Field	Description
Name	Unique reference name for the business calendar entry. To select a business calendar name or create a new one: **a.** To search for an existing business calendar entry name, click . **b.** If no business calendar entry name exists and you want to create a new business calendar entry name, click **New**. The Business Calendar Entry Name form appears. **Note:** To learn how to create business calendar entry names, see Create business calendar entry names.
Start	Start date and time for the business calendar period. It uses the time zone that you specified in the business calendar record that you associate with it. For example, if you are defining a fiscal period for the month of January, the starting date time is January 31, 2020 12:00 a.m.
End	End date and time for the business calendar period. It uses the time zone that you specified in the business calendar record that you associate with it. For example, if you are defining a fiscal period for the month of January, the ending date time is January 31, 2020 11:59 p.m.

4. Click **Submit**.

Create business calendar entry names

Create a name for an individual business calendar entry. You create this record so you can easily identify a specific business calendar entry when you use it in a business calendar.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **Business Calendar** > **Business Calendar Entry Names**.

   Note:
   Alternately, you can create a business calendar entry name at the time you create the associated business calendar entry. To learn more about creating a business calendar entry, see Define business calendar entries.

2. Click **New**.
3. On the form, fill in the fields.

Business calendar entry name form

Field	Description
Short name	Abbreviated name for the business calendar entry. For example, January.
Long name	Longer name for the business calendar entry. For example, January fiscal period.
Label	Unique label for identification of the business calendar entry when you perform filtering.
Calendar	Name of the business calendar that you want the business calendar entry to belong to. This entry creates the relationship between a business calendar and the business calendar entries that define each individual time span with it. For example, if you are creating a business calendar entry for the January fiscal period, you select FY20 to associate it with Fiscal Year 2020. To select an existing business calendar: **a.** Click or **b.** Click inside the **Calendar** field to view and select from a listing of recent business calendars that are created in **Business Calendar** > **Business Calendars**. **Note:** If you attempt to use the and see a dialog box with **Cancel** and **Leave** options, click **Cancel** and then click inside the **Calendar** field to select from a listing of recent business calendars.
Application	Application that is associated with this business calendar. If you are working on an application or are creating a business calendar from an application record, the field defaults to that application. Otherwise, the field defaults to **Global**. Any records that are created from the table record, such as modules and security rules, are assigned to this application by default. To learn more about application scoping, see Application scope

4. Click **Submit**.

Result

Records are created in the Business Calendar Entry (business_calendar_span) and the Business Calendar Entry Name (business_calendar_span_name) tables. If you create business calendar entries in the Business Calendar Entries related list in the Business Calendar form, these records are also associated with the specific business calendar.

Define business calendar filtering options

Define date and time filtering options for a business calendar. These filtering options are in addition to the standard date and time filtering in the ServiceNow AI Platform. You can provide your users with additional filtering options when they use the business calendar with ServiceNow AI Platform packages and functional areas.

Before you begin

Role required: admin

About this task

To learn more about standard filtering options, see Values for date/time fields .

Note:
The easiest way to create business calendar filtering options and associate them with a business calendar is to use the Business Calendar Filter Options related list in the Business Calendar form. To learn more about creating a business calendar, see Create a business calendar.

Procedure

1. Navigate to **All** > **Business Calendar** > **Business Calendar Filter Options**.
2. Click **New**.
3. On the form, fill in the fields.

Business calendar filter options form

Field	Description
Period	Number of periods that should be traversed forward, backward, or both, for the business calendar. For example, if the business calendar reflects a week: ◦ Provide a filter option for the current week by entering `0` for the base period. ◦ Provide a filter option for a previous period by entering a negative number that represents the number of periods that should be traversed from the base period. For example, enter `-1` for the last period, or enter `-2` for two periods ago. ◦ Provide a filter option for a future period by entering a positive number. For example, enter `1` for the next period, or enter `2` for the next two weeks. If the business calendar reflects a week: ◦ Provide a filter option for the current week by entering `0` for the base period. ◦ Provide a filter option for a previous period by entering a negative number that represents the number of periods that should be traversed from the base period. For example, enter `-1` for the last period, or enter `-2` for two periods ago. ◦ Provide a filter option for a future period by entering a positive number. For example, enter `1` for next period, or enter `2` for the next two weeks. If the business calendar reflects a fiscal year with 13 periods: ◦ Provide a filter option for the current fiscal year by entering `0` for the base period. ◦ Provide a filter option for a previous fiscal year by entering a negative number that represents the number of periods that should be traversed from the base period. For example, enter `-2` for two fiscal periods ago, or enter `-13` for the last fiscal year. ◦ Provide a filter option for a future fiscal year by entering a positive number. For example, enter `13` for the next fiscal year. If the business calendar reflects a fiscal year with four fiscal quarters: ◦ Provide a filter option for the current fiscal quarter by entering `0` for the base period. ◦ Provide a filter option for a previous fiscal quarter by entering a negative number that represents the number of periods that should be traversed from

Field	Description
	the base period. For example, enter `-1` for the last fiscal quarter, or enter `-2` for the last two fiscal quarters. ◦ Provide a filter option for a future fiscal quarter by entering a positive number. For example, enter `1` for the next fiscal quarter, or enter `2` for the next two fiscal quarters.
Calendar	Name of the business calendar that you want the filter options to belong to. To select a business calendar, or create a new one: **a.** Search for an existing business calendar by clicking . **b.** If no business calendar exists, or if you want to create a new business calendar, click **New**.
Application	Application that is associated with this business calendar. If you are working on an application or are creating a business calendar from an application record, the field defaults to that application. Otherwise, the field defaults to **Global**. Any records that are created from the table record, such as modules and security rules, are assigned to this application by default. To learn more about application scoping, see Application scope

Note:
Within a calendar filtering option record, you can only associate a business calendar with a single filter option. To create multiple filtering options, you have to create multiple filtering option records, each with its own period entry.

4. Click **Submit**.

Result
Records are created in the Business Calendar Filter Option (business_calendar_filter_option) table. If you create business calendar filtering options in the Business Calendar Filter Options related list in the Business Calendar form, these records are also associated with the specific business calendar.

Pair business calendars with packages

Pair a business calendar that you created with a package, or functional area, in the ServiceNow AI Platform. When you associate a business calendar with a package, it enables the use of the calendar in that specific functional area. Until you pair a business calendar definition with at least one package, it remains a definition record only.

Before you begin
Role required: admin

About this task
A package is an application, or generic set of functions, such as HR Service Delivery or Finance Close Automation, in the ServiceNow AI Platform. For example, if you created a fiscal calendar, you can associate it with the Finance Close Automation package.

Note:
Within a single calendar package record, you can only associate a business calendar with a single package. To associate business calendars with multiple packages, you have to create multiple calendar package records. For example, if you created a fiscal calendar and want to associate it with the HR and Fixed Assets packages:

- In the first record you create, in the **Calendar** field, select the fiscal year, and in the **Package** field, select **HR**.
- In the second record you create, you again select the same fiscal calendar, and then select **Fixed Asset** in the **Package** field.

Procedure

1. Navigate to **All** > **Business Calendar** > **Calendars for Package**.
2. Click **New**.
3. On the form, fill in the fields.

Package form

Field	Description
Calendar	Name of the business calendar that you want to associate with a ServiceNow AI Platform package. Click 🔍 to access Business Calendars where you can search for an existing business calendar. If there is no existing calendar, you can create a new one.
Package	Name of the ServiceNow AI Platform package that you want to associate with the selected business calendar. Click 🔍 to access Packages where you can search for an existing package. If there is no existing package, you can create a new one.

4. Click **Submit**.

Result

A record is created in the Calendars for Package (calendars_for_package) table that associates the selected business calendar with a ServiceNow AI Platform package.

Defining fiscal calendars

Create fiscal calendars with defined fiscal periods, for use with Financial Management applications, such as Finance Close Automation. You can define fiscal periods, such as one month long, four weeks long, Q1 2015, or July 2015.

To create a fiscal calendar, activate the Fiscal Calendar [com.snc.fiscal_calendar] plugin. You can then use the legacy functions on the **Fiscal Calendar** menu to generate and validate fiscal calendars and their associated fiscal periods. You can create one of these types:

Calendar type	Description
Regular	A calendar with 12 periods, each of which is one month long.
13 Period	A calendar with 13 periods, each of which is four weeks long. **Note:** This calendar is the only type that doesn't use quarters.

Calendar type	Description
445	A calendar that divides a year into subsequent 4-week, 4-week, 5-week periods. The first and second periods have four weeks, and the last has five weeks. The application creates these periods for the calendar duration.
454	A calendar that divides a year into subsequent 4-week, 5-week, 4-week periods. The first period has four weeks, the second has five weeks, and the third has four weeks. The application creates these periods for the calendar duration.
544	A calendar that divides a year into subsequent 5-week, 4-week, 4-week periods. The first period has five weeks, and the second and third have four weeks.

Important:
After you start working with a set of financial data with a fiscal calendar type, you can't change to another type of fiscal calendar. Verify that your fiscal calendar can be used with your financial data. If you import records into the Fiscal period [fiscal_period] table, you should validate that they're correct and don't contain any gaps.

After the fiscal calendar has been generated, don't change either **Start date time** or **End date time** fields of fiscal period records to avoid validation issues.

Note:

- The legacy functions that you may have used, on the **System Scheduler** > **Schedules** and Fiscal Calendar menus, to define schedules and fiscal calendars, continue to be supported. However, you should use the Business Calendar menu in their places to define calendars and schedules and then associate them with specific ServiceNow AI Platform packages. To learn more about defining and using business calendars, see Creating business calendars.
- For Strategic Portfolio Management (SPM) applications, you must define Fiscal Calendars using the Generate a fiscal calendar feature.

Activate Fiscal Calendar plugin

If you have the admin role, you can activate the Fiscal Calendar plugin (com.snc.fiscal_calendar). This plugin includes demo data and activates related plugins if they are not already active.

Before you begin

Role required: admin

Procedure

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

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

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

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

Related topics

List of plugins (Washington DC)

Generate a fiscal calendar

You can generate a fiscal calendar for a specified time period.

Before you begin

Role required: fiscal_calendar_admin

About this task

Important:
Once you start working with a set of financial data with a fiscal calendar type, you cannot change to another type of fiscal calendar. Verify that your fiscal calendar can be used with your financial data.

Note:
Instead of generating fiscal calendars, can also use the Business Calendar functions to create fiscal calendars. To learn more about creating business calendars and schedules, see Creating business calendars.

Procedure

1. Navigate to **All** > **Fiscal Calendar** > **Generate**
2. In the **Fiscal Calendar Types** section, select the type of fiscal year calendar.

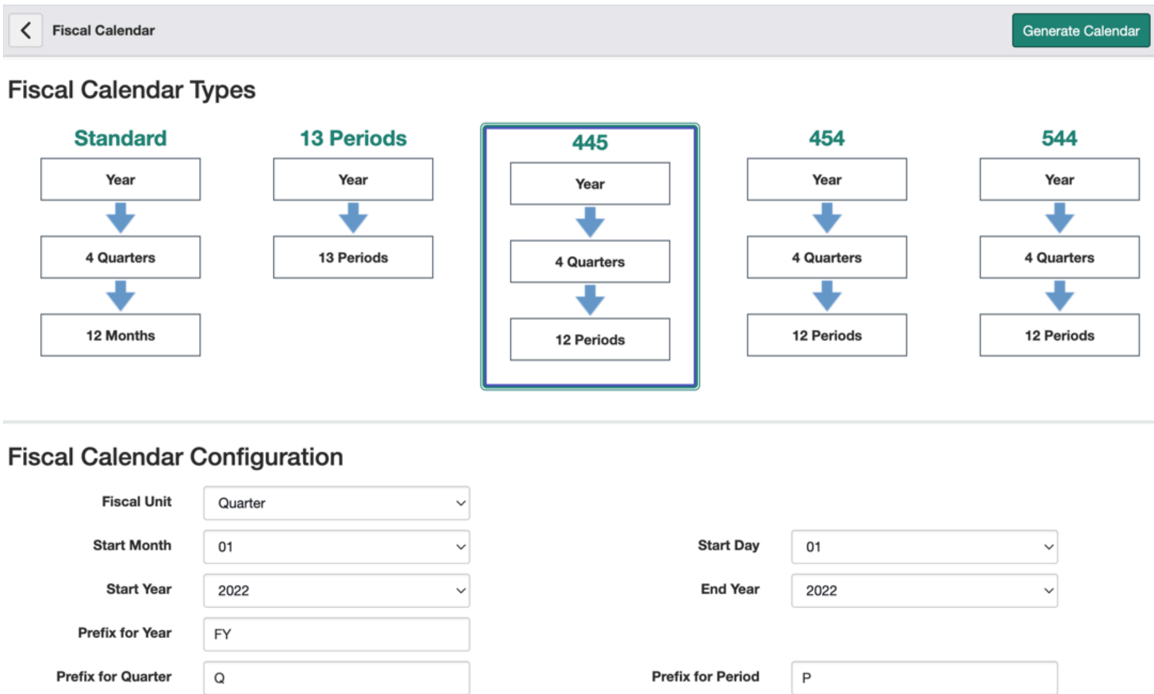

3. In the **Fiscal Calendar Configuration** section, fill in the fields.

Field	Description
Fiscal Unit	Base unit for the calendar. ○ **Month**: Generate fiscal calendar with monthly periods. ○ **Quarter**: Generate fiscal calendar with quarterly periods.
Start Month	Month that is the beginning of your fiscal year.
Start Year	Starting year for the calendar.
Prefix for Year	Prefix that the application uses in the name of the records that represent the fiscal year.
Prefix for Quarter/Period	Prefix that the application uses in the name of the records that represent the fiscal quarter or period.
Start Day	Day that represents the beginning of each month.
End Year	Ending year for the calendar.
Prefix for Month/ Period	Prefix that the application uses in the name of the records that represent the month or period.

4. Click **Generate Calendar**.
5. Navigate to **System Definition** > **Fiscal Periods** and verify that the monthly, quarterly, and annual fiscal period records exist.

View, modify, and validate fiscal periods

After you generate a fiscal calendar, you can view fiscal period records, modify the start and end date, deactivate a fiscal period if necessary, and validate.

Before you begin

Role required: fiscal_calendar_user

Procedure

1. Navigate to **All** > **System Definition** > **Fiscal Periods**.
2. Open any of the fiscal periods (see table for field descriptions) to view, modify, or deactivate.

Fiscal period form

Field	Description
Name	Name of the fiscal period.
Start date time	Date the fiscal period starts.
End date time	Date the fiscal period ends.
Fiscal Type	[Read-Only] Type of period, either **Year** or **Quarter**.

3. Validate the fiscal periods to ensure that there are no gaps and that they match a valid style of calendar.
   a. To validate, return to the list view of fiscal periods and click **Validate Periods**.

Set a system time zone

Set your system time zone with a system property.

Before you begin

Role required: admin

About this task

The ServiceNow AI Platform stores time values in Universal Coordinated Time. Times appear globally based on the system time zone (glide.sys.default.tz), or if specified, based on user settings.

Procedure

1. Navigate to **All** > **System Properties** > **System**.
2. Locate the property **System timezone for all users unless overridden in the user's record** [glide.sys.default.tz].

   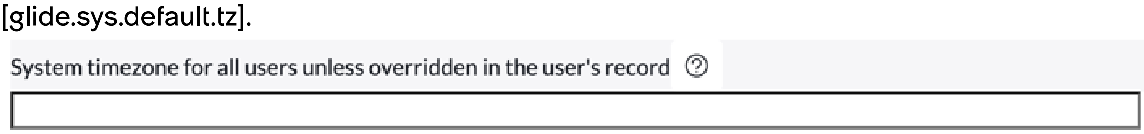

   By default, the field is empty. If you have not defined a time zone for this property, *America/Los Angeles* is the default.

   For more information on setting the default time, see: Configure Next Experience language and region preferences.

3. Add a time zone in the format Country/City or using the name of a time zone (for example, *GMT*) and select **Save**.

   For available values, see Time zone representation.

The new system time zone automatically cascades to all users who don’t already have a specified time zone. If a user selects a different time zone, or if the administrator selects a different time zone for them, the selected time zone is assigned to the user. The system time zone isn’t used.

Change the time zone choice list

Wherever users have a choice of time zone, the choices are populated using the **Time Zone** choice list on the `User [sys_user]` table. Not all time zones appear by default.

About this task

To add or remove time zones from the list of time zones:

Procedure

1. Navigate to **All** > **User Administration** > **Users** and open any user record, or click **New**.
   Notice that the default time zone is **System (America/Los_Angeles)**.
2. Right-click **Time zone**, and then select *Personalize Choices*.
3. Highlight the desired time zone from the **Available** or **Selected** lists, and then **Add** or **Remove** the time zones as needed.

Change a time zone in a scheduled report

By default, scheduled reports use the time zone of the user who runs them (the user in the **Run As** field) to evaluate the query parameters. You are able to change this time zone.

About this task

For example, a report on 'incidents open today' resolves 'today' based on the user's time zone. To manually specify the time zone used to resolve the parameters of the report:

Procedure

1. Navigate to **All** > **Reports** > **Scheduled Reports**.
2. Choose the report with the time zone you want to edit.
3. Add Run as fields to the layout.
   a. Configure the form layout to add the **Run as tz** field to the **Scheduled Report** form.
      See Configuring the form layout for instructions.

      The Run as tz field lets you choose either a the time zone of the user who configured the schedule, or a specific time zone.

   b. **Optional:** Configure the form layout to add the **Run as [+]** field to the **Scheduled Report** form.
      See Configuring the form layout for instructions.

      The Run as [+] field lets you choose a specific user to run the report as.
4. Select the time zone you would like to use.
5. **Optional:** Select the user to run the report as.
6. Click **Submit**.

Change a time zone in a scheduled data import

By default, scheduled data imports are run using the time zone of the user who creates them. However, there is a way to manually specify the time zone for the import.

About this task

To change the time zone of the scheduled data import:

Procedure

1. Configure the form to add the **Run As tz** field to the **Scheduled Data Import** form.
2. Select the appropriate time zone.
3. Click **Update**.

Define a relative duration

You can define a relative duration to work out duration time for SLAs.

Procedure

From the left navigation pane, select **System Scheduler** > **Relative Durations.**.

Four base Relative Durations are available:

- Two business days by 4pm
- Three business days by 4pm
- Next business day by 4pm
- End of next business day

Note:
The business schedule you use defines business day durations. If you define no schedule, by default the durations are 24 X 7.

Example:
Look at the End of next business day Relative Duration. From the Relative Durations list, select the **End of next business day** Relative Duration. The variable *days* is set to one, because the result of this calculation should land one day in the future. The rest of the script is as in the screenshot. If desired, you can customize the time at which you want the Relative Duration to end (currently set to 5pm).

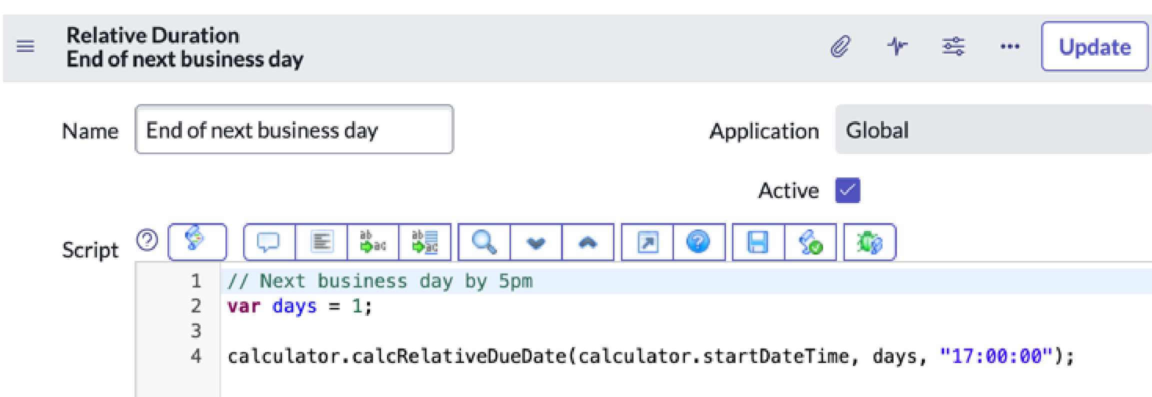

There is one more important Relative Duration design aspect that is used by the other three out-of-box Relative Durations. To illustrate this design lets look at **2 business days by 4pm**.

As you can see in the image, within the script there is an if-statement. This if-statement is checking to see if the calculated time is after 10am. If it is, then an extra day is added to the calculation. Hence the description of `2 business days by 4pm if before 10am`.

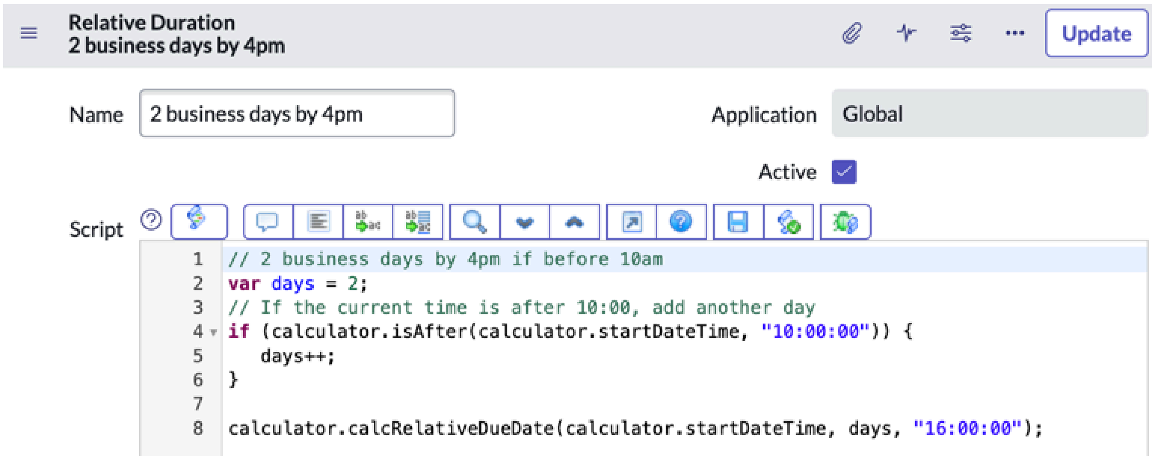

The "End of the business day" has nothing to do with the associated schedule. The end time of 17:00 is hard-coded into this Relative Duration script. If you want the time to be different than 5pm, you must change it in the script.

Use a relative duration

When you define an SLA, you can set the **Duration type** to be a relative duration.

About this task

When using relative durations, you can also use the **Relative duration works on** field. It calculates against the SLA record, or the task record that the SLA record is used for.

This example demonstrates how a relative duration of *End of next business day* works.

Procedure

1. Create an SLA that has a relative duration of **End of next business day**.
2. Complete the rest of the fields of this SLA with the values as shown, also setting a Timezone if you want.

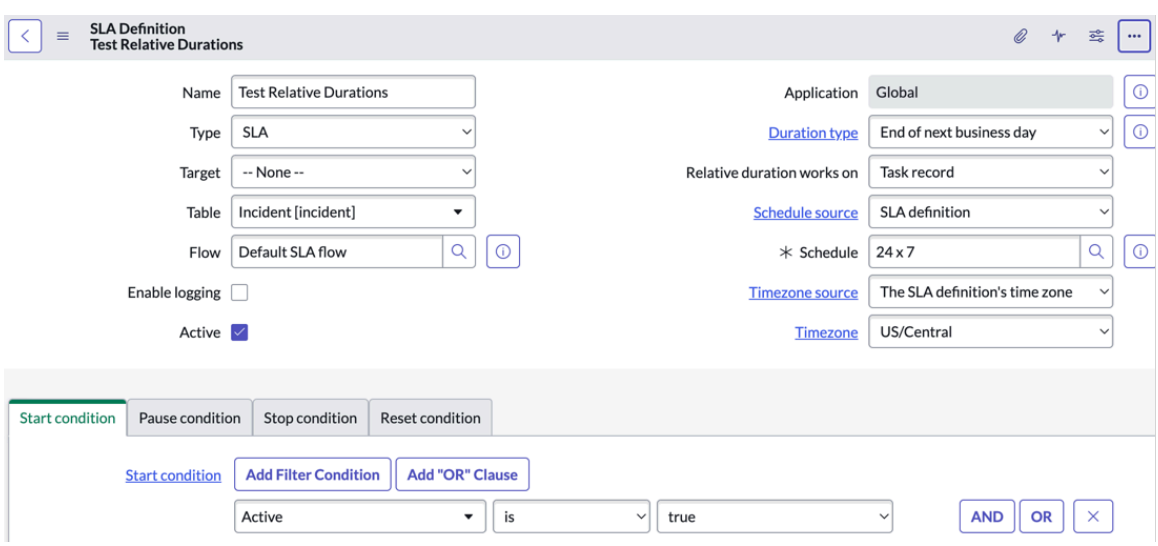

3. To show how this Relative Duration works, create a new incident.

The SLA starts for this incident. In the **Task SLAs** related list, if you look at the **Business time left** and **Start time** fields, notice that the time left is until the next business day at 5pm.

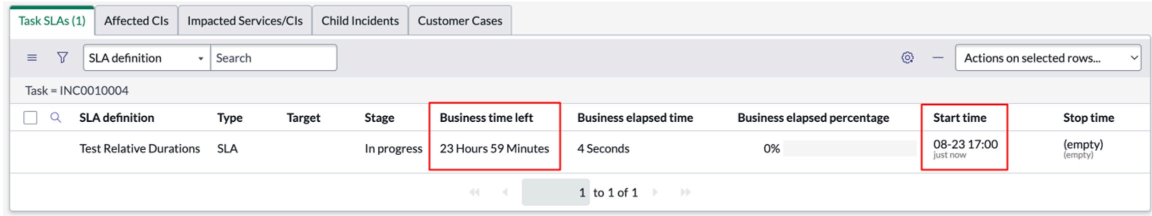

Note:
Pause conditions are not compatible with Relative Durations.

Create schedule entries in the Schedule calendar

Add events such as meeting, time-off, or appointment to display them in the schedule calendar.

Before you begin

Role required: schedule_admin, itil_admin, admin

About this task

You can add and analyze the events using a calendar view. You can also do this in a time line view for a given day, week, or month, which is defined as four weeks in the time line view.

A calendar view displays events for a specified time period. A time line view displays all events for the selected time line in the calendar.

Procedure

1. Navigate to **All** > **System Scheduler** > **Schedules** > **Schedules**.
2. Select a schedule.
3. Select **Show Schedule** related link.
   Using the Next Experience user interface, you can view the schedules in the schedule calendar for all schedule types except the roster type. The schedules of roster type display in the on-call schedule.
4. Create a schedule entry.
   a. Do any of the following to open the **Add event** form:
      - Select a day in the calendar.
      - Drag the span of days for which you want to create the event.

   b. In the **Name** field, enter a name for the event.

   c. From the **Type** menu, select the type of event.

   d. From the **Show as** field, select the status you want to display for this event.

   e. In the **Start** field, select the calendar icon and select the start date for the event.

   f. In the **End** field, select the calendar icon and select the end date for the event.
      If the event is for the entire day, select **All Day**.

g. From the **Repeats** menu, select how often you want this event repeated.

h. Select **Save**.

- The schedule entry displays in the calendar for that time period and is based on the user's time zone.
- You can reschedule an event by dragging and moving it to another date. You can also update that particular occurrence or all occurrences for an event.
- You can click on an event to view the details in a pop-up display and edit or delete the event by selecting the respective icon.
  - When you add or edit an event, the time zone in the event displays the **Timezone** field value from the schedule record.

    Note:
    If the **Timezone** field value is **Floating**, then the time zone isn't displayed.

  - When you add or edit an event, the calendar uses the start and end dates from the time zone of the schedule.

    Note:
    If the **Timezone** field value is **Floating**, then the calendar uses the start and end dates from the time zone of the logged-in user.

Time configuration references

This section provides information on properties of time configuration.

Time zones

All times are stored in Coordinated Universal Time (UTC) and appear globally based on the system time zone. However, times appear to users in their local time zone, according to their user preference settings.

Time zone representation

Time zones that have the Country/City format are primary time zone IDs. Other time zone IDs are links to the primary time zone.

- For example, US/Pacific is a link to the America/Los_Angeles time zone. Both America/Los_Angeles and US/Pacific represent Pacific Standard Time with the same zone offset and Daylight Savings Time (DST) schedule.
- Other than the representation, there is no impact on date and time functionality.

In the absence of a default time zone for the user or the system, JVM reads default time zone information from the server.

- Depending on how the server is configured, it might return the Country/City or link, for example, US/Pacific or America/Los_Angeles.
- Administrators should configure their system with a default time zone `glide.sys.default.tz` to avoid system dependencies. For more information, see Time zone representation.

Daylight Saving Time

In general, if you specify a time zone based on location (for example, **America/Los Angeles**), the system automatically adjusts for daylight saving time.

- If you specify a time zone based on the time zone name (for example, **GMT**), which is discouraged, it does not typically adjust for daylight saving time. There are however exceptions to these guidelines.
- For more detailed information on how daylight saving time is handled in Java, see https://www.iana.org/time-zones and https://en.wikipedia.org/wiki/Tz_database.

User preferences

Once the System Time Zone is defined, users can also select their own time zone from their user form, accessed through **Self-Service** > **My Profile**.

Note:

The System default appears as **System ([name of the default time zone])**. For example, if the System time zone is **America/Los_Angeles**, the user sees **System (America/Los Angeles)**.

Java time zone class

The ServiceNow AI Platform supports all Java time zone class time zones. To view a listing of these time zones:

1. In the **Type filter text** field, enter `sys_user.list`.
2. Open one of the user records in the listing.
3. Right-click the **time zone** field and select **Show Choices**.

   The resulting list is a complete listing of available time zone IDs.

Time zones in email notifications

The date and time stamp of a notification uses the **system time zone** and not the time zone of any recipient.

Note:

The *`glide.email.append.timezone`* property in **System Properties** > **Email** controls whether to append the time zone. If set to true, the system time zone of the instance is appended to any dates or date/times in outbound email messages (for example, 2010-07-02 04:01:14 PST).

Time zones in service level agreements

Service level agreements have different options for which time zone to use. To set a time-zone for SLAs, navigate to the SLA definition and locate the **Timezone source** field:

Time zone for service level agreements

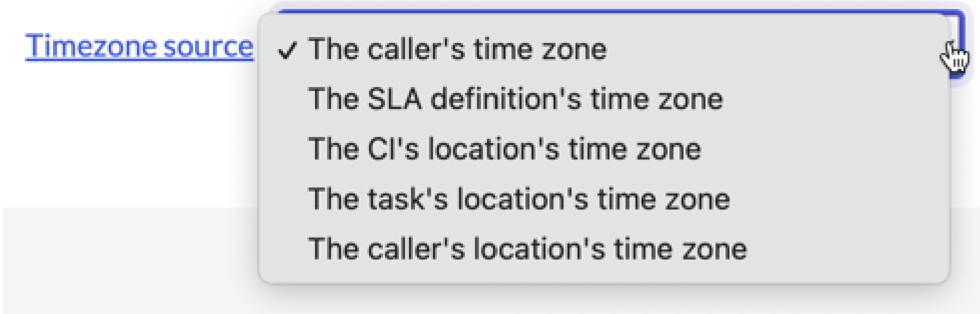

Some special considerations:

- If you select the time zone for the caller, unpredictable behavior can occur if the caller does not have a time zone defined.
- If you select the time zone for the SLA definition, you must manual specify the time zone on the form for the SLA.

Time zone in scripting

When you create scripting on the server, the script uses several `GlideSystem` date and time functions to obtain time values. For more information about specific methods and to learn the format in which each returns the requested time, see the GlideSystem API.

Enhancements

DST enhancement requires that reports and queries observe Daylight Saving Time rules. Changes affect trend charts, line charts, and filters using the "trend on" operation.

Time zone changer

The time zone changer is active by default. Users can change their time zone for the current session in the system settings (the gear on the top right of the banner). At the next login, the time zone setting reverts to the system default, or user profile setting for time zone.

Related topics

- Set a system time zone
- Change the time zone choice list
- Change a time zone in a scheduled report
- Change a time zone in a scheduled data import

Viewing record information over time

You can use three different interfaces to view record information over time: reports and performance analytics, timeline pages, and schedule pages.

Reports and performance analytics

Report example

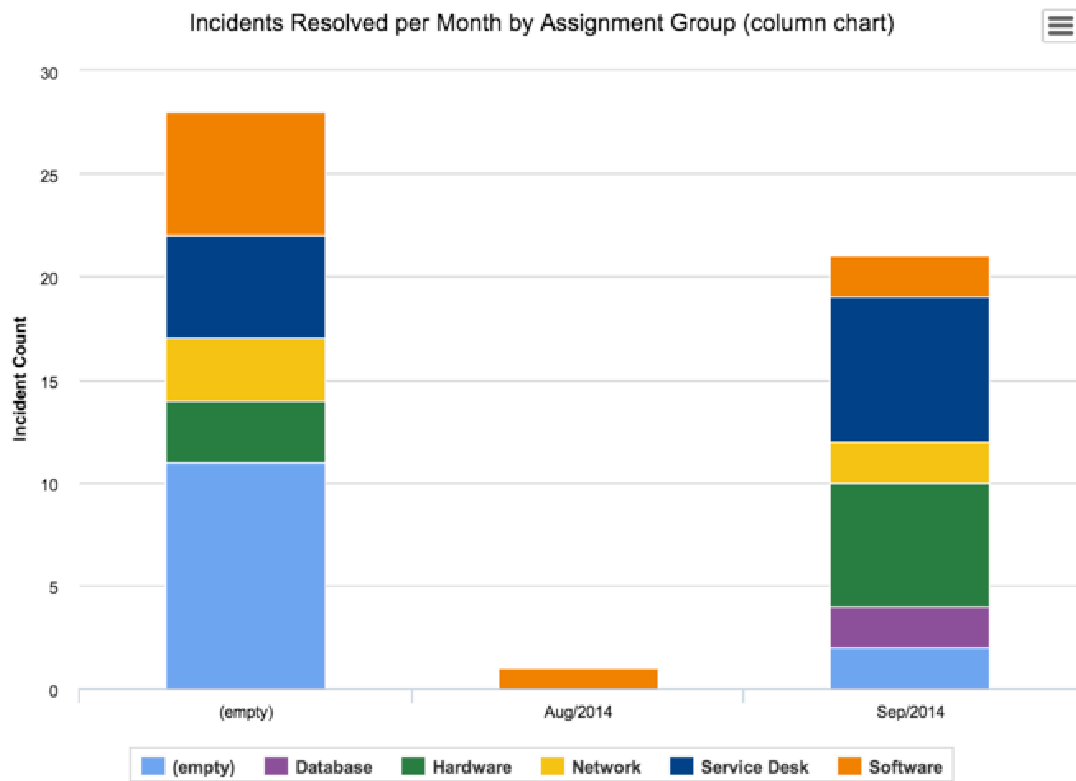

You can visualize data using Reporting.

servicenow

Timeline pages

Timeline pages example

This type of page enables you to define linear time lines from records with time information.

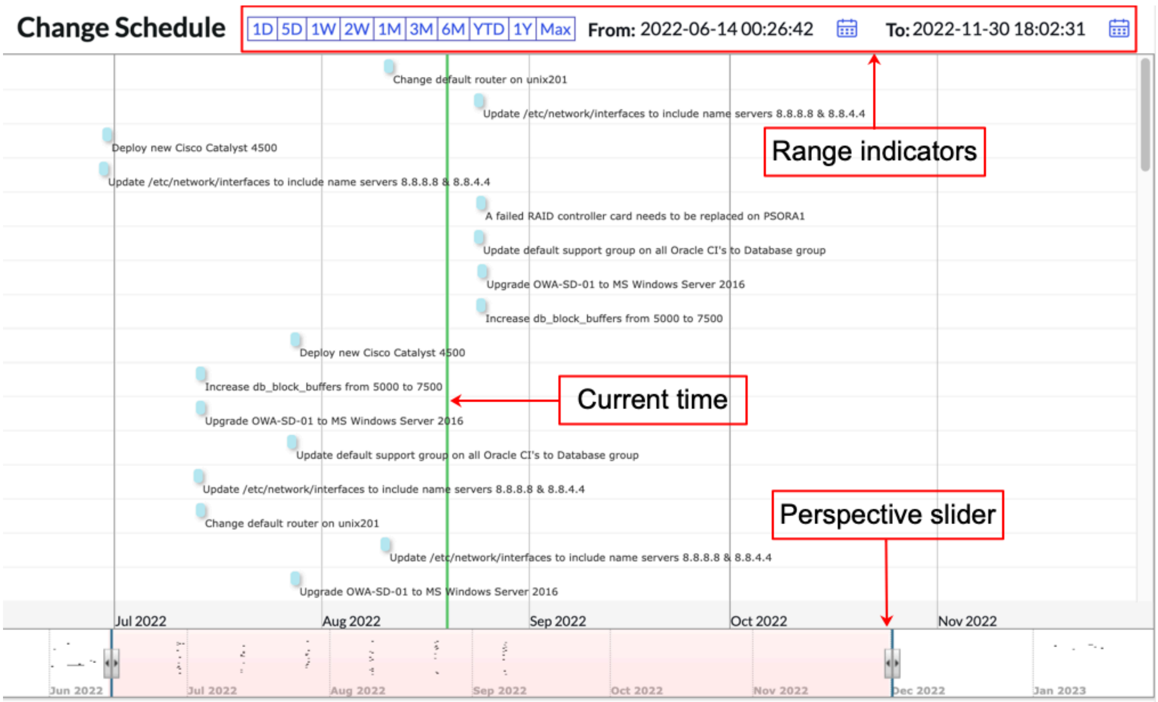

This type of page enables you to define linear time lines from records with time information.

Schedule pages

Schedule pages example

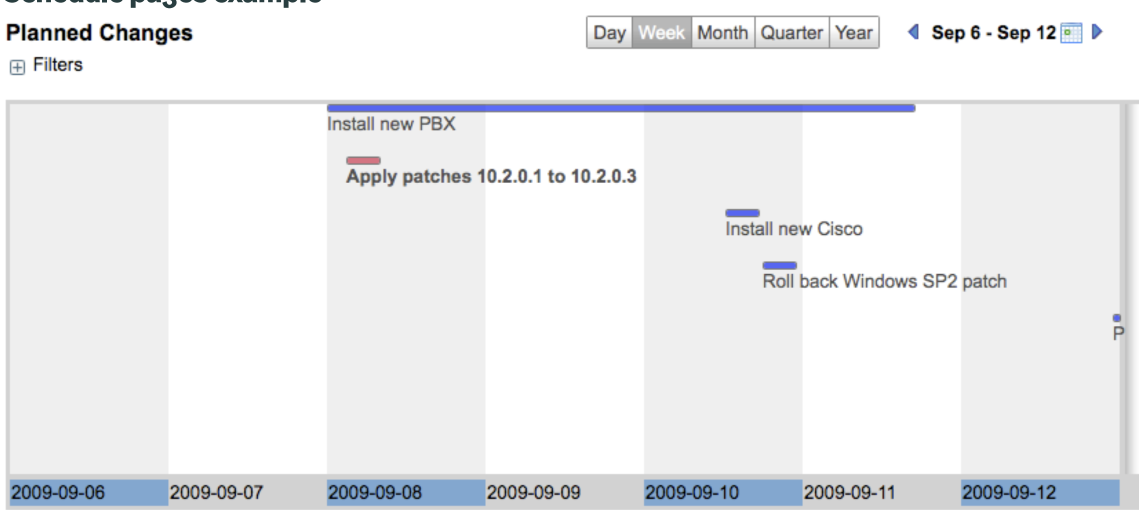

Timeline pages

Use timeline pages to track any activity bounded by two dates, such as change request start and end dates, or incident open and close dates. Administrators can create timeline pages, which are user interactive.

Features

- Make selected timelines available to users by role.
- Select perspective from daily to yearly views.
- Create dynamic labels for timeline spans.
- Configure tooltips for each span.
- Permit span dragging and resizing by users.
- Lock timelines to prevent editing.

Viewing timelines

By default, only administrators have a module to access timeline records. ITIL users can only view timelines by selecting a context menu item from an incident record.

Timeline View

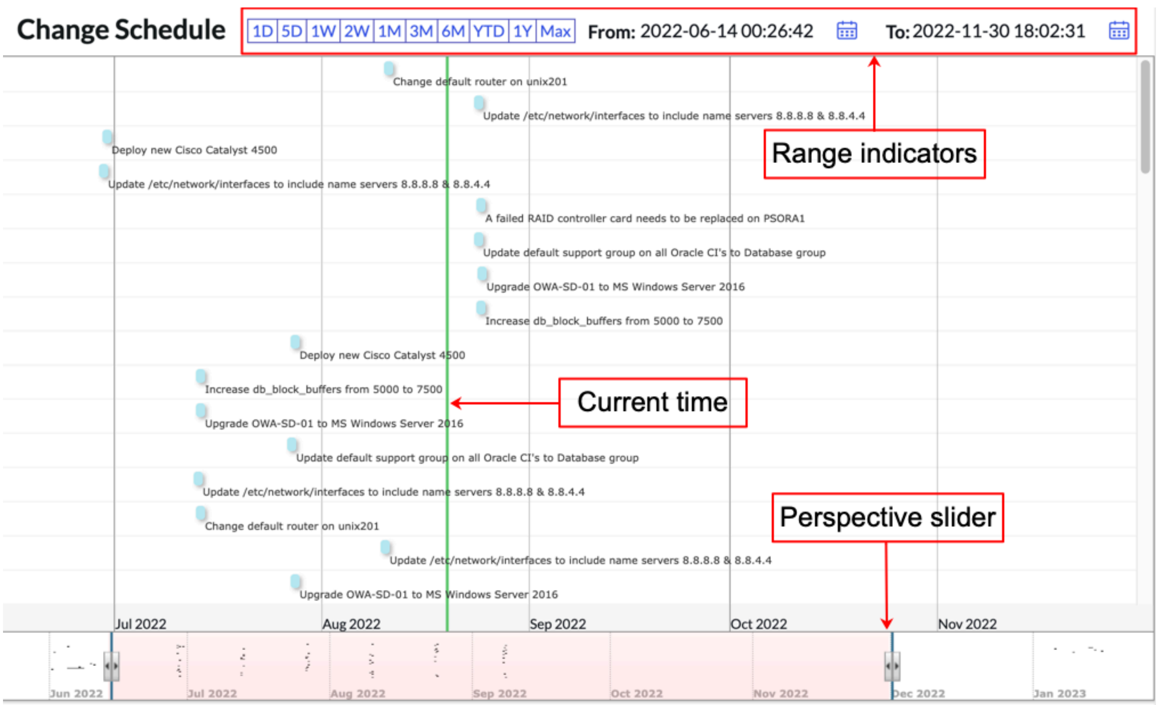

Administrator role users

Users with an administrator role can view timelines from the timeline page records.

1. Navigate to **System UI** > **Timeline Pages**.
2. Select a timeline record to view.
3. Click **View Timeline**.

Other users

By default, timeline metrics only appear for incident records. ITIL users can view timelines associated with any task record where the metrics are set to appear as timeline records. To view the Assigned to duration and the State duration metric timelines, ITIL users can select the **Metrics Timeline** context menu UI action.

1. Navigate to active task record, such as an incident.
2. Right-click the banner, and select **Metrics Timeline**.

To enable non-administrators to view other timelines, create a custom module.

Timing functionality

Timing functionality are tools that exist to answer the question "How Long?"

Related topics

- Time worked fields
- Time configuration SLAs
- Metric definitions

Metric definitions

Defined metrics can track how long an audited field holds a certain value.

For instance, a metric can track how long an incident is assigned to an individual, or how long an incident is in the state *Active*.

See Metrics for more information.

Time configuration SLAs

Service Level Agreements time how long a task meets a certain condition, and is primarily used to ensure that tasks are handled within a pre-determined time limit.

SLAs define the following conditions:

- Start Conditions
- Pause Conditions
- Stop Conditions

Once a task meets the Start Conditions, the SLA will time how long the task remains in that condition (unless it meets Pause Conditions). The timer will end if the Stop Conditions are met. If the time-limit is passed, the SLA will be marked *breached*.

Notifications can be driven off of the SLA to warn interested parties as the time limit approaches.

Time worked fields

The Task `[task]` table provides a time-tracking field called **Time worked**. This field measures how long a record has been viewed in order to measure work time on a ticket.

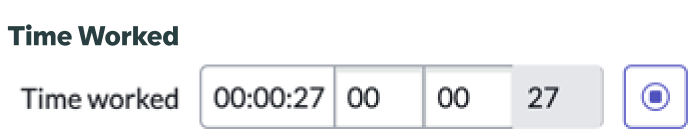

Any table that extends Task can use this field. To add the field, Personalize a form.

As the record is viewed, the timer counts upward. To pause the timer, click the stop icon ( )

To resume the timer, click the start icon ( ).

When the task is saved, the amount of new time in the timer is used to generate a record on the Time Worked [`task_time_worked`] table. This table can be viewed as a related list on the task form.

By default, the time displayed in the **Time worked** field displays a cumulative value stored in the task record. If you modify a Time Worked record, the changes will not be reflected in the task timer.

You can set the property `com.snc.time_worked.update_task_timer` to enable updating of the task timer value based on changes to the time worked records. This is accomplished through the **Update task timer** business rule.

Note:
Timer fields, such as Time Worked, do not display correctly in Service Portal. For more information and a workaround, see the Timer fields do not display the correct data in Service Portal [KB0752404] article in the Now Support Knowledge Base.

Related topics

Timing functionality

Create a timeline page

Create a timeline page to track any activity bounded by two dates.

Procedure

Navigate to **System UI** > **Timeline Pages** and create a new record (see table for field descriptions).

Creating a timeline page

Field	Description
Name	Unique name that describes the function of this timeline. For example, `High Priority Change Requests`.
Table	Name of the table associated with this timeline, such as Change Request `[change_request]`.
Start date field	Time-related field from the specified table to use as the start date for the timeline. The timeline begins with the span for the record with the earliest start date from this field, after it applies the filter and sort order. For example, you can select **Updated** as the start date field. It then starts the span for each active change request on the date it was updated to a high priority.
End date field	Time-related field from the specified table to use as the end date for the timeline. The timeline ends with the span for the record with the latest date from this field, after it applies the filter and sort order. For example, you can select **Closed** as the end date field and display all high priority change requests by the date on which they were closed.
Display Options	
Show grid lines	Select this check box to show horizontal background shading to highlight alternate spans.
Show left pane	Select this check box to show label text in a pane on the left of the timeline. The text that appears in this pane is defined in **Span text fields**.
Show summary pane	Select this check box to show the pink, perspective slider at the bottom of the timeline. Move the slider from right to left to scroll across the chart. Adjust the end points of the slider to change the magnification. A narrow slider zooms in on the spans and provides a more detailed view. A wide slider pulls the view out and makes more of the timeline visible on the screen.

Field	Description
Auto refresh	Select an automatic refresh interval or disable automatic refresh. When auto refresh is disabled, the timeline adjusts only when the browser is manually refreshed or when a start or end date field is updated in a record.
CSS span color	Enter a custom span color using any CSS color format, such as RGB or hexadecimal. If this field is blank, the default span color, light blue, is used. The dark blue color indicates tasks that have a planned start date only, with no planned end date. Adding a planned end date updates the color to light blue.
Show span text	Select this check box to display the content of the **Span text fields** as labels below each span.
Span text fields	Select fields from the specified table to have those values appear as span labels. For example, you might select **Number** and **Short description**. The span labels also appear in the left pane if the left pane is visible.
Show tooltips	Select this check box to display tooltips when the cursor rests on a span.
Tooltip text fields	Select from the specified table the fields whose values appear as tooltips. For example, you might select **Category**, **Assigned to**, and **Due date**.
Filtering and Sorting	
Condition	Create a condition to filter the results that appear in the timeline. For example, a condition that displays only active, high priority incidents. Add the condition count to a condition field to preview what records will be returned by this condition set.
Perform custom sort?	Select this check box to enable custom sorting. Configure the sort order by selecting fields in the **Sort by** and **Sort by order** fields.
Sort by	Select any field in the list for sorting the spans in the timeline. Common practice is to select either the **Start date field** or the **End date field** as the sorting field. If you select a different sorting field, also include that field in the list of **Tooltip text fields** to give users a way of discovering the sort criteria.
Sort by order	Select the sort order for the sorting fields selected.
Interactive Options	
Allow horizontal moving?	Select this check box to permit users to drag timeline spans horizontally. Dragging changes the start and end dates and updates the record.
Allow start time dragging?	Select this check box to permit users to update the record by dragging the start time of a span.
Allow end time dragging?	Select this check box to permit users to update the record by dragging the end time of a span.
Range calculator	Specify a script include that Using DurationCalculator to calculate a due date range restrictions and processes parent updates, if appropriate.

Customize the timeline page span style

The Timeline Page Span Styles related list allows you to define conditional span styles.

Procedure

1. Navigate to **All** > **System UI** > **Timeline Pages**.
2. Open the timeline page for which you want to define the span style.
3. Go to the Timeline Page Span Styles related list and click **New**.
4. Fill in the form and click **Submit**.

Customizing timeline page span styles

Field	Description
Condition	Condition Builder that filters the results that appear in the timeline.
Label color	Color for the text under each span.
Label decoration	Character style for the text under each span: **Bold**, **Italic**, **Underline**, or **Line-through**.
Order	Number to determine the sequence in which the style conditions are evaluated. Style conditions with a lower order are evaluated first.
Span color	Color of each span.
Timeline page	Timeline page to which the span style applies. By default, the span applies to the current timeline.

Timeline sub item

Use the Timeline Sub Items related list to define child spans for the timeline, based on records in a table that references the parent timeline's table.

About this task

You can create a timeline sub item generate a hierarchical relationship, starting from a timeline page to any number of levels. For example, if there is a timeline page for a release, the sub item might be sprints, and the sprint might have stories as a sub item.

Procedure

1. Navigate to **All** > **System UI** > **Timeline Pages**.
2. Open the timeline page for which you want to add a sub item.
3. Go to the Timeline Sub Items related list and click **New**.
4. Fill in the form (see table for field descriptions) and click **Submit**.

Timeline Sub Items

Field	Definition
Parent	[Read-only] Identifies the parent of the timeline sub item.
Name	Enter a unique name that describes the function of this timeline. For example, Sprints for late Priority Changes.
Table	Select the table called by this timeline. The selected table must contain at least one field that references the table you selected for the parent timeline page. For example, if the parent timeline page uses the Scrum Release [rm_release_scrum] table, you might choose the Sprint [rm_sprint] table for a timeline sub item.

Field	Definition
Start date field	Select a time-related field from the specified table to use as the start date for the timeline. For example, Planned start date.
End date field	Select a time-related field from the specified table to use as the end date for the timeline. For example, Planned end date.
Parent Reference Column	Select a reference field on which to base the timeline connection between the sub item records and the parent records. ◦ If multiple reference fields are available, choose the reference field that forms part of the hierarchy modeled by this timeline. ◦ If this list is blank, the sub item table contains no reference fields to the parent table. In this case, you must choose a different table for the sub item. ServiceNow uses the parent reference column to determine which records appear at each level of the timeline.
Display Options	
CSS span color	Enter a custom span color using any CSS color format, such as RGB or hexadecimal. If this field is blank, it uses the default span color, light blue. The dark blue color indicates tasks with a planned start date only, with no planned end date. Adding a planned end date updates the color to light blue.
Span text fields	Select fields from the specified table to have those values appear as span labels. For example, you might select **Number** and **Short description**. If the timeline displays the left pane, span labels also appear in the left pane
Tooltip text fields	Select fields from the specified table to have those values appear as tooltips. For example, you might select **Category**, **Assigned to**, and **Due date**.
Filtering and Sorting	
Condition builder	Create a condition to filter the results that appear in the sub item. For example, you might create a condition that displays only active, high priority incidents.
Perform custom sort	Select this check box to enable custom sorting. Configure the sort order by selecting fields in the **Sort by** and**Sort by order** fields.
Sort by	Select any field in the list for sorting the spans in the timeline. Common practice is to select either the **Start date field** or the **End date field** as the sorting field. If you select a different sorting field, also include that field in the list of tooltip text fields to give users a way of discovering the sort criteria.
Sort by order	Select the sort order for the sorting fields selected.
Interactive options	
Allow horizontal moving?	Select this check box to permit users to drag timeline spans horizontally. Dragging changes the start and end dates and updates the record.
Allow start time dragging?	Select this check box to permit users to update the record by dragging the start time of a span.
Allow end time dragging?	Select this check box to permit users to update the record by dragging the end time of a span.

Field	Definition
Restriction	Specify the behavior when you drag a child span. This option is available only if you did not specify **Range calculator** for the parent timeline page. ◦ **None**: No restriction is in place. ◦ **Restrict by parent**: Move the child span only within the time frame defined by the parent span. ◦ **Update parent**: Update the span when you move the child span outside the time frame defined by the parent span.

Display a metric as a timeline

Administrators can enable users to display any metric on a timeline by activating the *Timeline Metrics* UI action.

Procedure

1. Navigate to **All** > **Metrics** > **Definitions**.
2. Select the metric you want to display on a timeline.
   For example, Problem State Duration.
3. Select the **Timeline** check box.
4. Click **Update**.

Make a timeline visible to a selected user

Create a custom module within an application and defining the roles that can access it, to make selected timelines available to users.

About this task

To permit these users to update task records directly from the timeline, configure the timeline to enable span dragging.

> **Note:**
> Timelines delivered by a custom module are not entirely dynamic. The left pane, summary pane, auto-refresh feature, and the grid lines are not dynamic, and do not reflect updates to the timeline record after creating the module link. However, the data represented by the spans, labels, and tooltips display all updates in the custom module.

To create a timeline page module:

Procedure

1. Click the **Edit application** icon on an application (such as Incident) in the navigation pane.
2. In the application Modules related list, click **New**.
3. Configure the Module form to add the **Timeline page** field.
4. Fill in the form as described in the table and click **Submit**.

   The Module form provides the following fields.

Making timelines visible to selected users

Field	Description
Title	Name of the module as it appears in the navigation pane. For example, you might use Planning Timeline.

Field	Description
Order	Number to determine the sequence in which this condition should be evaluated if more than one matching condition exists. Conditions with a lower order are evaluated first.
Application	Application for the new module.
Hint	Brief description of the module that appears when the user places the cursor over the module name. For example, you might enter Weekly view of high priority changes.
Active	Check box that enables the module for the defined roles. Clear this check box to disable the module for all users.
Image	Icon that should appear with the module name.
Link type	Select **Timeline Page**. When you select this link, the **Timeline Page** field appears.
Timeline Page	Timeline page you want to appear in this module. For example, for the Change application select a change-related timeline, for the Incident application select an incident-related timeline.
Roles	Roles that can access this module.
Override application menu roles	Check box that indicates that this module should appear when the user has the specified roles. Otherwise, the user must have the roles specified by both the application menu and the module.

Range calculator scripts

You can specify a script include that calculates range restrictions and processes parent updates in a timeline page.

Range Calculator Scripts

Following are three examples of script includes that help specify range restrictions.

ExampleUpdateParentsRangeCalculator

Updates parent records when a child record span is moved or resized in the timeline.

```
var ExampleUpdateParentsRangeCalculator = Class.create();
ExampleUpdateParentsRangeCalculator.prototype = {
initialize: function() { },
updateParents: function(id, table, startDate, endDate){
 if (table == "rm_sprint"){
   var releaseID;
   var sprint = new GlideRecord(table);
   sprint.addQuery('sys_id', id);
   sprint.query();
   if (sprint.next())
     releaseID = sprint.release + "";
   if (releaseID) {
         var now_GR = new GlideRecord("rm_release_scrum");
         gr.addQuery("sys_id", releaseID);
         gr.query();
         if (gr.next()) {
           if (startDate  && startDate <
 this.getTimeMs(gr.start_date))
```

```
                gr.start_date = this.getTimeObject(startDate);
            if (endDate && endDate > this.getTimeMs(gr.end_date))
               gr.end_date = this.getTimeObject(endDate);
               gr.update();
         }
    }
  }
},
getMinRangeDetails: function(id, table){ return [-1, -1, "",
 ""]; },
getMaxRangeDetails: function(id, table){ return [-1, -1, ""]; },
getTimeMs: function(date){
  return new GlideScheduleDateTime(date).getMS(); },
getTimeObject: function(timeMS) {
  var gdt = new GlideDateTime();
  gdt.setNumericValue(timeMS);
  return gdt; },
logMessage: function(message){ gs.log(message); },
type: 'ExampleUpdateParentsRangeCalculator'
}
```

In this example, the span is identified based on the id and table from function(id, table, startDate, endDate).

ExampleMinRangeCalculator

Defines the earliest start date and the latest end date for a specified span.

```
var ExampleMinRangeCalculator = Class.create();
  ExampleMinRangeCalculator.prototype = {
      initialize: function() { },
      updateParents: function(id, table, startDate, endDate){ },
      getMinRangeDetails: function(id, table){
           var min = -1;
           var max = -1;
           var minID = "";
           var maxID = "";
           if (table == "rm_release_scrum"){
                    var now_GR = new GlideRecord("rm_sprint");
                    gr.addQuery("release", id);
                    gr.query();
                    while(gr.next()){
                         var start =
 this.getTimeMs(gr["start_date"]);
                         var end =
 this.getTimeMs(gr["end_date"]);
                         var id = gr["sys_id"];
                         if (min == - 1 || start <= min){
                                if (start != min)
                                     minID = "";
                         min = start;
                         minID += "," + id;
                         }
                         if (max == -1 || end >= max){
                                   if (end != max)
                                           maxID = "";
                                   max = end;
```

```
                                    maxID += "," + id;
                            }
                     }
            }
        return [min, max, minID, maxID];
},
getMaxRangeDetails: function(id, table){ return [-1, -1, ""]; },
getTimeMs: function(date){ return new
 GlideScheduleDateTime(date).getMS(); },
getTimeObject: function(timeMS) {
  var gdt = new GlideDateTime();
  gdt.setNumericValue(timeMS);
  return gdt; },
logMessage: function(message){ gs.log(message); },
type: 'ExampleUpdateParentsRangeCalculator'
}
```

ExampleMaxRangeCalculator

Defines the earliest start date and the latest end date for a specified span.

```
var ExampleMaxRangeCalculator = Class.create();
ExampleMaxRangeCalculator.prototype = {
   initialize: function() { },
   updateParents: function(id, table, startDate, endDate){ },
   getMinRangeDetails: function(id, table){ return [-1, -1, "",
 ""]; },
   getMaxRangeDetails: function(id, table){
      if (table == "rm_sprint"){
         var sprint = new GlideRecord(table);
         sprint.addQuery('sys_id', id);
         sprint.query();
         if (sprint.next())
            releaseID = sprint.release + "";
         if (releaseID) {
            var now_GR = new GlideRecord("rm_release_scrum");
            gr.addQuery("sys_id", releaseID);
            gr.query();
            if (gr.next())
               return [this.getTimeMs(gr.start_date),
                this.getTimeMs(gr.end_date), gr.sys_id];
         }
      }
      return [-1, -1, ""];
   },
   getTimeMs: function(date){ return new
 ScheduleDateTime(date).getMS(); },
   getTimeObject: function(timeMS) {
     var gdt = new GlideDateTime();
     gdt.setNumericValue(timeMS);
     return gdt; },
   logMessage: function(message){ gs.log(message); },
   type: 'ExampleUpdateParentsRangeCalculator'
}
```

Use the following two functions to obtain the correct start and end dates in the three example script includes provided for reference.

```
getTimeMs: function(date){
        return new ScheduleDateTime(date).getMS();
    }
```

```
getTimeObject: function(timeMS) {
        var gdt = new GlideDateTime();
        gdt.setNumericValue(timeMS);
        return gdt;
    }
```

Timelines

Timelines display a linear calendar of tasks, such as incidents and change requests, defined by their start and end dates.

A span represents each task on the timeline, which is displayed as a horizontal bar and may have distinctive color coding. Each span has a label and can have tooltip text containing additional information about the task. You can also change the perspective of the calendar for a more granular view of the data.

Changes in perspective

Use the range selectors at the top of the timeline to change the perspective.

The increments go from one day to one year. To limit the timeline to an increment between the start date of the first span and the end date of the last span, click **Max**. Use the starting and ending calendar fields to select the timeline perspective. These fields control the same perspective as the slider at the bottom of the timeline. The green, vertical line indicates the current date and time, and sweeps across the timeline automatically.

Timeline Perspective Bar

The pink slider at the bottom of the timeline offers another way to change the perspective. Move the slider from right to left to view all the spans on a long timeline. Adjust the end points of the slider to make arbitrary changes to the magnification. A narrow slider zooms in on the spans and provides a more detailed view of complex timelines. A wide slider pulls the view out and makes more of the timeline visible on the screen.

Timeline Perspective Slider

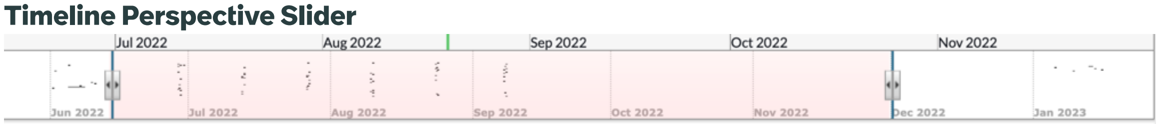

Managing spans

You can view span details, focus and zoom spans, and edit spans.

A span represents each task on the timeline, which is displayed as a horizontal bar and may have distinctive color coding. Each span has a label and can have tooltip text containing additional information about the task. You can also change the perspective of the calendar for a more granular view of the data.

Span details

If the left navigation pane is visible, the span label is visible in the left pane.

If the left pane is hidden, point the cursor over a span to view the tooltip.

Timeline Tooltip

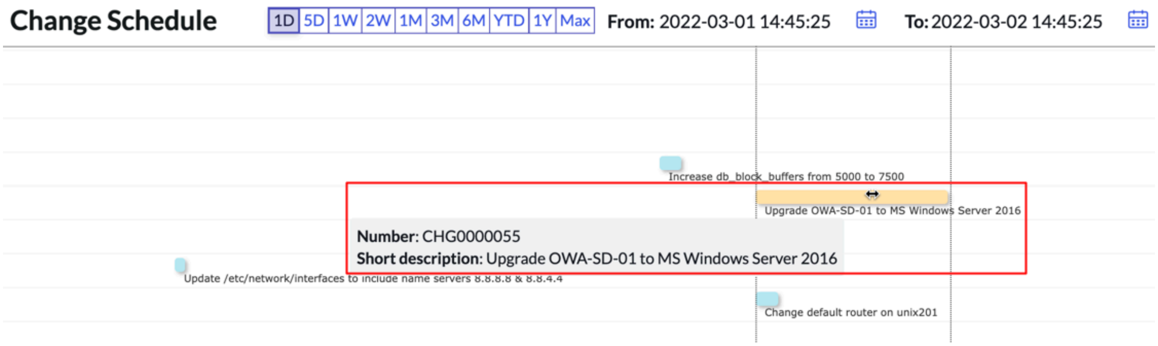

Span focus and zoom

You can use keyboard shortcuts when using timelines.

Focusing and Zooming on a Span

Operating System	Shortcut
Macintosh	Command + Click
Windows	Control + Click

Span edits

If the timeline has been configured for editing (horizontal moving, start and end date dragging), use the mouse to adjust the spans to suit your needs.

Move the start or end dates to change the duration of the task.

Timeline Drag Start

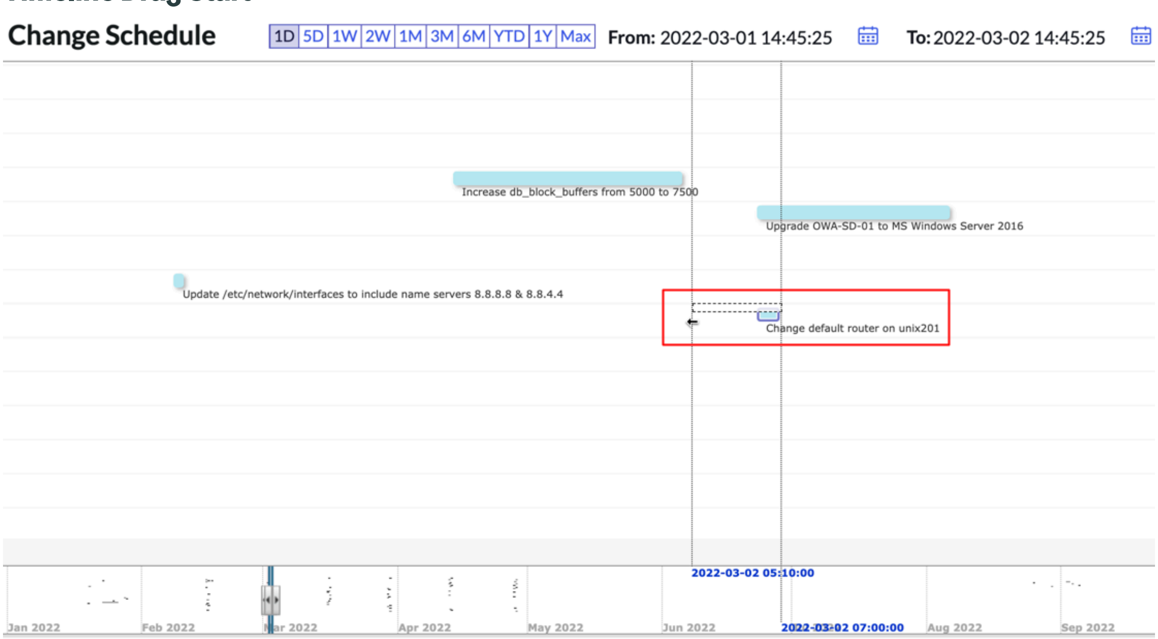

Move the entire span horizontally to adjust the start and end dates but not the duration.

Timeline Drag Span

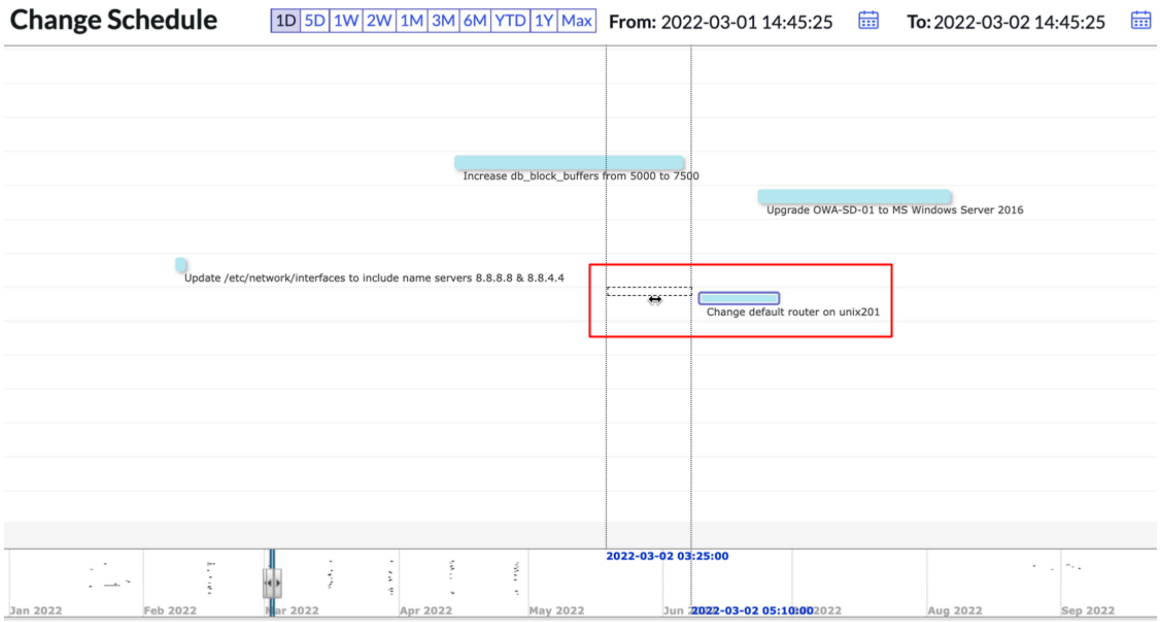

Client transaction timings

The Client Transaction Timings plugin enhances the system logs by providing additional information about the durations of transactions taking place between the client and the server.

You can track down performance issues to their source by viewing where the time is being consumed, and how time was spent during a transaction.

This plugin requires the Response Time Indicator Response Time Indicator to be enabled, and collects information from the following browsers:

- Firefox
- Internet Explorer
- Chrome

Client Transactions Information

Installing the plugin adds the *Client Transactions* module to the *System Logs* application. It provides a list of every logged transaction between client and server within the last day:

Client Transactions Information

Field	Description
Created	Moment the transaction was recorded.
Response Time	Number of ms spent by the server in fulfilling the transaction.
Business Rule Time	Number of ms spent by business rules triggered by the transaction.
SQL Time	Number of ms spent by the SQL database.
Client Response Time	(Load_completion_time) - (start_time). It is inclusive of server time.
Client Network Time	Number of ms spent by the network the client is connecting through.
Browser Time	Number of ms spent by the browser during the transaction.
Client Script Time	The number of ms spent executing client scripts.

Client Transactions Information (continued)

Field	Description
UI Policy Time	Number of ms spent executing ui policy.
Type	Type of transaction: • Form • List • Other
Table	Name of table that appeared. For example, incident, change_request.
View	View for this form/list.

Client Detailed Information

A more detailed breakdown of the client timings for all Form rendering (but not list rendering) is also tracked. To see details, drill into a particular client transaction record and observe the related list at the base of the screen.

Client Detailed Information

Field	Description
Order	Order during the load in which this operation occurred.
Type	Type of operation.
Name	Descriptive name of this particular operation
Duration	Number of ms this operation took to complete.

Components installed with Client Transaction Timings

The Client Transaction Timings plugin installs several components.

Database Table Structure

The plugin adds the table *syslog_client_transaction*.

Scripts

The plugin relies on the new script include *AJAXClientTiming*. This script gathers the information required and populates them on the *syslog_client_transaction* table.

Dependencies

This plugin does not require any other plugins, but does not gather information unless the Response Time Indicator Response Time Indicator is enabled.

Activating the Plugin

To activate the plugin, navigate to **System Definition** > **Plugins** and activate the plugin.

Note:
New instances have the plugin activated by default.

Timing values

Timing values are broken down into several sections.

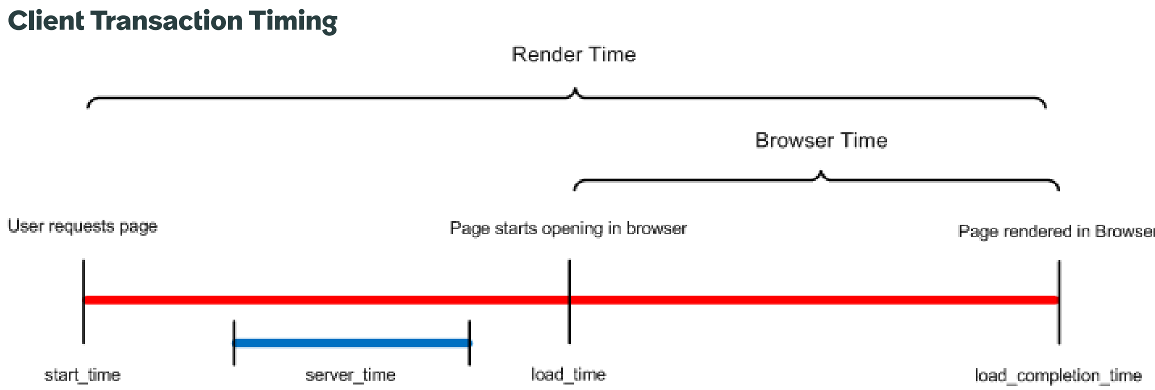

The variables in this diagram are defined as follows:

Timing Values

Variable	Description
start_time	Date and time the user requests a page (the user clicks on a link). This value is set by hooking into the *before unload* event of the previous page. WebKit browsers do not properly support the *beforeunload* event, which is why the client timings are not supported i Safari or Chrome.
load_time	Date and time that the current page starts loading in the browser. An inline javascript that runs as the first script in the HTML body sets this value.
server_time	Time in ms spent by the server processing the transaction. The server reports this value to the client.
load_completion_time	Date and time that the page is fully rendered in the browser. This operation is performed as the last script on the page and identifies the time the page completed loading.

The following timings are reported at the bottom right of many forms and lists:

Timing Values

Label	Element	Description	Calculation
Response Time	client_response_time	Overall time to deliver the page by subtracting the time the user requests the page from the time the page is fully rendered in the browser.	load_completion_time - start_time
Server Time	client_server_time	Time the server takes to process the transaction.	server_time
Network Time	client_network_time	Time the network takes to process the request. Calculated by subtracting the time of the user's request, from the time the page starts loading in the browser,	load_time - start_time - server_time

Timing Values (continued)

Label	Element	Description	Calculation
		and then subtracting the server processing time.	
Browser Time	browser_time	Time the browser takes to deliver the page by subtracting the time the page is fully rendered from the time the page starts loading in the browser.	load_completion_time - load_time

Currency administration

Currency fields provide features for handling the calculation, conversion, and display of currency values.

Get started

Choose one of these tiles to get started.

Explore	Configure	References
Learn about currency administration concepts and features.	Configure defaults for standard currency and FX currency.	Get details about currency administration properties.

Exploring currency administration

The ServiceNow AI Platform® supports two types of currency fields: default (or standard) currency and Foreign Exchange (FX) currency.

Default (standard) currency

The default, or standard, currency fields that are delivered with the ServiceNow AI Platform.

FX Currency

FX Currency is an alternate type of currency field that is highly configurable, supports custom rate tables, and provides for improved visibility into conversion calculations. You can add FX Currency fields to base and custom applications, and they work independently of existing standard currency fields.

Standard currency fields

The standard, or default, currency fields are installed with the base system and used in most non-financial applications.

A standard currency field stores a currency value, a currency code, and a reference currency value.

- Three-letter ISO code that identifies the currency that you select for the currency amount. For example, USD for the US Dollar, or EUR for the Euro.
- The reference currency value is the currency value, expressed in the reference currency. When you save a currency value, the reference currency value is calculated using a rate conversion.

A price field is similar to a currency field, but with special features for conversion and display. To learn more, see Price fields.

Locale settings

There are two locale settings, system and user. The system locale determines the reference currency, and the user locale determines the session currency.

System locale

You set the system locale using the *`glide.system.locale`* property. The value is in the format `Language.Country`, where the:

- `Language` is an ISO 639 language code.
- `Country` is an ISO 3166 language code.

Internally, this value is used as specified by Java. The system locale setting should be in the Java supported locales list. Set the system locale once after a fresh zBoot, because reference currency values that appear in currency fields are in the currency implied by the system locale. To set this property, see Control default currency field configuration and use in an instance.

Note:
Do not change the system locale after currency values have been entered into the instance. When you change the system locale, the reference currency values are not adjusted. There is no rate conversion. This persistence results in invalid aggregations and filtering.

User locale

The ServiceNow AI Platform® determines the user locale in the following sequence:

- User record in which the country, or both the country and language, are specified.
- System locale set using the *`glide.system.locale`* property.
- Browser locale.

Amounts in currency fields are composed of a currency code and amount. Amounts always appear in the session currency, and are formatted in the Java-specified format for the user locale.

Session and reference currency

The default, or standard, currency fields in the ServiceNow AI Platform® use two kinds of currency, Session and Reference.

Session currency

The session currency is defined for the user by the user's locale or single-currency mode.

Reference

The system local determines the reference currency, and is the standard used across the entire instance.

Each time you enter a value into a currency or price field, the system stores three pieces of information:

- Value as entered, in the user's locale.
- Currency code, in the user's locale.
- Value converted to the reference currency using the current exchange rate.

Note:
In multiple-currency mode, the currency code saved in the currency field may not be the same as the session currency code. For example, the session currency could be the Euro and the number entered could be the Japanese Yen.

Session currency

When users view a currency value, they can see the value as entered or in the session-currency format. The format contains the:

- Currency symbol
- Value converted to the session currency and shown in a localized number format.

The user's locale determines the session currency format.

The number format can differ in features such as the decimal separator based on the locale. For example, the US formatting is 1,234,567.89, while German formatting is 1.234.567,89. The ServiceNow AI Platform® determines the session currency in the following sequence:

- Single-currency mode setup using `glide.i18n.single_currency` and `glide.i18n.single_currency.code`.
- Default currency for the user's locale.

Reference currency

To perform calculations on heterogeneous currency values, the ServiceNow AI Platform® stores currency values converted to a system currency, referred to as the reference currency. Every currency field contains a reference currency value. The system determines the reference currency in the following sequence:

- System locale set using the property `glide.system.locale`
- Java default locale, typically en.US

The filtering and aggregation features use the reference currency value to perform calculations on default currency fields. These features can yield inaccurate results because of conversion rate changes.

Issues with currency fields

Users are often confused by the results of filtering, sorting, and displaying currency fields because the system works with at least two currencies for each value: the session currency and the reference currency.

Note:
Aggregations and filtering of currency fields use the reference currency, and the user sees the session currency. Because of changing conversion rates, the filtered reference currency values might not result in the same order as the session currency values would suggest. The same issue happens with aggregations.

The user might see the following issues:

- Lists filtered on currency fields might not be in the expected order. It uses the reference currency values for filtering but displays session currency values.
- Aggregation of currency fields may not produce the expected results because reference currency values are aggregated, and then converted to the session currency.
- Currency values may not appear as expected because currency values are formatted based on the user's locale and not on the currency code.

This confusion may be the result of the difference between session and reference currencies, changing conversion rates, and different session currencies used by different users.

Single-currency mode

Single-currency mode enables all users of the platform to view currency values in the same currency.

Before enabling single-currency mode, set the system locale. To configure single-currency mode, set the following properties:

- *`glide.i18n.single_currency`*: **true** or **false**
- *`glide.i18n.single_currency.code`*: the three-letter ISO currency code
- *`glide.system.locale`*: the system locale

Note:
For detailed information about valid locale formats, see Locale settings.

Using the single-currency mode has the following limitations:

- The single-currency mode changes the currency that appears in the user views, but does not change the number formatting. Even though users in different countries see currency values in one currency, the number formatting, as determined by the user's locale, might not be what they expect.
- Price fields, which have flexible conversion and currency features, can't be used because currency value input is constrained to the single currency.

You can avoid the effects of rate conversions by setting the system locale and the reference currency to be the single currency.

Price fields

A price field is a currency field that enables control over conversions and display.

You can select conversion and display options per price field, and change them at any time. There are three variations:

Calculated [Default]

Behaves in the same way as the default currency field type. Whenever conversions are performed, it uses the latest currency conversion rates. The price field appears in the session currency for the user.

Fixed

When the price field appears, it is shown in the currency code used when you entered the value. Whenever conversions are performed, it uses the latest currency conversion rates.

Multiple

Enables you to enter multiple price values for an item using a different currency for each price. The field value is the value entered in the user's session currency.

Otherwise, the first price entered converts to the user's session currency. Whenever conversions are performed, it uses the latest currency rates.

Note:
The first value entered is used during display. The additional values are not used during calculations.

You can change the currency code and numeric value for a price field. When you click the **Edit** icon that appears next to the price field, a form appears that you use to edit the details for the price field:

Price field currency details

Field	Description
Currency	List of currencies enabled in the system in the combo box. In single currency mode, the currency is a label and cannot be changed.
Amount	Numeric value formatted in the user's locale.
Type	Calculated, Fixed, Multiple. • When you change the price type to Multiple, the system creates child records for all currencies enabled in the platform. The child records are populated with values converted from the amount field, using latest currency conversion rates. • You cannot change the price type in single currency mode. • You can modify the price type at any time.

Default currency values in forms

In forms, currency values appear in the currency in which they were entered.

A combo box provides a list of available currencies, and the user's locale determines its format. When entering or changing the numeric currency value, enter it in the same format specified by the user's locale. In the form for a new record, the reference currency is pre-selected in the currency list, with the numeric value set to zero.

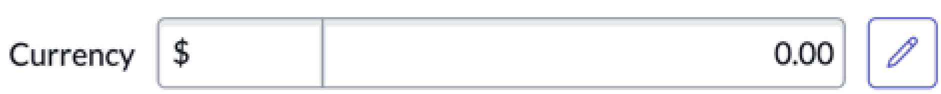

- If the record is read only, the currency value appears in the **Currency** field as it was entered, and is formatted in the user's locale. In the **Price** field, the session currency value appears.
- In Single currency mode, the currency is a label that cannot be changed, and the **Edit** icon does not appear.

Editing the currency instance record

If a currency instance record exists, an **Edit** icon appears next to the **Currency** field if you have edit rights for the fx_currency_instance table. Users with the financial_mgmt_user role can edit the values associated with the currency field.

Note:
Normally, you should not edit the fx_currency_instance table directly. The ServiceNow AI Platform maintains these tables, and your changes could have unintended consequences.

Default currency values in reports

Currency values in reports appear in the user's session currency, and are formatted in the user's locale with a currency symbol.

The user determines how the report is run.

- Shared report: The user who runs the report.
- Scheduled report: Generally run as the user who scheduled the report.

The two user-specific values in the report are:

- User session currency
- Converted value

Note:
A user that has a different session currency than the person who runs a report might receive unexpected results.

Default currency values in lists

In lists, default currency values appear in the user's session currency, formatted for display in the user's locale. Typically, a formatted number follows the currency symbol.

Different field types appear as follows in lists:

Field	Description
Currency field type	Value expressed in the user's session currency.
Price field type	• **Calculated**: Value in the user's session currency. • **Fixed**: Currency value, as entered by the user. • **Multiple**: Value associated with the user's session currency if this value exists. Otherwise, the first value entered converts to the user's session currency.

A Globe icon appears beside the currency value that enables you change to one of the following values:

- Value as entered by the user.
- Value in session currency.
- Value as entered and, in brackets, the value in reference currency.

The icon appears when the user's session currency is different from the currency entered. Clicking the icon cycles through the listed displays. In the preview for the record, currency values appear as entered, formatted in the user's locale.

Aggregation

Aggregation operations can be used on currency columns. Aggregation operations include total, group by, average, minimum, and maximum. Aggregation is done in two steps:

- Aggregate the reference currency values for all records.
- Convert this aggregate to the user's session currency for display.

Note:
The resulting value may not be what you expect. The conversion rate used for the value that appears for the currency field, and for its reference currency, that is used for the aggregation, may have changed.

This limitation extends to different price types.

- For price type fixed, the calculated reference value can be old.
- For price type multiple, the reference value of the first price entered is used. The other values are not used.

The aggregate value is shown formatted in user's locale with a currency symbol. Currency fields are stored with four decimal places, and aggregates have four decimal places. For upgrades, you must set the `glide.currency_price.use_all_fraction_digits` property. See Change default currency decimal places for more information.

Filtering

You can set up filters on currency fields. The currency value is entered as a currency code and numeric value. Filtering is done in two steps:

- The filter currency value is converted to the reference currency.
- The filter's calculated reference value is compared with the reference value in the records.

Matching records are shown in the list view.

Note:
At times, the filtering results might not provide the expected result. The conversion rate used when you run the filter might be different than the conversion rate used when calculating the reference values in the individual records. This limitation extends to the following price types:

- For a Fixed price field type, the calculated reference value can be old.
- For a Multiple price field type, it uses the reference value of the first price, and does not use the other values.

Default currency values in import and export

In general, currency values crossing the boundaries of the platform are represented in the user's session currency and formatted in the user's locale.

Import

Currency values are imported as strings, just like other fields. The default transform mapping to a currency field uses `setDisplayValue()`. The expected format for this function is:

- A number formatted in the user's locale. The number is taken as a value in the user's session currency, for example, 1,234.56.
- The number prefixed by the three-letter currency code separated by a semicolon, for example, EUR;1.234,56.

This behavior can be customized in transform map scripts.

Export

Currency values can be exported in several formats, including CSV, XLSX, PDF, and XML. Depending on the export format, currency values are exported in the user's session currency and formatted in the user's locale or they're exported in the reference currency value with no formatting.

Related topics

Exporting currency fields to Excel

Default currency values in scripts

You can use currency fields in scripts.

These methods are available on GlideElement objects.

To display currency values, use the *getDisplayValue()* display API. To work with currency values in any way other than display, use the APIs that return/accept unformatted numbers.

Note:
Before performing calculations on the value, do not use the *getDisplayValue()* methods and then process the string to remove formatting information.

Methods such as *getValue()* and *getCurrencyValue()* return unformatted numbers as strings. To obtain the floating point value, use the JavaScript function *parseFloat()*, and then use resulting value to perform calculations. To obtain the currency associated with these values, use the APIs that return the currency code. You can also use the *getCurrencyCode()* methods to determine a field's currency.

```
var rate = parseFloat(current.base_rate);
var currencyCode = current.base_rate.getCurrencyCode();
```

Use the *setValue()* method to set the value of a currency field. If this currency is the user's session currency, use a plain number or the floating point number of a string containing it. Otherwise, prefix the value with the 3-letter ISO currency code.

```
var totalCost = rate*current.hourly_rate;
current.total_cost.setValue(currencyCode + ";" + totalCost);
```

You are working with the reference currency value when you use GlideAggregate on currency or price fields. Be sure to convert the aggregate values to the user's session currency for display. The resulting value may not be what you expect. The conversion rate used for the currency or price field value, and for its reference currency, which is used for the aggregation, may have changed.

When you delete a record containing a currency value, the platform deletes any associated currency records.

Note:
Do not use *deleteMultiple()* on tables with currency fields. Always delete each record individually.

Currency values contain four decimal places.

- APIs that return values such as `getValue()` return up to four decimal places. Trailing zeros are always removed.
- APIs that return display values such as `getDisplayValue()` have at least two decimal places and up to four decimal places.
- GlideAggregate returns four decimal places.

You can have the system use two decimal places. When you set it to two decimal places, numeric values returned by the API contain two decimal places. Although currency conversion rates may have more decimal places, currency fields store only two decimal places. APIs that accept numeric values round decimal places to two places.

- APIs that return values such as `getValue()` return up to two decimal places. The trailing zeroes are removed for values read from the database, but if a value such as 00 is set later, 1.00 can be returned. The number of trailing zeros returned is not consistent.
- APIs that return display values such as getDisplayValue() contain up to two decimal places. It could sometimes return two places even for values such as 7.10, but could remove trailing zeros at other times. The number of trailing zeros returned is not consistent.
- GlideAggregate returns two decimal places.

Note:
To learn how to change the number of decimal places used by the system, see Change default currency decimal places.

The following example, the user's locale is set to German (de.DE), and the reference currency set to USD. The values use a currency value of 21345.67 in Japanese yen, 1563.72 in Euros, and 1152.48 in US dollars.

Methods to access currency fields

Method name	Description	Example
`getValue()`	Returns the currency value in the user's session currency as an unformatted number.	1563.72
`getReferenceValue()`	Returns the currency value in the reference currency as an unformatted number.	1152.48
`getSessionValue()`	Returns the currency value in the user's session currency as and unformatted number.	1563.72
`getCurrencyValue()`	Returns the currency value as entered as an unformatted number.	21345.67
`getDisplayValue()`	Returns the currency value in the user's session currency, formatted in the user's locale with a currency symbol.	€1.563,72
`getSessionDisplayValue()`	Returns the currency value in the user's session currency, formatted in the user's locale with a currency symbol.	€1.563,72

Methods to access currency fields (continued)

Method name	Description	Example
`getReferenceDisplayValue()`	Returns the currency value in the reference currency, formatted in the user's locale with a currency symbol.	$1,152.48
`getCurrencyDisplayValue()`	Returns the currency value as entered formatted in the user's locale with a currency symbol.	¥21.345,67
`getCurrencyString()`	Returns the currency value as entered as an unformatted number, prefixed by the 3-letter ISO currency code, separated by a semicolon.	JPY 21345.67
`getCurrencyCode()`	Returns the 3-letter ISO currency code for the currency value as entered.	JPY
`getSessionCurrencyCode()`	Returns the 3-letter ISO currency code for the user's session currency.	EUR
`getReferenceCurrencyCode()`	Returns the 3-letter ISO currency code for the reference currency.	USD
`setValue()`	Sets the currency value as: • An unformatted number taken as a value in the user's session currency. • An unformatted number prefixed by a 3-letter currency code separated by a semicolon.	4369.21 or JPY 4369.21
`setDisplayValue()`	Sets the currency value as: • A number formatted in the user's locale that is taken as a value in the user's session currency. • A number formatted in the user's locale prefixed by a 3-letter currency code separated by a semicolon.	4369.21 or JPY 4369.21

FX Currency fields

Use FX (Foreign Exchange) Currency fields as an alternative to the standard (default) currency fields found in most applications. They are highly configurable, support use of custom rate tables, report aggregations, and provide for improved visibility into conversion calculations.

Note:
You don't have to install a plugin to use FX Currency functionality. It's activated by default.

Use considerations

Use the FX Currency field type in place of the standard currency fields when you must:

- Support values expressed in multiple currencies, for example, EUR 100, JPY 200, USD 30, as values in a field. FX Currency fields provide better support for multi-currency applications than do standard currency fields.
- Control how the FX Currency field is configured, including specifying its internal behavior and currency value formatting.

  To learn more about currency configuration, see Configuring FX Currency field behavior.
- Create custom conversion rate tables.

  To learn more about custom table rates, see Add conversion rates using a custom rate table.
- Have more control over when the actual currency conversion occurs. FX Currency values convert on demand only, making the conversion results much more transparent.

  To learn more about currency conversions, see Understanding how FX Currency field conversions work.
- Provide greater transparency into currency conversion calculations. You can view and configure the calculation components, including conversion dates, reference currency, and number of displayed decimals.
- Build query conditions using FX Currency field values. For available query conditions, see Operators available for filters and queries.
- Access FX Currency fields through richer scriptable APIs.
- Aggregate currency values as entered by users, or aggregate them in selected reference currencies, for reporting purposes.

Important:
You can add FX Currency fields to base and custom applications, and they work independently of existing standard currency fields. However, you can't convert or change a standard currency field to an FX Currency type field. No functions are available to convert standard currency fields to FX Currency fields, then migrate the data from one to the other.

Related topics

Field types

Understanding FX Currency values in forms

In forms, the FX Currency field behaves like a dot walkable field in script. It consists of an entry field, and an accompanying list for selection of a currency code.

Specifically, FX Currency fields that appear on forms contain the following:

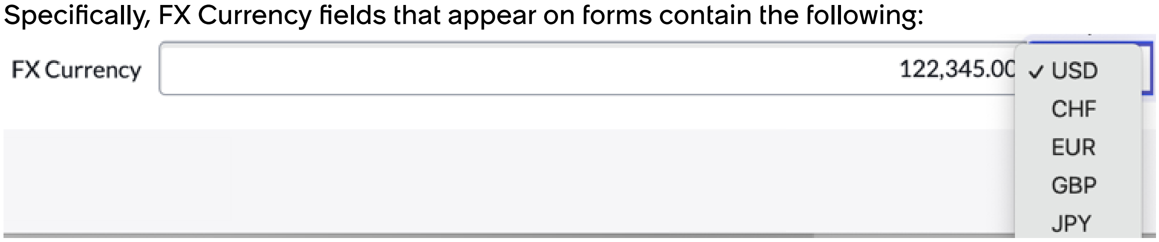

- An empty field for entry of a numeric currency value, formatted according to your assigned user locale. For example, in the US, it is formatted with two decimal places, for entry of cents.

  If this is a new record, the field is empty. If it is an existing record, the previous currency amount that you entered appears.

- A selection list of active currencies from the Currency [fx_currency] table. The default currency code is based on your user locale. You can select another currency for use during the entry session.

  For example, in the US, the default currency value, based on the user locale, is USD. You can select another currency code to use in the current entry session.

- An edit icon ( ). Users with an assigned currency_instance_admin role can click it to edit currency detail in the accompanying Currency Instance [fx_currency2_instance] record.

Understanding how FX Currency field conversions work

You can use existing system currency rates or define a custom conversion rate table for use in FX Currency processing.

A currency conversion rate represents the conversion factor for exchanging one currency for another. The conversion rate is used when converting a currency value, expressed in one currency, to another currency. Conversions can occur for storage (for example, conversion to a reference currency) or for display (for example, conversion to a session currency).

When conversions occur

FX Currency fields expose the values inside the Currency Instance [fx_currency2_instance] record as dot walkable fields. You can access to all data held in an FX Currency field without any transformations. Conversions can happen in these situations:

1. Calculation of a reference value, calculated when you enter a currency value, whether on insert or update. The reference currency value is saved as well as the currency value as entered. Once this calculation takes place, the reference value is not recalculated.
2. Calculation of the session display value returned from `getDisplayValue()`, which performs conversions to the required currency.

   The conversion rate is based on the settings of the **Conversion Rate Table** and **Conversion Date Source** fields in **System Localization** > **FX Currency Configuration.**

   Note:
   To learn more about these settings, see Specify the rate table and date source for currency conversions.

Conversion rates

FX Currency conversion rates are stored in the following ways:

- System rate [fx_system_rate] table: Contains system conversion rates. Each record contains the conversion rate from the Euro to a given currency. The rates are downloaded daily from the ECB website by a scheduled job called Retrieve System Rates. To learn more, see Schedule the rate update job.
- Custom rate table: Contains custom currency conversion rates, and custom rate tables extend the Currency Conversion Rate [fx_conversion_rate] table. To learn more about creating custom rate tables, see Add conversion rates using a custom rate table.

Understanding FX Currency values in lists and reports

By default, the currency values that appear in the FX Currency fields in lists and reports are the values that you directly enter into the FX Currency fields. For reports, the aggregations of FX Currency fields are, by default, reference currency values.

FX Currency values in list views

In lists, FX Currency values appear as a formatted currency string returned from `getDisplayValue()`.

- The settings in the **Display Digits** and **Display Value Currency** fields in **System Localization** > **FX Currency Configuration** determine how the currency values appear in lists.
- To learn more, see Identify the FX Currency field and its display parameters.

For an FX Currency field that depends on another FX Currency field, the currency value comes from the Currency Instance [fx_currency2_instance] record associated with the latter. The display value format is based on how you configured its display settings in **System Localization** > **FX Currency Configuration.**

> **Note:**
> An FX Currency field in a custom table is a reference field that stores the reference (sys_id) to the actual record. Because the sys_id is 32 characters, the **Max length** of the FX Currency column must be 32 (default).

Filtering FX Currency values in lists

In lists, FX Currency field filtering behaves differently than with standard currency fields. The filtering takes place using the currency type selected when a currency value is originally entered into an FX Currency field. It doesn't take reference currencies into account when you filter a list. So, if the unfiltered list contains FX Currency field entries in multiple currencies, a filtered list would only contain those records originally entered using the specified filtering currency.

- For example, if you entered a currency value in Euros (EUR) and the reference currency is US Dollars (USD), the reference currency value may appear in an unfiltered list.
- However, if you filter the list, and select USD as the currency, the currency value, originally entered in Euros, doesn't appear. In this case, only currency values originally entered in USD would appear.

Dependent FX Currency fields

Use of FX Currency view or dependent fields enable you to use alternate display configurations defined in **System Localization** > **FX Currency Configuration** for a particular FX Currency field. To learn more about FX Currency field configurations, see Configuring FX Currency field behavior.

When you create an FX Currency field, you can mark it as dependent on another FX Currency field. Marking an FX Currency field as dependent enables it to use the configuration settings for the original FX Currency field to display the amount values in different currency. Any changes that you make to the dependent field don't affect the original FX Currency field.

FX Currency reporting, aggregation, and group by

You can select FX Currency fields for inclusion on reports, and they display in the same manner as in list views.

- FX Currency fields appear as a formatted currency string returned from *getDisplayValue()*.
- If you set the **Display Value Currency** field in **System Localization** > **All Properties** to **Display in Session Currency**, the session currency assigned to the user running the report appears.

You can aggregate like reference currency values within a Currency Instance record. The following is the aggregated FX Currency data format in the ServiceNow AI Platform®:

Aggregate FX Currency value

Type	Description
String	Value of the aggregation. If the values being aggregated are FX Currency values, the returned value is in the format `currency_code:currency_value`, such as: `USD:134.980000` Otherwise, the returned value is just the currency value, such as: `134.980000`. **Note:** If the field contains FX Currency values of mixed currency types, the values cannot be aggregated. The value is returned as a semicolon (;).

For example, if the **Cost** field on a transaction record is an FX Currency field, you can group by *cost.amount* or *cost.currency*.

- Aggregation operations, such as SUM, translate as `SUM(amount)` or `SUM(reference_amount)`. The aggregation source determines the translation when the currency associated with the currency values in the grouping is the same.

  For example, the aggregate value returned by `GlideAggregate.getAggregate()` is a string of the format USD; 1234.56. The total currency value is expressed in the single reference currency that is associated with all currency values in the grouping.

- If there are different currencies associated with currency values in the grouping, then the aggregate value is empty. For example, if some of the currency values are in US Dollars, some are in Euros, and others in Yen, the aggregate value would be empty.

You can also use a Group By command on FX Currency fields.

- It is equivalent to grouping by currency and amount, or by reference_currency and reference_amount, as dictated by the aggregation source.
- The value of the grouping is available by dot-walking on the FX Currency field, or using an API such as `getDisplayValue()`.

Related topics

Configuring FX Currency global settings

Configuring FX Currency field behavior

FX Currency values in import and export

In general, currency values crossing the boundaries of the platform represent whatever is returned by `getDisplayValue`. Usually this currency value is the default as entered by a user into an FX Currency field for a transaction.

FX Currency import

If using the default `setDisplayValue()` import method, FX Currency works in manner similar to the standard currency field. It imports a value such as `USD;1,234.56` as `setDisplayValue("USD;1,234.56")`. It then parses the value as a three letter ISO currency code, followed by a semi-colon, followed by a number formatted in the locale of the user. This method should work for most use cases.

Defaults for other Currency Instance fields come from FX Currency field configuration you defined in **System Localization** > **FX Currency Configuration.** To explicitly set these fields, you can use a custom import script that sets those fields to desired values, possibly using values from another import column.

There are various ways to import FX Currency data.

Method	Description
Importing data into custom rate tables	Main method for importing custom rates, other than editing rates directly. When you modify existing rates, existing currency values that use that rate aren't updated. **Note:** This is specific for importing only.

FX Currency export

There are various ways to export FX Currency data.

Method	Description
Show XML	Shows an XML file with instance records for each FX Currency field bundled inside the XML payload.
Export to XML	Shows an unload XML with unloads for all instance records.
Export to Excel/PDF	Uses the default method of exporting data that extracts data from fields using `getDisplayValue()`. For FX Currency fields, using this method returns a formatted currency string of the form $1,234.56. To access specific information inside an FX Currency field directly without any transformations, simply export that column using dot walking.

Configuring currency fields

Configure defaults for standard currency and Foreign Exchange (FX) currency.

Setting up defaults required for standard currency use

Set up the defaults required for use of standard currency fields in your instance.

Control default currency field configuration and use in an instance

You can control how currency fields are configured and used in your instance.

These properties are available for currency fields.

Note:
To open the System Properties [sys_properties] table, enter `sys_properties.list` in the navigation filter.

Properties for default currency fields

Property	Description
glide.sys.audit_currency_value	When **true**, currency fields in audit records are the value entered by the user, in the format USD;1234.56. When **false**, the value is the numeric value in the session currency. • Type: Boolean • Default value: **false** • Location: System Property [sys_properties] table • Learn more: Configure default currency fields in audit records
glide.currency_price.use_all_fraction_digits	When **true**, currency fields have four decimal places. When **false**, two decimal places are used. • Type: Boolean • Default value: **false** • Location: System Property [sys_properties] table • Learn more: Change default currency decimal places
glide.currency_price_optimizer.enabled	When **true**, the optimizer is used. When **false**, the optimizer is not used.

Properties for default currency fields (continued)

Property	Description
	• Type: Boolean • Default value: **true** • Location: System Property [sys_properties] table • Learn more: Configure the currency optimizer
glide.currency_price_optimizer.min_rows	Minimum number of rows in parent table needed for the optimizer. When fewer than this number of rows are present, the optimizer is not used. • Type: Number • Default value: **4** • Location: System Property [sys_properties] table • Learn more: Configure the currency optimizer
glide.excel.convert_to_user_currency	When **true**, currency fields in a table are output in the user's session currency. • Type: Boolean • Default value: **false** • Location: System Property [sys_properties] table • Learn more: Exporting currency fields to Excel
glide.excel.fixed_currency_usd	When **true**, currency fields in a table are converted to US dollars. • Type: Boolean • Default value: **false** • Location: System Property [sys_properties] table • Learn more: Exporting currency fields to Excel
glide.system.locale	Value is of the format `Language.Country` where the language is an

Properties for default currency fields (continued)

Property	Description
	ISO 639 language code and the country is an ISO 3166 language code. Internally, this value is used as specified by Java. The system locale setting should be in the Java supported locales list. • Type: String • Default value: empty • Location: **System Properties** > **System Localization** • Learn more: Locale settings

Related topics

Exporting currency fields to Excel

Change default currency decimal places

You can specify the number of decimal places stored and used in currency fields and calculations.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **sys_properties.list**.
2. Find *glide.currency_price.use_all_fraction_digits*, and set to the desired value.

Value	Description
false	Uses two decimal places.
true	Uses four decimal places. This value is the default.

Configure default currency fields in audit records

You can control what currency value is stored in audit records.

Before you begin

Role required: admin

About this task

Procedure

1. Navigate to **All** > **sys_properties.list**.
2. Find *glide.sys.audit_currency_value*, and set to the desired value.

Value	Description
false	Value entered in the audit record is the numeric value in the session currency. It does not include the currency code. This value is the default.
true	Value as entered by the user in the format USD. For example, 1234.56.

Configure the currency optimizer

You can configure the system to use an optimizer for currency/price fields to speed up list view (reading currency/price values).

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **sys_properties.list**.
2. To turn on the currency optimizer, find *glide.currency_price_optimizer.enabled*, and set to the desired value.

Value	Description
false	Optimizer is not used.
true	Optimizer is used. This value is the default.

3. To set the minimum number of rows needed by the currency optimizer, find *glide.currency_price_optimizer.min_rows*, and set to the desired value. The default value is 4.

Default currency conversions

Currency values can be converted to other currencies when stored and accessed.

Conversions can happen in these situations.

- The currency value converts to the reference currency when stored, whether on insert or update. The reference currency value is saved as well as the currency value.
- The currency value converts to the user's session currency for display.
- The value entered into a filter converts to the reference currency using the currency specified in the filter.

Conversion rates are stored in the fx_rate table. Each record contains the conversion rate from a given currency to the Euro. The rates are updated daily from the ECB website by a scheduled job called ECB Exchange Rate Load.

A currency conversion from one currency to another involves two rates.

- Actual: In the fx_rate table, each record contains the conversion rate from a given currency to the Euro.
- Expected: In the fx_rate table, each record contains the conversion rate from the Euro to the given currency.

Whenever you perform a conversion, it uses the latest conversion rates. Therefore, calculations can potentially yield unexpected results. For example:

- Different currency values can have different rates applied to them while storing the reference currency value. Therefore, aggregation can combine values at different rates and convert back at another rate.
- A filter value is converted at current rates while the values it filters in the database can be converted at different rates. A filter for $100 at today's rate can match a value of $99 obtained at yesterday's rates.

Note:
For display purposes, the currency value used is what the user entered converted to session currency. However, for aggregation and filtering, the reference currency value is used. Using the reference currency enables currency values converted at different rates to be compared together.

Schedule the rate update job

Schedule the ECB Exchange Rate Load to perform a nightly download of currency-conversion tables from the European Central Bank.

Before you begin

Role required: admin

About this task

You can adjust the frequency of this behavior or disable it entirely. For information about turning off the regularly scheduled update and maintaining the Exchange Rate table manually, see Use your own currency-conversion rates and FX Currency fields.

Procedure

1. Navigate to **All** > **System Scheduler** > **Scheduled Jobs**.
2. Open the job named **Retrieve System Rates**.
3. Modify the schedule, as needed.

   After the job runs, it stores and loads the rates from the Exchange Rate [fx_rate] table. Navigate to **System Localization** > **Exchange Rates** to see them.

Use your own currency-conversion rates

All currency conversions are based on the rates stored in the Exchange Rate table. You can turn off the regularly scheduled update from the European Central Bank (ECB), and maintain the table manually.

Before you begin

Role required: admin

About this task

ECB Exchange Rate Load loads exchange rates from the ECB for the following currencies:

http://www.ecb.int/stats/eurofxref/eurofxref-daily.xml.

If ECB does not supply the daily rates for a specific currency, you can manually enter rates into the Exchange Rate table. Use an import set, or another service (for example, JSON or SOAP) that offers upload of more currency rates. You can then add a similar scheduled job to update these currencies.

Procedure

1. Navigate to **All** > **System Scheduler** > **Scheduled Jobs**.
2. Open the job named **ECB Exchange Rate Load**.
3. In the *Trigger type* field, select **-- None --**.
4. Enter new exchange rates either manually or with an Import Set.

Setting up defaults required for FX Currency use

Set up the defaults required for use of FX Currency fields in your instance.

Perform the following setup tasks to enable FX Currency processing:

Add conversion rates using a custom rate table

You can create a custom rate table to ensure repeatable and consistent calculations. In a custom rate table, you control the conversion rates and the frequency of rate conversions. If you don't define custom rate tables, FX Currency fields use the System Conversion Currency Rate [fx_system_rate] table, which updates rates daily from the ECB (European Central Bank) website.

Before you begin

Role required: admin, currency_admin

About this task

A conversion rate can be quoted in different ways.

- Direct conversion rate (A -> B) that may or not be invertible (that is, given rate R for A -> B, the rate for B -> A can be derived as 1/R).
- Indirect conversion rate through some intermediate currency (A -> C and C -> B).

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Click **New**.
3. In **Label** field, enter the name of the conversation rate.
   For example, Convert-US-Euro.
4. In the **Extends** field, select the Currency Conversion Rate [fx_conversion_rate] table.
5. Complete the remaining fields to create the new extended table.
   For more details, see Create a table for details.
6. Click **Submit**.
7. Navigate to the new table and fill in the fields on the form.

Field	Description
From Currency	Currency to convert the rate conversion table from. For example, USD if you are converting from US Dollars.
To Currency	Currency to convert the rate conversion table to. For example, EUR if you are converting to Euros.

Field	Description
Span start	Date/time from which this currency rate is valid, for example, January 1, 2020 00:00:00.
Span end	Date/time to which this currency rate is valid, for example, January 30, 2020 23:59:59.
Order	Order in which a currency conversion selects currency rates when their time spans overlap. For a date/time and a currency pair, if there are multiple rates with a Span start less than or equal to the date/time, and Span end greater than or equal to the date/time, the currency conversion selects the lowest order rate.
Rate	Actual conversion rate, stored as a decimal with six decimal places.

8. Click **Update**.

What to do next

After you create the custom rate table:

1. Manually populate the table, or import data to populate it.
2. To use this rate table on a global basis, change the *glide.currency2.system_rate_table* system property.

   To use it for a specific FX Currency field, change the **Conversion Rate Table** field setting in **System Localization** > **FX Currency Configuration.**

Related topics

FX Currency values in import and export

Configuring FX Currency global settings

Specify the rate table and date source for currency conversions

Configuring FX Currency global settings

Before setting up and using FX Currency, review the related system properties to confirm or modify its default behavior.

Before you begin

Role required: admin

About this task

The global defaults you establish in system properties apply the configuration option to all FX Currency fields. You can then customize them for specific dictionary attributes in **System Localization** > **FX Currency Configuration.** When you create a currency configuration, it applies those settings only to the FX Currency field that you add it to. To learn more about FX Currency configurations, see Configuring FX Currency field behavior.

FX Currency configuration settings for specific FX Currency fields take precedence over system properties. Administrators can set a global configuration for a property, and then apply exceptions on a field-by-field basis.

- If you change system defaults, except for *display_digits* and *display_value*, the changes apply to updated and new values only.
- Existing values and data don't change but rather display differently, depending on configuration.

> **Important:**
> Don't edit the Currency [fx_currency] table directly. The system maintains this table and your changes could have unintended consequences.

Procedure

1. Navigate to **All** > **System Localization** > **FX Currency Configuration.**
2. Review the system defaults and make any desired changes.

FX Currency field properties and attributes

System property	FX Currency configuration mapping	Description
glide.currency2.display_digits	currency_display_digits	Default number of digits displayed in FX Currency fields. **Note:** In a form, FX Currency fields display the number of digits you enter, up to 12 digits in length. ◦ **Type**: Number ◦ **Default value**: -1, or currency default ◦ **Available options**: 0–12 ◦ **Develop**: GlideCurrencyFormatter
glide.currency2.display_value	currency_display_value	Default currency display value returned from `getDisplayValue()`. ◦ **Type**: String ◦ **Default value**: as_entered ◦ **Available options**: ▪ as_entered: Currency value as entered by the user ▪ in_session_currency: Currency value as entered by the user and converted to the session currency ▪ in_reference_currency: Currency value as

servicenow

System property	FX Currency configuration mapping	Description
		entered by the user converted to the reference currency ◦ **Develop**: GlideCurrencyFormatter
glide.currency2.system_rate_table	conversion_rate_table	Conversion rate table used to perform currency conversions. By default, the source for currency conversions is the System Currency Conversion Rate [fx_system_rate] table, which stores the daily reference rates retrieved from the European Central Bank. If you have defined a custom rate table, you can change the rate table used in currency conversions. ◦ **Type**: String ◦ **Default value**: fx_system_rate ◦ **Available options**: Any table that extends the **Currency Conversion Rate** table [fx_conversion_rate] ◦ **Learn more**: Understanding how FX Currency field conversions workand Default currency conversions ◦ **Develop**: GlideCurrencyConverter GlideCurrencyExchangeValue
glide.currency2.default_reference_currency	reference_currency	Three-digit ISO currency code used for the reference currency. The default reference currency for FX Currency fields is the system reference currency. ◦ **Type**: String ◦ **Default value**: If you set *glide.system.reference_currenc* to true, the three-digit ISO currency code determined by *glide.system.reference_currenc*

System property	FX Currency configuration mapping	Description
		If set to false, the system locale is the default. ◦ **Learn more**: Session and reference currency ◦ **Develop**: GlideElementCurrency2

Add an FX Currency field to a table

Before using and configuring an FX Currency field, you first add it an existing table, or to a new one, in **System Definition** > **Tables.**

Before you begin

Role required: admin

About this task

For example, you create a new custom table, named `transactions`, that stores transactional information and contains an FX Currency field called `cost`.

Procedure

1. Follow the normal procedures to add a table, or access an existing table.
   To learn more about creating and modifying tables in an instance, see Create a table.
2. In the **Columns** section, use the **Table Columns** embedded list to add columns to the table.
3. In the **Column label** field, enter the column name.
   For example, `cost`.
4. In the **Type** field, select **FX Currency** as the field type.

   Example

   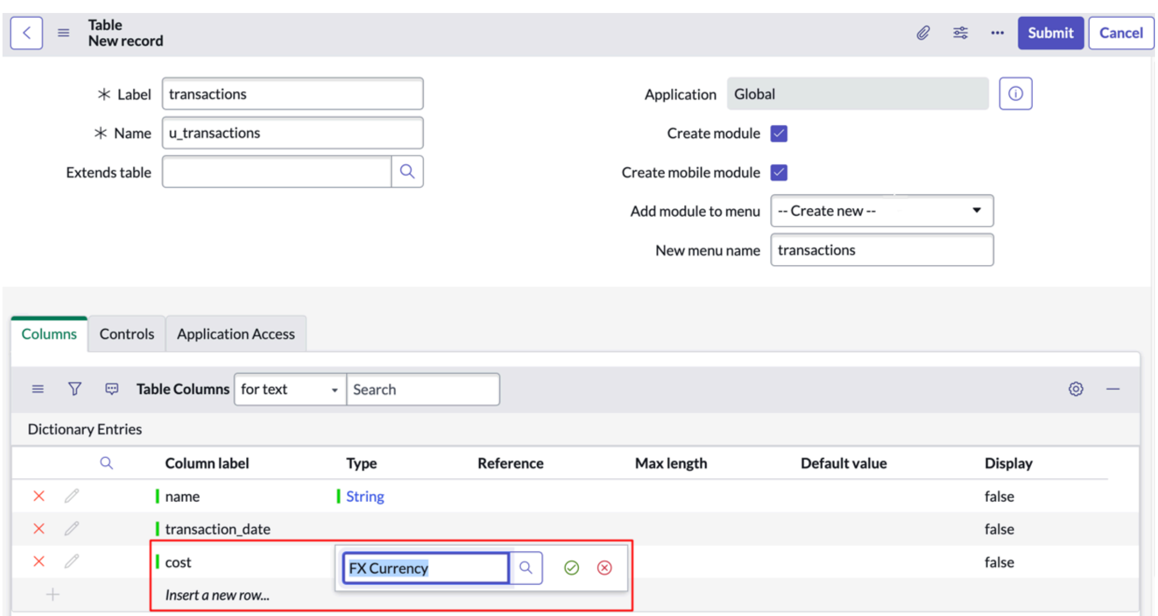

5. As needed, enter other required parameters.

Note:
Do not change the **Max length** value for **FX Currency** in custom tables. **Max length** must be 32 (default).

6. Click **Submit**.

Configuring FX Currency field behavior

After establishing global currency defaults, you can create custom FX Currency configurations for specific table fields. Identify the table fields that use FX Currency, and define how currency conversion calculations should be performed, displayed, filtered, and aggregated.

For each selected field that uses FX Currency, you configure settings that:

- Govern decimal precision and currency display in lists, forms, and reports.
- Identify the reference currency code for converting currency values entered into an FX Currency field into a common currency for report aggregations.
- Identify the rate table and data source to use for currency conversions.
- Specify the rate and target fields to use for filtering custom conversion rates.

Note:
When you create a currency configuration, it applies those settings only to the FX Currency field that you add it to. An FX Currency configuration takes precedence over global system properties that you established in **System Localization** > **All Properties.**

To learn more about global currency defaults, see Configuring FX Currency global settings.

You can change the defaults at any time. Doing so doesn't affect existing values, but it can make working with existing values difficult.

- For example, if you modify the **Reference Currency** field in the Currency tab, existing instance records have the old reference currency, but new instance records have the new reference currency. You can't perform an operation such as sum aggregation because the reference values are in different currencies.
- However, because the actual reference currency in the Currency Instance [fx_currency2_instance] table is accessible in an FX Currency field, you can use any script that performs this aggregation.

Identify the FX Currency field and its display parameters

Specify the table, and the field within the table (for example, cost), that the FX Currency configuration applies to. Then, select the decimal precision and currency display parameters for the field.

Before you begin

Role required: admin, currency_admin

Procedure

1. Navigate to **All** > **System Localization** > **FX Currency Configuration**.
2. Click **New**.
3. In the **Table** field, select the table that contains the FX Currency dictionary attribute.
   For example, if you created a custom table that stores transactional information and contains an FX Currency field, select it from the table listing. To learn more about creating tables in an instance, see Create a table.

4. In the **Field** field, select the column with the FX Currency field type.
   For example, if you created a custom table column for entry of the transaction cost, select it from the listing of fields for the table.
5. Click **Save**.
6. Click the Display tab.
7. In the **Display Digits** field, select the number of digits that should appear after the decimal point in lists and reports.

Option	Description
Currency Default	Number of digits that appear on lists and forms, based on what is ap propriate and customary for the specified currency itself. For exam ple: ◦ USD: Two digits appear after the decimal, for example, 12.33. ◦ Dinar: Three digits appear after the decimal, for example, 69.535. ◦ Yen: No digits appear after the decimal, because the Yen uses no cents and numbers round to the next latest value. For example, a calculated currency value of 69.535 rounds up to, and appears as, 67 Yen.
0–12	Specific number of digits that should appear after the decimal, with selectable values from 0 through 12. If you select more than 1, it pads the values with zeroes. ◦ For example, if you select 3, and the calculated value is 1.1, the cur rency value appears as 1.100. ◦ If you select 5, this value appears as 1.10000.

Note:
In a form, FX Currency fields display the number of digits as entered by the user, up to 12 digits in length.

8. In the **Display Value Currency** field, select what currency value appears in the designated FX Currency field on forms and lists.

Option	Description
Use Global Default	Use the global default you established in the `glide.currency2.display_value` system property in **Sys tem Localization** > **All Properties.**
Display As Entered	Currency value, as entered into the designated FX Currency field. For example, if the user locale is US Dollars, the value appears in Euros if they select the Euro as the currency during an entry session.
Display in Session Currency	Currency value entered into the designated FX Currency field, con verted to the session currency. ◦ For example, if the user locale is the United States, the session cur rency appears in US Dollars. ◦ If the user locale is Japan, the session currency appears in the Yen. To learn more about the user locale, see Locale settings.

Option	Description
Display in Reference Currency	Currency value entered into the designated FX Currency field, con verted to the reference currency designated in the **Reference Curren cy** field in the Reference tab. For example, if you enter the currency amount in Euros, it converts to and appears in US Dollars if that is the designated reference curren cy. To learn more the reference currency, see Set the reference cur rency.

9. In the **Aggregation Source** field, select the currency type to use when aggregating currency amounts for the designated field on reports.

Option	Description
Use Global Default	Use the global default you established in the *`glide.currency2.default_reference_currency`* proper ty in **System Localization** > **All Properties.**
As Entered Values	Currency values, as entered into the designated FX Currency field. **Note:** If you select this option, you may experience situations where amounts entered in dissimilar currencies can't be aggregated on reports. For example, if you enter currency amounts in US Dollars, Yen, and Euros, the amounts can't be properly aggregat ed, or may result in erroneous totals.
Reference values	Currency values, as entered into the designated FX Currency field, converted to the reference currency you designate in the **Reference Currency** field in the Reference tab. For example, if you designate USD as the reference currency, values entered into this field convert to the reference currency, regardless of the currency you entered them in. The currency values that appear on reports are the aggregated converted reference values.

10. Click **Save**.

Set the reference currency

Set the reference currency for the designated FX Currency field. The ServiceNow AI Platform uses the reference currency to convert values entered into the FX Currency field into a common currency for report aggregation purposes.

Before you begin

Roles required: admin

Procedure

1. Click the Currency tab.
2. In the **Reference Currency** field, select the three-digit ISO currency code used for the reference currency.
   For example, select **USD** if the currency values expressed in various currencies convert into US Dollars for report aggregation.

To use the global reference currency default you established in the *glide.currency2.default_reference_currency* property in **System Localization** > **All Properties.**

3. In the **Reference Currency Source** field, designate the source of the reference currency.
   - To use as the source of its reference currency when creating or updating a currency value for this FX Currency field, select the dot-walked field.
     - For example, `Cost` is an FX Currency field, and `my_currency` is a reference or string field in the Currency [fx_currency] table that contains a currency code such as USD.
     - When you create an FX Currency configuration for the `Cost` field, you select `my_currency` as its reference currency.
   - If you do not want to select the reference currency from another field, select **--None--**.
4. Click **Save**.

Specify the rate table and date source for currency conversions

Identify the rate table and data source to use for currency conversions for the designated FX Currency field.

Procedure

1. Click the Rate Table tab.
2. In the **Conversion Rate Table** field, select the rate table used for conversions of currency values in the FX Currency field:
   - To use the default system rate table you defined in the *glide.currency2.system_rate_table* property in **System Localization** > **All Properties.**, select **System Currency Conversion Rate**. This table stores daily reference rates retrieved from the European Central Bank.

     To learn more about the system rate table, see Default currency conversions.
   - If you have defined a custom rate table, you can select it to use in currency conversions for the FX Currency field. To learn more about how to define custom rate tables, see Add conversion rates using a custom rate table.
3. In the **Conversion Date Source** field, select the field to use as the source of the date and time for the currency conversion for the FX Currency field:
   - To use the date and time at which the user performs the currency conversion, select **--None--**.
   - To use as the source of the date and time for the currency conversion, select the dot-walked field.

     For example, to perform the currency conversion based on the currency rates in effect, select a transaction date field. So, if you entered the transaction on July 15, and you perform the currency conversion on July 31, it uses the currency conversion rates in effect on July 15.
4. Click **Save**.

Related topics

Understanding how FX Currency field conversions work

Add conversion rates using a custom rate table

Select the rate and target table fields used for filtering

Optionally specify the rate table and target table fields to use for filtering conversion rate information.

About this task

For example, you maintain interfaces to several ERP systems, have currency conversion rates that are unique to each ERP, but you maintain all rates in a single table.

- To identify rates for a specific ERP system, there is a **System** field in the rate table that contain entries identifying each ERP system, for example, ABC, XYZ. In the target table with the FX Currency field, you have a similarly named field for identifying the conversion rates required for specific ERP systems.
- When you perform currency conversions, and the rate and target table field entries match, it uses the rates that are unique to that specific ERP.

Procedure

1. Click the Rate Filter tab.
2. In the **Rate Table Field** field, select the field in the rate table to match against the specified target table field for conversion rate filtering.
   Select **--None--** skip use of a rate table field for conversion rate filtering.
3. In the **Target Table Field** field, select the field in the target table to match against the specified rate table field for rate filtering.
   Select **--None--** skip use of a target table field for rate filtering.
4. Click **Save**.

Related topics

Understanding FX Currency values in lists and reports

Understanding FX Currency values in forms

Updating the Currency Instance record

If you have an assigned role of currency_instance_admin, you can review and edit the details of a currency entry in an FX Currency field. You can't edit the reference currency and value because they are calculated automatically.

Before you begin

Role required: admin, currency_instance_admin, with ACL access to the FX Currency field.

Procedure

1. Access the Currency Instance record using one of the following methods:

2. On the form, fill in the fields.

Currency Instance

Field	Description
Amount	Currency amount that you entered into the FX Currency field, formatted according to the currency code in your user locale. The amount stores internally as a currency code and a numerical value of up to 12 digits in length. For example, `USD; 1234.56`.

Field	Description
	Note: The currency value has two parts: the currency amount and a currency code. The amount stores as a decimal value, for example, `1234.56`.
Currency Code	Three-letter ISO code that identifies the currency that you select for the currency amount. For example, USD for the US Dollar, or EUR for the Euro.
Conversion Rate	Conversion rate used to calculate the reference currency value. It converts a session currency value, expressed in one currency, to another, and stores as a reference to a record in the Currency Conversion Rate [fx_conversion_rate] table. The conversation rate is based on the **Conversion Rate Table** and **Conversion Date Source** settings for the FX Currency field in the Rate tab of FX Currency Configuration. To learn more about these settings, see Configuring FX Currency field behavior. You change or select the conversion rate as needed.
Reference Amount	Calculated reference currency amount. The ServiceNow AI Platform converts the currency value entered into the **Amount** field into a common reference currency. It uses the settings in the **Reference Currency Source** and **Reference Currency** fields in **System Localization** > **FX Currency Configuration** to do so. You can't change the calculated reference amount. **Note:** To learn more about reference currencies, see Set the reference currency.
Reference Currency Code	Three-letter ISO code that identifies the reference currency, for example, USD, to which the entered currency value converted. You can't change the reference currency code.

3. Click **Save**.

Currency administration references

The following sections show the properties associated with currency administration.

Life cycle of records containing FX Currency fields

The behavior of FX Currency fields varies during the processing that occurs during the lifetime of a record containing them.

Insert / Update

The FX Currency field points to a Currency Instance [fx_currency2_instance] record. When you change the currency value in an FX Currency field, it determines the conversion rate and calculates the reference currency before:

- The *before* business rules run.
- The *after* business rules run, and it includes any further changes you may have made to the *before* business rules.

Note:
If the FX Currency field contains an invalid currency code, an exception condition may appear before these two stages occur.

Auditing

Since an FX Currency field points to a Currency Instance record that stores multiple values, the audit string is a composite that contains this information. The string stored in the System Audit [sys_audit] table is in the format of `EUR;111.222;4555525f5553445f3230313931323033`, with the following values, separated by semicolons:

- Three letter ISO currency code. For example, `EUR`.
- Amount as an unformatted number. For example, `111.222`.
- System identifier (sys_id) for the conversion rate record in the Currency Conversion Rate [fx_conversion_rate] table. For example, `4555525f5553445f3230313931323033`.

When creating history lines for a user, the audited string is formatted, using the locale of the user. It's in the format of `€111.22;2019-12-03 17:00:00-3000-01-01 23:59:59;fx_system_rate`, with the following values, separated by semicolons:

- Formatted currency string in the user locale. For example, `€111.22`.
- Start span-End span, as expressed in the conversion rate record in the System Rate [fx_system_rate] table. For example, `2019-12-03 17:00:00-3000-01-01 23:59:59`.
- Name of the conversion rate table. For example, `fx_system_rate`.

Dot-walkable Currency Instance fields

You can dot-walk certain fields in the Currency Instance [fx_currency2_instance] record, and field values stored in the database are consistent with each other. However, in script, since fields can be changed individually, and you change only some of the fields, they can be inconsistent.

Currency code

The **Currency** field isn't a reference field, so you can set it to any value. However, if you enter an invalid currency code, an exception message is generated. An empty value is considered a session currency code.

Query conditions

You can add query conditions on dot-walked fields in the instance table. However, when the query condition value is in a special format with currency code and amount, it's treated as a compound condition on the dot walked fields. The value has to be in format `XYZ,abc` where:

- `XYZ` is a currency code, and
- `abc` is an amount.

For example, `USD, 12.34`.

A condition such as `cost>USD,12.34` is treated as `cost.currency>USD AND cost.amount12.34`.

When the second operand is another Currency2 field, the condition translates in a similar manner:

`cost1>cost2` is treated as `cost1.currency=cost2.currency AND cost1.amount>cost2.amount`.

ServiceNow AI Platform translation and localization

Provide service around the globe with language translation and localization for specified regions.

Localization Framework	System Localization	Dynamic Translation
Standardize the process of translating artifacts such as Service Catalog items, Knowledge Base articles, and Document Templates across the platform.	Set up location-based translation and currency support for users on your instance, wherever they are in the world.	Translate user-generated content live on your instance to streamline communication for your agents and end users.

For additional resources about globalization tools on the ServiceNow AI Platform, see the Localization page in the ServiceNow Customer Success Center.

Localization Framework

The Localization Framework app uses configurable workflows to standardize the process of translating artifacts across the platform. Translate Service Catalog items, Knowledge Base articles, Document Templates, and more.

System localization

Give your users specialized support with System Localization. Define locales and adjust localization settings to display specific time/date and currency to users. Multiple languages are available on the ServiceNow AI Platform, along with options to create custom translations for ServiceNow materials, as well as for languages not translated on the platform.

Dynamic Translation

Translate user-generated content live on your instance to streamline communication for your agents and end users. Integrate with the machine translation service provider of your choice to enable translation on forms, activity streams, and chats. Microsoft Azure and Google Cloud Translator Service are partially preconfigured with the activation of Dynamic Translation, or you can configure a custom service provider using Integration Hub.

Localization Framework

The ServiceNow® Localization Framework application standardizes and enhances the translation process of artifacts across the platform through configurable workflows. Such artifacts include ServiceNow® Service Catalog items, Virtual Agent topics, Natural Language Understanding models, and Document Templates of ServiceNow® HR Service Delivery.

Get started

Explore Localization Framework Learn how the Localization Framework handles translations and artifacts.	Configure Localization Framework Set up and implement localization for your instance.	Request translations Make requests for translations and localization.
Create translation projects Manage and approve projects, and carry out translations for single or bundled localization tasks and artifacts.	Monitoring Localization Framework Track Localization Framework performance and behavior.	Localization Framework reference Get familiar with specific terminology and other Localization Framework reference materials.

Troubleshoot and get help

- Localization in the ServiceNow Customer Success Center
- ServiceNow Community forums
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Explore Localization Framework

Learn more about the Localization Framework app that enables a systematic localization process, so you can provide a seamless user experience in multiple languages.

Localization Framework overview

The Localization Framework streamlines the process of adapting and translating ServiceNow instances to meet the needs of different regions and languages. With this application, users can create and manage translation projects, assign tasks to translators, and track progress. The Localization Framework also supports the creation of language packs, which contain translated content for specific languages. Organizations can ensure that their ServiceNow instances are accessible and user-friendly for users around the world by reducing the time and effort needed for localization.

The localization process

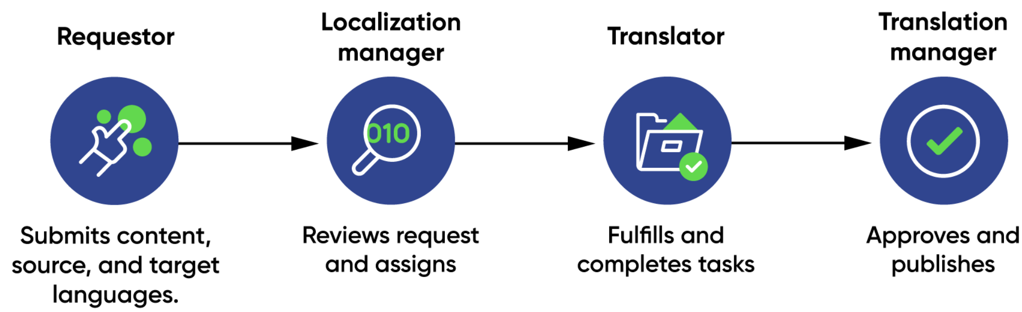

Localization Framework benefits

Benefit	Feature	Users
Define and use custom language artifacts, in addition to installing artifact-specific product plugins.	Artifact configurations	Administrators
Tailor Localization Framework workflows and translation preferences	Localization Framework settings	Administrators
Request translations	Request translations	Administrators, editors,managers, and requestors
Manage translation projects	Create a localization project	Managers
Bundle translation projects together	Add a localization request item to a project	Managers
Perform translation projects	Fulfill a localization task	Fulfillers
Approve and publish localizations and translations	Approve a localization task	Fulfillers
Use artifacts in localization, with support for Virtual Agent and several other ServiceNow products.	Supported artifacts in Localization Framework	Administrators, managers, and users

Note:
Localization Framework does not support Domain Separation.

Supported artifacts in Localization Framework

Localization Framework extends its framework capabilities to certain artifacts and translates the content of the artifact items.

With the activation of the Localization Framework Installer plugin (com.glide.localization_framework.installer), some artifacts are available by default. Some artifacts require the installation of the artifact-specific product plugins.

Artifacts available with Localization Framework installation

The following artifacts are available with the installation of Localization Framework:

- Catalog Item (Service Catalog)
- NLU Model (Natural Language Understanding Models)
- Surveys (Assessments and Surveys)

Artifacts available with artifact-specific plugins

The following artifacts are available with the installation of artifact-specific product plugins after the installation of Localization Framework.

- Virtual Agent Topic - available with the activation of the Glide Virtual Agent Lite plugin (com.glide.cs.chatbot.lite).
- Document Template Block Content - available with the activation of the Human Resources Scoped App: Lifecycle Events for Enterprise plugin [com.sn_hr_lifecycle_ent].
- HTML Document Template - available with the activation of the Human Resources Scoped App: Lifecycle Events for Enterprise plugin [com.sn_hr_lifecycle_ent].
- Knowledge Base article - available with the activation of Knowledge Management Advanced plugin [com.snc.knowledge_advanced] and set `glide.knowman.translation.enable_lf_article_translation` property to true.
- Email notifications, email templates and email layout - available with the activation of Glide Notification Translation plugin (com.glide.notification.translation).

Localization Framework support for Service Catalog items

Localization Framework enables translation of Service Catalog items when a catalog item is requested for translation into a language and a localization requested item (LRITM) is created.

Localizing Service Catalog items

You can do the following tasks in Localization Framework for translating the Service Catalog items.

- Select multiple languages from the available languages in the instance to configure settings with the Catalog Item artifact. To configure localization settings, see Localization Framework settings.
- Create LRITM's by requesting translations into one or more languages.
- Translate the content of the Service Catalog items directly using the edit translations option.
- Read, edit, and machine translate the content, send the content to TMS, send the content via email, export, and import the content of the Service Catalog item.

  Note:
  Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework.

- Approve and publish the translated content.
- Monitor the overall localization health of the Service Catalog items from the Localization Insights dashboard.

To request translations for a catalog item, see Request translations for Service Catalog items.

Localization Framework support for catalog builder in Service Catalog

Localization Framework enables translation of Service Catalog items directly from the catalog builder dashboard in Service Catalog.

You can do the following tasks on the Catalog Builder dashboard in Service Catalog when the Localization Framework installer plugin (`com.glide.localization_framework.installer`) is installed:

- Use the **Translation status** tab that is displayed when a Service Catalog item is selected, to perform the following tasks:
  - View the list of all active languages.
  - Request the translation of the Service Catalog items directly using the **Request Translation** option if the `glide.sc.catalog_builder.localization.enable_request_translation_mode` property is set to true.

    Note:
    You can request translations for the items in the active languages only if the Localization Framework setting is available for the particular language.
  - View the translation status of the Service Catalog items.

    Note:
    The translation status of the Service Catalog items is based on the status and completion of the Localization Framework scheduler jobs.
  - Use the **Edit Translation** option in case the `glide.sc.catalog_builder.localization.enable_request_translation_mode` property is set to false, so that you can machine translate and publish the translations on the ad hoc UI.

    Note:
    Machine translate option is only available if the Localization Framework setting is available for that particular language.
- Request the translations of the Service Catalog items when creating the items.

For more information on the Localization Framework integration with catalog builder, see Localization Framework integration with Catalog Builder.

Virtual Agent Designer integration with Localization Framework

The integration of Virtual Agent Designer with Localization Framework enables you to translate Virtual Agent topics using the translation modes supported by Localization Framework with or without a localization task.

Localizing Virtual Agent topics

The integration of Virtual Agent Designer with Localization Framework enables you to do the following tasks.

- Select multiple languages from the available languages in the instance to configure settings with the Virtual Agent artifact. To configure localization settings, see Localization Framework settings.
- Translate the content directly from the Virtual Agent Designer using the edit option.

- Read, edit, machine translate, export, and import the content of the Virtual Agent topic from the Virtual Agent Designer if the Virtual Agent artifact is configured in the localization settings.

  Note:

  - **Note:**
    You need the localization_editor role to edit the translated content from the comparison UI of the Virtual Agent Designer directly. For more information about the localization_editor role, see Localization Framework Roles.
  - Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework.
- Request translations from Virtual Agent Designer.

  Note:
  Requesting translations creates localization tasks. You need the localization_requester role to request translations.
- Approve and publish the translated content.
- Monitor the overall localization health of the Virtual Agent topics from the Localization Insights dashboard.

For more information about localizing the Virtual Agent topics in the Virtual Agent Designer, see Virtual Agent Localization.

Localization Framework support for Natural Language Understanding models

Localization Framework integrates with NLU Workbench to translate the Natural Language Understanding models along with their associated intents, utterances, and entities into multiple languages.

Localizing Natural Language Understanding models

The integration of Localization Framework with NLU Workbench enables you to do the following tasks.

- Auto-configure the NLU Model artifact in the Artifacts Configuration [sn_lf_config] table with the activation of the Localization Framework plugin (com.glide.localization_framework.installer).
- Select multiple languages from the available languages in the instance to configure settings with Natural Language Understanding models artifacts. To configure localization settings, see Localization Framework settings.
- Request translations for Natural Language Understanding Models and translate the content of the localization task using the translation modes.

  Note:
  Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework.
- Translate the content of the Natural Language Understanding models directly from the NLU Workbench.

Note:
Only users with the localization_editor role can translate content on the comparison UI of the Natural Language Understanding model in the NLU Workbench. For more information about the localization_editor role, see Localization Framework Roles.

- Approve and publish the translated content.

Note:
The localization of Natural Language Understanding models requires mapping between the Natural Language Understanding models in one language and their counterpart in other languages to use the localization framework support.

For more information about localizing Natural Language Understanding models, see Translate NLU Models.

Localization Framework support for Surveys

Localization Framework enables the translation of surveys items when a translation request for a survey is created and when a localization requested item (LRITM) is created.

Localizing Surveys

The extension of the Localization Framework capabilities to surveys enables you to do the following tasks:

- Select multiple languages from the available languages in the instance to configure settings with the survey artifact. To configure localization settings, see Localization Framework settings.
- Create LRITMs by requesting translations into one or more languages.
- Translate the content of the surveys directly using the edit translations option.
- Read, edit, and machine translate the content. You can also send the content to TMS or via email. You can export and import the content of the survey elements.

  Note:
  Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework. For more information, see Exploring Dynamic Translation.

- Approve and publish the translated content.
- Monitor the overall localization health of the surveys from the Localization Insights dashboard.

For more information on requesting and editing translation for surveys, see Request translations for surveys and Edit translations for surveys.

Localization Framework support for Knowledge Base

Localization Framework enables the translation of Knowledge Base (KB) articles when a KB article is requested for translation into a language and when a localization requested item (LRITM) is created.

Localizing Knowledge Base

The extension of the Localization Framework capabilities to KB articles enables you to do the following tasks:

- Select multiple languages from the available languages in the instance to configure settings with the knowledge base artifact. To configure localization settings, see Localization Framework settings.
- Create LRITMs by requesting translations into one or more languages.
- Translate the content of the KB articles directly using the request translations option.
- Read, edit, and machine translate the content. You can also send the content to TMS or via email. You can export and import the content of the KB articles.

  Note:
  Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework.

- Approve and publish the translated content.
- Translate Knowledge Base articles in bulk in multiple languages.

For more information on translating the KB articles, see Knowledge base article localization.

Localization Framework support for email notifications

Localization Framework enables the translation of emails (notification, template, and layout) when an email artifact is requested for translation into a language and when a localization requested item (LRITM) is created.

Localizing Email Notifications

The extension of the Localization Framework capabilities to email enables you to do the following tasks:

- Select multiple languages from the available languages in the instance to configure settings with the email notifications, email template and email layout artifacts. To configure localization settings, see Localization Framework settings.
- Create LRITMs by requesting translations into one or more languages.
- Translate the content of the email notification directly using the edit translations option.
- Read, edit, and machine translate the content. You can also send the content to TMS or via email. You can export and import the content of the email notification.

  Note:
  Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework.

- Approve and publish the translated content.

For more information about localizing email notification, see Email notification localization.

Localization Framework support for HR Service Delivery

Localization Framework extends its framework capabilities to Document Templates and Document Template Block Contents in HR Service Delivery.

Localizing Document Templates

The extension of the Localization Framework capabilities to HR Service Delivery enables you to do the following tasks:

- Select multiple languages from the available languages in the instance to configure settings either with the Document Template Block Content or HTML Document Template artifacts. To configure localization settings, see Localization Framework settings.
- Translate the template content manually, through machine translation, or by using export and import.

  Note:
  Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to use machine translation in Localization Framework.

- Publish the translation.

For more information about translating the Document Templates, see Use Localization Framework for Document Templates.

Localization Framework Roles

Localization Framework uses different roles to manage different aspects of the application.

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

Localization administrator [localization_admin]

Manage the following aspects of the Localization Framework application.

- Configures the supported artifacts and custom artifacts for Localization Framework.
- Configures localization settings for the activated languages.
- Can request translations into multiple languages for Service Catalog items, Virtual Agent topics, or for any custom artifacts from their respective artifact tables.
- Can request translations for Service Catalog items, Virtual Agent topics, or for any custom artifact for a particular language from their respective artifact reports in the insights dashboard.
- Can configure RWS and XTM TMS and can create and configure a custom TMS by inheriting the following roles:
  - flow_designer: Can edit the TMS subflows and can create a new TMS configuration record.
  - connection_admin: Can access the connections and credentials tables and records.
- Owns the localization insights dashboard and reports. Can execute the scheduled jobs to display the reports on the dashboard.
- Views the localization health of the system.
- localization_admin role inherits the lcoalization_requestor, localization_editor, and localization_fulfiller role.

Contains Roles

List of roles contained within the role.

- localization_fulfiller
- localization_manager
- localization_requestor
- localization_editor

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Localization requester [localization_requestor]

Request translations for all the configured artifacts into one or more languages.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Note:
Localization requester can fulfill translations by inheriting the localization_fulfiller role.

Localization fulfiller [localization_fulfiller]

Provide translations for the requested artifact items and approve and publishe the translations.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Fulfillers can translate a localization task in the following ways:

- Use manual translation if the fulfiller understands both the source and the target language.
- Use machine translation by using the translator configured in the translation preferences.
- Exports data to a third-party translator and imports it back after translations.
- Sends the data via email for translations and receives the translated data back via email.
- Sends translatable data to a translation management system (TMS) for translations.

- Verifies and approves the translations.
- Publishes the translations if the language setting workflow has no approval flow.

Localization editor [localization_editor]

Provides the ability for the users to translate the content without the need of a localization task.

Contains Roles

The Localization editor [localization_editor] role contains the Localization requestor [localization_requestor] role.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Users assigned to this role can translate content from Virtual Agent Designer NLU Workbench, and Document Templates of HR Service Delivery.

A localization editor can translate the content using machine translation, and also by exporting and importing the content.

Localization manager [localization_manager]

Owns the localization projects. Adds all the localization requested items to the projects. Only a Localization manager has the authority to start the project.

Contains Roles

The Localization manager [localization_manager] role contains the Localization fulfiller [localization_fulfiller] role.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

If the Localization Framework Hub and Spoke architecture is set up, the localization_manager role has access to the following tables on the hub instance:

- Localization Framework Hub Translation Request [sn_lf_hub_translation_request]
- Localization Framework Hub Translation Item [sn_lf_hub_translation_item]

For more information, see Localization Framework Hub and Spoke architecture

Workflows in the Localization Framework

Workflows define actions for completing a localization task. Each localization task undergoes the workflow configured for its target language to complete the task.

Workflows in the Localization Framework

Workflows	Description
Translation → Publish	The localization task undergoes translation manually or by using one of the translation modes. After translation, the content gets published manually. **Note:** Ensure that you assign the **Localization Fulfiller Group** for the workflow to localize the content.
Translation → Translation Approval → Auto Publish	The localization task undergoes translation manually or by using one of the translation modes. After translation, the task undergoes translation approval and gets auto-published. **Note:** Ensure that you assign the **Localization Fulfiller Group** and **Translation Approver Group** for the workflow to localize the content.
Auto Translation → Translation Approval → Auto Publish	The localization task undergoes auto-translation. After translation, the task undergoes translation approval and gets auto-published. **Note:** Ensure that you assign the **Localization Fulfiller Group**, **Auto Translation Mode**, and **Translation Approver Group** for the workflow to localize the content.
Auto Translation → Publish	The localization task undergoes auto-translation and gets published manually. **Note:** Ensure that you assign the **Localization Fulfiller Group** and **Auto Translation Mode** for the workflow to localize the content.
Auto Translation → Auto Publish	The localization task undergoes auto-translation and gets auto-published. **Note:** Ensure that you assign the **Localization Fulfiller Group** and **Auto Translation Mode** for the workflow to localize the content.

Workflows in the Localization Framework (continued)

Workflows	Description
Business Approval → Translation → Publish	The localization task needs business approval for localization. The content gets translated manually or translated by using one of the translation modes. The content gets manually published. **Note:** Ensure that you assign the **Business Approver Group** and **Localization Fulfiller Group** for the workflow to localize the content.
Business Approval → Translation → Translation Approval → Auto Publish	The localization task needs business approval for localization, and the content gets translated manually or by using one of the translation modes. After translation, the task undergoes translation approval and gets auto-published. **Note:** Ensure that you assign the **Business Approver Group**,**Localization Fulfiller Group**, and **Translation Approver Group** for the workflow to localize the content.
Business Approval → Auto Translation → Translation Approval → Auto Publish	The localization task needs business approval for localization, and the content gets auto-translated. After translation, the task undergoes translation approval and gets auto-published. **Note:** Ensure that you assign the **Business Approver Group**,**Localization Fulfiller Group**, **Auto Translation Mode**, and **Translation Approver Group** for the workflow to localize the content.
Business Approval → Auto Translation → Publish	The localization task needs business approval for localization. The content gets auto-translated but gets manually published. **Note:** Ensure that you assign the **Business Approver Group**,**Localization Fulfiller Group**, and **Auto Translation Mode** for the workflow to localize the content.
Business Approval → Auto Translation → Auto Publish	The localization task needs business approval for localization. The content gets auto-translated and auto-published.

Workflows in the Localization Framework (continued)

Workflows	Description
	Note: Ensure that you assign the **Business Approver Group**, **Localization Fulfiller Group**, and **Auto Translation Mode** for the workflow to localize the content.

Configure the Localization Framework

Plan and configure your Localization Framework implementation.

Configuration overview

After activating Localization Framework, you can create custom artifacts and processing scripts, and adjust preferences for email, workflows, and other settings

Activating the Localization Framework

When you activate the Localization Framework Installer plugin (com.glide.localization_framework.installer), Localization Framework is available by default for the Service Catalog items and Virtual Agent topics, Natural Language Understanding models, and Document Templates for HR Service Delivery.

After you activate the Localization Framework Installer plugin (com.glide.localization_framework.installer), you can:

- Create configurations to onboard any other application to Localization Framework.
- Request translations of the artifacts into one or more languages.
- Translate content of the artifacts using manual translation or any of the supported translation modes.
- Approve translations by proofreading the translated content.
- Publish the translated content.
- Analyze the localization health of the system using the Localization Insights dashboard.
- Install optional plugins to set up Localization Framework Hub and Spoke architecture.

For more information on activating this plugin, see Activate Localization Framework.

Activate Localization Framework

You can activate the Localization Framework Installer plugin (com.glide.localization_framework.installer) for Localization Framework if you have the admin role. The application installs related plugins if they are not already installed. Optionally, you can activate plugins for the Localization Framework Hub and Spoke architecture on a set of instances.

Before you begin

Role required: admin

About this task

The Localization Framework Installer plugin (com.glide.localization_framework.installer) can be activated in two ways:

- Manual installation of the plugin.
- Auto-installation of the plugin.

  When a language plugin is installed, the Localization Framework Installer plugin (com.glide.localization_framework.installer) is automatically installed with the language plugin. For more information, see language plugins.

The following items are installed with Localization Framework :

- Plugins
- Roles
- Scheduled jobs
- Tables

For more information, see Components installed with Localization Framework.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Localization Framework Installer plugin (com.glide.localization_framework.installer) using the filter criteria and search bar.

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

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

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

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

What to do next

Optionally, next you can install plugins for the Localization Framework Hub and Spoke architecture. This enables you to designate one instance as a hub to fulfill localization requests, and one or more instances as spokes that submit requests to the hub.

Decide which instance should function as the hub. On that instance, follow the above procedure to search for and install the Localization Framework Hub plugin [com.sn.localization_framework.hub].

Decide which instance(s) should function as spokes. On each of those instances, follow the above procedure to search for and install the Localization Framework Spoke plugin [com.sn.localization_framework.spoke].

Components installed with Localization Framework

The Localization Framework Installer plugin (com.glide.localization_framework.installer) activates several related plugins that are not already active. With installation of the core plugins, optional plugins for the Localization Framework Hub and Spoke architecture are available.

Plugins available with Localization Framework Installer plugin (com.glide.localization_framework.installer)

Plugin	Description
Localization Framework [com.glide.localization_framework]	Manage localization workflows across the platform and products.
Localization Framework for Service Catalog [com.glide.localization_framework.service_catalog]	Enables translation of Service Catalog items using Localization Framework.
Localization Framework for Virtual Agent Topic [com.glide.localization_framework.va]	Enables translation of Virtual Agent topics using Localization Framework.
RWS Translation Management System Spoke for Localization Framework [com.glide.localization_framework.sdl_spoke]	Integrates Localization Framework with the RWS Translation Management System.
XTM Translation Management System Spoke for Localization Framework [com.glide.localization_framework.xtm_spoke]	Integrates Localization Framework with the XTM Translation Management System.

Optional plugins for the Localization Framework Hub and Spoke architecture

Localization Framework Hub [com.sn.localization_framework.hub] Scope: sn_lf_hub	It's a framework that enables to fulfill the translation requests received from Localization Framework Spoke. After installation and configuration, the hub instance manages incoming translation requests. For more information, see Localization Framework Hub and Spoke architecture.
Localization Framework Spoke [com.sn.localization_framework.spoke] Scope: sn_lf_spoke	It's a framework that enables to fulfill the translation requests using Localization Framework Hub. After installation and configuration, the spoke instance submits its translation requests to the hub.

Artifact configurations

You can configure the required artifacts and use the Localization Framework to fulfill the translation requests.

Any artifact configured in the Localization Framework application has its own record in the artifact configurations [sn_lf_config] table.

With the activation of the Localization Framework Installer plugin (com.glide.localization_framework.installer), a record for Service Catalog item and Natural Language Understanding Model are automatically created and configured in Localization Framework.

A record for the Virtual Agent topic is also created and configured in the framework, since Virtual Agent is automatically enabled with the base system. For more information, see Virtual Agent Lite.

Records for Document Template Block Content and HTML Document Template are created in the artifact configurations [sn_lf_config] table when you activate the Human Resources Scoped App: Lifecycle Events for Enterprise plugin [com.sn_hr_lifecycle_ent]. For more information about activating this plugin, see Activate Lifecycle Events for Enterprise.

You can also create and configure an artifact of your choice and a record for the same is created in the artifact configurations [sn_lf_config] table. For more information, see Create and configure a custom artifact.

Create and configure a custom artifact

Create and configure custom artifacts of your choice to use the Localization Framework functionality.

Before you begin

Ensure that you have activated the Localization Framework Installer plugin (com.glide.localization_framework.installer) before configuring a custom artifact.

Role required: localization_admin and script_include_admin.

Procedure

1. Navigate to **All** > **Localization Framework** > **Artifact Configurations**.
2. Select **New**.
3. In the Artifact Configuration form, fill in the fields.

Artifact Configuration form fields

Field	Description
Name	Identifier of the artifact record.
Internal name	An internal name for the artifact.
Table name	Identifier of the table where information of the artifact is stored. For example, Catalog Items [sc_cat_item] table.
Processor script	Extracts the translatable content associated with the translation request and saves the translations into the database.You can select a script include using the search icon, or you can create a new script include. The processor script contains the following functions: ◦ getTranslatableContent: Extracts the translatable content. ◦ saveTranslatedContent: Saves the translated content back to the system. For more information, see Processor script functions.To create a new processor script for the artifact, see Create a processor script. **Note:** Under **Related Links** on an existing artifact record, you can switch to **Legacy View** for script editing.

Field	Description
Application	Application scope. The value is the current application selected for the logged-in user session. For more information, see .
Active	Option to activate the configuration record.
Generate insights	Option to turn on or turn off the localization insights for the artifact. For more information, see Localization Insights dashboard.

4. Select **Submit**.

What to do next

For examples of Localization Framework artifacts that you can set up, see the LF Artifact blog posts list on https://www.servicenow.com/community/international-localization/need-to-translate-a-par-dashboard-check-this/ta-p/2839751 in the ServiceNow Community. These examples include Processor scripts.

Create a processor script

You can create a processor script to include in your custom artifact configuration to extract the translatable content and save the translated content.

Before you begin

Role required: localization_admin and script_include_admin

Procedure

1. Navigate to **All** > **Localization Framework** > **Artifact Configurations**.
2. Click **New** or open an artifact that you want to update the processor script for.
3. Click the **Create Processor Script** related link.
4. In the **Script include name** field, provide the script include name to integrate the artifact with Localization Framework.
5. Click **Submit**.

Processor script functions

The processor script contains functions that help you extract the translatable content and save the translated content which is included when you configure an artifact.

getTranslatedContent(params)

This function is responsible for extracting the translatable content for the Localization Framework to use and consume that content and contains the following parameters:

getTranslatedContent(params)

Name	Type	Description
tableName	String	The table name of the artifact record.
sysId	String	The sys_id of the artifact record.
language	String	The language into which the artifact has to be translated (target language).

Note:
This script function is available by default when you configure the processor script. This script function requires that you configure its parameters to build the LFDocumentContent object.

saveTranslatedContent(documentContent)

This function is responsible for saving the translations back to the system and has a parameter which is an object of LFDocumentContent.

All the processor scripts extend LFArtifactProcessorSNC which defines the default implementation of this function, that is, saving the translations to standard tables. Defining this function in the processor script overrides the default behavior.

Note:
This script function is commented by default when a new processor script is created. For more information, see Read script helper functions for LFDocumentContentBuilder.

Create a UI action for the custom artifact

Create a UI action to request translations for the custom artifact.

Before you begin

Role required: localization_admin

Procedure

1. Navigate to the custom artifact table.
   Find the table using its table name. For example, **[sys_cs_topic]** is the table name for Virtual Agent topics.
2. Click the hamburger icon (≡) at any of the columns and navigate to **Configure UI Actions**.
3. Click **New**.
4. In the UI Actions form, fill in the fields.

UI Action form fields

Field	Description
Name	Name for the UI action. For example, **Request Translations**.
Client	Option to execute the script in the user's browser. When enabled, the **Onclick** field appears above the Condition field.
Condition	Condition to decide when to show this UI action.
Onclick	Provide the JavaScript function name as: `renderLanguagePickerModal();`
Script	Option to execute the script to run the request translations function. The script varies for **Form button** and **List choice**.

Field	Description
	○ Form Button `function renderLanguagePickerModal() {` `    var dlg = new GlideModal("sn_lf_language_picker");` `    dlg.setTitle(getMessage("Request Translations"));` `    dlg.setPreference("sys_id", g_form.getUniqueValue());` `    dlg.setPreference("artifact_config_internal_name", "catalog_item");` `    dlg.setPreference("request_type", "form");` `    dlg.setPreference("focusTrap", true);` `    dlg.render();` `}` ○ List choice `function renderLanguagePickerModal() {` `    var sys_ids = g_list.getChecked();` `    if (!sys_ids || sys_ids.length == 0)` `        return;` `    var dlg = new GlideModal("sn_lf_language_picker");` `    dlg.setTitle(getMessage("Request Translations"));` `    dlg.setPreference("sys_id", sys_ids);` `    dlg.setPreference("artifact_config_internal_name", "catalog_item");` `    dlg.setPreference("request_type", "list");`

Field	Description
	`` dlg.setPreference("focusTrap", true); dlg.render(); } ```
Form button	Option to enable the UI action as a form button.
List choice	Option to enable the UI action in the choice list.

Note:
Configure the form to see all the fields. For more information, see Create a UI action.

5. Click **Submit**.

Read script helper functions for LFDocumentContentBuilder

The `LFDocumentContentBuilder` provides utility functions to build document content. Use these functions to write the logic to extract the translatable content from any artifacts regardless of them using standard tables or not for translations.

initialize (version, targetLanguage, sys_id, tableName)

Use this function as a constructor for `LFDocumentContentBuilder` and includes the following parameters:

Name	Type	Description
Version	String	The version type of the document content to be generated. For Washington DC, this value is v1.
targetLanguage	String	The target language of the localization task document content. The target language is the same as the language available from the read script function argument.
Sys_id	String	The sys_id of the artifact record.
tableName	String	The table name of the artifact record.

processString (string, groupName, label)

Use this function to add the LFDocumentContentElement in the document content for the given string and includes the following parameters.

Name	Type	Description
string	String	The value for which a document content entry is created.
groupName	String	The name of the group to which element should belong.
label	String	The unique label identifier for the created element.

processStringArray (stringArray, groupName, label)

Use this function to add elements (object of LFDocumentContentElement) in the document content for the given strings and includes the following parameters.

Name	Type	Description
stringArray	List of Strings	The list of values for which document content entry is created.
groupName	String	The name of the group these strings belongs to.
label	String	The unique label identifier for these strings.

processScript (scriptContent, groupName, label)

Use this function to add an element in the document content for the strings present as input of getMessage, gs.getMessage, or gs.getMessageLang in the provided Script. It has the following parameters:

Name	Type	Description
scriptContent	String	The Script to be processed to get the strings present as input of getMessage, gs.getMessage, or gs.getMessageLang.
groupName	String	The name of the group for this input.
label	String	The unique label identifier for this input.

For example,

```
var documentContentBuilder = new LFDocumentContentBuilder("v1",
 language, sysId, tableName);
Var name = "name";
documentContentBuilder .processScript("gs.getMessage('Hello {0},
 How are you', name)", "Script_Group", "Script_Label");
```

processTranslatableFieldsForSingleRecord (glideRecord, groupName)

Use this function to create elements for the given glideRecord object, in document content for all its translatable fields. These translatable fields can be of translated_field, translated_text or translated_html types. It includes the following parameters:

Name	Type	Description
glideRecord	GlideRecord	The GlideRecord object to be processed for translatable fields. The GlideRecord should be pointing to the valid table record.
groupName	String	The name of the group for this input.

processTranslatableFieldsForMultipleRecords (tableName, encodedQuery, groupName)

Use this function to fetch all the qualified records for the given tableName and encodedQuery and adds elements (object of LFDocumentContentElement) in the document content for all the translatable fields. It includes the following parameters:

Name	Type	Description
tableName	String	The table to fetch records.
encodedQuery	String	The encoded query for filtering records from the processTranslatableFieldsForSingleRecord table.
groupName	String	The name of the group for this input.

processElement()

Use this function to add an element (object of LFDocumentContentElement) to LFDocumentContent object.

element	LFDocumentContentElement	Element to be added to the document content.

build()

Returns the LFDocumentContent object.

Note:
The getFinalJSON function has been deprecated in the Washington DC release and it is recommended to use the `build()` function.

A sample script to build the document content.

```
getTranslatableContent: function(params) {
    var tableName = params.tableName;
    var sysId = params.sysId;
    var language = params.language;
    var lfDocumentContentBuilder = new
 global.LFDocumentContentBuilder("v1", language, sysId,
 tableName);
    // This will create a new object of LFDocumentContent and
 stores it internally

 lfDocumentContentBuilder.processTranslatableFieldsForSingleReco
rd(gr, "Basic Info");
    var field =
 LFDocumentContentHelper.createField(originalValue,
 translatedValue);
    field.setTextType(LFDocumentContentHelper.PLAIN_TEXT_TYPE);
    // All the other relevant attributes can be set
    var element =
 LFDocumentContentHelper.createElement(groupName, label);
    element.addField(field);
    lfDocumentContentBuilder.processElement(element); // Adds
 element to the LFDocumentContent object
    return lfDocumentContentBuilder.build(); // return the
 LFDocumentContent object
}
```

Sample script to extract data from the LFDocumentContent

```
saveTranslatedContent: function(documentContent) {
    // LFDocumentContent object is passed as an argument to the
 saveTranslatedContent
    var targetLanguage = documentContent.getLanguage();
    var version = documentContent.getVersion();
    var elements = documentContent.getElements(); // Array of
 LFDocumentContentElement objects
    for (var idx = 0; idx < elements.length; idx++) {
        var element = elements[idx];
        var groupName = element.getGroupName();
        var fields = element.getFields(); // Array of
 LFDocumentContentField objects
        for (var fieldIdx = 0; fieldIdx < fields.length;
 fieldIdx++) {
            var field = fields[fieldIdx];
            var originalValue = field.getOriginalValue();
            var translatedValue = field.getTranslatedValue();
            // Get all the other required members and process
 them appropriately
        }
    }
}
```

Localization Framework settings

You can define translation and workflow preferences for one or more languages and for one or all the available artifacts with Localization Framework settings.

Before you define the translation and workflow preferences, ensure that you have the required artifacts to define the localization settings. For more information, see Artifact configurations.

Note:
Languages with no localization settings can't be selected for translation requests.

You can enable one or more of the following translation modes to define the translation preferences. For more information, see Translation modes.

- Activate the required translator if you want to use machine translation.

  Note:
  Machine translation in the Localization Framework application is powered by Dynamic Translation. Therefore, activating the Dynamic Translation plugin (com.glide.dynamic_translation) enables you to use the machine translation capability in Localization Framework

  .
- Configure a TMS service of your choice if you want to send the content to a translation management system for translation. For more information, see Translation Management System configurations.
- Configure an email subflow if you want to use a subflow other than the default subflow for sending the translatable content file for translation via email. For more information, see Create a custom email subflow.

Create a custom email subflow

Create an email subflow according to your requirements. You can then select this subflow to define translation preferences for a language.

Before you begin

Roles required: localization_admin and flow_designer

About this task

The Localization Framework provides a default email subflow that you can customize.

Copy the existing email subflow configured for the Localization Framework and then customize it. Publish the customized subflow so that it is available for selection when setting up the translation preference for a language. For more information, see Create a subflow.

Procedure

1. Navigate to **All** > **Process Automation** > **Flow Designer**.
2. In the Flow Designer window, click **Subflows**.
3. Click **Send LF Email for Translation** subflow.
   The Send LF Email for Translation subflow opens in the Flow Designer in read-only format. The inputs in this subflow are described in the following table:

Send LF Email for Translation input parameters

Name	Description	Type
lftask_record	Associated Localization Framework task record.	Reference.Localization Task [sn_lf_insight]

Name	Description	Type
attachment_record	Contains a file related to the LFTASK that can be sent via Email for translation.	Reference.Attachment
send_to	Contains email addresses of the recipients to send the LF task file.	String
source_language	Language from which the localization task was created.	String
target_language	Language into which the content of the localization task gets translated into.	String

4. In the subflow header, click the more actions icon and select **Copy subflow.**
5. In the **New subflow name** field, enter the name of the new subflow.
6. To limit the application scope of the subflow, in the **Application** field, select the required application, otherwise select **Global**.
7. Click **Copy**.

   The new subflow is created and opens in the subflow designer editor. The inputs from the copied subflow appear.

   Note:
   Inputs are fixed and cannot be edited.

8. In the Actions section, click **Send Email**.
   This expands the Send Email section. The fields in the Send Email section are described below:

Send Email fields

Field name	Description
Action	Name of the action. Default: **Send Email**.
Target Record	Associated Localization Framework task record
Table	Name of the table where the localization task resides.
Include Watermark	Adds a watermark in the file associated with the localization task.
To	Contains email addresses of the recipients who should be addressed directly.
CC	Contains email addresses of the recipients who should see the list of recipients directly addressed in the email.

Field name	Description
BCC	Contains email addresses of the recipients who are copied on the email but not displayed in the email.
Subject	The text about the email that its recipients can see.
Body	The text message about the localization task being sent for translation.

9. Click **Done**, when you have modified the fields as per your requirement.
10. Click **Copy Attachment**.
    The Copy Attachment action expands. The fields in the Copy Attachment section are described below:

Copy Attachment fields

Field name	Description
Action	Name of the action. Copy Attachment is displayed by default.
Source Attachment Record [Attachment]	Translatable content file of the localization task.
Target Record	Represents the target record in the email.
Table	Represents the email table.

11. Click **Save**, to save the subflow in draft status.
12. Click **Publish**.

Result
Use the published subflow to define the translation preference for a language.

Translation Management System configurations

The Localization Framework supports integration with a translation management system to localize the content.

The TMS configuration table lists all the TMS services configured in the instance.

RWS and XTM TMS services are automatically enabled in Localization Framework with the activation of the Localization Framework Installer plugin (com.glide.localization_framework.installer) and appear in the TMS configuration [sn_lf_tms_config] table.

To use the RWS TMS service for localization of the artifacts, you must configure the connection and credentials. For more information about configuring RWS TMS, see RWS Translation Management System spoke.

To use the XTM TMS service for localization of the artifacts, you must configure the connection and credentials. For more information about configuring XTM TMS, see XTM Translation Management System spoke.

You can also sign up and create an account with a TMS service of your choice. To create and configure a custom TMS, see Create a custom translation management system.

RWS Translation Management System spoke

The RWS Translation Management System spoke provides subflows and actions to localize the artifacts using the RWS translation management system service.

Activate the Localization Framework Installer plugin (com.glide.localization_framework.installer) to activate the RWS Translation Management System Spoke for Localization Framework (com.glide.localization_framework.rws_spoke).

Subscription

The RWS TMS spoke requires any ServiceNow subscription that is Standard or above and includes the Localization Framework.

Supported versions

Version v1.

Spoke Subflows

The spoke provides subflows in the Published state to integrate with the RWS TMS service. These subflows can be used as part of other subflows and flows.

Spoke subflows

Subflow	Description
Request Translations [request_translations]	Contains set of inputs and outputs to request translations of the artifacts.
Retrieve Translations [retrieve_translations]	Contains set of inputs and outputs to retrieve translations of the translated artifacts.

Spoke actions

The spoke provides actions to integrate with the RWS TMS service.

Spoke actions

Action	Description
Create Project [create_project_on_sdl]	Creates a project on RWS for translation.
Get Project Details [get_sdl_project_details]	Fetches the project's details for the translation request.
Get Translated File [get_translated_file_from_sdl]	Fetches the translated file from RWS TMS after translations.
Mark File As Complete [complete_file_on_sdl]	Marks the file as complete, after getting the translated content.
Upload File [upload_file_on_sdl_tms]	Sends a file to RWS for translations.

RWS Translation Management System account requirements

Ensure that you are registered with RWS TMS and obtain the credential information. For more information on RWS Managed Translation API, see RWS documentation.

Connection and Credential alias requirements

The IntegrationHub uses aliases to manage connection and credential information. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you do not need to update any actions that use the connection. For more information, see Connections and Credentials Information.

To use the spoke connection aliases, use an associated connection record and a credential record created with the RWS spoke installation.

This spoke uses the RWS TMS alias (sn_sdl_spoke.SDL_TMS) to authorize actions.

The following are the connection alias requirements.

- Connection type: HTTP
- Connection URL: For example, `https://languagecloud.sdl.com`. This is a sample URL. Users can have a different URL based on their geographical preference.

See Configure RWS TMS in the Localization Framework to connect to the RWS TMS services and use them to localize the artifacts.

Configure RWS TMS in the Localization Framework

Configure RWS TMS to use the RWS TMS service in the Localization Framework.

Before you begin
Roles required: connection_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   A credential record for RWS TMS is created when the RWS TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
2. Open the RWS TMS credential record.
3. In the Connections related list, open the RWS connection record.
   A connection record for RWS is created when the RWS TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.

   All the form fields are populated with the information that is fetched through the RWS spoke installation during the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
4. In the Attributes section of the HTTP/HTTPS Connection form, provide the *`project_option_id`* connection attribute, and save the form.
   RWS provides an option to create automated or reviewed projects and provides the respective option IDs during RWS account creation.
5. In the **Credential** field, click the information icon ( ) and open the record.
   A record for OAuth 2.0 credentials is created when the RWS TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
6. In the **OAuth Entity Profile** field, click the information icon ( ) and open the record.

An OAuth Entity profile is created when the RWS TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.

7. In the **OAuth provider** field, click the information icon (ⓘ) and open the record.
   An Application Registries record is created when the RWS TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
8. Provide the Client ID and Client Secret provided by RWS at the time of account creation, and save the form.
9. Open the OAuth 2.0 Credentials form again, and click the **Get OAuth Token** related link.
10. Provide the user name and password.
    RWS supports resource owner password credentials as the grant type. Provide these credentials to generate the token for the first time.

    Note:
    The user name and password do not persist in the ServiceNow instance.

11. Click **Get OAuth Token**.
    Once the OAuth token is created, the ServiceNow® platform regenerates the token until the refresh token expires.

XTM Translation Management System spoke

The XTM Translation Management System spoke provides subflows and actions to localize the artifacts using the XTM translation management system service.

Activate the Localization Framework Installer plugin (com.glide.localization_framework.installer) to activate the XTM Translation Management System Spoke for localization Framework (com.glide.localization_framework.xtm_spoke).

Subscription

The XTM TMS spoke requires any ServiceNow subscription that is Standard or above and includes the Localization Framework.

Supported versions

Version v1.

Spoke Subflows

The spoke provides subflows in the Published state to integrate with the XTM TMS service. These subflows can be used as part of other subflows and flows.

Spoke subflows

Subflow	Description
Request Translations [request_translations]	Contains set of inputs and outputs to request translations of the artifacts.
Retrieve Translations [retrieve_translations]	Contains set of inputs and outputs to retrieve translations of the translated artifacts.

Spoke actions

The spoke provides actions to integrate with the XTM TMS service.

Spoke actions

Action	Description
Create Project [create_project]	Creates a project on XTM for translation.
Get Project Details [get_xtm_tms_project_details]	Fetches the project's details for the translation request.
Get Translated File [get_translated_file_from_xtm_tms]	Fetches the translated file from XTM TMS after translations.

XTM Translation Management System account requirements

Ensure that you are registered with XTM TMS and obtain the credential information. For more information, see XTM documentation.

Connection and Credential alias requirements

The IntegrationHub uses aliases to manage connection and credential information. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you do not need to update any actions that use the connection. For more information, see Connections and Credentials Information.

To use the spoke connection aliases, use an associated connection record and a credential record created with the XTM TMS spoke installation.

This spoke uses the XTM TMS alias (sn_xtm_spoke.XTM_TMS) to authorize actions.

The following are the connection alias requirements.

- Connection type: HTTP
- Connection URL: For example, `https://languagecloud.sdl.com`. This is a sample URL. Users can have a different URL based on their geographical preference.

See Configure XTM TMS in the Localization Framework to connect to the XTM TMS services and use them to localize the artifacts.

Configure XTM TMS in the Localization Framework

Configure XTM TMS to use the XTM TMS service in the Localization Framework.

Before you begin

Roles required: connection_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   A credential record for XTM TMS is created when the XTM TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
2. Open the **XTM TMS** (ID = sn_xtm_spoke.XTM_TMS) credential record.
3. In the Connection Attributes tab of the form, provide the *`template_id`* and *`customer_id`* attributes, and save the form.
4. In the Connections tab, open the XTM connection record by selecting the information icon (ⓘ).

A connection record for XTM is created when the XTM TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation. All the form fields are populated with the information fetched through the XTM spoke installation.

> **Note:**
> In the **Connection URL** field, enter the URL of the XTM connection.

5. In the **Credential** column, select the value **XTM**to open the **OAuth 2.0 Credentials** [oauth_2_0_credentials] record.
   A record for OAuth 2.0 credentials is created when the XTM TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
6. In the **OAuth Entity Profile** field, select the information icon (ⓘ) to open the record.
   An OAuth Entity Profile is created when the XTM TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
7. In the **OAuth provider** field, select the information icon (ⓘ) to open the **Application Registries** [oauth_entity] record.
   An Application Registries record is created when the XTM TMS spoke is installed with the Localization Framework Installer plugin (com.glide.localization_framework.installer) activation.
8. Provide the Client ID and Client Secret provided by XTM at the time of account creation, and save the form.
9. Open the OAuth 2.0 Credentials form again, and select the **Get OAuth Token** related link.
10. Provide the username and password.
    XTM supports resource owner password credentials as the grant type. Provide these credentials to generate the token for the first time.

    > **Note:**
    > The username and password do not persist in the ServiceNow instance.
11. Select **Get OAuth Token**.
    Once the OAuth token is created, the ServiceNow® platform regenerates the token until the refresh token expires.

Integrate with a translation management system

Integrate with any translation management system of your choice to localize the artifacts requested for translations.

Before you begin

Role required: localization_admin

Procedure

1. Create subflows to request translations and to retrieve the translated content using Integration Hub subflows, actions, and connections and credentials.

   a. Ensure that the error messages and codes from the TMS provider are mapped correctly with the error messages and codes supported the ServiceNow platform.

      Standard error codes and messages

Error code	Error message
40510	Error while requesting translations.

Error code	Error message
40520	Error while retrieving translations.
40500	Authentication failed.

b. Ensure that the input and output fields are in the following format for the subflows:

Request translation subflow

Field Name	Description	Type
Inputs		
lftask_number	Associated localization task number.	String
xliff_1_2_content	Content of an artifact which must be translated in XLIFF 1.2 format. For example, `'csv', 'xliff 1.2'`	String
source_language	Language code of the source text. For example, `'en'`	String
target_language	Language code of the target text. For example, `'fr', 'it'`	String
Outputs		
project_info	The ID for the project created for an artifact configuration.	Object
status	Status of the request. For example, Success or Failure.	Choice
error_message	The reason for the error that populates when an error occurs.	String
error_code	The error code that populates when an error occurs.	String

Retrieve translation subflow

Field Name	Description	Type
Inputs		
lftask_number	Associated localization task number.	String
project_info	The ID for the project created for an artifact configuration.	Object
Outputs		
translated_xliff_1_2	Translated XLIFF file received from the TMS.	String
status	Status of the request. For example, Success or Failure.	Choice
error_message	The reason for the error that populates when an error occurs.	String
error_code	The error code that populates when an error occurs.	String

2. Create and configure a custom translation management system and activate it.

   See Create a custom translation management system.

   Note:
   Localization Framework depends on Integration Hub to build integrations with third-party systems.

Create a custom translation management system

Configure Localization Framework with a third-party translation management system of your choice to localize an artifact.

Before you begin

Roles required: localization_admin and flow_designer

Ensure that the Localization Framework Installer plugin (com.glide.localization_framework.installer) is activated.

Procedure

1. Navigate to **All** > **Localization Framework** > **TMS Configuration**.

   Note:
   By default, RWS, and XTM translation management system configurations are available.

2. Click **New**.
3. In the TMS Configuration form, fill in the fields.

TMS Configuration form fields

Field	Description
Name	Identifier of the TMS configuration.
Request translation subflow	IntegrationHub subflow that interacts with the third-party translation management system to request translations.
Retrieve translation subflow	IntegrationHub subflow that interacts with the third-party translation management system to retrieve translated content.
Application	The scope of the application. The default value is Global.
Version	Version of the contract used in subflows. The default value is **v1.**
Enable auto retrieval	Option to retrieve the translated content automatically. When selected, the translations are retrieved as per the scheduled jobs which run daily.

4. Click **Submit**.

Configure the Localization Framework preferences

Configure the preferences to request the translation in the activated languages.

Before you begin

Role required: localization_admin or localization_manager

Procedure

1. Navigate to **All** > **Localization Framework** > **Settings**.
2. Select **New**.
3. On the form, fill in the fields.

Settings form

Field	Description
Name	Identifier of the translation setting.
Artifact	Artifact for the settings being defined. You can select either **All Artifacts** or an artifact of your choice, such as **catalog_item** or **nlu_model**. **Note:** To select Virtual Agent topics, Document Templates, Document Template Block Content, HTML Document Template, Knowledge Base or Email notification artifacts, you must activate specific plugins. For more information, see Supported artifacts in Localization Framework.
Languages	Target records of the languages available in the instance. All the available languages are populated based on the artifact selected.

4. Select the **Translation Preferences** tab.
5. On the form, fill in the fields.

Translation Preferences form

Field	Description
Machine Translation	
Enable machine translate	Option to enable the machine translation service for selected languages. **Note:** Dynamic Translation must be installed to use machine translation. For more information, see Activate Dynamic Translation.
Translator	List of the available translator services for machine translation. **Note:** Ensure that at least one third-party translator configuration record is integrated and active. For more information, see Integration with other translation services.
Translation Management System	
Enable TMS	Option to enable Translation Management System (TMS) for selected languages.
TMS	List of all the TMS services available for machine translation. By default, the **RWS** and the **XTM** TMS are available. You can also choose a configured custom TMS. For more information, see Translation Management System configurations. If you are configuring a spoke instance per the Localization Framework Hub and Spoke architecture, select `Localization Framework Hub` in this field. For more information, see Localization Framework Hub and Spoke architecture .
Send via Email	
Enable send via email	Option to enable email service for content localization for selected languages.
Email subflow	Available subflows for your email. The default subflow for Localization Framework is `Send LF Email For Translation`. To choose another subflow, create a subflow before setting up the translation preference.
Default file format	Option to select a file format to send and receive the content. By default, XLIFF 1.2 and CSV file formats are supported.
Default recipient(s) list	List of recipient email addresses through which the content is sent to the translator for translation.
Export/Import	

Field	Description
Enable export/ import	Option to enable the export and import of content for the translation of selected languages.
Default file format	Available options of file formats to send and receive the content. By default, XLIFF 1.2 and CSV file formats are supported.

6. Select the **Workflow Preferences** tab.
7. On the form, fill in the fields.

Workflow Preferences form

Field	Description
Workflow	List of all the workflows available for the translation process. For more information on the available workflows, see Workflows in the Localization Framework.
Business approver group	List of all the available business approver groups. This field appears only when a workflow that requires business approvals is selected from the **Workflow** field. **Note:** Ensure that the business approver group contains the localization_fulfiller role.
Localization fulfiller group	List of all the available localization fulfiller groups. **Note:** Ensure that the localization fulfiller group contains the localization_fulfiller role.
Auto translation mode	Option to select the auto-translation mode. The available options are as follows: ◦ **Machine Translate** ◦ **Send to TMS**. If the instance has been set up as a spoke in the Localization Framework Hub and Spoke architecture, this setting submits the translation request to the hub. ◦ **Send via Email** For more information, see Translation Modes. **Note:** Ensure that you have configured the Translation Preferences for the selected auto translation mode.
Translation approver group	List of all the available translation approver groups. This field appears only when a workflow that requires translation approvals is selected from the **Workflow** field. **Note:** Ensure that the translation approver group contains the localization_fulfiller role.

8. Select the **Project Preferences** tab.
9. On the form, fill in the fields.

Project Preferences form

Field	Description
Enable projects	Option to enable the creation of the localization projects. If enabled, localization requested items (LRITMs) are created when translation of the artifact is requested. You can then create a project and add the corresponding requested items to start the translation. You can use this option for bulk translation, so multiple artifacts are translated at the same time. The progress of translation of the LRITMs is monitored at the project level. **Note:** You need the localization_manager role to create a project. For more information, see Create translation projects.
Enable auto creation	Option to enable the auto creation of the localization projects. If enabled, the following occurs when translation is requested: ◦ A localization project is auto created. The corresponding localization requested items (LRITMs) are created, bundled, and added to the localization project. ◦ The project starts automatically based on the localization framework settings.
Localization manager group	List of all the available user groups. **Note:** Ensure that the group you designate contains the localization_manager role.

10. Select **Submit**.

Localization Framework Properties: Update Set Strategies

Use update sets to migrate your translations to another instance. Configure properties for update sets according to your business requirements.

Update sets enable you to transfer your artifact translations to other instances. For background information on update sets, see System update sets.

For localization tasks, the default has been to create a dedicated update set per task in the scope of the artifact. In the Washington DC release, you can specify your update set strategy to control whether translations are bundled into one update set, or distributed into granular update sets.

The settings for update set strategies vary according to the type of translation, defined as follows:

localization tasks

Part of workflows that often include approval processes. These are based on tasks.

adhoc

Edited and published directly by the localization editor, from the artifact record (such as a catalog item). The **Edit translation** button is visible to the editor, and it opens a UI page for translation. Adhoc translations are processed in the user context.

Update set strategy to use when translations are published from localization task

Save translation in update set created for each localization task	The system will create a dedicated update set for each task, using the task name. This strategy is most appropriate when there are a limited number of tasks.
Save translations in 'LF: Translations' update set	Granular update sets are not created. Instead, the system creates or reuses one update set named 'LF: Translations', and saves all tasks to it. This strategy can be more convenient when there are many localization tasks.

Update set strategy to use when translations are published in adhoc way

Save translations in user preferred update set	The translations are written to the update set that the user has specified for a scope, because the processing happens in the user's context. When artifacts are in an application scope, their translations are saved to the user's preferred update set for that scope.
Save translations in 'LF: Translations' update set	The system creates or reuses one update set named 'LF: Translations', and all adhoc translations are written to it.

To set these properties, navigate to **Localization Framework** > **Properties**. Choose the appropriate strategy and select **Save**.

Localization Framework Hub and Spoke architecture

Enable one central instance to fulfill translation requests submitted from other instances using the Localization Framework Hub and Spoke architecture. Install optional Localization Framework plugins to configure one of your instances as a hub and other instances as spokes.

By default, translations are requested and fulfilled on the same instance. However, in the Washington DC release, there is an available hub-spoke model for translation requests and fulfillment. With this architecture, translation requests are a two-step process:

1. The spoke instance submits a translation request to the hub instance.
2. The hub instance manages the fulfillment of the request, such as by sending the content to a translation management system (TMS). The hub returns the translated contents to the spoke instance.

Consider the Localization Framework Hub and Spoke architecture when you want to limit fulfillers' access to one instance only.

To set up this architecture, use the admin role to install the following plugins:

- Localization Framework Hub [com.sn.localization_framework.hub]. Install this plugin on the instance you designate as a hub.
- Localization Framework Spoke [com.sn.localization_framework.spoke]. Install this plugin on the instances you designate as spokes.

For more information on installation, see Activate Localization Framework.

When both plugins are installed, use the admin or localization_admin role to configure the connection between hub and spoke instances. On the hub instance:

1. Navigate to **Localization Framework** > **Hub** > **Spoke Settings.**
2. In the **Name** field, enter a unique name for each entry.
3. In the tab for **Translation Preferences**, find the **Translation Management System** (TMS) section. Enter the TMS service your organization uses.
4. In the tab **Workflow Preferences**, configure according to your requirements. Depending on the workflow you choose, different subfields are populated.
5. The tab **Project Preferences** is optional. If no project is entered, any request for translations into more than one target language is broken into separate tasks, one for each language.

Note:
If Dynamic Translation is not installed, the option for machine translation is not available.

On a spoke instance:

1. Navigate to the **Settings All Artifacts** table [sn_lf_setting].
2. In the tab **Translation Preferences**, find the **Translation Management System** (TMS) section.
    - **Enable TMS**: select.
    - **TMS**: Set to `Localization Framework Hub`.
    - **TMS Additional Info**: for **spokeSettingsName**, enter the name of the connection setting that has been configured on the hub.
3. Navigate to the **Connection & Credential Aliases** table. Find and open the record named *`Localization Framework`* **Hub TMS**.
4. Add the connection to the hub in the **Connections** tab. You can enter the URL to your hub in the **Connection URL** field. Set authentication and other attributes appropriate to your environment.

From the perspective of the spoke instance, the hub acts as a TMS.

Request translations

The Localization Framework provides the means to select one or more activated languages while requesting translations for artifacts from English into different languages.

Request translations in the Localization Framework

- Define localization preferences for the languages available in the instance to request translations.
- Lists only the activated languages in the instance to choose for localization of the content.

Note:
Before you request for translations, ensure to activate all the required languages in the instance. For more information, see Activate a language.

- Displays the translation status for each activated language only if translations have been requested for insights in those languages. The translation status is updated when the insights dashboard is generated.
- Enables you to request translations from English into the selected language from the localization dashboard for the selected insights records.
- Enables you to request translations for insights from English into one or more activated languages from their respective tables.

Request translations from an artifact table

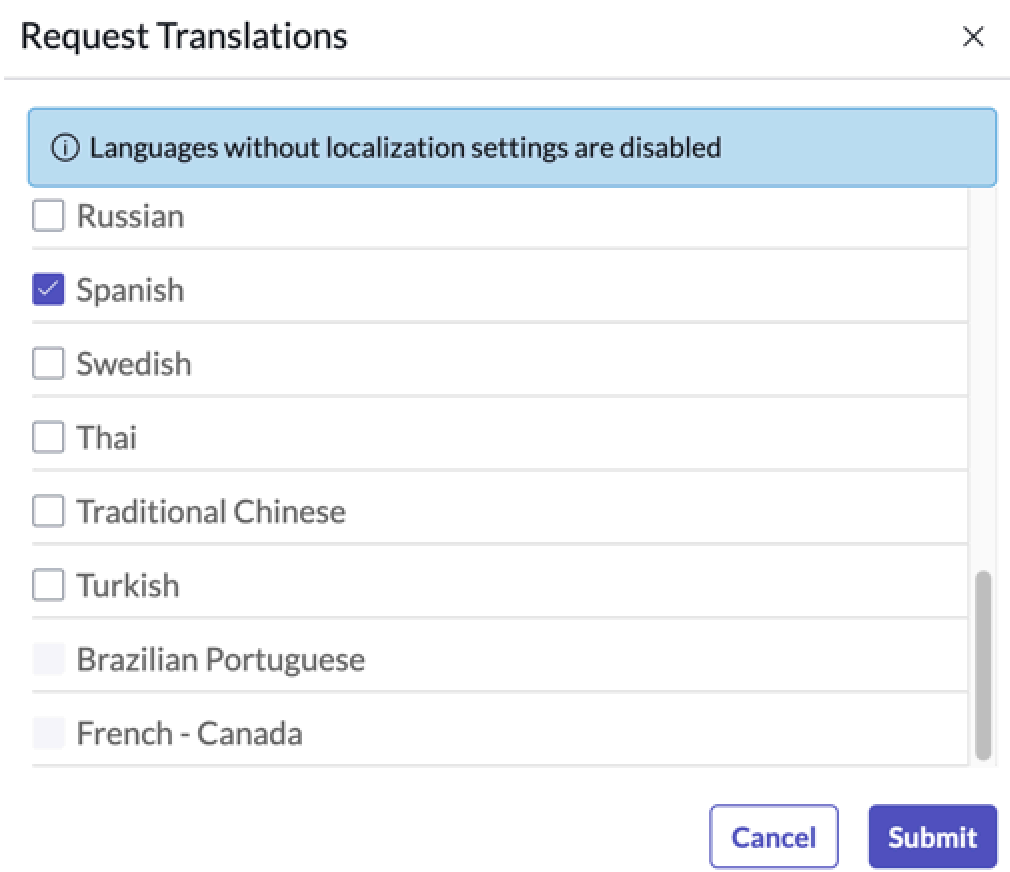

Note:
Languages without localization settings are disabled for selection while requesting translations.

Request translations

The Localization Framework enables you to select one or more activated languages while requesting translations for artifacts.

Request translations in the Localization Framework

- Define localization preferences for the languages available in the instance to request translations.
- Lists only the activated languages in the instance to choose for localization of the content.

Note:
Before you request for translations, ensure to activate all the required languages in the instance. For more information, see Activate a language.

- Displays the translation status for each activated language only if translations have been requested for insights in those languages. The translation status is updated when the insights dashboard is generated.
- Enables you to request translations in the selected language from the localization dashboard for the selected insights records.
- Enables you to request translations for insights into one or more activated languages from their respective tables.

Request translations from an artifact table

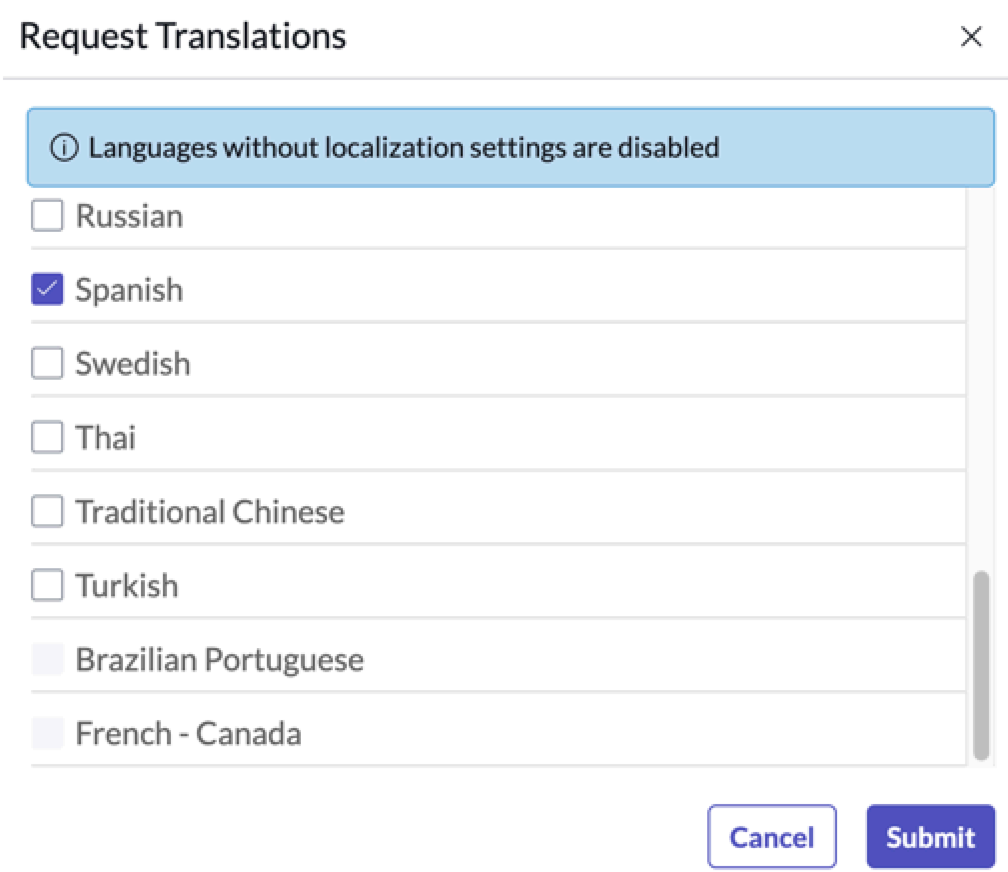

Note:
Languages without localization settings are disabled for selection while requesting translations.

Related topics

Fulfill a localization task

Approve a localization task

Request translations for Service Catalog items

Request translations for multiple Service Catalog items to localize them from English into one or more languages.

Before you begin

Roles required:

- localization_requestor: To request translations
- catalog_admin: To access the Service Catalog items [sc_cat_item] table.
- catalog_editor (or catalog_manager) if you are the editor or manager of the item that is being translated: To access the Service Catalog items [sc_cat_item] table.

Procedure

1. Navigate to **All** > **Service Catalog** > **Catalog Definitions** > **Maintain Items**.
2. Select single or multiple items from the list or select the **Select All** check box.
3. From the **Actions on selected rows** drop-down list, select **Request Translations**.
4. In the **Request Translations** dialog box, select the languages.
5. Select **Submit**.
   Localization requested items (LRITMs) are created for the selected items in all the selected languages. A success message with a link to the list of requested items is displayed. Click the link to see the list of requested items.

Request adhoc translation for Service Catalog items

Request adhoc translations for Service Catalog items to localize them from English into one or more languages using the edit translations option. After you make the changes, you can publish the translations.

Before you begin

The Localization Framework for Service Catalog plugin [com.glide.localization_framework.service_catalog] must be installed for the Edit Translations button to appear on the service catalog form.

Roles required:

- localization_editor: To request translations
- catalog_admin: To access the Service Catalog items [sc_cat_item] table.
- catalog_editor (or catalog_manager) if you are the editor or manager of the item that is being translated: To access the Service Catalog items [sc_cat_item] table.

Procedure

1. Navigate to **All** > **Service Catalog** > **Catalog Definitions** > **Maintain Items**.
2. Select the item for which you want to request adhoc translations.
3. Click **Edit Translations**.
   The adhoc comparison UI of Localization Framework is displayed.
4. On the adhoc comparison UI of Localization Framework, you can use the followings options.

   Note:

   The options displayed on the adhoc comparison UI of Localization Framework are based on the Localization Framework **Settings**.
5. After your changes are completed, click **Publish Translations**.

Create translation projects

The Localization Framework provides the means to assemble, manage, and complete localization and translation tasks, from single item assignments to large batch projects.

The **Localization Projects** feature enable users to bundle multiple localization request items (LRITMs) under one project. This allows bulk translation of requested items, thereby increasing the efficiency of the process. To work with Localization Projects, the feature needs to be enabled in Localization settings. A user with localization_manager role can enable the feature in Localization settings.

Localization tasks are not created automatically. The localization manager adds the requested items to a project and starts the project. All the requested items are then bundled under the localization tasks.

You can also automatically create a localization projects by enabling the **Enable auto creation** option in the Localization Framework settings. A localization project is auto created and the corresponding LRITMs are created, bundled, and added to the localization project. For more information, see Configure the Localization Framework preferences.

Each localization task has one or more requested items and is grouped in the following sequence:

1. Localization Framework **settings**: All the requested items are first grouped based on LF settings.
2. Source language: After the requested items are grouped according to the settings, then same source languages are bundled under one localization task.
3. Group of 150 requested items: If the number of requested items for a source language is more than 150, then for every group of 150 requested items, one localization task is created.

Create a localization project

Create localization projects to bundle multiple requested items under one localization task.

Before you begin

Enable Localization Projects in Localization Framework settings. For more information, see Configure the Localization Framework preferences.

Role required: localization_manager

Procedure

1. Navigate to **All** > **Localization Framework** > **My Projects**.
2. Select **New**.
3. On the form, fill in the fields.

Localization Project

Field	Description
Number	Unique project number.
State	Current state of the project. Initially this is set to **Draft**.
Assigned to	Person to whom the project is assigned.
Assignment group	Group of the assigned person.
Short Description	Brief description of the project.

Field	Description
Work Notes	Notes regarding the project.
Additional comments (Customer Visible)	Notes that are visible to the customer.

4. Select **Submit**.

Add a localization request item to a project

Add multiple localization request items (LRITMs) to a project for the fulfillers to bulk translate the items.

Before you begin

Role required: localization_manager

Procedure

1. Navigate to **All** > **Localization Framework** > **My Projects**.
2. Select the project from the list to add the requested items.
3. On the Localization Project record, from the related list, select the Localization Requested Items tab.
4. Select **Edit**.

   Note:
   The Edit button is visible only when a project is in **Draft** state.

5. Select one or more items for translations from the **Collection** window.
6. Add the selected items to **Localization Requested Items List** window using the right arrow icon.
7. Save the changes.
8. On the form header, select **Start Project**.
   The state of the project is changed to **In progress** and one or more Localization tasks are created based on settings and requested items.

   Note:
   You cannot add items to the project once you start the project.

Fulfill a localization task

Fulfill a localization task created for each selected language in the translation request.

Before you begin

Role required: localization_fulfiller

Procedure

1. Navigate to **All** > **Localization Framework** > **My Tasks**.
2. Select a localization task in the Open state.
3. On the Localization Task page, select **Translate** in the form header.

Note:
The **Translate** UI action is visible in all states except Awaiting Business Approval, Close Complete, and Incomplete.

A two-column comparison UI page opens with the Source Language column on the left and the Target Language column on the right.

Each Localization requested item has the following options:

- **Expand all**: Expand all the groups on the comparison UI.
- **Collapse all**: Collapse all the groups on the comparison UI.
- **Show groups with missing translations**: See only the groups that have missing translations when the content is translated.

4. Translate the content manually or select one of the following translation modes from the task header.

<table>
<tr><th>Translation mode</th><th>Description</th></tr>
<tr><td>Machine Translate</td><td>Translates content using the translation service provider configured in the translation preferences. If there is more than one LRITM, then the operation is asynchronous.</td></tr>
<tr><td>Send to TMS</td><td>Sends the artifact to the translation management system configured in translation preferences for translation. If there are more than is LRITM, then the operation is asynchronous.</td></tr>
<tr><td>Export Source File</td><td>Exports the source file for translation in the file format configured in the translation preferences or as per the selected file format in the comparison UI.</td></tr>
<tr><td>Import Translated File</td><td>Imports the translated file in the file format configured in the translation preferences from the third-party translator.

A merge conflict occurs when you have certain fields pre-populated with translations, but those fields get updated when you import the translated file. These conflicts are shown right below the fields updated with the latest translations.

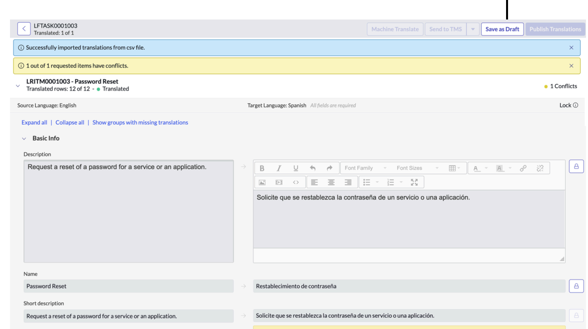
</td></tr>
</table>

Translation mode	Description
	To resolve the merge conflict, do the following: **a.** Click **See versions**. **b.** Choose a version that you want to retain. **c.** Click **Save**.
Send Source File via Email	Sends the source file to the email addresses configured in translation preferences for translation. If there is more than one LRITM, then the operation is asynchronous.

Note:
When any translation mode is selected to translate the content, the content is saved as a draft to retain changes done by the user and then the translation process is initiated.

5. **Optional:** If you want to make any changes to the verified translations, click the lock icon ( ) beside the field with verified translations.
   Verified translations are pre-existing translations available in the database for the given text. The framework automatically fetches the existing translations and locks the fields.
6. **Optional:** Save the translated content by clicking **Save as Draft**.
   You can later revisit the content or send the content for translation approval.

   (Optional) This option moves the task state to Translation in Progress.
7. Click **Submit for Approval**.

   Note:
   This option is available only for tasks where the workflow configured for the target language needs translation approval.

   The task is sent for translation approval.
8. Click **Publish Translations**.

   Note:
   This option is available only for tasks where the workflow configured for the target language has the **Publish** option.

Result
The translations of the localization task are published and the task is updated with the Closed Complete state.

You can move translation records across instances by using update sets. The translation-related XML files are saved to the update set specified in the settings. For more information, see Localization Framework Properties: Update Set Strategies.

Approve a localization task

Proofread the translated content and approve a localization task.

Before you begin
Role required: localization_fulfiller

Procedure

1. Navigate to **All** > **Localization Framework** > **My Tasks**.
2. Select a localization task in the **Open**state.
3. Select **Translate**.
   A list of all the requested items with the translation status is displayed. You can expand each requested item to see the details.
4. Select **Submit For Approval**.
   The state of the localization task is changed to **Under Review**. An approval localization task is created and assigned to an approver according to the approver group assigned in localization settings.
5. Select to open the approval localization task record.
   The state of the approval localization task is set to **Requested**.
6. From the form header, select **Verify Translations**.
   A list of all the requested items is displayed. Each requested expands to display the comparison UI.
7. On the form header, select **Approve**.
   All the translations are published and the states of the requested items, localization task, and the localization project is changed to **Closed Complete**.
8. On the form header, select **Request Changes** if you want to reject the translations.
   The approval localization task state is changed to **Rejected**. The requester needs to create another approval task and assign to the approval group.

Monitor the Localization Framework

Track localization performance and request translations with Localization Framework Dashboards.

Localization Insights dashboard

The Localization Insights dashboard provides the localization health of the system with reports that are generated from the localization insights records as per the scheduled jobs. Localization insights contain the list of artifacts with their translation statuses which, when generated, display a dashboard with information.

Localization Insights dashboard

- The localization Insights dashboard shows the last updated time of the localization insights using which the reports were generated.
- You can opt in to or opt out of an artifact from the insights scheduled job.
- Based on the artifacts, this dashboard contains these following tabs:
  - General: Displays general translation information.
  - Catalog Items: Displays translation information related to Service Catalog items.
  - VA Topics: Displays translation information related to Virtual Agent (VA) topics.
  - Surveys: Displays translation information related to Surveys.

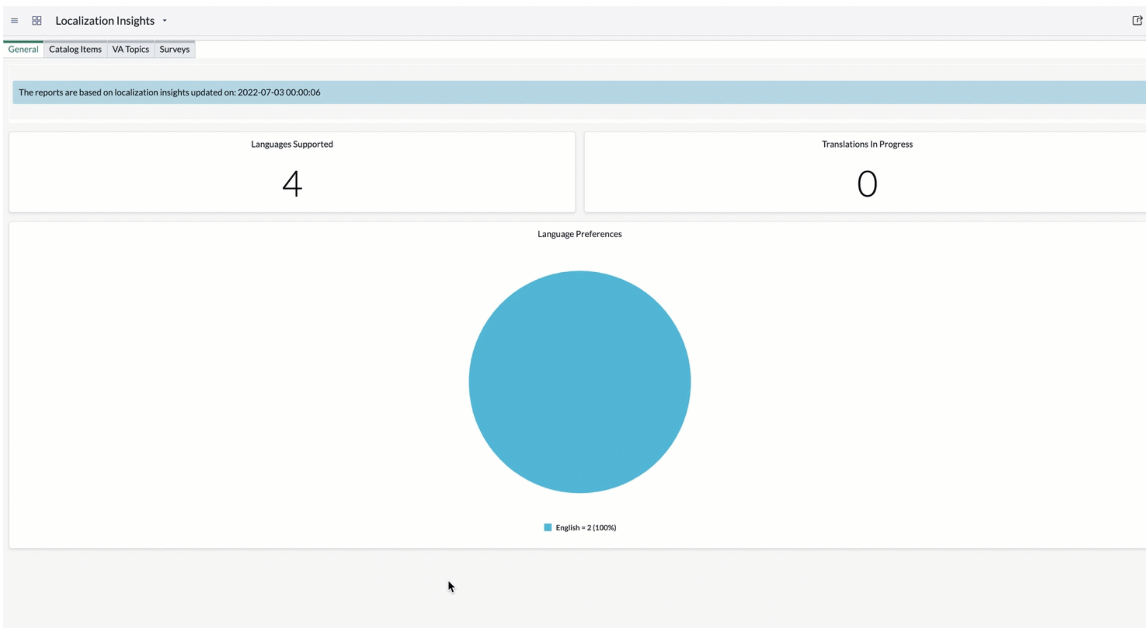

Required ServiceNow AI Platform role

You need the localization_admin or localization_requestor role to access the localization insights records for executing the scheduled jobs to generate the Localization Insights dashboard.

Use case

User	Dashboard use
Users with the localization_requester or localization_admin roles	View the localization health of the system. Request translations for Service Catalog items, Surveys and Virtual Agent topics for a particular language from the Localization Insights dashboard.

Data visualizations

Title	Type	Source table	Description
Languages Supported	Single Score	sys_language	The total number of activated languages in the instance.
Translations In Progress	Single Score	sn_lf_insight	The total number of translations in progress in the instance.
Language Preferences	Horizontal Bar	localization_insights_view	Provides information about the number of users logged in to the instance along with their preferred languages.
Catalog Items	Single Score	sn_lf_insight	The total number of Service Catalog items in the instance.
Catalog Items Translation In Progress	Single Score	sn_lf_insight	The total number of Service Catalog items translations in progress in the instance.
Catalog Item Translations	Vertical Bar	sn_lf_insight	The total number of Service Catalog items requested for translations per language. It also indicates the following translation statuses with defined color legends: • **Completely Translated** • **Translation in Progress** • **Partially Translated** • **Not Translated**
Virtual Agent Topics	Single Score	sn_lf_insight	The total number of published Virtual Agent topics in the instance.
Virtual Agent Topic Translation In Progress	Single Score	sn_lf_insight	The total number of Virtual Agent topic translations in progress in the instance.

Title	Type	Source table	Description
Virtual Agent Topic Translations	Vertical Bar	sn_lf_insight	The total number of published Virtual Agent topics requested for translations per language. It also indicates the following translation statuses with defined color legends: • **Completely Translated** • **Translation in Progress** • **Partially Translated** • **Not Translated**
Surveys	Single Score		The total number of surveys in the instance.
Surveys Translation In Progress	Single Score		The total number of survey translations in progress in the instance.
Survey Translations	Vertical Bar		The total number of surveys requested for translations per language. It also indicates the following translation statuses with defined color legends: • **Completely Translated** • **Translation in Progress** • **Partially Translated** • **Not Translated**

Request translations from Insights Dashboard

Request translations for Service Catalog items or Virtual Agent topics from the insights dashboard to localize them into the selected language based on their translation status.

Before you begin

Role required: localization_admin or localization_requestor

Procedure

1. Navigate to **All** > **Localization Framework** > **Insights Dashboard**.

   The dashboard provides **Catalog Items Translations** report and **Virtual Agent Topic Translations** report by default using which you can request translations.

   If you have created a report for your custom artifact to display in the dashboard, you can also request translations for the custom artifacts. For more information, see Create a report for custom artifact.

2. In the **Catalog Items Translations** report or **Virtual Agent Topic Translations** report, click the vertical bar of the language into which you want to request translations.
   For example:
   - Select a language from the **Catalog Items Translations** report to request translations for the Service Catalog items in that language.
   - Select a language from the **Virtual Agent Topic Translations** report to request translations for the Virtual Agent topics in that language.
   - If you have a custom report created for an artifact, select a language from that report to request translations into the selected language.
3. Select one or more documents from the localization insights list.
   - You can click the **Select All** check box to request translations for all the items in the list.
4. From the **Actions on selected rows** list, select **Request Translations**.
5. Click **OK** in the Confirmation dialog box to submit the request.

Localization Framework reference

Reference topics provide information on components and technical aspects of the Localization Framework.

Translation modes

The Localization Framework offers various translation modes for content translation. Localization fulfillers can use one of the translation modes to translate content and fulfill localization requests.

Machine Translate

Machine translation in the Localization Framework is powered by Dynamic Translation.

Activate the required translation service provider in the instance to machine translate the content. For information, see Create a custom translator configuration.

Send to Translation Management Systems

The Localization Framework integrates with third-party Translation Management Systems (TMS) to localize the content.

RWS and XTM Translation Management Systems are supported by default and requires configuration. For more information see, Configure RWS TMS in the Localization Framework and Configure XTM TMS in the Localization Framework.

You can configure a custom Translation Management System of your choice. For more information, see Create a custom translation management system.

When the Localization Framework Hub and Spoke architecture is installed, the TMS must be configured accordingly. For more information, see Localization Framework Hub and Spoke architecture.

Zip attachment should be supported on the instance. For more details on attachment configuration, see Configure system attachment properties

Send via Email

The Localization Framework provides **Send via Email** as a translation mode that can be used to send the translatable content for translation to a TMS service via email.

If you have more than one requested item in a task, then the requested CSV or XLIFF items are compressed into one ZIP file and attached to the email. Either CSV or XLIFF attachment should be supported on the instance. For more details on attachment configuration, see Configure system attachment properties

You can configure the **Send via Email** mode as a translation preference for one or more activated languages. For more information see, Configure the Localization Framework preferences.

Export/Import

Localization Framework supports the following translation modes:

- **Export**: Export the translatable content for translation to a TMS service. If you have more than one requested item in a task, the individual CSV or XLIFF files are compressed and downloaded as one zip file.
- **Import**: Import the translated content from a TMS service. If you have more than one requested item in the task, you can import individual CSV or XLIFF files or in a ZIP format. Any files unrelated to the task are skipped.

Either CSV or XLIFF attachment should be supported on the instance. For more details on attachment configuration, see Configure system attachment properties

States of localization projects and tasks

The state of a localization project or of a task defines the status of the localization request. The states of the localization tasks change based on the workflow that is defined for translations.

Project States

State	Definition
Draft	Initial state when the localization project is created.
In progress	State when the localization project starts.
Closed Complete	State when all the tasks are closed complete.
Closed Incomplete	State when at least one of the tasks is closed incomplete.

Task States

State	Definition
Awaiting Business Approval	Localization task is awaiting business approval.

Task States (continued)

State	Definition
Error in Translation	Localization task has encountered an error while translating the content.
Closed Complete	Localization task has been fulfilled and the translations are published.
Closed Incomplete	Localization task has been closed without fulfilling the request.
Open	Localization tasks are not sent for translation.
In Progress	Translation of localization tasks is in progress.
Under Review	Approver is reviewing the translations.

Error messages in the Localization Framework

Error messages are displayed when the appropriate settings are not configured for the translation and workflow preferences.

Error messages in Artifact Configurations

Error message	Scenario
Internal name cannot be modified for the existing record.	When someone tries to modify the internal name of the existing artifact record.
Either a Processor script or a Read script is required for an active artifact configuration	When there is no processor script or read script available for the artifact configuration record.
A configuration record exists for the table: {table_name}	When you provide a table name for an artifact configuration record that exists for another artifact.
A configuration record exists with internal name: {internal_name}	When you provide an internal name for an artifact configuration record that exists for another artifact.

Error messages in Settings

Error message	Scenario
Contains invalid Email Address(es): {comma-separated email addresses}.	When some of the email addresses fail the validation.
Selected Localization Fulfiller group does not have localization_fulfiller role.	When the selected user group does not have the localization_fulfiller role.
Selected Business Approver group does not have localization_fulfiller role.	When the selected user group doesn't have the localization_fulfiller role.
Selected Translation Approver group does not have localization_fulfiller role.	When the selected user group does not have the localization_fulfiller role.
Ensure that the corresponding translation preference for your auto translation option is enabled.	When auto translate is selected in the workflow preference.

Error messages in Settings (continued)

Error message	Scenario
{Translator_name} translator is not of version v3.	When the corresponding translator version is not v3.
{Translator_name} translator has no translate subflow configured.	When the corresponding translator's translate subflow is empty.
Selected translator is not active.	When the corresponding translator configuration is inactive.
There are no languages available for the selected artifact. All the possible languages were configured in other localization setting records.	A warning message when there are no languages available for the selected artifact.
A setting record exists for the artifact - language pair(s): {pair of artifact and language}. Resolve the conflict(s) to proceed.	When a user tries to add a language or artifact in settings and that record exists in another record.
The Default recipient(s) list contains invalid email address(es): {comma_separated_email_adddresses}.	When the provided email addresses of the recipients are invalid.
{0} translator is not active.	When the translator is not active.

Error messages in Request Translations

Error message	Scenario
Your request for translations failed. Please reinitiate the request.	When there is an error in submitting the translation request.
Error in requesting translations for the following artifact item and language pair(s): {pair of artifact and language}. Reinitiate the request for the listed pairs.	Error is seen only on the list view when there is an error requesting translations for certain language pair(s): {pair of artifact and language}.

Error messages in the Comparison UI

Error message	Scenario
Machine Translate	
Failed to machine translate using {translator name}. {Reason for failure}. **Note:** The reason is fetched from Dynamic Translation.	This message appears in two scenarios: • Error message shown in the comparison UI for machine translation. • Message logged in the activity stream when failed to machine translate the content.
Failed to machine translate using {translator name}. Translations for one or more fields have failed.	This message appears in two scenarios: • Error message shown in the comparison UI when translations for one or more fields fail. • Message logged in the activity stream when translations for one or more fields fail.

Error messages in the Comparison UI (continued)

Error message	Scenario
Machine Translate	
Failed to machine translate.	Error message shown in the comparison UI in case the call for machine translate fails.
Error while saving the draft. Contact your system administrator.	Error message shown in the comparison UI when content that is requested for translation, is not saved as a draft message successfully. For any translation mode selected to translate the content, the content is saved as a draft to retain changes done by user and then the translation process is initiated.
Send via Email	
Failed to send source file via email to: {email addresses}.	Message logged in the activity stream when sending source file via email fails.
Failed to send email.	This error message appears in two scenarios: • Error message persisted and shown in the comparison UI when sending the source file via email fails. • Error message shown in the comparison UI in case the call for sending email fails.
Following are the invalid email addresses: {comma separated invalid email addresses}, please try again.	When one of the email addresses provided are incorrect while sending via email.
Send to TMS	
Error while sending the task to {tms name} TMS for translation. {Reason for failure}.	This error message appears in two scenarios: • Message logged in the activity stream when sending translatable content to TMS fails. • Error message persisted and shown in the comparison UI when sending translatable content to TMS fails.
Failed to send to TMS.	Error message shown in the comparison UI in case the call for sending translatable content to TMS fails.
A record related to {0} exists.	when you try to create another TMS with the same name.
Retrieval from TMS	
Failed to retrieve translations from {tms name} TMS. {Reason for failure}.	This error appears in two scenarios: • Message logged in the activity stream when retrieving translations from XTM TMS fails. • Error message persisted and shown in the comparison UI when retrieving translations from XTM TMS fails.
Failed to retrieve translations from {tms name} TMS.	Error message shown in the comparison UI in case the call for retrieving translations from XTM TMS fails.

Error messages in the Comparison UI (continued)

Error message	Scenario
Machine Translate	
Export Source File	
Failed to export the {CSV/XLIFF} file. {Reason for failure}.	This error message appears in two scenarios: • Error message persisted and shown in the comparison UI when exporting the source file for translations fails. • Message logged in the activity stream when exporting the source file for translations fails.
Failed to export file.	Error message shown in the comparison UI in case the call for exporting the source file for translation fails.
Import Translated File	
Failed to import translations from {CSV/XLIFF} file. {Reason for failure}.	This message appears in two scenarios: • Message logged in the activity stream when importing the translated file fails. • Error message persisted and shown in the comparison UI when importing the translated file fails.
Import failed.	Error message shown in the comparison UI in case the call for importing the translated file fails.
Failed to publish translations.	When failed to publish the translations.
Save as Draft	
Error while creating the draft.	When there is an error after clicking Save as draft.
Approve Translations	
Error while requesting for changes.	When an error occurs while an approver requests changes for the translations.
Publish Translations	
Error in publishing the translations.	When attempting to translate content that has already been translated.
Error while submitting the translations.	When publishing translations or when sending translations for approval.

Error messages in Insights

Error message	Scenario
Insights generation job is already running, please try again later	When you try to generate the insights record multiple times.
Your request for translations failed for the selected insights. Click here to see the list of insights for failed requests and reinitiate the request.	Lists the insights records for which it failed to request translations.

Error messages in Insights (continued)

Error message	Scenario
Localization setting does not exist for the following table and language pair(s): {pair of table name and language}. Uncheck the record(s) of the mentioned pair(s) and resubmit.	Lists the language pair(s): {pair of table name and language} for which a settings record does not exist.

System Localization

Localization allows administrators to accommodate users from a variety of different countries, using different languages and currencies, within the same instance.

Get started

Explore

Learn about localizing an instance for different users.

Configure

Configure languages, locales, regions, and locations for an instance.

Translate custom content

Translate content created or modified on an instance.

Translate to a new language

Translate an instance to a language without a language plugin.

Reference

Get details about properties, tables, and more.

Troubleshoot and get help

- Localization in the ServiceNow Customer Success Center
- Localization Expert Series in ServiceNow University

- Digital guidebook for localization on the ServiceNow® Developer Site: ServiceNow Application Localization
- Internationalization & Localization in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring System Localization

Learn about support for localizing ServiceNow instances as part of your globalization strategy.

System Localization overview

Localization is the process of making an instance specific to a given region or locale without impacting functionality.

System localization includes the language and locale settings for an instance and is a ServiceNow AI Platform feature that is active by default (com.glide.i18n.core). Translations for several languages are available through additional internationalization (I18N) language plugins (for example, com.snc.i18n.spanish). These language plugins install the I18N: Internationalization plugin (com.glide.i18n), which provides the elements necessary for translating an instance.

Languages and translation

The I18N language plugins contain translations of the UI string content that ServiceNow provides by default in your instance. Whenever you have customized UI strings that you want to translate, you have to identify and translate these strings manually. In addition, you can create translations for languages for which ServiceNow doesn’t provide a plugin.

Note:
Using the Localization Framework application is recommend for translating to an unsupported language or translating large amounts of custom UI string content in a supported language.

Locales, locations, and regions

Locale settings support displaying dates, times, and currencies properly based on your location. In addition, you can configure locations to identify users, assets, incidents, and more on the instance. Configure regions to organize languages in language selectors.

System Localization benefits

Benefit	Feature	Users
Display translations of base system UI string content.	• Language internationalization support • Activate a language	Administrators
Display dates, times, and currencies in an instance based on users' locations.	Configuring locales	
Translate UI string content that you create or modify.	Translating custom content	

Related applications

Localization Framework

Using the Localization Framework application is recommend for translating to an unsupported language or translating large amounts of custom UI string content in a supported language. Localization Framework also supports using Dynamic Translation for live translation of form fields, translating knowledge articles, and integrating with third-party translation services.

Translation management for Knowledge Management

Knowledge article translation uses a separate system that relies on the I18N: Knowledge Management Internationalization plugin. If this plugin isn't activated, knowledge article translation falls back on System Localization, with translations kept in the Translated Text [sys_translated_text] table.

Language internationalization support

The ServiceNow platform supports multiple languages.

Language support is available when at least one language plugin has been activated, as described in Activate a language.

When a user logs in, the language for the instance session is determined by the following logic:

1. If the language selection at login is enabled, that language is used.
2. If not, the language preference selected using the language picker in the header bar is used.
3. If not, the language setting in the User [sys_user] table is used.
4. If none of the above are true, the system default language is used.

You can also enable a fallback language to use if a translation isn't available in the preferred language, as described in Set a fallback language.

Some parts of the system aren't translated and remain in the language used to create them:

- Journal fields
- Report titles
- Any field that stores free-form text

Additionally, plugin names and spoke names aren't translated. The English version of plugins and spoke names is always kept, even when they're used in other contexts. Feature names are translated, however.

Beginning in the Washington DC release, new product and application names are localized in all supported languages. Existing product and application names are being localized incrementally.

FAQ

Question	Answer
Will fields that I created or edited myself be translated?	The ServiceNow AI Platform does not perform translations. ServiceNow only provides translations of strings that the company itself has created. When you create a field, a message, or other textual UI content, translate this content manually.

FAQ (continued)

Question	Answer
I am not seeing journal fields translated.	Translation tables are not provided for user-defined string fields, such as short description, additional comments, and report titles. These fields can appear only in the language they were created in.

Related topics

Locating translatable UI strings

User specific language

The user has several choices for identifying the language to be used for the instance, as described here.

To enable user-specific languages, first activate the I18N Translation plugins for the desired languages.

- *Language picker at login:* If user-specific language is enabled, users see a choice list on the login page to select their language. To control display of the choice list on the login page, navigate to **All** > **System Properties** > **UI Properties** and use the property *Show the language select box on the login page to allow the user to specify the language they would like to be logged in with*.
- *Language picker in the user menu (Next Experience UI):* Users can select their language from the user menu by navigating to **Preferences** > **Language & Region**. The user menu is located in the Next Experience Unified Navigation.
- *Language picker in the System Settings window (Core UI):* Users can select their language in the General tab of the System settings window. Access this window by clicking the gear icon on the right edge of the banner.
- *Language setting on the user table:* If you have users that require a different language in addition to the global language, you can specify a language for them in **All** > **User Administration** > **Users**. If the **Language** field isn't already on the form, you can configure the form and add it.

The **Language** field in the User [sys_user] table overrides the default global language for that user's sessions.

User language

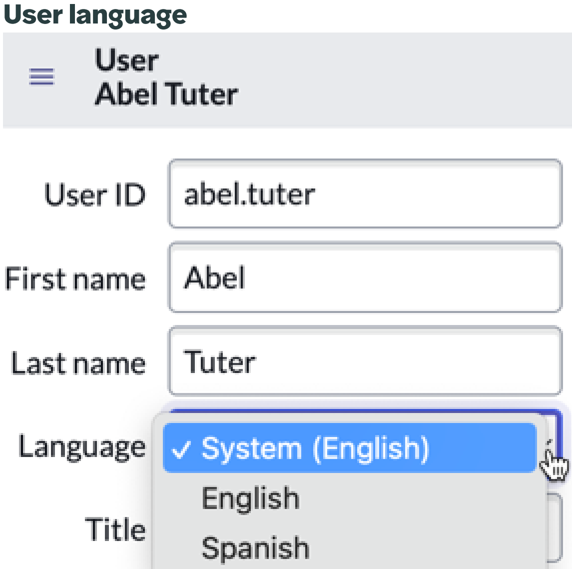

Note:
Setting the language for the system guest user sets the language for both the login page and all users without a user role.

Localizing currency and price fields

You can localize currencies for item prices and options.

Standard currency fields display the same price in different currencies based on the user's language. A price field is similar to a currency field, but with special features for conversion and display.

The system has the concept of a reference currency based on the instance locale. All currency values are automatically converted to the reference currency before aggregation or conversion. Do not change the global locale setting after you have data in the system. Changing the global locale setting after adding data to the system can cause aggregations to calculate and display incorrectly. For information about configuring the instance locale, see Configuring locales.

Related topics

- Currency administration
- Price fields
- Default currency values in scripts

Translating the knowledge base

The knowledge base has its own Translation Management system. However, when this is disabled, knowledge base articles use translated_html fields for translated content.

When any of the internationalization (I18N) plugins are enabled, the Knowledge Management Internationalization plugin is enabled. This plugin allows knowledge article authors to create language-specific articles and to translate and access articles. For more information about Knowledge Management internationalization, see Translation Management in the Knowledge Management documentation.

If the Knowledge Management Internationalization plugin is disabled, knowledge base articles use translated_html fields for article content. This type of field displays the translation based on the language of the user, if multiple translations are stored. To learn about using translated_html fields, see Translating text fields.

Translating Service Catalog cart labels

You can specify language-specific labels in these Service Catalog screens: Cart, Edit Cart, and Check Out (including workflows and approvals).

The text for the labels is stored on the Messages [sys_ui_message] table. Localized Message records for these labels are included in the I18N Internationalization language plugins. The application is Global. You can add a record for an unsupported language, or edit an existing record. If you edit an existing record, export the record to another application, edit it in that application, and re-import it. Otherwise, your edits may be overwritten when your instance is zbooted or updated. For more information about exporting translation records, see Export and edit translation records.

Service catalog cart labels

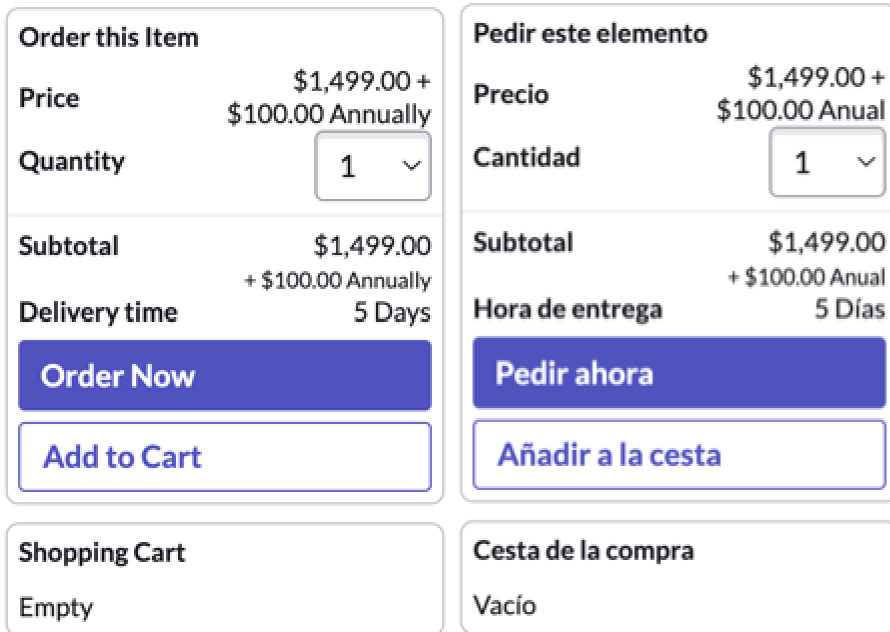

Configuring System Localization

Configure the locales, languages, regions, and locations needed to accommodate your users.

Configuration overview

The following list is a high-level overview of the workflow to get started localizing an instance. However, there may be use cases where performing these steps in a different order is preferred.

1. Set the instance locale

   By default, the instance locale is the United States (en.US). You should set the instance locale before a system has gone into production and not change it after.

2. Activate a language

   Activate any of the ServiceNow® internationalization (I18N) language plugins to get translations of base system UI strings in the languages your users need. American English is the default language and doesn't need to be activated.

3. Set the default language for an instance

   Select a default language for an instance if it should be a language other than American English.

4. Set a fallback language

   Specify the fallback language for another language so that if UI text isn’t translated in a language, the text displays in the fallback language instead. By default, if a UI string isn’t translated in the preferred language of an instance session, it displays in English.

5. Set up locations

   Configure the locations that your organization supports to help identify users, assets, incidents, and more by location.

6. Create regions for language selection in portals

Map languages to regions to organize the language selector for guest users.

7. Ensure that translated choices appear for applications added by plugins

   If you activate other application plugins after having activated language plugins, translated choices in lists might not appear for the applications or features installed by their plugins. Ensure that translated choices appear for those applications by activating translated choice records after corresponding English choice records have been added by plugins.

After you configure your instance using the localization options and supported languages, you might need to translate custom content or translate the instance to an unsupported language. For more information, see Translating custom content and Translating to an unsupported language.

Configuring locales

Specify a locale for an instance so information such as dates, times, and currencies display properly based on your location.

By default, the base system uses US standard formatting. For example, the currency default is the US dollar sign ($) displayed with two decimal places: $100.00. By customizing your locale, you can make things such as currency appear as you expect. For example, in France, you may want to see 100,00 € instead of $100.00.

Set the instance locale

Set the instance locale using a locale code.

Before you begin

Role required: admin

About this task

You should set the instance locale before a system has gone into production and not change it after it goes into production. If a user's locale must be changed, update the **Country code** field on the user record.

> **Important:**
> The value of this property determines the system's default currency, or reference currency, into which all prices are automatically converted before other sums or conversions are performed. Changing this property after any price or currency fields have been given a value (for Service Catalog items, Assets, Project Tasks, etc.) may result in improper conversion or prices that sum incorrectly.

Procedure

1. Navigate to **All** > **System Properties** > **System Localization**.
2. Enter the locale code to use under **Locale code to use for localization. Format is [language code].[country code] (e.g. en.GB for Britain fr.FR for France, de.DE for Germany, or ja.JP for Japan).**

   Common locales

Country	Locale code
United States	en.US
Great Britain	en.GB
France	fr.FR

Country	Locale code
Germany	de.DE
Japan	ja.JP

3. Select **Save**.

Activate a language

By default, the language supported by the platform is American English. You can activate plugins for other supported languages if you have the admin role.

Before you begin

Role required: admin

About this task

The following plugins are currently available:

- I18N: Arabic Translations
- I18N: Brazilian Portuguese Translations
- I18N: Czech Translations
- I18N: Dutch Translations
- I18N: Finnish Translations
- I18N: French Canada Translations
- I18N: French Translations
- I18N: German Translations
- I18N: Hebrew Translations
- I18N: Hungarian Translations
- I18N: Italian Translations
- I18N: Japanese Translations
- I18N: Korean Translations
- I18N: Norwegian Translations
- I18N: Polish Translations
- I18N: Portuguese Translations
- I18N: Russian Translations
- I18N: Simplified Chinese Translations
- I18N: Spanish Translations
- I18N: Swedish Translations
- I18N: Thai Translations
- I18N: Traditional Chinese Translations
- I18N: Turkish Translations

In addition, the I18N: Internationalization plugin (com.glide.i18n) provides the elements necessary for translating an instance without any translation preloaded. This plugin is useful for translating an instance to a language other than those listed above. For more information on using the I18N: Internationalization plugin to translate an instance, see Translating to an unsupported language.

Note:

- Activating internationalization plugins for any of the available languages automatically activates the I18N: Knowledge Management Internationalization Plugin v2 plugin (com.glideapp.knowledge.i18n2).
- Product and Application Names are kept in English in localization plugins. Feature names are translated.

Procedure

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

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

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

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

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

Set the default language for an instance

Change the language that appears by default for an instance.

Before you begin

Activate the languages that your users need. For more information, see Activate a language for supported languages or Translating to an unsupported language for custom translations.

Role required: admin

About this task

This property defines the language that users with a role see if a language is not specified in their user record. Users without a role see the default guest language, as described in User specific language.

Procedure

1. Navigate to **All** > **System Properties** > **System Localization**.
2. In the **Default language for the system (two character values)** field, enter a two-character language code.

   You can set the default language to any language that is active on the instance. The following language codes and languages are examples of possible values.

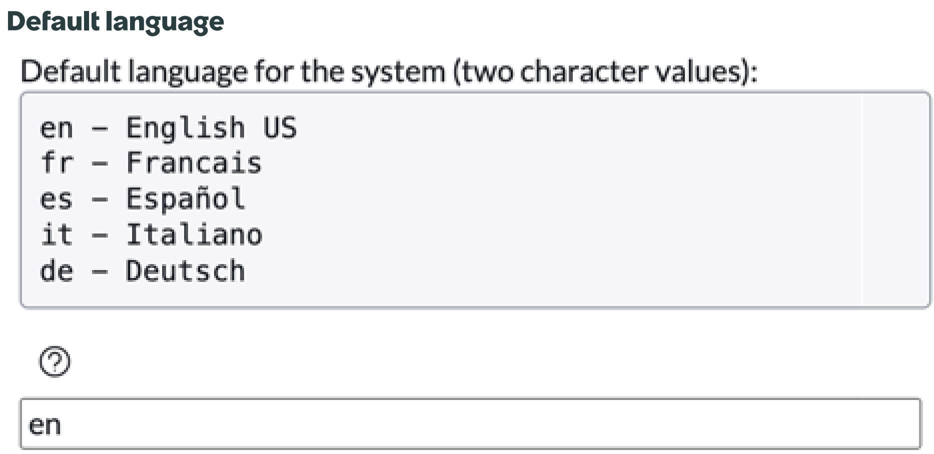

3. Select **Save**.

Set a fallback language

Set a fallback language to be used when user interface text is not translated in the user's preferred language to accommodate the language needs of your users.

Before you begin

Activate the languages that your users need. For more information, see Activate a language for supported languages or Translating to an unsupported language for custom translations.

Role required: admin

About this task

By default, if a UI string is not translated in the preferred language of an instance session, it displays in English. Setting fallback languages creates a three-level hierarchy in which an intermediate language is used before defaulting to English. Setting a fallback language also means you don't need to add custom translations for UI strings that are suitable in the fallback language and users get UI strings in the languages most appropriate for them.

For example, you could set the fallback language for Mexican Spanish to Spanish. If a UI string is not translated in Mexican Spanish, it displays in Spanish instead. If a Spanish translation is not available either, the UI string displays in English. As an admin, you would need to provide custom translations of UI strings in Mexican Spanish only when the translation differs from the Spanish translation.

> **Note:**
> To turn off fallback languages and always fall back to English, add the *`glide_i18n.language_fallback_enabled`* property to the System Properties [sys_properties] table and set the value to `false`.

Procedure

1. Navigate to **All** > **System Localization** > **Languages**.
2. In the Languages [sys_language] table, select the name of an active language.
3. In the **Fallback** field of the language record, select the name of the fallback language.
   You can select only an active language as the fallback language. You cannot set English as the fallback language because it is the default fallback language.
4. Select **Update**.

Set up locations

If your organization supports more than one distinct location, you can configure these locations in the platform to help further identify users, assets, and incidents.

Before you begin
Role required: admin

Procedure

1. In the application navigator, navigate to **All** > **User Administration** > **Locations**.
2. Select **New**.
3. On the Location form, enter the location name, address, contact, and other location details.
4. Select **Submit**.

Result
The new location is available as a reference anywhere you can specify one.

Create regions for language selection in portals

Create regions to group languages in the language selector in portal headers.

Before you begin
Role required: admin

About this task
In portal headers, the language selector displays the current language selection of the user and a menu of languages that the user can choose from. If languages are assigned to regions, languages display according to the configured region groups.

Languages grouped by region in a portal header

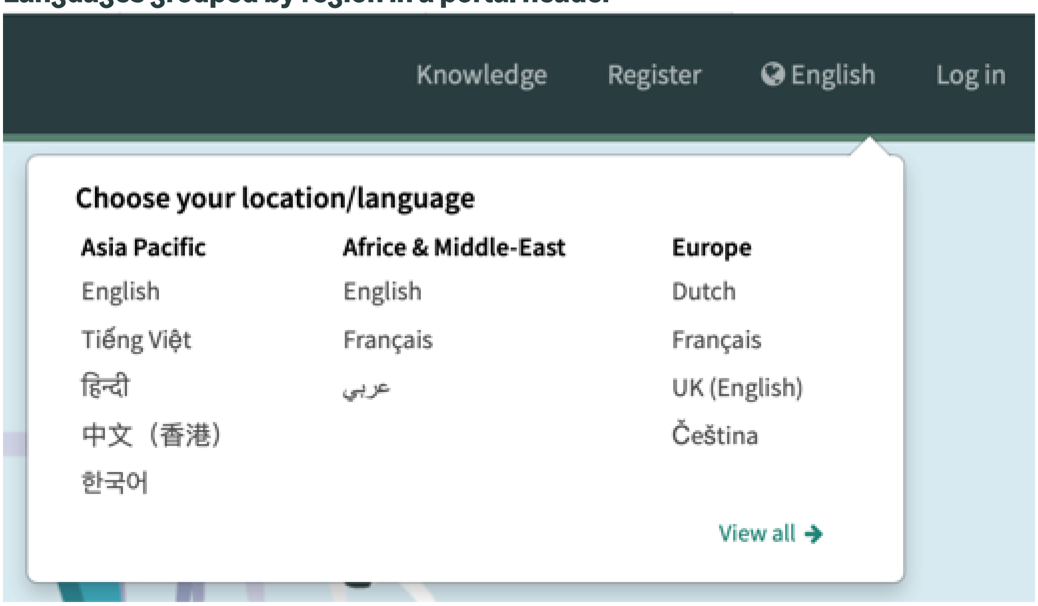

For more information about the language selector in portal headers, see Enable a language selector in portal headers.

Procedure

1. Create regions.
    a. Navigate to **System Localization** > **Regions**.
    b. Select **New**.
    c. In the **Region** field, add a name for the region.
    d. In the **Order** field, assign a numerical value in multiples of hundred for the region.
    The order controls where the regions appear in the language selector, with the lowest ordered region on the left and the highest ordered region on the right.
    e. Select **Submit**.
    f. Repeat these steps to create additional regions.
2. Assign languages to regions.
    a. Navigate to **System Localization** > **Language Regions**.
    b. Select **New**.
    c. In the **Region** field, select a region from the list.
    d. In the **Language** field, select a language to assign to the region from the list.
    e. Select **Submit**.
    f. Repeat these steps to assign additional languages to regions.

Ensure that translated choices appear for applications added by plugins

Ensure that translated choices appear for applications by activating translated choice records after corresponding English choice records have been added by plugins.

Before you begin

You must have activated a language and then activated a plugin that includes the corresponding English choice records. For more information, see Activate a language and Activate a plugin.

Role required: admin

About this task

When you activate plugins after having activated languages, translated choices in lists might not appear for applications or features installed by the plugins. Translated choice records installed by a language plugin that don't already have corresponding English records in the Choices [sys_choice] table are set to **Inactive** by default. If you activate plugins that include those corresponding English choice records later, you need to make these records active to display translated choices for any applications or features installed with the plugin.

Procedure

1. Navigate to **All** > **System Localization** > **Languages**.
2. In the Languages [sys_language] table, select the name of any active language.

    You can select any language because this process updates choice records for all languages.
3. In the Related Links section, select **Repair non-english sys_choice records**.

4. In the Repair Non-English Choices? dialog box, select **Yes**.
   A job that makes the **Inactive** field for all translated choice records match their corresponding English choice records runs 15 minutes after you make this choice.

Result

Translations for all choices with corresponding English choice records appear for applications or features installed with plugins.

Translating custom content

Translate content you create or modify to a supported language or add an unsupported language.

> **Note:**
> Using the Localization Framework application is recommend for translating to an unsupported language or translating large amounts of custom UI string content in a supported language. For more information, see Localization Framework.

ServiceNow provides translations of the base system UI strings in supported languages through the internationalization (I18N) language plugins. For applications you create and modifications you make to the ServiceNow AI Platform® content, you must provide your own translations. You can also provide any translations to unsupported languages that ServiceNow does not provide translations for. The ServiceNow AI Platform does not provide any automatic translation capability.

The translation method depends on the material that you are translating. Repeat the translation process whenever you edit or update content.

Translation methods and use cases

What needs translation	Approach to take	Resources
Field labels in a table for a customer-created application	1. Create English versions for all fields in the table. 2. Locate the records for the table in the application. 3. Open the Field Labels translation table. 4. Export the records, translate the appropriate fields, set the language field, and reimport through an import set.	• Export and edit translation records • Alternatively, follow the procedure for translating individual labels in Translate a field label.
HTML in a UI page in a customer-created application	1. Create the English version. When creating the HTML in the UI Page form, wrap any translatable text in `gs.getMessage` calls. 2. Create a Message [sys_ui_message] record for each English message with	Translate a client script message

Translation methods and use cases (continued)

What needs translation	Approach to take	Resources
	a corresponding key/value pair. 3. Create additional Message records for each message for each other language.	
Service Catalog items	1. Navigate to **All** > **System Properties** > **System Localization** and set **Display translation prefix on translatable strings** to **Yes**. 2. Create the item in English. 3. Change the preferred language to a different language. 4. Translate all the translatable items, now identified by translation prefix. 5. Save the record.	• To activate translation prefixes, see Debug translations. • For translatable text fields, which are common in Service Catalog items, see Translating text fields.
Service Portal widget	1. Create the English version. Externalize any translatable text by surrounding the text with `${}` wrappers. 2. Create a Message [sys_ui_message] record for each English translatable text with a corresponding key/value pair. 3. Create additional Message records for each piece of translatable text for each other language.	• Translate a client script message • Internalize a widget in Service Portal documentation • For unofficial advice, see Building multi-lingual service portals
Edit a translation provided in a plugin	1. Export the records from the translation table. 2. Edit them outside the translation table and reimport them. Do not edit the records directly, because your changes could be overwritten	Export and edit translation records

Translation methods and use cases (continued)

What needs translation	Approach to take	Resources
	when you upgrade or zBoot your instance.	
Translate an instance to an unsupported language	1. Create a Choice [sys_choice] record for that language. 2. Add other translation records through the Export and import set approach.	Translating to an unsupported language
Support Chat texts	Support Chat is mostly translated through the I18N language plugins. If you customize any Support Chat texts, follow the procedure for Service Portal widgets.	Internationalize a widget in Service Portal documentation
Notifications and surveys	System Localization does not cover platform notifications or surveys. The usual approach is to create a separate version of each survey or notification in each language and use business rules to send the right language version to the right end user.	—

Locating translatable UI strings

Locate individual untranslated strings to provide your own translation of customer-created and customer-edited strings, such as field names and messages.

The internationalization (I18N) language plugins contain only translations of the UI content that ServiceNow provides by default in your instance. If you have customized UI strings that you want to translate, you can identify and translate these strings manually. To perform these translations, you must first take the following actions:

- Identify the strings that need to be translated.
- Identify which translation tables contain the strings.

To identify strings that have not been translated into any language and their translation tables, see List untranslated items.

If you know where a string appears in the UI and only have a small amount of material to translate, consider using the debug functionality to identify which strings are translatable and their translation tables. This functionality consists of prefixes that appear before all translatable strings and identify the translation table that the strings are in. For more information, see Debug translations.

List untranslated items

Get a list of all custom UI strings that have not been translated into any language (English only) to provide your own translations.

Before you begin

Role required: admin

About this task

This procedure activates modules corresponding to the following translation tables that list records with only an English version.

- Translated Name / Fields
- Messages
- Field Labels
- Choices

This method of listing untranslated items does not identify strings that have been translated into at least one non-English language but are missing translations in other languages.

> **Note:**
> This process works best with a small number of untranslated items. For information about a workaround for large numbers of untranslated items, see the Inefficient load of customizations for System Localization -> Non-Translated Items [KB0813304] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **System Definition** > **Application Menus**.
2. Select the **System Localization** application menu.
3. From the **Modules** list in the **System Localization** record, update the **Customizations**, **Translated Name/Fields**, **Messages**, **Field Labels**, and **Choice** modules to set the **Active** field to **True** using the list editor.

   When the list is sorted by **Order**, the modules appear in order from 500 to 900.

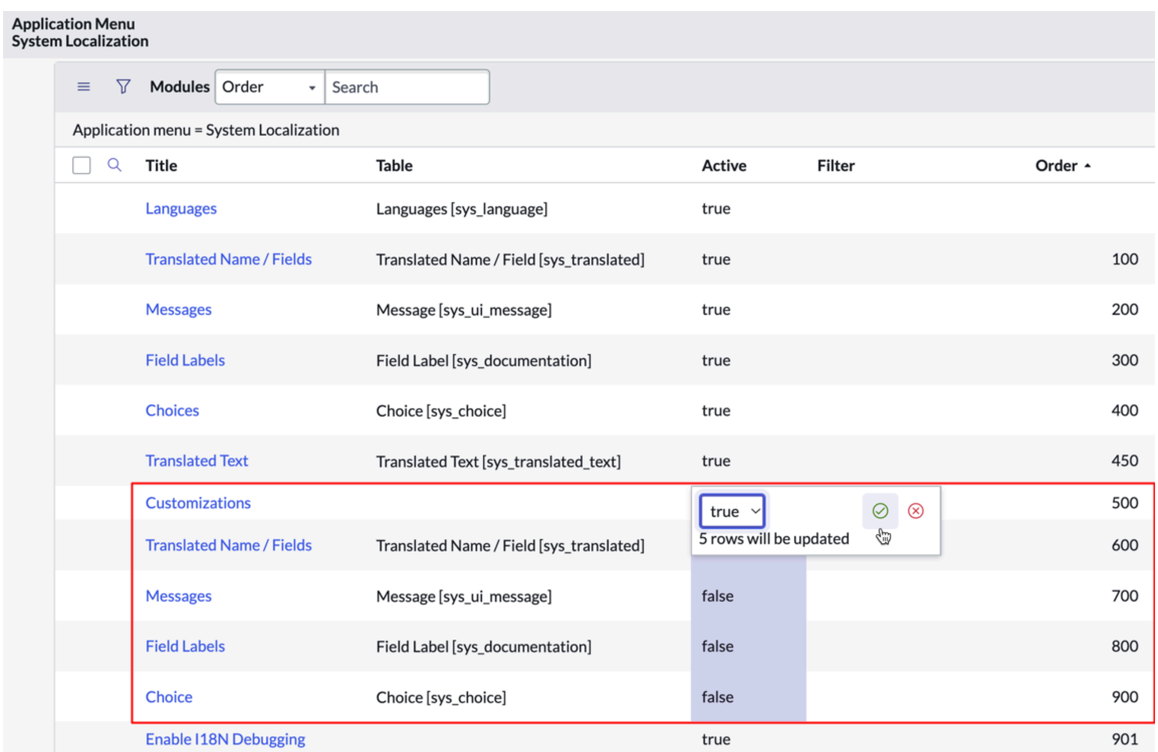

Note:
Instead of **Customizations**, your list might show a record titled **Non-translated Items**. When you activate the **Customizations** module, the title changes to **Non-translated Items**.

After the application navigator refreshes, the modules appear in the filter navigator beneath **Non-translated Items**.

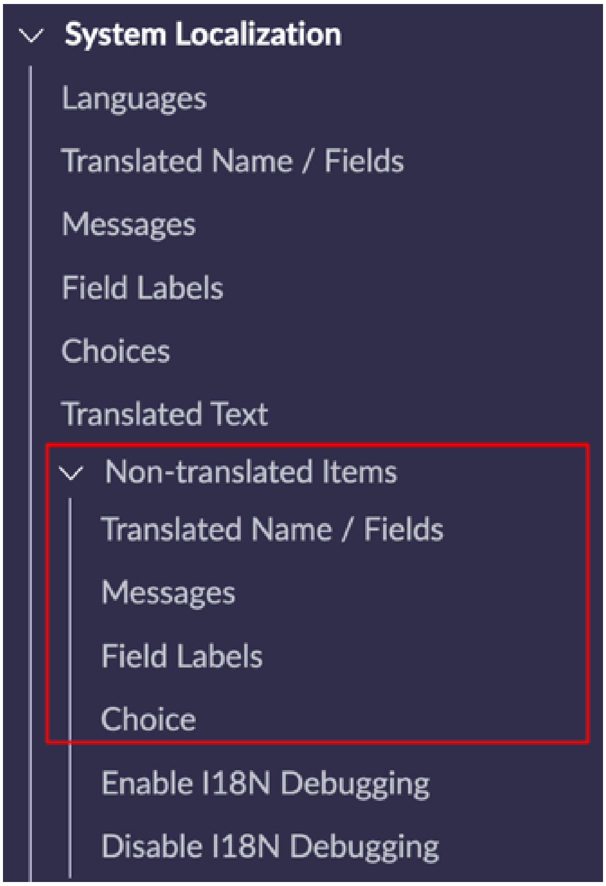

4. From the application menu, open each module and export their lists to any supported format, such as Microsoft Excel.

Example
For example, select **Non-translated Items** > **Translated Name/Fields**, right click any column, and export the list.

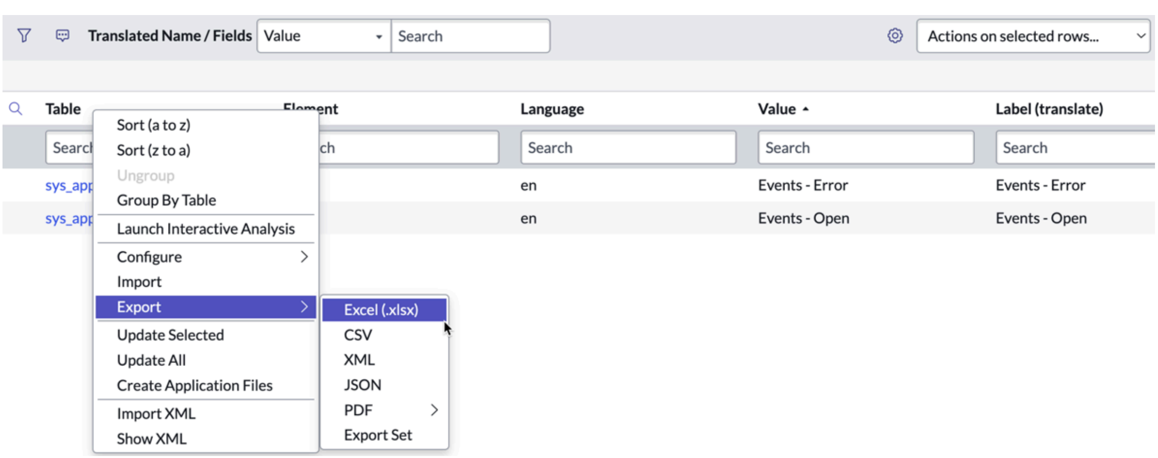

What to do next

You can translate the fields in an external application and reload the translations as an update set. For more information see, Export and edit translation records. If you have a small number of strings to translate, you can translate them individually. For more information, see Translating individual UI strings.

Debug translations

Show which UI strings are translatable and their translation table to debug translations.

Before you begin

Role required: admin

About this task

You can set translation prefixes to appear only in the current user session or to appear in all sessions.

Translation prefixes map to translation tables as follows.

Translation prefixes mapped to tables

Prefix	Table	Navigation
TRF	Translated Name / Fields [sys_translated]	**System Localization** > **Translated Names / Fields**
MSG	Messages [sys_ui_message]	**System Localization** > **Messages**
GMLD	Field Label [sys_documentation]	**System Localization** > **Field Labels**
TRT	Translated Text [sys_translated_text]	**System Localization** > **Translated Text**
CHC	Choice [sys_choice]	**System Localization** > **Choices**

Note:
A few strings may not display translatable prefixes. In this case, the string is not stored on any of these tables. Examples of this behavior include:

- Text embedded in images, such as the buttons in the Service Catalog.
- Text defined by properties, such as the text which follows the banner.

Procedure

1. Enable translation prefixes on field labels.
    - To enable prefixes for the current user session, navigate to **All** > **System Localization** > **Enable I18N Debugging**.
    - To enable prefixes for all sessions, navigate to **All** > **System Properties** > **System Localization** and set **Display translation prefix on translatable strings** to **Yes**.

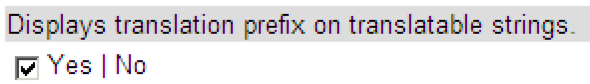

2. Refresh the page or change your user language to one of the non-English languages installed on your instance to see the prefixes.

Result

When enabled, translation prefixes appear before translatable UI strings.

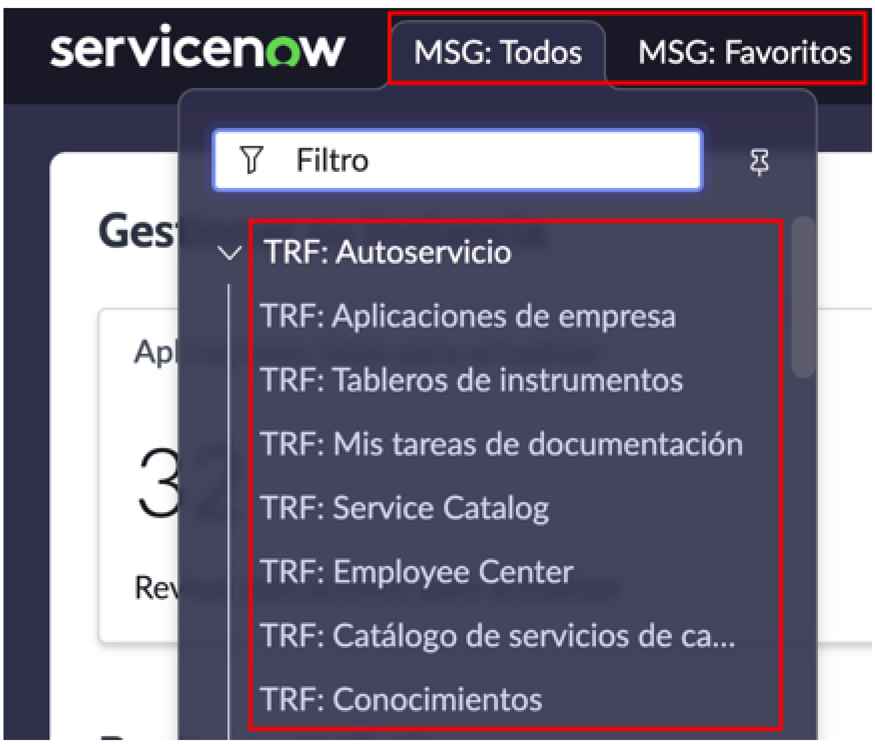

What to do next

When you are done debugging, disable the translation prefixes.

- To disable prefixes for the current user session, either log out of the session or navigate to **All** > **System Localization** > **Disable I18N Debugging**.
- To disable prefixes for all sessions, navigate to **All** > **System Properties** > **System Localization** and set **Display translation prefix on translatable strings** to **No**.

Translating individual UI strings

When translating just a few UI strings, such as when you add customizations to a translated instance, use the procedure that applies to the type of text being translated.

UI strings include field names and values, related list names, and messages.

Fields and related lists

Three types of fields support translated strings for fields and related list names:

Translated_field

Stores field labels, related list names, and certain field values. The value of the translated_field replaces the label, list name, or field value when the user selects the matching language. Translated_field values have a one-to-many relationship with their associated keys. As a result, multiple records can reference one translated_field value. For more information, see the following topics:

- Translate a field label
- Translate a related list name
- Translate a field value for the Classic Environment

Translated_text

Stores long text values in plain text. The value of the translated_text replaces the plain text when the user selects the matching language. Translated_text values have a one-to-one relationship with their associated keys. As a result, only one record can reference a translated_text value. For more information, see Translating text fields.

Translated_html

Stores long text values in HTML. The value of the translated_html replaces the HTML when the user selects the matching language. Translated_html values have a one-to-one relationship with their associated keys. As a result, only one record can reference a translated_html value. For more information, see Translating text fields.

The translated values are stored as separate records and display the proper value according to the end user's language. To determine the field type, right-click the field on the form, select **Configure Dictionary**, and check the **Type** field.

Messages

Informational messages, confirmation messages, error messages, and other types of system messages across the platform can be translated. Messages that support translation come from HTML in a UI page in a customer-created application or Service Portal widget messages when these texts are wrapped correctly. When a client script contains a getMessage call or a server script contains a `gs.getMessage` call, the system checks the Messages [sys_ui_message] table for a translation. Fore more information, see Translate a client script message.

Translate a field label

Field labels are the names that appear on forms and lists to describe the type of information the field contains.

Before you begin

Role required: admin

About this task

The following procedure works best for translating individual field labels, such as those added with a customization. To translate large numbers of field labels, use the procedure described in Export and edit translation records.

Procedure

1. Navigate to the field on the form.
2. Right-click the field label and select **Configure Label**.
3. In the *Field Label* form, replace the English text with the text of the target language in the *Label*, *Plural*, and *Hint* fields.
4. Enter the BCP 47 code of the target language.
5. Right-click the header bar and select **Insert**.
   Selecting **Insert** creates a new record in the Field Label [sys_documentation] table for this field label in the selected language. For more information about this table, see Field Label table.

Translate a field value for the Classic Environment

Field values are the text entries that are used for fields with the type translated_field, such as the *Title* or *Hint* field in the Module [sys_app_module] table.

Before you begin

Role required: admin

About this task

The following procedure works best for translating values for individual fields, such as those added with a customization. To translate large numbers of field values, use the procedure described in Export and edit translation records.

Procedure

1. Use the language picker in the header bar to switch to the desired target language.
2. Navigate to the field on the form.
3. Update the field value with the text translated into the target language (Spanish).

4. Select **Update**.

This creates a new record in the Translated Name / Field [sys_translated] table for the selected language or updates the existing record if a translation already exists. For more information about this table, see Translated Name / Field table.

Translated field

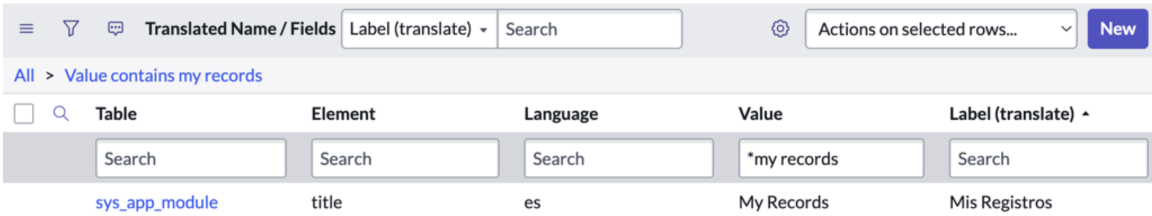

Translate a choice label

Translate the labels of choices in a list.

Before you begin

Role required: admin

About this task

To translate a choice in a list, you must create a choice record for each language that you want to support in the Choice [sys_choice] table. For more information about this table, see Choice table.

Procedure

1. Navigate to **All** > **System Localization** > **Choices**.
2. Select **New**.
3. On the form, fill in the fields.

Choice form

Field	Description
Table	The name of the table that contains the choice.
Element	The column name for the field that contains the choice.
Language	The BCP 47 code of the language for the translated choice. The identifier can contain a language code or language code followed by a country or region code. For example, *tr* for Turkish or *es-MX* for Spanish Mexico.
Label	The label for the choice in the specified language.
Value	A key that is the same for the choice in each language. For example, you could use the English label or the sequence number as the value for all languages to map them to each other.
Sequence	A number that determines in what order the option appears in the list if you don't want to list choices alphabetically.

Example
The following record translates the "Email" choice in the subcategory field on the Incident table into Spanish.

- **Table**: Incident [incident]
- **Element**: subcategory
- **Language**: es
- **Label**: Correo electrónico
- **Value**: email

4. Select **Submit**.
5. **Optional:** From the context menu, select **Insert and Stay** and update the form to translate the choice into another language.

Translating text fields

Translate long pieces of text into the language of the user.

Translated text fields allow the same field to display different content based on the language of the user.

There are two different translated text fields:

- Translated Text: Stores long text values in plain text.
- Translated HTML: Store long text values in HTML.

These fields operate the same as normal text and HTML fields respectively, except that they can store multiple inputs in multiple languages. The most frequent use of translated text fields is Service Catalog names, descriptions, or variables.

Note:
There is a slight performance penalty associated with changing a normal HTML or text field into a translated HTML or text field. It is best only to use translated fields if the translated capability is required.

For translated text and HTML fields, the English language text is stored in the main table, but the values of other languages are stored in the Translated Text [sys_translated_text] table. Each translated field on every row has one or more entries in the Translated Text table: one per language for which ServiceNow provides a translation. For more information, see Translated text table.

Note:
Knowledge article translation uses a separate system that relies on the I18N: Knowledge Management Internationalization Plugin. If this plugin is not activated, knowledge article translation falls back on the system described here, with translations kept in the Translated Text [sys_translated_text] table. For more information about Knowledge Management internationalization, see Translation management in the Knowledge Management documentation.

Translate long text content

Translate the content of translated text and HTML fields directly in the form where the content is specified.

Before you begin
Role required: admin

About this task

Use the following procedure to translate the content for customer-created text fields. To edit existing text translations, export the records from the Translated Text table and follow the procedure in Export and edit translation records.

Procedure

1. Open the record with the field that you want to translate.
2. Verify that the field whose content you want to translate is a translated_text or translated_html field.
   The following steps are one possible way to verify a field type:

   a. Right-click the label of the field and select **Configure dictionary**.

      Example
      In this example, the context menu has been opened for the **Name** field for a record in the Record Producer [sc_cat_item_producer] table.

      Context menu for field

      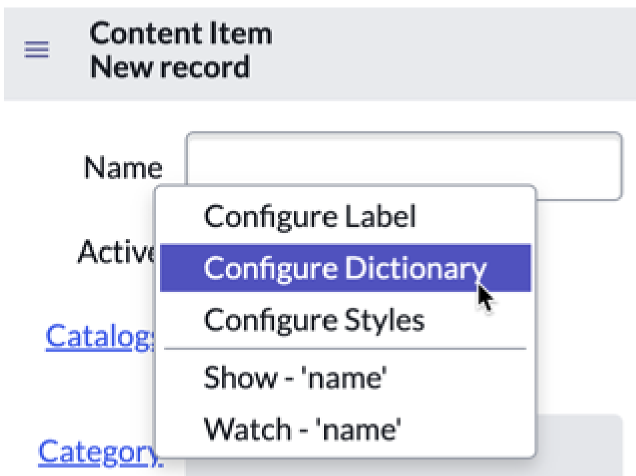

   b. In the **Dictionary Entry** record that opens, confirm that the **Type** field is **Translated Text** or **Translated HTML**.

      Example
      In this example, you see that the **Name** field is of the type **Translated Text**.

Dictionary entry for field

c. Navigate back to the form.

3. Use the language picker to switch to the desired language.
4. In the relevant field, replace the English text with the text of the target language.
5. Select **Submit**.
   A new record in the Translated Text [sys_translated_text] table for the active language is added. For more information about this table, see Translated text table.

Translate a related list name

Related lists appear at the bottom of forms. You can translate a related list name by configuring the list control.

Before you begin

Role required: personalize_control

Procedure

1. Use the language picker to switch to the desired language.
2. Navigate to the related list on the form.
3. Open the list control for the related list by performing the appropriate action for the list version.
4. On the List Control form, replace the existing *Label* with the text of the target language.
5. Select **Submit** or **Update**.
   The system creates a new entry in the Translated Name / Field [sys_translated] table or updates the existing entry for this language. For more information about this table, see Translated Name / Field table.

Translate a client script message

Translate the messages used by client scripts.

Before you begin

Create a client script that contains messages. For information about creating client scripts, see Client scripts.

Role required: admin

About this task

Client scripts include a multi-line *Messages* field, which is for message strings that the client script can use as a key to look up a localized message alternative from the Message [sys_ui_message] table. Each message key in a client script is on a separate line in the *Messages* field. The instance looks for a localized message string anytime the client script makes a getMessage(msg) call where the msg string matches a key in the *Messages* field in the client script.

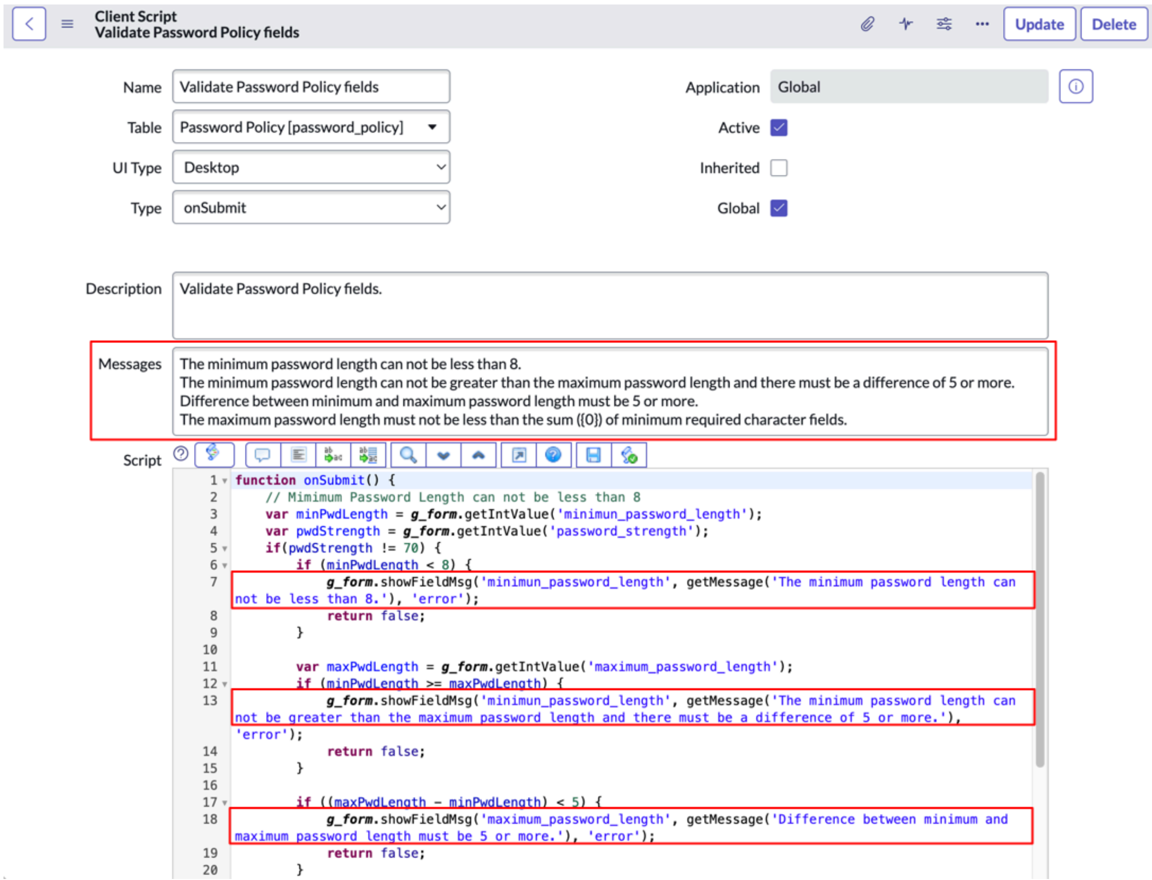

For example, if you add the string `Please populate the Reason` field to the *Messages* field in the client script, then the instance will look for a localized string from the Message `[sys_ui_message]` table any time the client script calls:

```
getMessage("Please populate the Reason field")
```

Note:
If you create a record in the Message [sys_ui_message] table that uses an existing key and the same language, the message for existing references using the key can change. For example, if a record exists with "Refresh" for the key, English for the language, and "Refresh" for the message, and you create another record with the same key and language but a different message of "Refresh the screen", existing UI strings using the "Refresh" key will display the new "Refresh the screen" message. Records in this table are created in the global scope and aren’t scope-restricted.

Procedure

1. Navigate to **All** > **System Localization** > **Messages**.
2. Select **New**.
3. On the form, fill in the required fields.

Message form

Field	Description
Key	The internal unique identifier for this message.
Language	The language the message is translated into.
Message	The translated text that users see.

Example

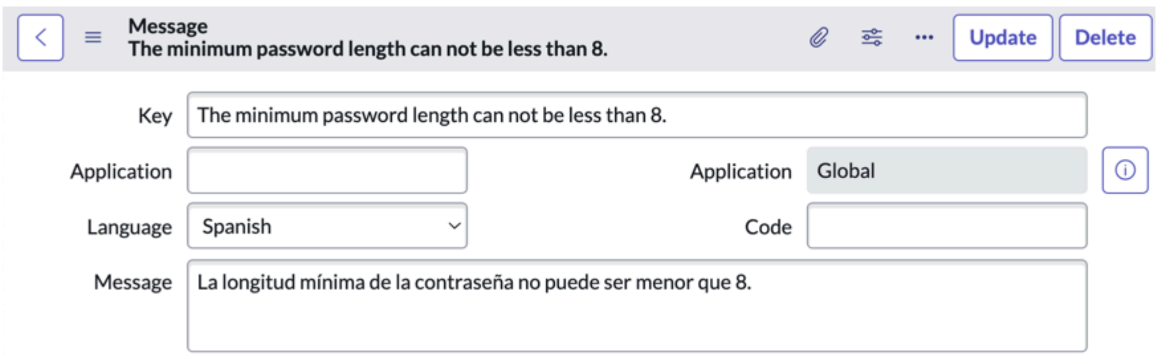

4. Select **Submit**.
   A new record is added to the Message `[sys_ui_message]` table for the localized message key string. For more information about this table, see Message table.

Export and edit translation records

Export translation records to translate to additional languages in your preferred translation environment and import the new or edited translation records through import sets.

Before you begin

Note:
Using the Localization Framework application is recommend for translating to an unsupported language or translating large amounts of custom UI string content in a supported language. For more information, see Localization Framework.

Role required: admin

About this task

In the following cases, translate strings in the UI outside of the ServiceNow AI Platform and import them through import sets:

- You have a significant amount of customer-created UI components and you want to translate these components into other languages.
- You want to modify translations that ServiceNow provides in internationalization plugins.
- You want to translate an instance into a language that ServiceNow does not support with internationalization plugins. For more information, see Translating to an unsupported language.

All translatable UI strings items are stored in translation table records. The only time you would translate UI directly in translation records is when you have a small amount of customer-created UI components and you know exactly where they are. For more information about translation tables, see Translation tables and Translating individual UI strings.

Procedure

1. If you are translating customer-created UI components, locate the translatable strings.
   For more information, Locating translatable UI strings.
2. Export the relevant translation table records into a format (such as Excel) that can be easily manipulated.
   For more information about exporting a list of records, see List export.
3. Within the exported document, translate the following fields in each row, depending on the translation table:

Translation table	Fields to translate
Field Labels	Label, Plural, Hint
Translated Name/Field	Label (translate)
Message	Message
Choices	Label
Translated Text	Value

4. Within each row, change the "Language" column to the BCP 47 code of the new language.
   In Translated Text and Message records, type the full name of the language instead of the BCP 47 code.
5. Save your changes.

What to do next

Import the translated document back into the instance as an import set.

Import a translation from an Excel spreadsheet

The System Import Sets application contains import tables and corresponding transform maps to assist with importing translations from an Excel spreadsheet.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Import Sets** > **Load Data**.
2. Select **Use Existing** and the **Table name** that matches the type of data being imported, as follows.
   - For choices, select the `[u_sys_choice]` table.
   - For field labels, select the `[u_sys_documentation]` table.
   - For translated names and fields, select the `[u_sys_translated]` table.
   - For messages, select the `[u_sys_ui_message]` table.
   - For translated text, select the `[u_sys_translated_text]` table.
3. Select **Upload an Excel file**, and then click **Browse** to select the source Excel spreadsheet.
4. If appropriate, specify the *Work sheet* and *Header row number*.
5. Select **Go**.

The translations are now available in the appropriate *Import Set Table*.

6. Navigate to **All** > **System Import Sets** > **Table Name** and review the imported information to verify that the import was successful.
7. To transform the imported data into the corresponding table, navigate to **All** > **System Import Sets** > **Run Transform**.
8. Select the appropriate transform map, as follows.

   Note:
   Make sure you choose a transform map that has the **Run Business Rule** option selected. If the transform map does not have this option selected, any customized translations you have may be overwritten during the next upgrade.

   - For choices, select the **Sys Choice Translation Map**.
   - For field labels, select the **Sys Documentation Translation Map**.
   - For translated names and fields, select the **Sys Translated Translation Map**.
   - For messages, select the **Sys UI Message Translation Map**.
   - For translated text, select the **Sys Translated Text Translation Map**.
9. Select **Transform**.

Translating to an unsupported language

Translate the interface of an instance to a language for which ServiceNow does not provide an internationalization (I18N) plugin.

Before you begin

Note:
Using the Localization Framework application is recommend for translating to an unsupported language or translating large amounts of custom UI string content in a supported language. For more information, see Localization Framework.

Contact Customer Service and Support to request a file containing strings from the interface in the English language.

Role required: admin

About this task

All interface items are stored as translatable strings in the translation tables. By default, these translation tables only contain English strings. To populate these tables with translated strings in an unsupported language, follow this workflow. For more information about translation tables, see Translation tables.

Procedure

1. Create a language record for the unsupported language.
2. Export and translate the strings that you want to translate, as described in Export and edit translation records.
   There is a very large number of translatable strings on the platform. Exporting all of them and translating them manually is probably not a realistic possibility. Consider exporting filtered subsets of the strings from each table and then translating them.

3. Import the translated document back into the instance as an import set, as described in Import a translation from an Excel spreadsheet.
4. Create a language choice for users to select the new language.

Create a language record

Add a record for a language not provided by an internationalization (I18N) plugin to associate new translations with.

Before you begin

Role required: admin

About this task

To translate the interface of an instance to a language for which ServiceNow does not provide an internationalization plugin, you must create a language record for a new translation in the Languages [sys_language] table. For more information about this table, see Languages table.

Procedure

1. Navigate to **All** > **System Localization** > **Languages**.
2. Select **New**.
3. Provide the language information.
   - **Name**: The name of the language, for example, Turkish.
   - **ID**: The BCP 47 code for the language, for example, tr.
   - **Text Direction**: The direction that the instance should use to display the language: Left-to-Right or Right-to-Left.
   - **Active**: Option to activate the language for the instance.
   - **Fallback**: The name of another active language to use as a fallback language. For more information about fallback languages, see Set a fallback language.
4. Select **Submit**.

Create a language choice for users

Allow users to select the language as a valid option in the language picker or a user record.

Before you begin

Role required: admin

About this task

You must create a choice record for a new translation in the Choice [sys_choice] table. For more information about this table, see Choice table.

Procedure

1. Navigate to **All** > **System Localization** > **Choices**.
2. Select **New**.
3. Enter the following fields.

Choice record fields to select the new language

Field	What to enter
Table	`sys_user`

Field	What to enter
Element	`preferred_language`
Language	Enter an identifier that complies with BCP 47. The identifier can contain a language code or language code followed by a country or region code. For example, *tr* for Turkish or *es-MX* for Spanish Mexico.
Label	Enter the name of the new language selection as you want it to appear in the language picker. For example, Mongolian.
Value	Enter an identifier that complies with BCP 47. The identifier can contain a language code or language code followed by a country or region code. For example, *tr* for Turkish or *es-MX* for Spanish Mexico.
Sequence	Enter a number to determine what order the option appears in the choice list if you do not want to list choices alphabetically. For example, 5.

Example
The following record allows a user who is using the Dutch localization of their instance to change languages to Mongolian:

- **Table:** User [sys_user]
- **Element:** preferred_language
- **Language:** nl
- **Label:** Mongoolse
- **Value:** mn

4. Select **Submit**.

Configure a language as reading from right to left

Use the Text Direction field to configure a language that reads from right to left, such as Hebrew.

Before you begin

Activate the languages that your users need. For more information, see Activate a language.

Role required: admin

About this task

Right-to-left language support is available only in the main user interface and on live feed. Other user interfaces and applications, such as the graphical Workflow Editor, reporting, CMS, chat, and the ServiceNow documentation sites, are not supported.

Procedure

1. Navigate to **All** > **System Localization** > **Languages**.
2. Select **New**.
3. Enter the **Name** of the language, such as Hebrew.
4. Enter the BCP 47 code for the language.

Example
For example, Hebrew is `he`.

5. In the Text Direction field, select **Right-to-Left**.
6. Select **Submit**.

System Localization reference

Reference topics provide additional information about configuration properties, tables, and more.

System Localization properties

System Localization properties configure translation, currency, and locale settings on the instance for users from multiple countries.

Some localization properties only appear after the I18N: Internationalization plugin is activated. Some properties are available on the System Localization properties page, but some lesser-used properties are available only from the System Properties [sys_properties] table.

Access the following properties by navigating to **All** > **System Properties** > **System Localization**.

Property	Description
Show the language select box on the login page to allow the user to specify the language they would like to be logged in with. *glide.ui.login.language.select*	Enables language selection upon login. • **Options**: Yes | No • **Default value**: Yes • **Dependency**: The I18N: Internationalization plugin must be active
Text Search stemming language. *glide.ts.stemming_language*	Selects the language to match derived words in text search. • **Options**: English, German, and French • **Default value**: English • **Dependency**: The I18N: Internationalization plugin must be active • **Learn more**: Zing matches derived words with stemming
Displays translation prefix on translatable strings. *glide.ui.i18n_test*	Adds a translation prefix on translatable fields to indicate where to find the string for translation. • **Options**: Yes | No • **Default value**: No • **Learn more**: Debug translations
Use a single currency model. Display all currencies in the same currency code, regardless of a user's locale, country, or language code.	Enable single currency mode. • **Options**: Yes | No • **Default value**: No • **Learn more**: Currency administration

Property	Description
glide.i18n.single_currency	
Default language for the system (two character values) *glide.sys.language*	• **Options**: en, fr, es, it, de • **Default value**: en
If using the single currency model, display all currencies using this currency code. Currency codes use the ISO 4217 three letter format. *glide.i18n.single_currency.code*	• **Options**: ISO 4217 three-letter currency code • **Default value**: USD • **Learn more**: Currency administration
Locale code to use for localization. *glide.system.locale*	Determines the system's default currency into which all prices are automatically converted before other sums or conversions are performed. Changing this property after any price or currency fields have been given a value may result in improper conversion or prices that sum incorrectly. • **Options**: [language code].[country code] (for example, en.GB for Britain, fr.FR for France, de.DE for Germany, or ja.JP for Japan) • **Default value**: None • **Learn more**: Locale settings **Note:** You have already set a value on this system property, do not change this value once a system has gone into production. If a user's locale must be changed, update the "Country code" field on the user record.
Spell checker dictionary to use for English users. *glide.spell.dictionary.en*	Sets the spell checker dictionary used in the system for English users. • **Options**: Brazilian Portuguese, Dutch, English US, English UK, French, German, Italian, Portuguese, Russian, Spanish, and Thai • **Default value**: English US • **Learn more**: Locale settings

Access the following properties by entering `sys_properties.list` in the navigation filter to open the System Properties [sys_properties] table.

Property	Description
glide_i18n.language_fallback_enabled	Option to allow setting fallback languages for the instance. Setting fallback languages creates a three-level hierarchy in which an intermediate language is used before defaulting to English when user interface text is not translated in the user's preferred language. • Type: true | false • Default value: true • **Learn more**: Set a fallback language

Translation tables

ServiceNow stores translation information in these tables.

- Languages [sys_language]
- Translated Name / Field [sys_translated]
- Messages [sys_ui_message]
- Field label [sys_documentation]
- Choice [sys_choice]
- Translated Text [sys_translated_text]

To see which table contains a specific translated label, enable the display of translation prefixes. For more information, see Debug translations.

Note:
The Languages table is available only after at least one I18N:Internationalization plugin has been activated.

Related topics

Activate a language

Languages table

The Languages [sys_language] table contains a list of the languages for which translated text is available.

To enable translation to a new language, add a record to the Languages table. The main fields for this table are:

- Name: Language name
- ID: BCP 47 code for the language
- Text Direction: Direction of text in this language
- Active: Indicator that shows whether the language has been activated (true) for this instance or not (false)
- Fallback: Name of the language that the instance uses if a translation is not available

Choice table

The Choice [sys_choice] table contains translated text for options that appear in lists.

The main fields for this table are:

- Table: The name of the table this translation applies to.
- Element: The name of the field this translation is used for.
- Language: The BCP 47 code of the language for the translated choice.
- Value: A key that is the same for the choice in each language. For example, you could use the English label or the sequence number as the value for all languages to map them to each other.
- Label: The label for the choice in the specified language.
- Sequence: A number that determines in what order the option appears in the list if you don't want to list choices alphabetically.

Choice list

Choices | Table | Search

Table ▲	Element	Language	Value	Label
change_request	priority	es	1	1 - Crítica
change_request	priority	en	1	1 - Critical
change_request	priority	es	2	2 - Alta
change_request	priority	en	2	2 - High
change_request	priority	es	3	3 - Moderada
change_request	priority	en	3	3 - Moderate
change_request	priority	es	4	4 - Baja
change_request	priority	en	4	4 - Low

Translated choices

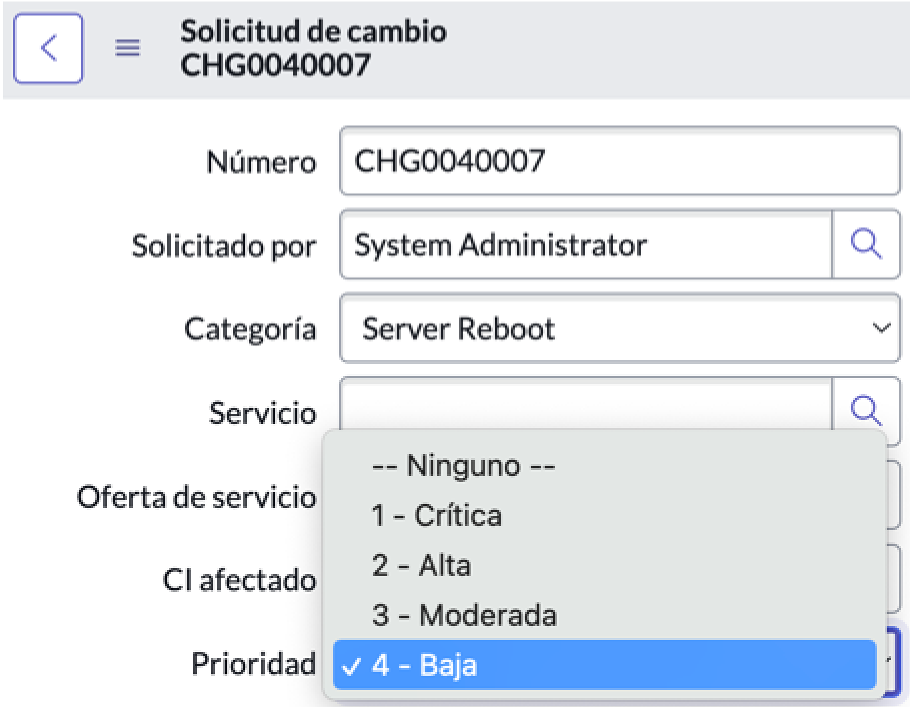

Translated Name / Field table

The Translated Name / Fields [sys_translated] table stores translated values for text fields where the field type is translated_field (see the dictionary entry).

This option is available for text fields up to 255 characters in length. Some examples are names, titles, and short descriptions. The main fields for this table are:

- Table: name of the table this translation applies to.
- Element: name of the field this translation applies to.
- Label (translate): translated text that users see on forms and lists.
- Language: BCP 47 code of the language for this translated text.
- Value: English value that causes this translated text to be displayed.

Translated Name/Fields table

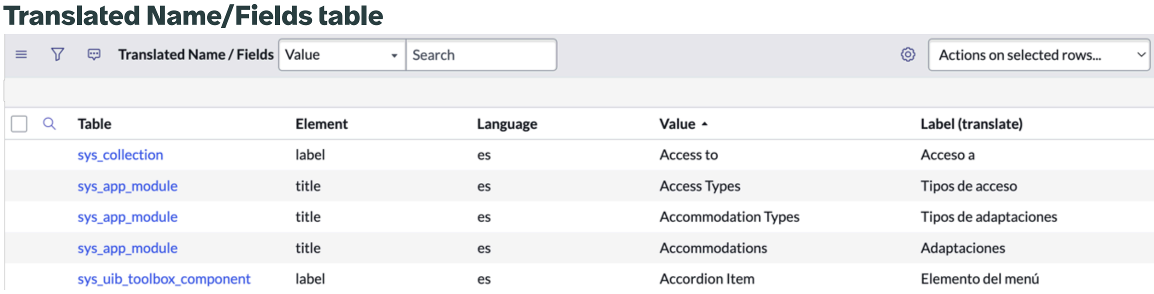
Translated Name / Fields | Value | Search | Actions on selected rows...

Table	Element	Language	Value ▲	Label (translate)
sys_collection	label	es	Access to	Acceso a
sys_app_module	title	es	Access Types	Tipos de acceso
sys_app_module	title	es	Accommodation Types	Tipos de adaptaciones
sys_app_module	title	es	Accommodations	Adaptaciones
sys_uib_toolbox_component	label	es	Accordion Item	Elemento del menú

Field Label table

The Field Label [sys_documentation] table stores the text of table names along with the singular and plural labels for each field in the table.

For each table name and field label, the Field Label table contains a record for each installed language. ServiceNow uses the table and field names from this table to display lists and forms in the proper language. The main fields for this table are:

- Table: name of the table this translation applies to.
- Element: name of the field this translation applies to.
- Language: BCP 47 code of the language for the translated text.
- Label: translated text that users see.
- Plural: plural of the label.
- Help: reserved for future use. A question mark (?) icon appears next to hint labels when accessibility mode is enabled.
- Hint: text that pops up when the cursor rests on the field.
- URL: URL for a web page that provides information about the field. When a URL is provided, the text appears blue and underlined to indicate the label is clickable.

Field label


Field Labels Element Search

Table	Element	Language	Label	Plural
incident	close_code	en	Resolution code	Resolution codes
incident	close_code	es	Código de resolución	Códigos de resolución
incident_fact_table	close_count	en	Close count	Close counts
incident_fact_table	close_count	es	Número de cierre	Número de cierres

Translated field

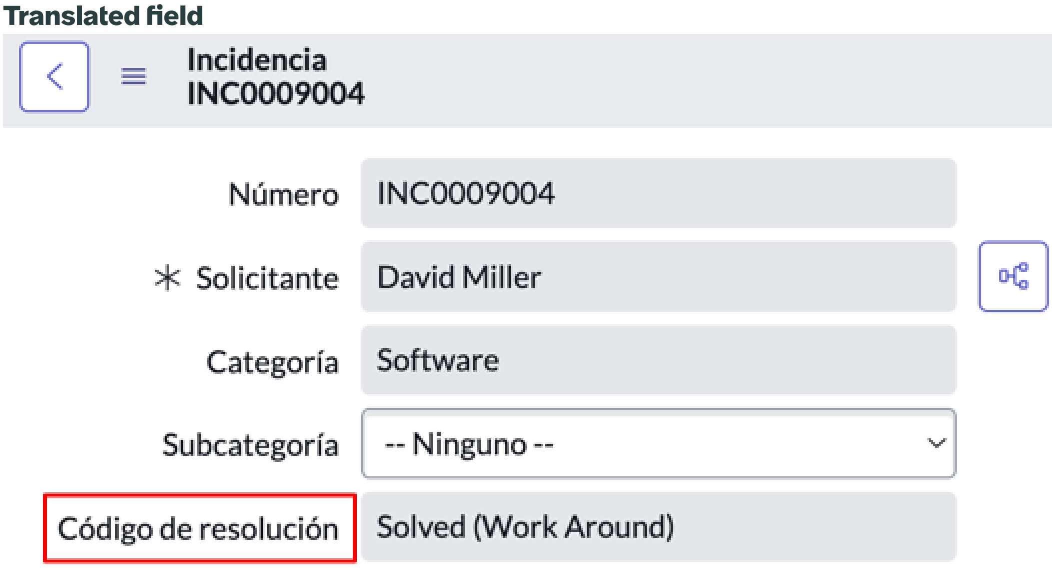

Translated text table

The Translated Text [sys_translated_text] table stores translations for fields with the field type translated_text or translated_html (see the dictionary entry).

This field type is typically used for long text fields, up to 65,000 characters in length, such as survey name and introduction. The main fields for this table are:

- Document: internal identifier of the record this translation applies to.
- Field name: field this translated text appears in, for example, Close notes.
- Language: language the text is translated into.

- Table Name: table this translation applies to.
- Value: translated text that the user sees.

Translated text list

Related topics

Translating text fields

Message table

The Message [sys_ui_message] table contains the translations for informational messages, confirmation messages, error messages, and other types of system messages across the platform.

The Message table contains custom translatable text from the following sources when these texts are wrapped correctly:

- HTML in a UI page in a customer-created application
- Service Portal widget messages (including in surveys)

ServiceNow checks this table for translated text when a client script contains a getMessage call or a server script contains a `gs.getMessage` call. The main fields for this table are:

- Key: internal, case-insensitive, unique identifier of this message.
- Language: language the message is translated into.
- Message: translated text that users see.

Note:
If you create a record in the Message [sys_ui_message] table that uses an existing key and the same language, the message for existing references using the key can change. For example, if a record exists with "Refresh" for the key, English for the language, and "Refresh" for the message, and you create another record with the same key and language but a different message of "Refresh the screen", existing UI strings using the "Refresh" key will display the new "Refresh the screen" message. Records in this table are created in the global scope and aren’t scope-restricted.

For more information about translatable custom content in the Message table, see Translating custom content.

Message list

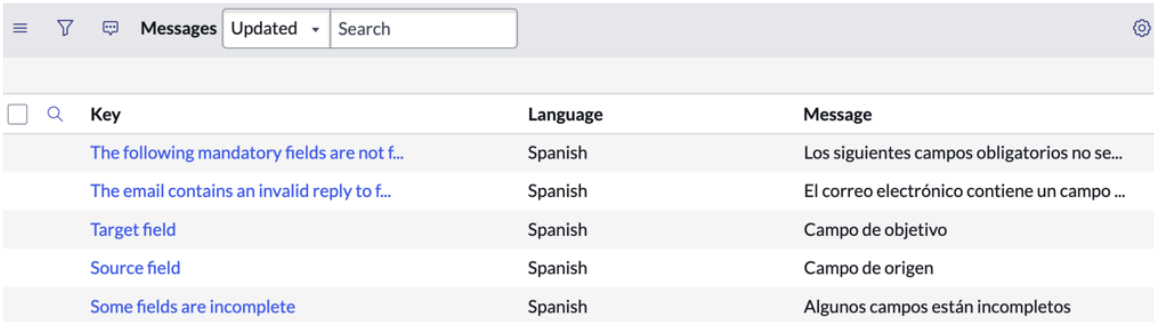

Key	Language	Message
The following mandatory fields are not f...	Spanish	Los siguientes campos obligatorios no se...
The email contains an invalid reply to f...	Spanish	El correo electrónico contiene un campo ...
Target field	Spanish	Campo de objetivo
Source field	Spanish	Campo de origen
Some fields are incomplete	Spanish	Algunos campos están incompletos

Translated message

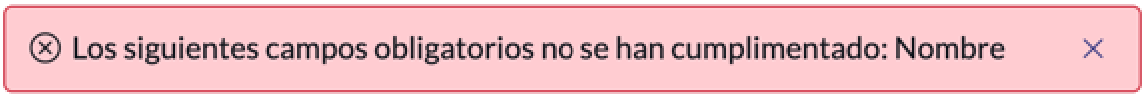

Dynamic Translation

You can dynamically translate text entered in an application, such as a user-entered text on forms, for a seamless localization experience.

An overview of Dynamic Translation

Dynamic Translation enables you to configure localization support in real-time across your instance. Once configured, your agents will be able to translate form fields, Agent Chat messages, and knowledge articles quickly and efficiently.

Get started

Explore

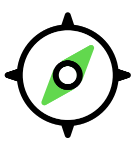

Explore Dynamic Translation

Configure

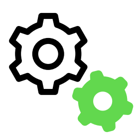

Configure Dynamic Translation

Integrate

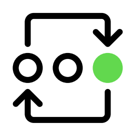

Integrate with third-party translation services

Use

Use Dynamic Translation across your instance

Reference

Reference for Dynamic Translation properties, actions, and error messages

Troubleshoot and get help

- Localization in the ServiceNow Customer Success Center
- Read articles on the Now Community site
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Dynamic Translation

Learn more about Dynamic Translation, including enabling live translation of form fields, translating knowledge articles, and integrating with third-party translation services.

Dynamic Translation overview

When you activate the Dynamic Translation plugin (com.glide.dynamic_translation), you make machine translation possible across your instance, including forms, activity streams, and chats. You can use a machine translation service provider of your choice. You can configure which tables and features enable dynamic translation for your users.

Dynamic Translation workflow

1. Procure any ServiceNow, Inc. subscription that is Professional or above, and includes Dynamic Translation.
2. Activate your choice of language or languages. For more information, see Activate a language. The language plugin also activates Localization Framework.
3. Activate the Dynamic Translation plugin (com.glide.dynamic_translation). For information, see Activate Dynamic Translation.
4. Sign up and create an account with a translation service provider. By default, the following translation service providers are supported:
    - Microsoft. To sign up, see the Microsoft documentation. For more information, see Microsoft Azure Translator Service spoke.
    - Google. To sign up, see the Google website. For more information, see Google Cloud Translator Service Spoke and Set up Google Cloud Translator Service spoke.

      Note:
      The Google Cloud Translator Service spoke is installed from the ServiceNow Store.

   You can also configure any translation service provider of your choice. For information about custom configurations, see Integrate with a translation service provider.
5. Configure your translation service provider in Dynamic Translation.
    - After you have signed up with a translation service provider, do the following in your instance: create a connection, specify the credential information, then activate the translator configuration.
    - Set at least one translator configuration as the default for detection, to detect the language of the user-entered text.
    - Set at least one translator configuration as the default for translation request calls.

- Dynamic Translation uses the ServiceNow® language codes and supports all the languages that the translation service providers support.

  Note:
  For languages where there is a difference in language codes, see Create a language code mapping to create a language code mapping record and associate it with the translator configuration record.

6. Configure your use cases as needed.
   - To enable dynamic translation in a form field, configure the dictionary attribute. For more information, see Enable dynamic translation for a field.
   - To enable dynamic translation in an activity stream, include the table name in the allow list property. For more information, see Dynamic translation on forms and activity streams.
   - For other use cases, see the links in the Features column of the following Dynamic Translation benefits table.

Dynamic Translation benefits

Benefit	Feature	Users
Live translation on forms	Dynamic translation on forms and activity streams	Agent
Knowledge article translation	Together with Localization Framework, enable the translation of Knowledge Base articles.	Agent
Set your preferred translation services	Dynamic Translation spoke and integrating with third-party services	Administrator
Enable live translation for multiple languages	DynamicTranslation API	Administrator
Translate text in conversations between end users and agents	Dynamic Translation for Agent Chat	End user, agent
Support multi-language interactions in certain Now Assist applications	Enable Dynamic Translation for Now Assist applications and Enable Dynamic Translation for capabilities in Generative AI Controller	Agent

Configuring Dynamic Translation

Activate and configure Dynamic Translation to integrate with a third-party machine translation provider to enable live translation of user-generated content across the ServiceNow AI Platform.

Configuration overview

Activate dynamic translation and integrate with your preferred service provider.

Activate Dynamic Translation

Activate the Dynamic Translation plugin (com.glide.dynamic_translation) to enable all the required APIs.

Before you begin

You must have the Glide Notification Translation (com.glide.notification.translation) plugin installed and enabled.

Role required: admin

About this task

Dynamic Translation plugin (com.glide.dynamic_translation) activates these related plugins if they are not already active.

Plugins for Dynamic Translation plugin (com.glide.dynamic_translation)

Plugin	Description
Flow Designer - Installer [com.glide.hub]	Supports the Flow Designer application.
ServiceNow IntegrationHub Runtime [com.glide.hub.integration.runtime]	Enables execution of IntegrationHub actions and flows.
ServiceNow IntegrationHub Action Step - REST [com.glide.hub.action_step.rest]	Action Step - REST
Microsoft Azure Translator Service Spoke [com.glide.microsoft_translation_spoke]	Provides subflows and actions to dynamically translate the user-entered text, and to detect the language of the text using Microsoft translation services.
IBM Watson Translator Service Spoke [com.glide.ibm_translation_spoke]	Provides subflows and actions to dynamically translate the user-entered text, and to detect the language of the text using IBM translation services.
ServiceNow Language Detection Service Spoke [com.glide.language_detection_spoke]	Uses ServiceNow® Language Detection service to provide the ability to detect the language of given text.
Dynamic Translation Spoke [com.glide.dynamic_translation.spoke]	Provides actions that enable the ability to translate text(s) from one language to other language(s) and to detect the language of given text.
Translation Commons [com.glide.translation.commons]	Common utilities for translation features.

Note:

- If you activate the Microsoft Azure Translator Service Spoke plugin when the Dynamic Translation plugin (com.glide.dynamic_translation) is inactive, the subflows and actions of the Microsoft Azure Translator Service Spoke plugin are enabled. These subflows can be used in other flows and subflows. However, you cannot call Dynamic Translation APIs.
- If you activate the IBM Watson Translator Service Spoke plugin when the Dynamic Translation plugin is inactive, the subflows and actions of the IBM Watson Translator Service Spoke plugin are enabled. These subflows can be used in other flows and subflows. However, you cannot call Dynamic Translation APIs.

  Important:
  IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

Procedure

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

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

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

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

Related topics

List of plugins (Washington DC)

Request for domain separation in Dynamic Translation

Enable the support for the Dynamic Translation framework in domain-separated instances.

Before you begin

Role required: admin

About this task

There are two ways to request a plugin.

- Access the Now Support Service Catalog directly by selecting **All** > **Service Catalog** > **Activate Plugin** on Now Support.
- Access the Now Support Service Catalog through the All Applications page on your instance by following these steps.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. On the All Applications page, select **Request Plugin** to open the **Activate Plugin** form on Now Support.

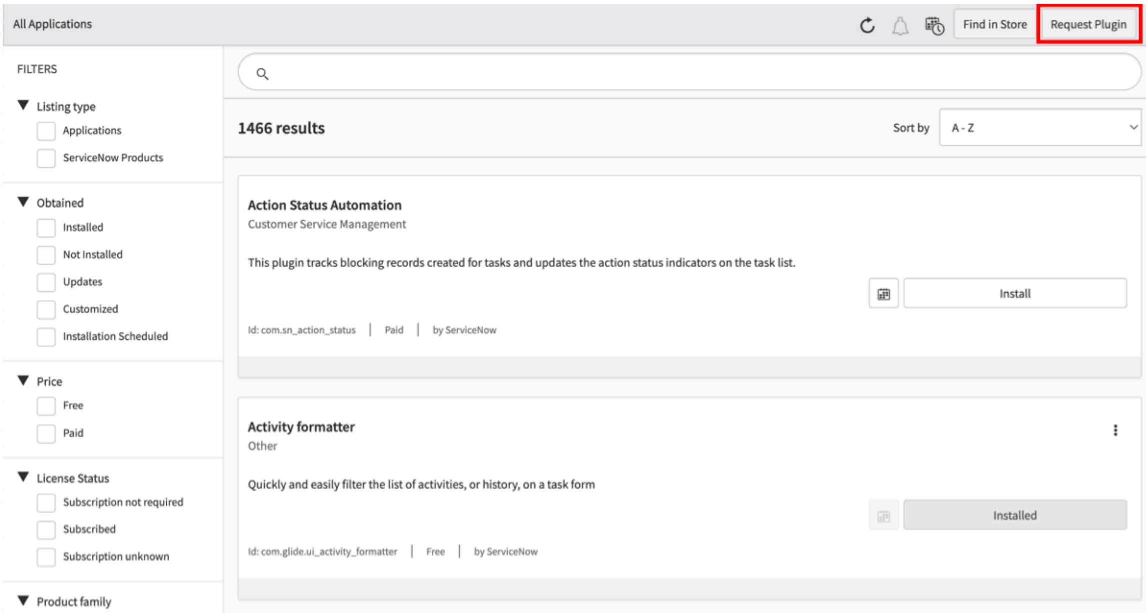

3. On Now Support, select the link to access the Now Support Service Portal Service Catalog.

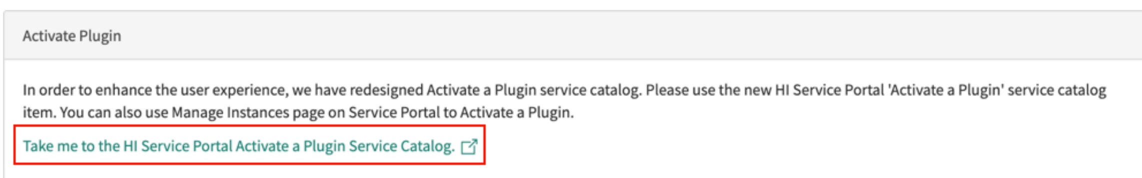

4. Select your instance.
5. Select **Actions > Activate Plugin**.
6. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Instance on which to activate the plugin.
Which plugin would you like to activate	Name of the plugin to activate. **Note:** If the system does not list the plugin you want or if you are activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.

Field	Description
Select Maintenance Date and Time	The date and time to activate the plugin. **Note:** Plugins are activated in two batches, once in the morning and once in the evening, on every business day in the US Pacific time zone. If the plugin must be activated at a specific time, enter the request in the **Reason/Comments** field.

Example

For example, see the following form to activate the CSM Workspace plugin on an instance named My Instance.

Activate Plugin form

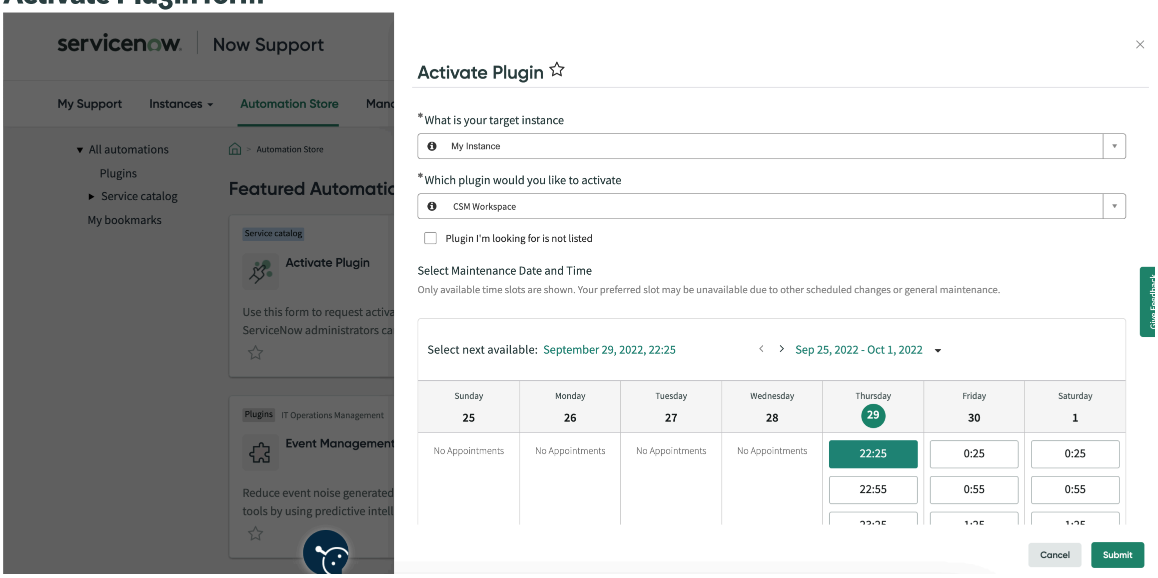

7. Select **Submit**.
   For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

What to do next

For information about configuring domain separation in Dynamic Translation, see Domain separation and Dynamic Translation.

Related topics

List of plugins (Washington DC)

Migrate to version v3 of a translator configuration

Implement the contract with subflows using complex objects in the input and output fields for dynamic translation of text.

Before you begin

Role required: admin

For a translator configuration that you added to integrate with a translation service provider of your choice, implement Integration Hub subflows with complex objects in the input and output fields. These subflows are required to interact with the translation service provider for the following tasks:

- Translate the input text
- Detect the language of the input text

For information on implementing these subflows, see Integrate with a translation service provider.

About this task

After an upgrade to Washington DC, the version of all translator configurations that are available by default but not customized is automatically updated to v3. The corresponding subflows have complex objects in the input and output fields.

However, you should migrate to version v3 for the following translator configurations:

- Translator configurations that are available by default but customized in New York
- Translator configuration that you added to integrate with a translation service provider of your choice

Procedure

1. Navigate to **All** > **Dynamic Translation** > **Translator Configurations**.
2. Select a translator configuration.
3. In the Translator Configuration form, update these fields.

Translator Configuration form fields

Field	Description
Active	Activates the translator configuration to integrate with Dynamic Translation.
Version	Version of the contract used in subflows. Select **v3**.
Choose a translate subflow	Integration Hub subflow with complex objects in the input and output fields to translate the input text. ◦ For the translator configurations that are available by default but customized, select Translate Text [translate_text_v3]. ◦ For the translator configuration that you added to integrate with a translation service provider of your choice, select the subflow that you implemented with complex objects. The contract now supports multiple texts in a single method call.
Choose a detect subflow	Integration Hub subflow with complex objects in the input and output fields to detect the language of the input text. ◦ For the translator configurations that are available by default but customized, select Detect Language [detect_language_v3]. ◦ For the translator configuration that you added to integrate with a translation service provider of your choice, select the subflow that you implemented with complex objects. The contract now supports multiple texts in a single method call.

Field	Description
Mark as default for translation	Marks the translator as the default configuration for dynamic translation of text.
Mark as default for detection	Marks the translator as the default configuration to detect language of the input text.

4. Click **Update**.

Integration with other translation services

Integrate third-party translator services to enable machine translation of user-generated text using the provider of your choice.

After you have activated Dynamic Translation, you can integrate your preferred translator service provider directly with your instance.

Translator services are listed in the Translator Configurations [sn_dt_translator_configuration] table. In this table, several third-party translators are partially preconfigured by default. You can complete the configuration of a default translator by providing your account credentials and connections. Alternatively, you can also create a custom spoke.

Note:
If you activate dynamic translation in the Generative AI Controller, a fully preconfigured spoke to Microsoft Azure OEM Translator Service is added as a row in Translator Configurations.

Also listed in the Translator Configurations table is a default record for the **ServiceNow** translator. This translator can be used for language detection only, not for translation, so its **Mark as default for translation** check box is not selectable. This optional language detection service is included in your Dynamic Translation subscription.

Integrate with a translation service provider

Integrate with a third-party translation provider of your choice for machine translation of user-generated text.

Before you begin

Role required: admin

Procedure

1. Create subflows to detect the language of the input text, and translate the input text by using Integration Hub subflows, actions, and connections and credentials.
   For more information, see Flow Designer.

   a. Map error codes from your machine translation provider to the error codes and standardized error messages supported by ServiceNow.

      Standard error codes and message

Error code	Error message
40050	Error code to populate error messages from the machine translation provider instead of standardized error messages.

Error code	Error message
40051	Unknown error occurred.
40052	Text ('text' field) has exceeded its maximum length.
40053	Source language is invalid.
40054	Target language is invalid.
40055	Request is not authorized because credentials are missing or invalid.
40056	Text cannot be translated to target languages.

b. Ensure that the inputs and output fields are in the following format for subflows.

Note:
While you create subflows, you can refer to the subflows that are available by default if required. For example, Detect Language [detect_language_v3] and Translate Text To Multiple Languages [translate_text_to_multiple_languages_v3].

Detect subflow

Field name	Description	Type
Inputs		
texts	Input texts whose language should be detected. In the Washington DC release, version v3 supports bulk detection- detection of multiple texts input in a single string.	Array.String
Outputs		
detections	Information about the code and confidence of the languages detected by the translator and their respective alternatives, error codes, and error messages. ▪ Code: Language code of the detected language as specified by the translation service provider. ▪ Confidence: Value that indicates the confidence in the result. The value is between zero and one. A low value indicates a low confidence. For example, French text with confidence of 0.93 `{ 'code' : 'fr',  'confidence' : '0.93'  }` For more information refer to the getDetectedLanguages section of DynamicTranslation - Scoped.	Array.Object. Each element of the Object is a String value.
status	Status of the request.	Choice

Translate subflow

Field name	Description	Type
Inputs		
texts	Texts to be translated. In the Washington DC release, version v3 supports bulk translation- translation of multiple texts input in a single string.	Array.String
target_languages	Array of language codes to which the text needs to be translated. For example, `['fr','it']`	Array String. Each element of the array is a String value.
source_language	Language code of the source text. For example, `en`	String
additional_parameters	Array of additional parameters. For example, `[{'parameter_name': 'textType', 'parameter_value': 'html'}]`	Array.Object. Each element of the array is an Object with String values.
Outputs		
translations	Translated texts in the specified target languages. Each element of the array is an Object containing target language and the translated text. For example, `[ {'target_language':'fr', 'translated_text':'Bonjour'}, {'target_language':'it', 'translated_text':'ciao'} ]` For more information refer to the getTranslations section of DynamicTranslation - Scoped.	Array.Object. Each element of the array is an Object with String values.
status	Status of the request.	Choice

Note:
Ensure that the **Run As** field is set to **User who initiates session** in the flow properties.

2. Create a translator configuration and activate it.
   For more information, see Create a custom translator configuration.

Create a custom translator configuration

Configure Dynamic Translation for a third-party machine translation provider of your choice.

Before you begin

Ensure that the Dynamic Translation plugin (com.glide.dynamic_translation) is activated.

Create subflows for your custom translation provider. See Integrate with a translation service provider.

Create language code mappings as needed. For more information see Language Code Mapping in Dynamic Translation.

Role required: admin

About this task

By default, Google and Microsoft translator configurations are available and ready for you to add your account details. Also by default the ServiceNow translator configuration, for language detection only, is available and active.

However, you can also configure a custom translator with the following procedure.

Procedure

1. Complete the prerequisites and confirm your application scope, then navigate to **All** > **Dynamic Translation** > **Translator Configurations**.
2. Select **New**.
3. On the form, fill in the fields.

Translator Configuration form fields

Field	Description
Name	Identifier of the translator configuration.
Active	Option to activate the translator configuration.
Version	Version of the contract used in subflows. Version v3 is used in releases before Xanadu Patch 3. Version v4 is used in releases from Xanadu Patch 3.
Preferences	
Choose a translate subflow	Workflow Studio subflow that interacts with the translation service provider to translate the input text. Search for the subflow that you created in the previous procedure (Integrate with a translation service provider).
Mark as default for translation	Option to mark the translator configuration as the default configuration for translation. When no translator configuration is specified while translating the text, the default configuration is used. ◦ The configuration must be active to be marked as default. ◦ You can select only one configuration as default for translation.

Field	Description
Choose a detect subflow	Workflow Studio subflow that interacts with the machine translation provider to detect the language of the input text. Search for the subflow that you created in the previous procedure (Integrate with a translation service provider).
Mark as default for detection	Option to mark the translator configuration as the default configuration for detection (determining the language of the submitted text). When no translator configuration is specified while detecting the text, the default configuration is used. ◦ The configuration must be active to be marked as default. ◦ You can select only one configuration as default for detection.
Language Code Mappings	
Language code mappings	Mapping codes for languages to associate with the translator configuration record. Associate with the translator configuration record by selecting the magnifying glass icon ( ), then selecting the language code mapping. ◦ Only languages for which the language code mapping was created are available for selection. ◦ You can select only one language code mapping per language. For more information, see Create a language code mapping.

4. Select **Submit**.

What to do next

If you are on Xanadu Patch 3 or above, and want to use Exclusion Framework in Dynamic Translation, add your translation provider's exclusion tags. For more information, see .

Google Cloud Translator Service Spoke

Provides subflows and actions to dynamically translate the user-entered text, and to detect the language of the text using the Google translation service.

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.

Subscription

This spoke requires one of the following subscriptions:

- Any ServiceNow subscription that is Professional or above, and includes Dynamic Translation.
- Integration Hub subscription. For more information, see Legal schedules - IntegrationHub overview.

Supported versions

Google Translate API version v3.

Scope of the spoke in the Dynamic Translation framework

To use Google Cloud Translator Service spoke in the Dynamic Translation framework, ensure that the spoke is installed in ServiceNow Store and the Dynamic Translation plugin is active. For more information, see Dynamic Translation.

Spoke flows

This spoke has no sample flows.

Spoke subflows

The Google Cloud Translator Service spoke provides subflows in the Published state to integrate with the Google translation service. These subflows can be used as part of other subflows and flows.

Subflow	Description
Detect Language	Contains a set of inputs, actions, and outputs to detect the language of the input text.
Translate Text To Multiple Languages	Contains a set of inputs, actions, and outputs to translate the input text to multiple languages.

Spoke actions

The Google Cloud Translator Service spoke provides actions to integrate with the Google translation service. Available actions include:

Action	Description
Detect Language	Detects the language of the input text.
Translate Text To Single Language	Translates the input text to a single language.
Translate Text To Multiple Languages	Translates the input text to multiple languages.

Google account requirements

The Google Cloud Translator Service spoke requires configuring your Google Cloud account to generate an OAuth 2.0 JWT Bearer Grant token.

Connection and credential alias requirements

Integration Hub uses aliases to manage connection and credential information, and OAuth credentials. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you don't need to update any actions that use the connection.

This spoke uses the GoogleTranslation alias record (sn_google_trans.GoogleTranslation) to authorize actions. For information about setting up the spoke, see Set up Google Cloud Translator Service spoke.

Set up Google Cloud Translator Service spoke

Integrate your Google account with your ServiceNow instance.

Before you begin

- Role required: admin
- Create a service account on Google Cloud, generate a service account key of JSON type and download it. For information on creating the service account key, see the Google documentation.
- Enable the Cloud Translation API service. For information on enabling a service, see the Google documentation.

Create a Java KeyStore certificate

Encrypt the security certificates obtained from Google by creating a Java KeyStore (JKS) file.

Before you begin

- Role required: admin
- Download a JSON file containing the service account key from Google. For information on downloading the JSON file, see the Google documentation. Following is a sample JSON file:


```
{
  "type": "service_account",
  "project_id": "primeval-nectar-242610",
  "private_key_id": "0c6c7b1511f1c236c1300c5933d528bbf45314e5",
  "private_key": "-----BEGIN PRIVATE
 KEY-----\nMIIEvQKBADANBgkqhkiG9w0BAQEFAASCBKcwggSjAgEAAoIBAQDg
97zOxfOenyu2\niXC8ArQHaAgOiy27sjG3jJXohUdHkawuX+nQLWZkIbZYlZX+c
MjNfpUFpzeHRkPm\n5OQ2bjHCgCTVFCbToKIH+VkMpS7LO+TFKW2TkWJHzWO2Xv
3ObT+2HkSb4tTZisMg\nV/sI6dGJaOko4yTIBwZc7vlr28nGJdw945+9NVkE
Pm/G8drYJRTlrImyfnnDORGR\npyKctHAC3wZ/rCE5b/VIHK5ZKspRT/qXVwZsH
5pW7a8NDrB1z+YQIGGOjlfINxBE\nQCE5GhiUF+XscfGdq3OculFoILO9xAk/41
yQb8ziDr4Ru5H1UFOljp4t6fLw3k2o\noDsr+R
+XAgMBAAECggEAA03QIY9JsSowR6mEE+YsQ7GU8LW9kbSfwOzWxMfOUIpE\np5e
OBOEtOEmodmuX/NkmMdJqmN8oUx3GkIULDvuWUn9OSwbVPSVuS8SvOJ7SbZyv\n
EvA1UkX/1gMftEV34FecaG3QXyO5MYq2q+pu3uYkNCrsxbNOTlXAE6xUOG81a
uoL\nYIl7pzVfcxDW5Of5VI5AuMc8+OrgadvOPu3JThjjauMvnDPKi5sbKDsDCf
TqqDys\nHMny8+uhfCOr2JtTgpc4HUgqGKpALAUQvRZYmmlKqfBNgXo+dDBKJsd
7pD9NnxQn\nM7l8wqQOkS2kfcJsHNcV4ZUORabz1q1MFlVdxxeijQKBgQD5r
7J/MgN4LLekaimE\n3T27KTA12aikQlOLORQHLbHZRexVD41WoRrQySSzUvQKXS
iwZwRQLO/hptDRnUyh\n6WVQLwtSiNvItLu+8Ak2mb3F+NDEFb/R1FPdTR3OG7w
XEeau1+cm8+8N63ehFAW4\n7IXTyDhtO5c3deXN7SDrRFjrnQKBgQDmqAdRkKPw
au6q+zSGPxqcFqlXS+ghpQkT\nUPxe45aZdRuFUOUCWA9BIbMnlVHZUDnQpiIbf
BJ9Tcs2/vBwLt2QjziZPbnIlmkp\nqIF4t5edar8Z6vnOlWmeupAmlFqKN7Goo
pUD+RIto68BfSKIPczfGLtu3PxLEmgj\nRYzY5HUTwwKBgCZO5zssCtjBmm9aYp
ayNMXU7DX/FjhmeEo4Olt4sEHUwTfAs3Y6\nThUGRf7QsgG+o3u4AjQPF8tblCI
U5i6x8gbNmCLYLXHWVGxuMBOWxOHvFsh8yRSa\nbWhSbmCgvPGYsjOPx+x5+cHd
GInYuaDn3RznY7l/SiUipYh4E2/pEQEJAoGAdu1G\nMMESNkD8ZD5324wn7TkmA
TPLMaXFYydLqKVSHjeqg/eszKOY4eO9UXiFJjZeSP2P\n8nnrkp4M3INgd4dCiG
nANgsEgq9C887FSvfmfazvca6dSIXNWqE4+Btf/4ot2RRT\nHyRKQiv2bR8XMgY
jXxiCc+wPTank9eLDd4V79D8CgYEA6Eywngil5bVO5ksW9gyp\nOu5eieKMNYnd
Pq5lyw1POyN9nx/iYQ3m2FI7d2K7aBpKOAXkLZQqv5tbqrF2fvdO\nbm4eLISLg
8kdZO9OohAsRuZ6JG2aICw4vTWgQyRWLe4RxdoY9gGKh8N844tUU7uN\nPCs
Wu/rdB8dT8wqerTFo81U=\n-----END PRIVATE KEY-----\n",
  "client_email":
 "service-now-google-trans-v3@primeval-nectar-242610.iam.gservi
ceaccount.com",
  "client_id": "113255944370425109391",
  "auth_uri": "https://accounts.google.com/o/oauth2/auth",
```

```
  "token_uri": "https://oauth2.googleapis.com/token",
  "auth_provider_x509_cert_url":
 "https://www.googleapis.com/oauth2/v1/certs",
  "client_x509_cert_url":
 "https://www.googleapis.com/robot/v1/metadata/x509/service-now
-google-trans-v3%40primeval-nectar-242610.iam.gserviceaccount.c
om"
}
```

Procedure

1. In the JSON file, perform the following for the `private_key` value.

   a. Unescape the `\n` characters.

      Tip:
      To unescape characters in a JSON file, you can parse a JSON string into a JavaScript object.

   b. Store the private key in a file with .pem extension.

2. In the JSON file, perform the following for the `client_x509_cert_url` value.

   a. Open the link specified for `client_x509_cert_url`.

   b. Unescape the `\n` characters for the certificate associated with the private_key_id mentioned in the JSON file.

   c. Store the certificate value in a file with .pem extension.

3. Use the openssl command to create a PKCS 12 file using the recently created private key and certificate files.

   ```
   openssl pkcs12 -export -in [path to certificate] -inkey
    [path to private key] -certfile [path to certificate ] -out
    testkeystore.p12
   ```

   For example, if certificate.pem is the certficate and privatekey.pem is the private key:

   ```
   openssl pkcs12 -export -in certificate.pem -inkey
    privatekey.pem -certfile certificate.pem -out testkeystore.p12
   ```

   A PKCS 12 file, testkeystore.p12, is created.

4. Specify an export password or source keystore password.

   Note:
   You should specify this password when creating a JWT key for Google Cloud Translator Service spoke.

5. Use the keytool command to create a JKS file from the PKCS 12 file.

   ```
   keytool -importkeystore -srckeystore testkeystore.p12
    -srcstoretype pkcs12 -destkeystore wso2carbon.jks
    -deststoretype JKS
   ```

   Note:
   testKeyStore.p12 is the PKCS 12 file and wso2carbon.jks is the JKS file.

6. Specify a destination keystore password.

Note:
You should specify this password when attaching a JKS certificate to Google Cloud Translator Service spoke.

Attach a Java KeyStore certificate to Google Cloud Translator Service spoke

Enable the JWT client authentication by attaching a valid Java KeyStore (JKS) certificate to Google Cloud Translator Service spoke.

Before you begin

- Role required: admin
- Valid Java KeyStore certificate

Procedure

1. Navigate to **All** > **System Definition** > **Certificates**.
2. Click **New**.
3. On the form, fill in the fields.

X.509 Certificate form fields

Field	Description
Name	Identifier of the certificate.
Notify on expiration	Users to be notified when the certificate expires.
Warn in days to expire	Number of days to send a notification before the certificate expires.
Active	Option to activate the certificate.
Type	Type of the certificate. Select **Java Key Store**.
Expires in days	Number of days until the certificate expires.
Key store password	Password to access the certificate. Use the destination keystore password specified when creating the JKS certificate. For more information on this password, see Create a Java KeyStore certificate.
Short description	Summary about the certificate.

4. Click the manage attachments icon ( ) and attach a JKS certificate.
5. To validate the JKS certificate, click **Validate Stores/Certificates**.
6. Click **Submit**.

Create a JWT signing key for Google Cloud Translator Service spoke

Assign a JSON Web Token (JWT) signing key to your Java KeyStore certificate.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **JWT Keys**.
2. Click **New**.
3. On the form, fill in the fields.

JWT Keys form fields

Field	Description
Name	Identifier of the JWT signing key.
Signing Keystore	Valid JKS certificate for which you want to assign the key.
Key Id	Key ID to identify which key is used when multiple keys are used to sign tokens.
Signing Algorithm	Algorithm to sign with the key.
Signing Key Password	Password associated with the key. Use the export password or the source keystore password specified when creating the JKS certificate. For more information on this password, see Create a Java KeyStore certificate.
Active	Option to activate the key.

4. Click **Submit**.

Create a JWT provider for Google Cloud Translator Service spoke

Add a JSON Web Token (JWT) provider to Google Cloud Translator Service spoke.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **JWT Providers**.
2. Click **New**.
3. On the form, fill in the fields.

JWT Provider form fields

Field	Description
Name	Identifier of the JWT provider.
Expiry Interval (sec)	Number of seconds that indicate the lifespan of the JWT provider token. Specify 3600.
Signing Configuration	JWT signing key.

4. Right-click the form header, and click **Save**.
5. In the Standard Claims related list, enter the values for these claims.

Claim

Claim name	Claim value
aud	https://www.googleapis.com/oauth2/v4/token
iss	*`client_email`* value from the JSON file.

6. In the Custom Claims related list, create the **scope** claim and enter its value as `https://www.googleapis.com/auth/cloud-translation`.
7. Click **Update**.

Configure the credential for the GoogleTranslation alias

Authorize actions of Google Cloud Translator Service spoke by configuring the Google OAuth 2.0 credential.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select the Google OAuth 2.0 credential.
3. Open the record specified in the **OAuth Entity Profile** field.
4. In the **JWT Provider** field, specify the JWT provider that you want to use.
5. Click **Update**.
6. To verify if an OAuth Access token is generated to connect to Google's translation services, click the **Get OAuth Token** related link of the Google OAuth 2.0 credential.

Configure the connection attributes for the GoogleTranslation alias

Connect to the Google's translation service by configuring the Google connection. Provide information that is used by HTTP(s) actions or activities to connect to that service.

Before you begin

Role required: connection_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections**.
2. Open for the record for **Google**.
3. In the Attributes related list of the HTTP(S) Connection form, fill in the fields.

HTTP(S) Connection form fields

Field	Value required
location	Location of the customer. If the value is not specified when creating an account on Google Cloud, then specify `global`.
project_id	Identifier of a project. Specify the *`project_id`* value from the JSON file.
version	API version that the related spokes are built for. The default value is `v3beta1`.

Field	Value required
	Because Google now supports the v3 version, you can override this value as v3.

4. Right-click the form header and click **Save**.

Microsoft Azure Translator Service spoke

Microsoft Azure Translator Service spoke provides subflows and actions to dynamically translate the user-entered texts, and to detect the language of the user-entered texts using the Microsoft translation service. Dynamic Translation supports detection and translation of single text and multiple texts.

https://player.vimeo.com/video/1019025021?h=e0b4f4b24b&badge=0&autopause=0&player_id=0&app_id=58479

Activation

Activate the Dynamic Translation plugin (com.glide.dynamic_translation) or the Microsoft Azure Translator Service Spoke plugin (com.glide.microsoft_translation_spoke).

Subscription

This spoke requires any ServiceNow subscription that is Professional or above, and includes Dynamic Translation.

Supported versions

API version v3.

Spoke subflows

The spoke provides subflows in the Published state to integrate with Microsoft translation service. These subflows can be used as part of other subflows and flows.

Spoke subflows

Subflow	Description
Detect Language [detect_language_v3]	Contains a set of inputs, actions, and output to detect the language of the input text.
Translate Text [translate_text_v3]	Contains a set of inputs, actions, and outputs to translate the input text to multiple languages.

Spoke actions

The spoke provides actions to integrate with Microsoft translation services.

Spoke actions

Action	Description
Detect Language [detect_language_v3]	Detects the language of the input text.
Translate Text [translate_text_v3]	Translates the input text to multiple languages.

Microsoft account requirements

Ensure that you are registered with Microsoft and obtain the credential information. For more information, see Microsoft documentation.

For information on Microsoft Translator Text API, see the Microsoft documentation.

Connection and credential alias requirements

IntegrationHub uses aliases to manage connection and credential information. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you do not need to update any actions that use the connection. For more information, see Credentials and connection information.

To use the spoke connection aliases, create an associated connection record and a credential record for each alias.

This spoke uses the MicrosoftTranslation alias (sn_ms_trans_spoke.MicrosoftTranslation) to authorize actions.

The following are the connection alias requirements.

- Connection type: HTTP
- Connection URL: For example, `https://api.cognitive.microsofttranslator.com`.

> **Note:**
> The URL provided here is a sample only. Microsoft translation customers may need to use custom API URLs or virtual network type API URLs, even when they do not subscribe to a virtual network. Microsoft customers who use multi-service or regional translation services also need to know their Ocp-Apim-Subscription-Region attribute, and their API URL may contain a regional reference. Check your Azure portal for this information. Refer to the Microsoft documentation for further details.

Perform the following tasks to connect to the Microsoft translation services and use them to dynamically translate your text.

Create a credential for the MicrosoftTranslation alias

Authorize actions of Microsoft Azure Translator Service spoke by creating a credential for the MicrosoftTranslation connection and credential alias.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Click **New**.
3. Select **API Key Credentials**.
4. In the API key Credentials form, fill the fields.

API key Credentials form fields

Field	Description
Name	Identifier of the credential.
Active	Option to activate the credential for use.
Applies to	MID servers for which the credentials are applicable. Possible options are: ◦ All MID servers in your network. ◦ Specific MID servers. In this case specify the servers in the **MID servers** field.
MID servers	Specific MID servers for which the credentials are applicable. This field is displayed only when you select **Specific MID servers** for the **Applies to** field.
Order	Sequence in which the credential is applied.
API Key	Unique code that is passed to an API to identify the connection.
Credential alias	Alias to use the credential.

5. Select **Submit**.

Create a connection for the MicrosoftTranslation alias

Connect to the Microsoft translation service by creating a connection. Provide information that is used by HTTP(s) actions or activities to connect to that service.

Before you begin

Ensure that you are in the application scope that you want for your connection. The default application scope is **Microsoft Azure Translator Service Spoke**.

Role required: connection_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
2. Select the MicrosoftTranslation alias (sn_ms_trans_spoke.MicrosoftTranslation).
3. In the **Connections** tab, click **New**.
4. In the HTTP(S) Connection form, fill the fields.

HTTP(S) Connection form fields

Field	Description
Name	Name of the HTTP(s) connection.
Credential	Credential record used to authorize the connection.
Connection alias	Alias record to associate with the connection. Using an alias enables you to update the connection record without having to reconfigure any actions or activities that use the alias.
Active	Option to activate the connection.
Domain	Domain where the connection is applicable.

Field	Description
URL builder	Option to enable the system to build the connection URL. Clear this check box.
Connection URL	Base URL to connect to the Microsoft's translation service. For example, `https://api.cognitive.microsofttranslator.com`. **Note:** This URL is a sample only. Microsoft translation customers may need to use custom API URLs or virtual network type API URLs, even when they do not subscribe to a virtual network. Check your Azure portal to obtain your API URL. For further information, refer to Microsoft documentation.
Use MID server	Option to use a MID Server for this connection.
Connection timeout	Timeout value in milliseconds. Overrides the global property for this usage.
Attributes	
Version	API version that the related spokes are built for. The default value is v3. Do not edit this value.
Ocp-Apim-Subscription-Region	Value of the regional translator resource. This value is optional when using a global translator resource. **Note:** This value is mandatory for customers using Microsoft's multi-service or regional translator resources. For information on supported regions, see the Microsoft documentation.

5. Click **Submit**.

Activate the Microsoft translator configuration

Make the Microsoft translation service available for use by activating the Microsoft translator configuration.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Dynamic Translation** > **Translator Configurations**.
2. Select **Microsoft**.
3. Select the **Active** check box.
4. In the **Preferences** section, choose a translator as default either for translation or detection, or for both.

Preferences

Field	Description
Mark as default for translation	Option to mark the translator as default for translation.

Field	Description
Mark as default for detection	Option to mark the translator as default for detection.

Note:
Make at least one translator configuration as the default configuration either for translation to enable the Dynamic Translation experience on forms and activity streams.

5. Select **Update**.

IBM Watson Language Translator Service spoke

The IBM Watson Language Translator Service Spoke has been deprecated by IBM. This information is provided for legacy context. Customers with existing spoke installations can continue to use the translator service until it is withdrawn by IBM.

Integrate Dynamic Translation with IBM translation service.

Important:
IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

Activation

Activate the Dynamic Translation plugin (com.glide.dynamic_translation) or the IBM Watson Language Translator Service Spoke plugin (com.glide.ibm_translation_spoke).

Subscription

This spoke requires any ServiceNow subscription that is Professional or above, and includes Dynamic Translation.

Supported versions

API version v3.

Spoke subflows

The spoke provides subflows in the Published state to integrate with the IBM Watson Language translation services. These subflows can be used as part of other subflows and flows.

Spoke subflows

Subflow	Description
Detect Language [detect_language_v3]	Contains a set of inputs, actions, and output to detect the language of the input text.
Translate Text [translate_text_v3]	Contains a set of inputs, actions, and outputs to translate the input text to multiple languages.

Spoke actions

The spoke provides actions to integrate with IBM translation services.

Spoke actions

Action	Description
Detect Language [detect_language_v3]	Detects the language of the input text.
Translate Text [translate_text_v3]	Translates the input text to multiple languages.

IBM account requirements

For information on the IBM Language Translator API and the notice of deprecation, see the IBM documentation.

Connection and credential alias requirements

IntegrationHub uses aliases to manage connection and credential information. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you do not need to update any actions that use the connection. For more information, see Credentials and connection information.

To use the spoke connection aliases, create an associated connection record and a credential record for each alias.

This spoke uses the IBMTranslation alias (sn_ibm_trans_spoke.IBMTranslation) to authorize actions.

The following are the connection alias requirements.

- Connection type: HTTP
- Connection URL: Base URL provided by IBM to connect to the language translator service. Users can have a different URL based on their geographical preference.

Perform the following tasks to connect to the IBM Watson Language translation services and use them to translate the dynamically generated text.

Create a credential for the IBMTranslation alias

Authorize actions of the IBM Watson Language Translator Service spoke by creating a credential for the IBM connection and credential alias.

Before you begin
Role required: admin

About this task

Important:
IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Click **New**.
3. Select **Basic Auth Credentials**.
4. In the Basic Auth Credentials form, fill the fields.

Basic Auth Credentials form fields

Field	Description
Name	Identifier of the credential.
Order	Sequence in which the credential is applied.
User name	User name to identify the connection. Enter `apikey`.
Password	Unique key provided by IBM to confirm the connection.
Active	Option to activate the credential for use.
Credential alias	Alias to use the credential.

5. Select **Submit**.

Create a connection for the IBMTranslation alias

Connect to the IBM translation service by creating a connection. Provide information that is used by HTTP(s) actions or activities to connect to that service.

Before you begin

Role required: connection_admin

About this task

Important:
IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
2. Select the IBMTranslation alias (sn_ibm_trans_spoke.IBMTranslation).
3. In the **Connections** tab, click **New**.
4. In the HTTP(S) Connection form, fill the fields.

HTTP(S) Connection form fields

Field	Description
Name	Name of the HTTP(s) connection.
Credential	Credential record used to authorize the connection.

Field	Description
Connection alias	Alias record to associate with the connection. Using an alias enables you to update the connection record without having to reconfigure any actions or activities that use the alias.
Active	Option to activate the connection.
Domain	Domain where the connection is applicable.
URL builder	Option to enable the system to build the connection URL. Clear this check box.
Connection URL	Base URL provided by IBM to connect to the language translator service. Users can have a different URL based on their geographical preference.
Use MID server	Option to use a MID Server for this connection. Clear this check box.
Attributes	
Learning Opt-Out	Attribute to specify if your data can be used as test data by IBM to train its models. If set to `true`, IBM does not use your data for any API request. By default, IBM services log requests and their results.
API Version Date	Release date of the API version you want to use. The default value is 2018-05-01. Do not edit this value.
Version	API version that the related spokes are built for. The default value is v3. Do not edit this value.

5. Select **Submit**.

Activate the IBM translator configuration

Make the IBM translation service available for use by activating the IBM translator configuration.

Before you begin

Role required: admin

About this task

> **Important:**
> IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

Procedure

1. Navigate to **All** > **Dynamic Translation** > **Translator Configurations**.
2. Select **IBM**.
3. Select the **Active** check box.
4. In the **Preferences** section, choose a translator as default either for translation or detection, or for both.

Preferences

Field	Description
Mark as default for translation	Option to mark the translator as default for translation.
Mark as default for detection	Option to mark the translator as default for detection.

Note:
Make at least one translator configuration as the default for translation to enable the Dynamic Translation experience on forms and activity streams.

5. Select **Update**.

Translating with Dynamic Translation

Translate user-generated content live on your instance to streamline communication for your agents and end users.

A common use case for Dynamic Translation

After you activate the Dynamic Translation plugin (com.glide.dynamic_translation), you can do the following tasks using a translation service provider:

- Detect the language of the input text
- Translate user-entered text from a source language to a target language

For information on activating this plugin, see Activate Dynamic Translation.

Default translation service providers for Dynamic Translation

By default, Dynamic Translation supports the following Microsoft and Google service providers for machine translation.

- Connections and credential aliases and the corresponding connection attributes.

  Note:
  You should create a connection for each translation service provider and specify the credential information.

- IntegrationHub subflows and actions for these spokes.
  - Microsoft Azure Translator Service spoke
  - Google Cloud Translator Service spoke
- Translator configurations

Note:
The Google Cloud Translator Service spoke for Dynamic Translation must be installed from ServiceNow Store.

Using Dynamic Translation

When your admin has configured Dynamic Translation, look for the translate icon ( ) next to fields on forms and in workspaces.

Dynamic translation on forms and activity streams

For a seamless localization experience on forms and activity streams, you can dynamically translate text based on the preferred language of a user.

Dynamic translation on forms

You can translate the field values on forms on the ServiceNow AI Platform and the configurable workspaces. For more information on Dynamic Translation support for configurable workspaces, see Dynamic Translation support for configurable workspace.

To enable dynamic translation at the field-level on Platform forms, see Enable dynamic translation for a field. After dynamic translation is enabled, select the translate icon ( ) to show or hide the translated content.

Dynamic translation in activity streams

You can enable dynamic translation in activity streams on the ServiceNow AI Platform, configurable workspace, and Service Portal. For more information on Dynamic Translation support for configurable workspaces, see Dynamic Translation support for configurable workspace.

Use the dynamic translation properties to enable or prevent the option for dynamic translation in the activity stream of the tables you list. When enabled, select the translate icon ( ) to show or hide the translated content for each comment or work note in the activity stream. For more information, see Dynamic Translation properties.

Length limitations for translation service providers

The following are the length limitations for translation service providers:

- For Microsoft, the input text for translation cannot have more than 10,000 characters including spaces. And the limit for the input text for detection is 50,000. For more information, see the Microsoft Documentation.
- For IBM, the size of the input text cannot exceed 50KB. For more information, see the IBM Documentation.

  Important:
  IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

- For Google, the recommended maximum length of each request is 5,000 characters. The maximum size of the input text cannot exceed 30,000 code points. For more information, see the Google Documentation.

Note:
Starting with the Rome release, you can use Dynamic Translation without the snc-internal role.

Enable dynamic translation for a field

Translate the dynamically generated text for a field on Platform forms by enabling the field-level dynamic translation. You can translate the text based on the preferred language of a user.

Before you begin

Role required: admin
Activate the Dynamic Translation plugin (com.glide.dynamic_translation).

About this task

Dynamic Translation supports only the following field types:

- String
- String (Full UTF-8)
- Multi Line Small Text Area
- Wide Text
- HTML

Note:

- The following are the length limitations for translation service providers:
  - For Microsoft, the input text for translation cannot have more than 10,000 characters including spaces. And the limit for the input text for detection is 50,000. For more information, see the Microsoft Documentation.
  - For IBM, the size of the input text cannot exceed 50KB. For more information, see the IBM Documentation.

    Important:
    IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.
  - For Google, the recommended maximum length of each request is 5,000 characters. The maximum size of the input text cannot exceed 30,000 code points. For more information, see the Google Documentation.
- You cannot translate a field on Agent Workspace forms.

Procedure

1. Navigate to any form, for example, an Incident form.
2. Right-click the required field and select **Configure Dictionary**.
3. In the Attributes related list, click **New**.
4. In the Dictionary Attribute form, fill the fields.

Dictionary Attribute form fields

Field	Value
Attribute	Dynamic Translation Enabled
Value	true

5. Click **Submit**.
6. Click **Update**.

Dynamic translation for a field

A translate icon appears next to the field. The logged-in user can click this icon to translate the text to the user's preferred language using the default translation service provider specified in the translator configurations. For information about error messages during translation, see Error messages in Dynamic Translation.

Note:
If you do not have any default translator configuration, you cannot view the translate icon .

Dynamic Translation for Agent Chat overview

Dynamic Translation for Agent Chat (DTAC) translates chat conversations from one language to another. This translating enables agents and requesters to have chat conversations with each other without the need of a shared language.

Dynamic translation for Agent Chat overview video

For example, a requester who speaks only French starts a chat conversation with a live agent who speaks only English. From the requester's view, the entire conversation is conducted in French, while from the agent's view, the entire conversation is conducted in English.

DTAC enables cross-locale support teams and merger and acquisition projects to save time and money by handling larger volumes of chat conversations from other regions.

Note:

- DTAC is supported on Agent Workspace, not on Connect Support.
- DTAC is currently not supported when autopilot is in progress. When the requester responds to prompts from autopilot topics, these responses are not translated for the agent, even if DTAC is turned on for that conversation. See Conversation Autopilot for more information.

How DTAC works

DTAC uses Dynamic translation and a translation service provider to translate chat conversations in real time. Available translation service providers include:

- Set up Google cloud translator service spoke
- IBM Translator Service
- Microsoft Azure Translator Service spoke

Mobile clients

DTAC is available for requesters on Agent Chat in mobile for iOS and Android.

Language limitations

Even though the corresponding I18N plugins are activated, DTAC does not support the following language translations:

- Brazilian Portuguese
- Canadian French

Translation support limited to one language

DTAC can only support translation from one language to another therefore additional agents or managers joining chat sessions should speak the same language as the primary agent.

- Agent chat transfers – Agents transferring a chat conversation should both speak the same language.
- Managers joining chat conversation – Managers joining chat conversation to monitor or assist agents should speak the same language as the primary agent.

Agent presence states

DTAC does not translate agent presence states.

Attachments within DTAC

When inserting an attachment into a chat, DTAC does not translate the file name or the contents of the attachment.

Translate a knowledge article from a translation task

Start translating articles once they are published and assigned to you for translation.

Before you begin

- Activate the Dynamic Translation plugin (com.glide.dynamic_translation).
- Enable the `glide.knowman.translation.enable_translation_task` property. For more information, see Translation management.

Role required: admin, knowledge_admin, or knowledge

About this task

Note:
This procedure has been superseded by the Localization Framework. For more information see Localization Framework support for Knowledge Base. This information is provided for legacy context.

For information about how administrators configure the translation API, see DynamicTranslation API.

Procedure

1. Navigate to **All** > **Knowledge** > **Translation Management** > **My Assigned Tasks**.
2. In the Knowledge Translation Tasks page, select one of the listed translation tasks assigned to you.
3. Select **Translate**.
   The **Translate from** and **Translate to** panes are side-by-side for easy comparison, as shown in the following figure.

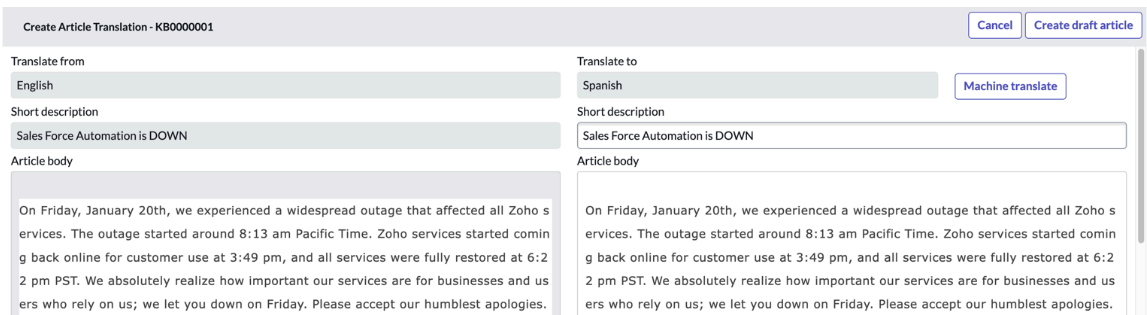

4. Select **Machine translate** to automatically translate to the targeted languages.
5. Select **Create draft article** to create a draft version of the translated article.
   If a translated article is already present for that language, a new draft version overwrites the existing translated content.
6. Select **Publish**.

   Note:
   By default, the translated version is always created as v1.0 so that the parent version remain same for all the translated versions. For more information, see Translated KB articles always show Parent version 1.0 regardless of version used for translation.

Result

The newly created draft is listed in the **Translated Versions** tab.

DynamicTranslation API

You can use the `DynamicTranslation` API to translate text in real time into multiple languages using translation service providers.

This API is available as part of Server APIs, Client APIs, and Service Portal APIs to interact with translation service providers. See DynamicTranslation - Scoped and DynamicTranslation - Client.

No matter what translation service provider you choose, this API works on a standard set of inputs and outputs, and displays standard error codes and messages.

There are a few limitations when you integrate with a translation service provider to enable the Dynamic Translation experience. For more information, see Limitations in Dynamic Translation.

Dynamic Translation spoke

The Dynamic Translation spoke provides actions to enable the ability to translate multiple texts from one language into one or more languages and detect the language of multiple texts seamlessly.

The Dynamic Translation APIs are accessible across the ServiceNow platform through the following access points:

- Server
- Platform Client
- Portal Client

In the Washington DC release, the Dynamic Translation spoke enables you to access the Dynamic Translation APIs from the Flow Designer.

The Dynamic Translation Spoke is a plugin (com.glide.dynamic_translation.spoke) that is available with activation of the Dynamic Translation plugin (com.glide.dynamic_translation).

The Dynamic Translation spoke contains actions that are built corresponding to the Dynamic Translation APIs.

The following actions are built in the Washington DC release in Dynamic Translation spoke:

Dynamic Translation Spoke actions and APIs

Action Name	API
Detect Language	getDetectedLanguage()
Detect Languages	getDetectedLanguages()
Translate Text	getTranslation()
Translate Texts	getTranslations()
Is Dynamic Translation Enabled	isDynamicTranslationEnabled()

For more information about accessing the actions from the Flow Designer, see Access Dynamic Translation spoke actions from the Flow Designer.

Access Dynamic Translation spoke actions from the Flow Designer

Access Dynamic Translation spoke actions from the Flow Designer to dynamically detect and translate the given texts.

Before you begin

- Ensure that the Dynamic Translation spoke is activated with Dynamic Translation plugin (com.glide.dynamic_translation) activation and configure the Dynamic Translation framework for a translation service provider.
- Role required: admin

Procedure

1. Navigate to **All** > **Process Automation** > **Flow Designer**.
2. Click the **Actions** tab.

3. Search the actions using the filter criteria and search bar.
   You can search for the actions with the spoke's name. For example, Dynamic Translation. All the actions built in the Dynamic Translation spoke are displayed.

   For more information, see Actions in Dynamic Translation spoke.

Language Detection spoke

ServiceNow Language Detection Service spoke provides a single subflow and action to detect the language of the given text using the ServiceNow language detection service.

The ServiceNow Language Detection Service spoke is an in-house language detection service offered by ServiceNow platform in the Washington DC release.

The ServiceNow Language Detection Service spoke plugin (com.glide.language_detection_spoke) is available with the activation of Dynamic Translation plugin (com.glide.dynamic_translation).

Subscription

This spoke requires any ServiceNow subscription that is Professional or above, and includes Dynamic Translation.

Supported versions

API version v3.

Spoke subflow

The spoke provides the subflow in the Published state to integrate with the ServiceNow language detection service.

Spoke subflow

Subflow	Description
Inputs	
Detect Language [detect_language_v3]	Detects the languages of texts using ServiceNow Language Detection Service.

Spoke action

The spoke provides the action in the Published state to integrate with the ServiceNow language detection service.

Spoke action

Subflow	Description
Inputs	
Detect Language [detect_language_v3]	Detects the languages of texts using ServiceNow Language Detection Service.

The ServiceNow Language Detection service uses the following Dynamic Translation APIs to detect the language of the input texts:

- getDetectedLanguage: Provides text to identify the languages.
- getDetectedLanguages: Provides texts to identify the languages.

Language Code Mapping in Dynamic Translation

Dynamic Translation enables you to add mapping between the ServiceNow® language codes and the language codes used by your third-party translation service provider.

- You can create a language code to map the ServiceNow® language code.
- You can view the default language code mapping records in the Language Code Mapping [sn_trans_commons_language_code_mapping] table.
- You can associate the language code mapping to a translator configuration record. For more information, see Create a custom translator configuration.

Note:
By default, language code mappings are included in the records of Microsoft Azure Translator Service and Google Cloud Translator Service. These preconfigured records for third-party providers are available in the Translator Configurations [sn_dt_translator_configuration] table.

To add a new language code mapping, see Create a language code mapping.

Create a language code mapping

Create a language code mapping to map the ServiceNow® language codes with the language codes of the translation service providers.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Dynamic Translation** > **Language Code Mappings**.
2. Click **New**.
3. On the form, fill in the fields.

Language Code Mapping form fields

Field	Description
Language code	Language for which a ServiceNow® language code mapping record is being created.
Mapped language code	Translation service provider's language code with which the ServiceNow® language code is being mapped.
Name	Identifier for the language and the mapping code. The name is auto-filled using the language code and the mapped language code in the "language code -> mapped language code" format. For example, fq -> fr-CA.

Field	Description
	Users can edit the fields if required.

4. Click **Submit**.

Note:
After creating the language code mapping, associate it with a corresponding translator configuration to map the existing ServiceNow® language code.

Reference for Dynamic Translation

Get details about Dynamic Translation domain separation, properties, actions, limitations, and error messages.

Domain separation and Dynamic Translation

Domain separation is supported in Dynamic Translation and is configured to apply to translator configurations. 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.

Activation information

You should activate the Domain Support - Domain Extensions Installer plugin (com.glide.domain.msp_extensions.installer). For information on how you can request for the plugin activation, see Request for domain separation in Dynamic Translation.

How domain separation works in Dynamic Translation

A service provider with domain separated instances can implement the Dynamic Translation framework so that users can use the translation service providers configured in the translator configurations specific to their domain. Translator configurations are process-separated in Dynamic Translation. All translation service providers configured in the translator configurations of the parent domain are available in child domains.

A translator configuration in the global domain is visible in all other domains. In any domain, a user can only view the translator configurations configured in the following domains:

- Current domain
- Parent domains
- Global domain

Note:
If there are multiple translator configurations in a domain for the same translation service provider, the translator configuration of the current domain or the closest parent domain is considered for dynamic translation.

Also, different connections can be configured for the same connection and credential alias of a translation service provider in multiple domains. However, Credentials and Connections are data-separated. So, a connection configured in a child domain is visible in parent domains. For information on domain separation for Connections and Credentials, see Domain separation for Credentials and Connections.

For example, consider the following scenario:

Connections for a translation service provider in multiple domains

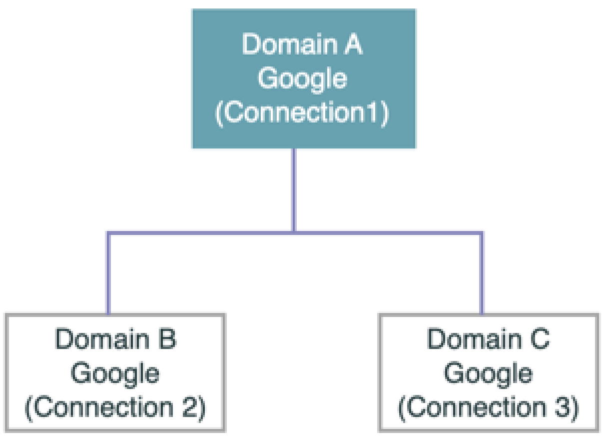

Different connections are configured for the same connection and credential alias of Google in multiple domains. When in Domain A, all the following connections are visible:

- Connection1
- Connection2
- Connection3

Note:
When multiple active connections are visible for the same connection and credential alias in a domain, the connection with the highest order is given priority when the **Support Multiple Active Connections** check box is selected.

Domain-separated table

Translator Configuration [sn_dt_translator_configuration].

Default translator configuration for a domain

The default translator configuration of the current domain is always considered for dynamic translation. If the current domain does not have any default translator configuration, then the available default translator configuration of the nearest parent is considered.

A domain can have multiple default translator configurations. In this case also, the default translator configuration of the current domain is considered for dynamic translation. For example, let us consider the following scenario:

Domain separation in Dynamic Translation

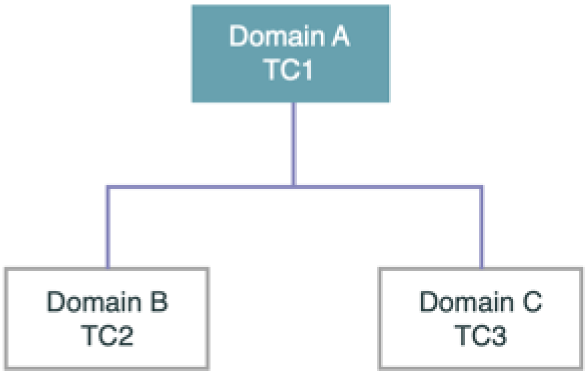

In Domain B, both TC1 and TC2 are visible. From Domain B, TC2 is first set as the default translator configuration. From Domain A, TC1 is then overridden and set as the default translator configuration. This results in multiple default translator configurations in Domain B. In this case, when in Domain B, TC2 is used as the default translator configuration for dynamic translation.

Overriding a translator configuration

In any domain, you can override the translator configuration of that domain or the parent domain. The overridden translator configuration of a parent domain is also visible in child domains. However, the overridden translator configuration of a child domain is not visible in the parent domain.

After you override a translator configuration of the same domain, only the overridden translator configuration is visible in that domain.

After you override a translator configuration of a parent domain, the following scenarios happen:

- Only the overridden translator configuration is visible in the child domain.
- The **Overrides** field of overridden translator configuration refers to the original translator configuration of the parent domain.

For example, consider the following scenario:

Domain separation in Dynamic Translation

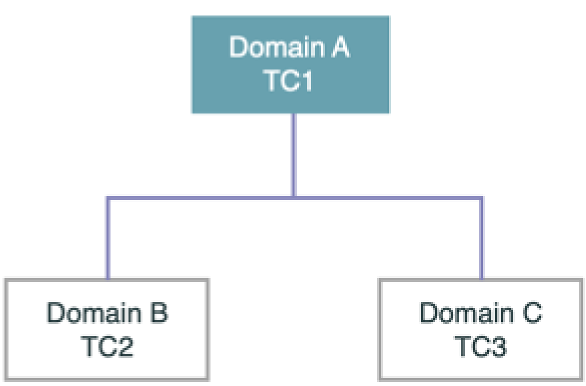

You can override a translator configuration TC1 from Domain B. After overriding, only the overridden configuration TC1 is available in Domain B and the **Overrides** field of TC1 refers to **TC1** of the parent domain.

Related topics

Domain separation for service providers

Dynamic Translation properties

Use the Dynamic Translation properties to customize, configure, and control the translation features.

Navigate to **Dynamic Translation** > **Properties** to view and edit these properties.

Dynamic Translation properties

Property	Description
Activity Stream	
Allow list tables sn.dt.activity_stream.allow_list	List of tables (comma-separated) for which the option to translate the content is enabled on the activity streams (work notes and comments). Example: `sn_hr_core_case,task`. • Type: string • Default value: none
Block list tables sn.dt.activity_stream.block_list	List of tables (comma-separated) for which the option to translate the content is inactive on the activity stream (work notes and comments). Example: `change_request,sn_hr_core_case_payroll`. • Type: string • Default value: none

Actions in Dynamic Translation spoke

Actions in Dynamic Translation spoke are built to access the Dynamic Translation APIs from the Flow Designer.

Dynamic Translation Spoke actions

Action	Description
Detect Language [detect_language]	Detects the language of the text using the configured translator.
Detect Languages [detect_languages]	Detects the languages of multiple texts using the configured translator.
Is Dynamic Translation Enabled [is_dynamic_translation_enabled]	Checks if the Dynamic Translation plugin (com.glide.dynamic_translation) is installed

Dynamic Translation Spoke actions (continued)

Action	Description
	and configured for a Translation Service Provider. • If the input is empty, checks if a default Translation Service Provider is configured. • If a translator is provided, checks if that Translation Service Provider is configured.
Translate Text [translate_text]	Translates text to multiple languages using the configured translation service provider. The action can translate a text string into multiple target languages.
Translate Texts [translate_texts]	Translate texts to multiple languages using the configured translation service provider. The action can translate an array of texts into multiple target languages.

Standard Error Codes

When an error occurs, error codes populate with the reason for the error. Refer to the following information for the standard error codes that appear while you work on the Dynamic Translation spoke actions corresponding to the Dynamic Translation APIs.

- Dynamic Translation - Scoped
- Dynamic Translation - Client

Limitations in Dynamic Translation

You must be aware of a few limitations when you integrate with a translation service provider to enable the Dynamic Translation experience.

Length limitations for translation service providers

The following are the length limitations for translation service providers:

- For Microsoft, the input text for translation cannot have more than 10,000 characters including spaces. And the limit for the input text for detection is 50,000. For more information, see the Microsoft Documentation.
- For IBM, the size of the input text cannot exceed 50KB. For more information, see the IBM Documentation.

Important:
IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

- For Google, the recommended maximum length of each request is 5,000 characters. The maximum size of the input text cannot exceed 30,000 code points. For more information, see the Google Documentation.

Language support limitations for translation service providers

Not all source language translations are supported by translation service providers.

For information about language translations supported by Microsoft, see the Microsoft documentation.

For information about language translations supported by IBM, see the IBM documentation for identifiable languages, and the IBM documentation for translation models.

Important:
IBM has announced the deprecation of the IBM Watson Translator Service for IBM Cloud in all regions. As of June 10, 2023, the Language Translator tile will be removed from the IBM Cloud Platform for new customers; only existing customers will be able to access the product. As of June 10, 2024, the service will reach its End of Support date. As of December 10, 2024, the service will be withdrawn entirely and will no longer be available to any customers. For more information, see https://cloud.ibm.com/docs/language-translator?topic=language-translator-release-notes.

For information about language translations supported by Google, see the Google documentation.

Error messages in Dynamic Translation

You must be aware of a few error scenarios while using Dynamic Translation.

Error messages on forms and activity streams

Error messages are displayed when the text cannot be dynamically translated on forms activity streams.

Forms and activity streams

Scenario	Error message
Field has no value. **Note:** This scenario is applicable only for fields on the form.	No text to translate.
Text to be translated is already in the logged-in user's preferred language.	This content is written in your preferred language. No need to translate.

Forms and activity streams (continued)

Scenario	Error message
Text cannot be translated to the user's preferred language by the default translation service provider.	Text cannot be translated to your preferred language.
Credentials configured for the translation service provider under **Connections & Credentials** are not valid.	Credentials are missing or invalid. Contact your administrator.
Text to be translated has exceeded the maximum length supported by the default translation service provider.	Text has exceeded the maximum length.
Error during translation. Possible reasons are: • Credentials configured for the translation service provider under the **Connections & Credentials** sub module are missing. • Translation got timed out. • Any other issue with the translation service provider. For example, the exhaustion of credentials for more transactions.	Unable to translate.

Error messages on translator configuration

Translator configuration

Scenario	Error message
When there is no valid subflow associated with the translator configuration.	Either a detect or a translate subflow is required for a valid translator configuration.
When multiple language code mappings are selected for the same language code	Multiple language code mappings cannot be selected for the same language code(s): xx.

ServiceNow AI Platform integrations

Integrate third-party applications into the ServiceNow AI Platform.

Integration with third-party applications and data sources

Connect with external services and configure your enterprise perfectly through the ServiceNow AI Platform. ServiceNow systems integrate easily with third-party apps, using industry standard technologies such as SOAP, REST, WSDL, and API, as well as several others such as Web Services, JDBC, and more. Load external data using XML, CSV, XLS, and other file types, all through multiple data sources including HTTPS, FTPS, and SCP. Integration between instances is even easier, due to shared integration points on the platform.

Integration with third-party applications and data sources

ServiceNow integrates with many third party applications and data sources.

The most common integrations are with CMDB, Incident Management, Problem Management, Change Management, User Administration, and Single Sign-on. A variety of techniques can

be used, most notably Web Services, JDBC, LDAP, Excel, CSV, and Email, as well as any industry standard technologies that use SOAP, REST, or WSDL. Additionally, API and command-line integrations can be done using a MID Server. ServiceNow has performed the following integrations with enterprise systems and platforms.

Technologies

The ServiceNow platform is based on service-oriented architecture (SOA), in which all data objects can use web services to access bi-directional data-level integration. The interface is also direct and dynamic because all modifications to existing objects and all new objects are automatically published as a Direct Web Service. A more indirect web service creation and usage can be achieved through Mapped Web Service where a transform map is used to gather incoming web service data into the final targeted tables. Finally, an advanced Scripted Web Service technique is available for defining process-based web services, where data is irrelevant, but serves more as a trigger for a process or a composite of actions that execute at the server.

Additionally the platform offers a rich interface for loading external data using import sets. Using this feature, you can load from various data sources such as HTTPS, FTPS, and SCP using file formats such as XML, CSV, and Microsoft Excel XLS files. Information can also be pulled from a data source using a direct JDBC connection, provided the network connectivity allows.

Information can be pulled from the platform to an external platform using an ODBC driver.

Forms, lists, and reports on the platform can be accessed directly using a URL, which facilitates integration on the UI level between two or more web applications.

A handful of single sign-on technologies is identified and implemented out of the box to allow fast integration with your portal, however, the technique is customizable in a script to allow for flexibility in the different SSO environments our customers have.

Integration between ServiceNow Instances

There are times when you find you need to perform a specific integration between your instance and another ServiceNow instance. Instance-to-Instance integrations are a snap because all of the integration points exist between the two instances.

Integration options

Nearly all ServiceNow customers obtain additional value by integrating with third-party applications.

ServiceNow customers have the following options for integrating with third-party applications.

- Activate a ServiceNow provided integration.
- Install a certified integration from the ServiceNow Store.
- Install an integration from Share.
- Contact your sales representative for available custom-built integrations.
- Build your own custom integration using the platform's integration interfaces.

Note:
Certified integrations have passed a set of interoperability, security, and performance test criteria defined by ServiceNow.

ServiceNow provided integrations

ServiceNow provides many integrations as part of the platform.

These integrations are considered part of the platform and are provided at no additional charge.

Provided Integrations

Integration	Type	Integration Point
Altiris (version 6.5)	MID Server	CMDB
Google Maps	Web services	Varies
Google Custom Search	Web services	Knowledge
Integration — IBM Watson Conversation Service	Web services	Varies
Microsoft SMS / SCCM	MID Server	CMDB
Verizon eBonding	Web services	Incident

Set up Google Maps API

You enable the geolocation feature using the appropriate entries from a Google Maps API for Business license, or a Google Maps JavaScript API Key.

Before you begin

Role required: admin

About this task

You can generate maps using basic JavaScript, and they are flexible enough to display even the most complicated of queries. These maps use standard Google Maps API for Business or Google Maps JavaScript API Key mapping features, including various link types to records in your instance. Use of this feature requires installation of the Google Maps plugin.

Map page

Procedure

1. Google Maps APIs for Business licenses are no longer available from Google, so, depending on whether have an existing license, use one of the following procedures.

   Note:
   You are required to use your own Google Maps JavaScript API Key from Google and not use any ServiceNow AI Platform® generated key.

   To learn more about Google Maps API for Business, and Google Maps JavaScript API Key, see https://developers.google.com/maps/gmp-get-started.
2. After you complete and save the configuration, create Map pages, or use the default pages included with the plugin.

Map pages define what data appears on the map, and the appearance of the links. For a tutorial on how to view all of the markers that link to your records, see the Display Map Markers on Google Maps blog post on the ServiceNow Community.

Property	Description
google.maps.auto_close	If true, automatically closes a map information window before opening a new one. ◦ Type: true/false ◦ Default value: true ◦ Location: **System Properties** > **Google Maps**
google.maps.client	Client ID for Google Maps API for Business. ◦ Type: string ◦ Default value: gme-servicenow ◦ Location: **System Properties** > **Google Maps**
google.maps.private.key	Private key for Google Maps API for Business. This key activates the geolocation feature, which locates users in the system precisely, using data from their mobile devices. ◦ Type: string ◦ Default value: empty ◦ Location: **System Properties** > **Google Maps**
google.maps.version	Version number of the current installation of Google Maps API for Business. ◦ Type: string ◦ Default value: current version number ◦ Location: **System Properties** > **Google Maps**
google.maps.key	Google Maps API for Business or Google Maps JavaScript API Key that ties to the URL of the server. This key authorizes development use of Google Maps API. ◦ Type: string ◦ Default value: empty ◦ Location: **System Properties** > **Google Maps**
google.maps.latitude	Starting latitude of the map. This value determines the starting position displayed in Google Maps. ◦ Type: string ◦ Default value: 36.008522 ◦ Location: **System Properties** > **Google Maps**
google.maps.longitude	Starting longitude of the map. This value determines the starting position displayed in Google Maps pages. ◦ Type: string ◦ Default value: -95.221764 ◦ Location: **System Properties** > **Google Maps**

Property	Description
google.maps.max_items	Maximum number of items to display on the map. ◦ Type: integer ◦ Default value: 500 ◦ Location: **System Properties** > **Google Maps**
google.maps.table	Table used by the map. The table needs the following fields: name, longitude, latitude. ◦ Type: string ◦ Default value: cmn_location ◦ Location: **System Properties** > **Google Maps**
google.maps.zoom	Starting zoom level of the map (1 is the lowest) ◦ Type: string ◦ Default value: 4 ◦ Location: **System Properties** > **Google Maps**

3. Perform these configurations:

a. Navigate to **System Properties** > **All**.

b. Set values for these system properties:

System property	Description
google.maps.refresh.latlong.duration	Number of days after which the location details are updated. **Note:** The maximum value is 30.
google.maps.refresh.latlong.record_count_per_table	Maximum number of records in a table that can be processed using the scheduled job in one refresh.

c. Specify the tables that store the location details in the Map Coordinates Refresh Configuration [cmn_coordinate_refresh_config] table.

d. If there are no fields to track when the coordinates are last retrieved in your tables:

i. Create the **coordinates_retrieved_on** field of the type **Date/Time** in your tables.

ii. Ensure that the value of the field that tracks time when the coordinates are last retrieved, is populated correctly in the Map Coordinates Refresh Configuration [cmn_coordinate_refresh_config] table. Refer to the sample fix script `Update coordinates updated on field` to update values in the **coordinates_retrieved_on** field.

e. Navigate to **System Definition** > **Scheduled Jobs**.

f. Open the record, **Refresh Latitude Longitude info of Records**.

g. Select the **Active** check box.

h. Specify the frequency at which the location details should be refreshed by providing frequency details in **Run** and **Repeat Interval**.

> **Note:**
> Ensure that you refresh the location details at least once every 30 days.

4. **Optional:** To access Google Maps from the ServiceNow instance based on your region, perform these configurations.

a. Navigate to **User Administration** > **Users**.

b. Select and assign the relevant **Country Code** for the required user.

> **Note:**
> If the required **Country Code** isn't available, add the dictionary entry. Ensure that you use the ISO 3166-1 country codes. For example, the country code of Morocco is `MA`.

c. Navigate to **System Properties** > **Google Maps**.

d. Select the **Yes** check box for **Enable this property only when you want to alter the Google map's behavior based on a region. The region is populated from the current User's 'Country code' field value in sys_user record. The Country code value on the sys_user record should be consistent with the expectations from the Google Map's region code. Most Unicode region identifiers are identical to ISO 3166-1 codes, with some notable exceptions. Check the codes properly before making any changes.**

e. Click **Save**.

Microsoft SCCM integration

The Microsoft SCCM integration is a one direction import of SCCM data into the ServiceNow Configuration Management Database (CMDB).

Scheduled imports bring relevant SCCM data into the ServiceNow instance from an SQL Server database and map it to tables in the CMDB. You can configure either a full or incremental data import. The import is achieved using a JDBC connection via the MID Server. The integration keeps the ServiceNow CMDB synchronized with the SCCM SQLServer database, so that only relevant data is imported from the SCCM database to the CMDB. The SCCM database is considered an authoritative source and is not written to.

Supported versions

The Microsoft SCCM legacy plugin is deprecated and no longer supported or available for new activation. For details, see the Deprecation Process [KB0867184] article in the Now Support knowledge base. Recommendation for all customers to migrate to the new Service Graph connector for Microsoft SCCM.

These are the plugins for the currently supported SCCM versions, all of which support Asset Intelligence and Incremental Software Reconciliation:

- Microsoft SCCM 2007
- Microsoft SCCM 2012 v2
- Microsoft SCCM 2016

Note:
If you upgrade your instance to a version of the Now Platform that does not provide support for your version of SCCM, you can continue to use that version. However, all new instances require the use of a supported SCCM version.

The Integration - Microsoft SCCM 2016 (com.snc.integration.sccm2016) plugin is compatible with SCCM versions 1606, 1906, 1910, and 2002.

Service Graph connector for Microsoft SCCM

To import SCCM data into the CMDB, you can use the Service Graph connector for Microsoft SCCM instead of using the legacy SCCM connector.

Using the Service Graph SCCM connector provides these advantages:

- Adherence to the Common Service Data Model (CSDM) data model.
- Use of Identification and Reconciliation engine (IRE) to avoid duplicate records.
- Support for removed software detection and delta detection.
- Use of the IntegrationHub ETL (2.2.1) which allows for easy data mapping and customizations.
- Higher performance during data processing.

For information on how to update from using the Microsoft SCCM connector plugin to using the Service Graph connector for Microsoft SCCM, see Update Microsoft SCCM 2016 plugin for Service Graph connector for Microsoft SCCM.

Available modules

- **Setup**:
  - Configure the data sources.
  - Specify database server settings and the MID Server.
  - Test the configuration.
- **Scheduled Import**: Schedule the import or execute the import immediately.
- **Data Sources**: A list of the pre-configured data sources defining the external CMDB database.
- **Progress**: The progress log for scheduled imports.
- **Transform History**: A log of transformations performed by scheduled imports.
- **CI Identification**: Identifiers used by SCCM to match import set data with the CMDB.
- **Import Set Data**: Tables that store imported SCCM data.
- **Web Services**: List of web services used to add or remove fields in the transform maps.

Asset intelligence

Asset Intelligence is an SCCM feature that can filter, normalize and clean up software records. When enabled, it populates normalized software data in a separate table within the SCCM database. You can target this separate table instead of the raw software data table. When you import software records into the CMDB with this feature enabled, you get a cleaner set of

software data. To use this feature, you must Activate SCCM Asset Intelligence scheduled imports in the target ServiceNow instance, as well as in the source SCCM database.

Activate a supported Microsoft SCCM plugin

The SCCM plugins can be activated by an administrator and requires the Integration - JDBC and CI Identification plugins. These dependent plugins are activated automatically with SCCM activation.

Before you begin

For a complete listing of the plugins that need to be activated, see Activate SCCM Asset Intelligence scheduled imports.

Role required: admin

Procedure

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

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

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

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

SCCM data import process and source tables

The Microsoft SCCM versions supported in the ServiceNow® platform offer identical features and the same imported data.

Data import process

1. A schedule called **SCCM System <version> Import** determines when the SCCM tables are imported into the ServiceNow® instance. Imports can be executed immediately or scheduled to run at defined intervals.
2. A MID Server retrieves the SCCM data and imports it into staging tables on the instance.
3. Transforms run on the data in the staging tables and map the SCCM data to existing fields in the CMDB.

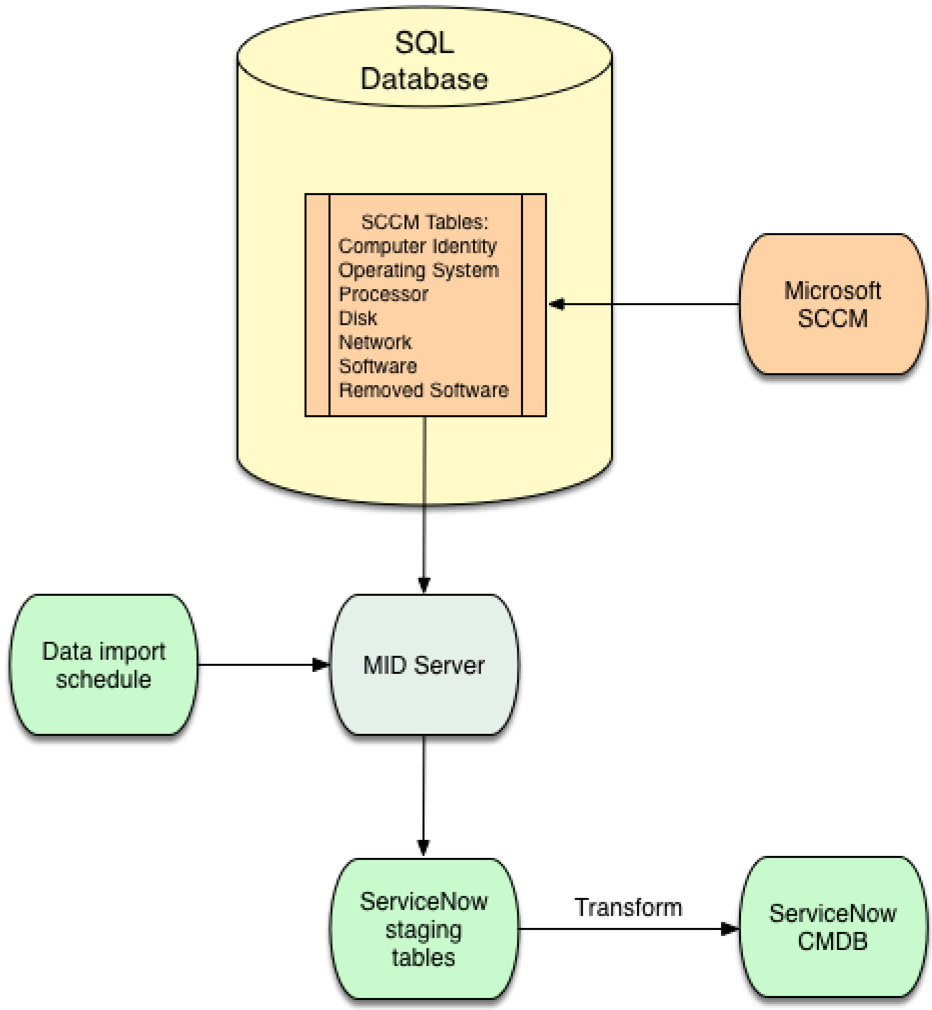

SCCM data imported

This table shows the SCCM source tables, the corresponding import set staging tables in the ServiceNow® instance, and the target tables in the CMDB.

SCCM table	Staging table	CMDB table
• v_GS_Computer_System • v_GS_Workstation_Status • v_GS_System • v_GS_PC_Bios • v_GS_Operating_System • v_GS_Computer_System_Product • v_GS_System_Enclosure • v_GS_Baseboard	SCCM <version> Computer Identity [imp_sccm<version>_computer_id]	Computer [cmdb_ci_computer]
v_GS_Disk	SCCM <version> Disk [imp_sccm<version>_disk]	Disk [cmdb_ci_disk]

SCCM table	Staging table	CMDB table
v_GS_Network_Adapter_Configuration	SCCM <version> Network [imp_sccm<version>_network]	Network Adapter [cmdb_ci_network_adapter]
v_GS_Operating_System	SCCM <version> Operating System [imp_sccm<version>_os]	Computer [cmdb_ci_computer]
v_GS_Processor	SCCM <version> Processor [imp_sccm<version>_processor]	Computer [cmdb_ci_computer]
• v_GS_Add_Remove_Programs • v_GS_Add_Remove_Programs_64	SCCM <version> Software [imp_sccm<version>_software]	• Software [cmdb_ci_spkg] • Software Installation [cmdb_sam_sw_install] (when Software Asset Management is enabled*)

*For more information about Software Asset Management and how to enable it, see Request Software Asset Management.

SCCM data sources

The ServiceNow® SCCM integration uses JDBC data sources to import software data from the SCCM database. Each data source contains the connection specifics for the SCCM database and names the MID Server the instance will use to import the data. The transforms that map the SCCM fields to the CMDB are defined in a related list in each data source record.

Transform maps

Transform maps are accessed from the **Transforms** related list in each data source record. The source fields in SCCM and the target fields in the CMDB are listed in the **Field Maps** related list in each Table Transform Map record. The SCCM integration uses the transform map utility provided with the ServiceNow® platform. For instructions on editing or creating a transform map, see Transform maps.

The SCCM integration provides two transform maps for incremental software imports. Only one transform map can be enabled (**Active**) at a time.

- **Incremental Import**: Enabled by default. This map should be configured as **Active** when ServiceNow® Software Asset Management is not enabled on the instance.
- **Incremental Import (SAM enabled)**: If the Software Asset Management plugin is activated, set this transform to **Active**

Note:
To force a one-time full import of all software data from the SCCM database, clear the value in the**Last run datetime** field on the data source. This operation can take a long time to execute, so the most efficient method is to use an incremental transform after the first full import.

Transforming the assigned user

The SCCM <version> Computer Identity transform script attempts to set the **Assigned to** field in the CMDB record by looking up the name of the user in the SCCM source table and comparing the value with the matching field in the ServiceNow sys_user table. If a match is found, that user is assigned to the record. If no match is found, the **Assigned to** field is left blank. The matching

field is controlled by the *glide.discovery.assigned_user_match_ field* system property, which is set to **user_name** by default.

Identifiers

The SCCM integration uses CI identification to update CIs created from data imported from SCCM with a resource ID. The Hardware Rule identifier returns the resource ID of a computer from SCCM and stores it in the Source [sys_object_source] table. When resource IDs are first imported, either from SCCM or Discovery, the [sys_object_source] table is populated with IDs for each CI it identifies. In subsequent imports, if an incoming ID matches that of an existing CI, IRE (Identification and Reconciliation Engine) updates the information for that CI in the CMDB. If the incoming resource ID does not match that of an existing CI, IRE creates a new CI and populates it with the resource ID.

For more information about CMDB Identification and Reconciliation and IRE, see CMDB Identification and Reconciliation.

Upgrades from pre-Geneva versions still preserve the legacy identifiers, but you can switch to the new identifiers using a property: `glide.discovery.use_cmdb_identifiers`. If you upgraded from a pre-Geneva version, you must manually add this property and set it to `true` to use the new identifiers. If you upgraded from Geneva, this property is available in the System Properties [sys_properties] table. To preserve functionality in custom legacy identifiers, convert them to the new CMDB identifier rules format before enabling this property. The system does not reconfigure your custom identifiers to the new framework automatically.

Scripts

Data population scripts populate the related data in the CMDB for each target CI discovered by the Hardware Rule identifier.

Software

The Microsoft SCCM integration reconciles the software package count in the records for a CI and removes a software instance from the Software Instance [cmdb_software_instance] table if the software package is uninstalled from the CI. For a full software import, the transform populates the Software [cmdb_ci_spkg] and Software Instance [cmdb_software_instance] tables. If the Software Asset Management plugin is enabled, the transform populates the Software Installation [cmdb_sam_sw_install] table.

Note:
The table data imported from SCCM must contain complete data for the CI. The instance assumes that the import represents all relationships that exist and adjusts the CMDB accordingly. Partial data received from SCCM tables can cause the deletion of active relationships.

Configure the SCCM integration and schedule an import

Importing data from Microsoft SCCM requires a connection to the SCCM database via a JDBC data source and a schedule that tells the MID Server when to fetch the data.

Before you begin
You must have sufficient credentials to query the SQL Server that contains the SCCM database. Create a SQL Server account in order to connect to the SCCM database on the MID Server to use the data sources.
Connect to your SCCM SQL Server and configure the following:

1. Create a new login user name and password for SQL Server authentication.
2. Choose the SCCM database that ServiceNow data sources will connect to.
3. Assign the 'db_datareader' role membership to the new user.

Role required: admin

About this task

The system uses the details of the database connection you create in this procedure to configure the default SCCM data sources automatically.

> **Note:**
> If you want to configure the JDBC data source to authenticate on your SQL Server database using the Windows MID Server service user, select the **Use integrated authentication** box on each of the existing data sources.

Procedure

1. Navigate to **All** > **Integration - Microsoft SCCM <version>** > **Setup**.
2. Complete the required fields to configure connection to the SCCM SQL database.
   Make sure to include a valid MID Server that has access to the SQL database server.

   SCCM data import setup

   

3. Under Related Links, click **Test data source connections** to test the connection to the SCCM database.
4. If the data source connects to the SCCM database, navigate to **Integration - Microsoft SCCM** > **Scheduled Import**.
   The Scheduled Data Import form appears for your version of the SCCM integration. The Computer Identity data source runs first. The other data sources listed on the form run in sequence, based on their default order.

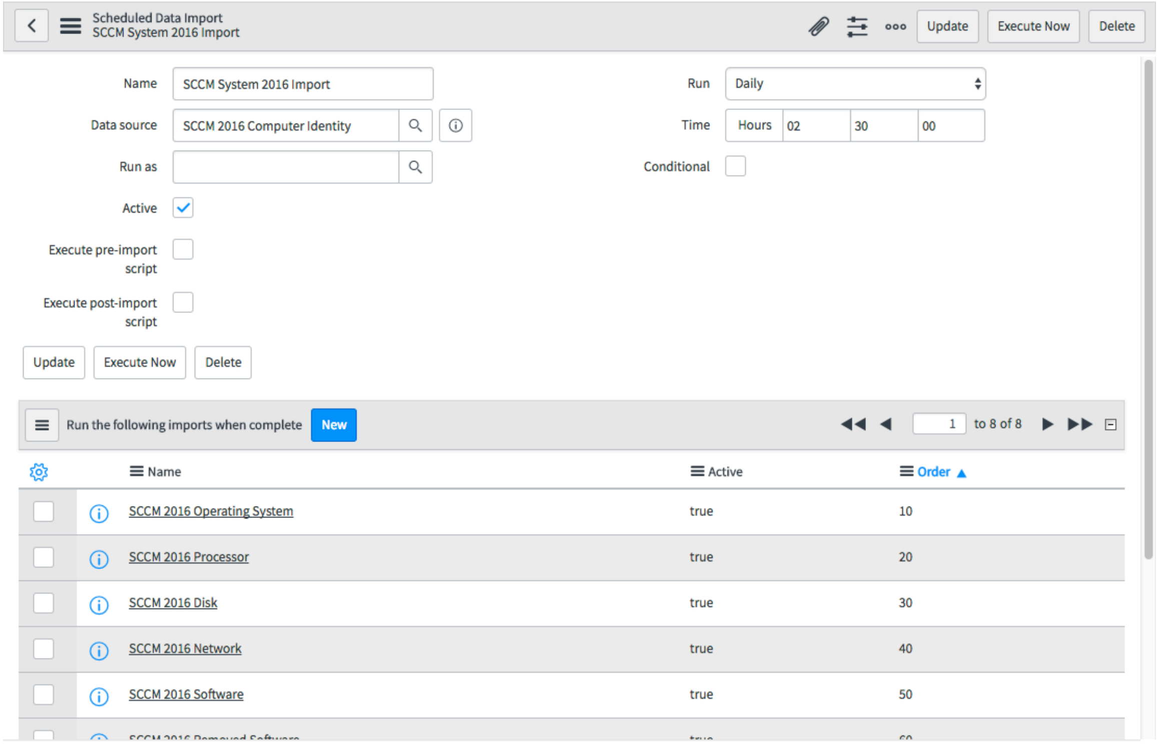

5. Select a calendar interval (Day, Week, etc.) in the **Run** field and specify the time of day the import should run.
   Your selection displays additional fields required to complete the schedule.
6. Click **Update** to schedule the import or click **Execute Now** to perform the import immediately.

   After the instance imports data using the **SCCM <version> Computer Identity** data source, the other data sources run in the order configured to retrieve the remainder of the SCCM data.

Activate SCCM Asset Intelligence scheduled imports

To prevent duplicate software imports, activate either the Asset Intelligence (AI) or the non-AI scheduled import.

Before you begin

Role required: admin

About this task

Make sure to limit the software import frequency to be no greater than the frequency of the table cleanup that tracks deleted software in SCCM.
Asset Intelligence is supported in these SCCM plugins:

- Integration - Microsoft SCCM 2007
- Integration - Microsoft SCCM 2012 v2
- Integration - Microsoft SCCM 2016

> **Important:**
> To improve the performance of your initial SCCM import, you can prevent the system from checking against deleted software prior to the import date. Navigate to **Integration - Microsoft SCCM <version>** > **Data Sources** > **SCCM <version> Removed Software** and enter the current date in the **Last run datetime** field, using the format **yyyy-mm-dd 00:00:00.0**. This field is populated automatically for each subsequent run of the removed software data source, so it is not necessary to alter the value after the initial import.

Procedure

1. Navigate to the **System Import Sets** > **Administration** > **Scheduled Imports**.
2. Set **Active** to **false** for these data import schedules:
   - SCCM <version> Software
   - SCCM <version> Removed Software
3. Set **Active** to **true** for these data import schedules:
   - SCCM <version> Software (with AI)
   - SCCM <version> Removed Software (with AI)

Related topics

Configure the SCCM integration and schedule an import

Upgrade the SCCM integration version

If you are using an earlier version of an System Center Configuration Manager(SCCM) plugin, you can switch over to a later version to take advantage of new features.

Before you begin

Role required: admin

About this task

The ServiceNow SCCM integrations are self-contained and can exist independently. They each use their own import set tables, data sources and transform maps. However, all SCCM integrations will transform data into the same tables within the ServiceNow CMDB. To avoid the data being overwritten by another source:

- Use one SCCM integration and disable all other SCCM scheduled imports.
- Perform a full import to clear the cmdb_software_instance table, the cmdb_sam_sw_install table, and other tables of old SCCM data.

Note:
It is possible to configure each plugin to integrate with SCCM 2007 or 2012 because the mechanism of the integration is actually the same, which is to leverage Java Database Connectivity (JDBC) imports. However, a data source must be modified if it is used for an SCCM version for which it was not written. Use the plugin version that corresponds to the SCCM version for which the data source is intended.

To change the SCCM integration:

- Disable the current integration by deactivating the SCCM import schedule.
- Activate the new SCCM plugin.
- Reimport all the software records when you are switching to an integration that supports incremental imports of removed software.

To disable the SCCM import schedule:

Procedure

1. Navigate to **All** > **Integration - Microsoft SCCM 20xx** > **Scheduled Import**.
2. Clear the **Active** check box.

Scheduled Import form

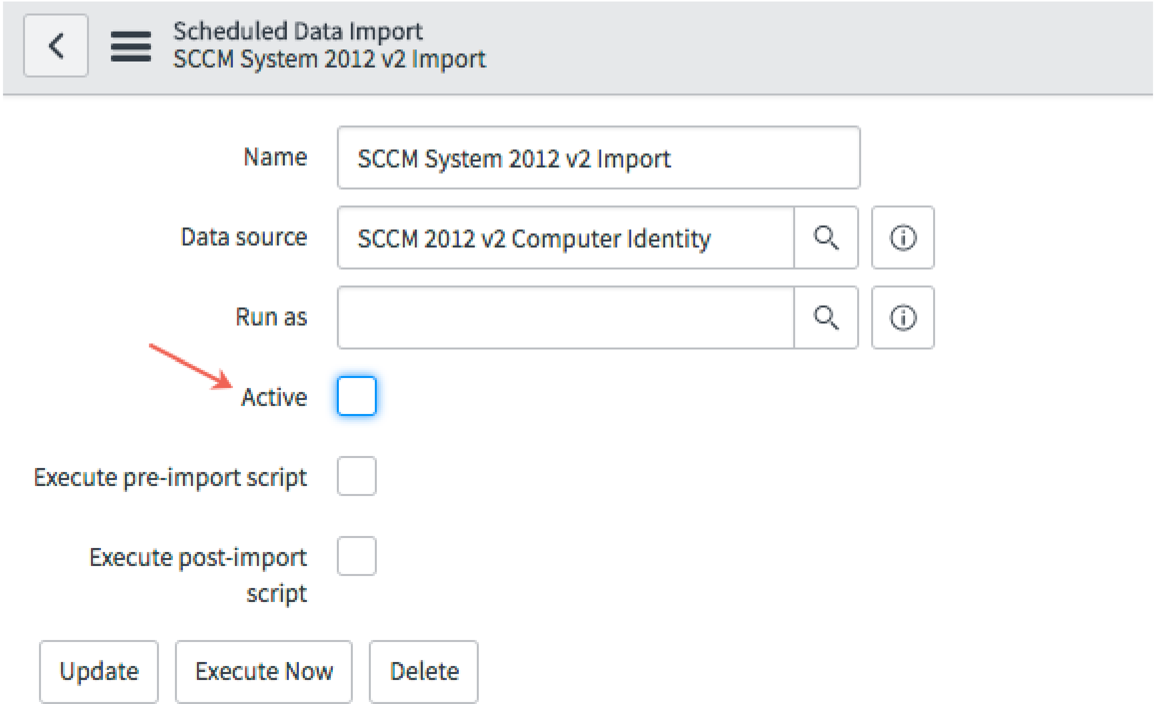

3. Click **Save** or **Update**.
4. To activate the new SCCM plugin, navigate to **System Definition** > **Plugins**.
5. Search on the name ***SCCM** to see all the available SCCM integrations plugins.
6. Activate the plugin.

Related topics

SCCM data import process and source tables

Migrate the Verizon eBonding Integration to a Production System

This topic outlines the tasks required for moving the Verizon eBonding integration from a Dev/Test environment to a Production environment.

This process requires coordination with your assigned Verizon implementation manager.

General Tasks

The following work can be performed any time:

- Obtain a production SOAP password from Verizon. Your Verizon SOAP user ID is the same as that used against the Verizon test broker.
- Create an integration user for Verizon in the ServiceNow Production instance. Use the same user name and password that was created for the development instance, so that the communication between Verizon and ServiceNow is properly authenticated.

Migrating to a Production Environment

Perform these tasks in order.

- Task 1: Request the Verizon eBonding plugin for the production instance. When the plugin is installed, the only immediate user interface change is an additional **Category** field named **VZ eBonding**. All of the triggers for the integration are dependent on the category being set to **VZ eBonding**. Enabling the plugin on the production instance should not affect your users. If you wish to take additional precautions, disable the **VZ eBonding** category from the list when the plugin is activated for the instance.
- Task 2: Set up the Verizon certificate and keystore. Copy these directly over from the Dev instance or regenerate them so that you have separate versions for the Dev and Production instances.
- Task 3: Configure the integration properties to work with the Production Verizon eBonding system. This is userID, password, and other properties provided by your Verizon representative.
- Task 4: Change all the Verizon eBroker SOAP Message functions to use Verizon's production SOAP Endpoint. In most cases, the endpoint URL will be: https://pubwebsvc.vzbi.com:443/Trouble_Management_v3r0. Verify this URL with the Verizon implementation specialist.
  - In the **SOAP Message Functions** related list, click each function, clear the **Lock** check box, and enter the production endpoint.
  - Configure the **SOAP Message Functions** related list and add the **SOAP endpoint** column. In the **SOAP Message Functions** list view, unlock and update the endpoint for all records by Edit multiple records in a list using the list editor.
- Task 5: Apply any customizations. If customizations to the integration on the Dev system were captured in an update set, apply that update set to the Production system.
- Task 6: Provide Verizon with the production SOAP Endpoint. The SOAP endpoint should look something like this: https://yourproductioninstance.service-now.com/ETMSPublish.do?WSDL.
- Task 7: Coordinate a Production turn up with your Verizon representative to do basic testing.

> **Warning:**
> Do not create test tickets against the production system without working with the Verizon team. Any tickets created are treated as real tickets.

Legacy: Altiris integration 2.0

The Altiris integration is deprecated in the Istanbul release.

The Altiris integration a one direction import of the Altiris data into ServiceNow CMDB (Configuration management database). The integration keeps the ServiceNow CMDB up to date with Altiris SQLServer database.

Data Import

Relevant data is imported from the Altiris database to the CMDB. The Altiris database is not written to, it is considered an authoritative source. The import is achieved using a JDBC connection via the MID Server.

Altiris integration application

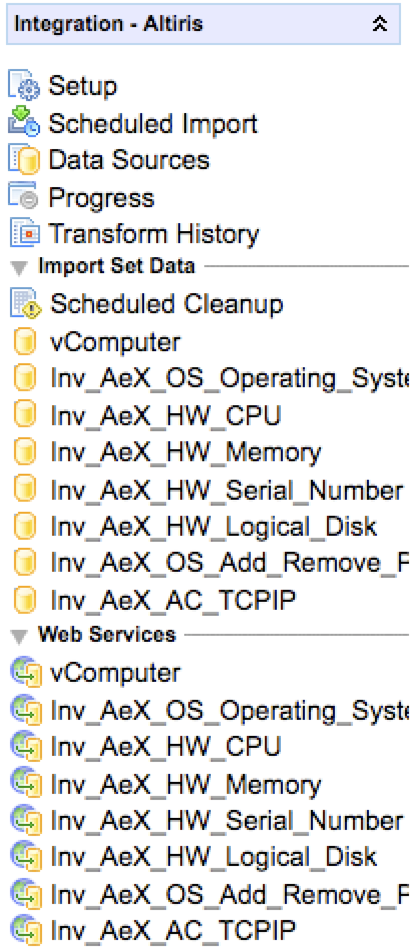

Information pulled from Altiris includes:

- PC Hardware (Model, Manufacturer, memory, clock speed, number of CPUs, etc.)
- Operating System Information (Name, Service Pack)
- Printers
- Disk information (physical, network, and logical)
- Network (IP Address, Netmask)

Configuration and Operational Modules

Enabling this integration will create the Integration - Altiris application.

The following are the configuration and operational modules for this integration.

Module	Description
Setup	• Configure the data sources from one form ◦ Provide the Database Server IP Address ◦ Provide the Database Name ◦ Provide the Database User ID and Password (this will need to be created on SQL DB side, see this article for help with

Module	Description
	this step: http://technet.microsoft.com/en-us/library/aa337562.aspx ◦ Find and select the MID Server • Specify Database server settings and MID server • Test configuration
Scheduled Import	Schedule the execution of the import or import immediately
Data Sources	A list of the pre-configured data sources defining the external CMDB database
Progress	A historical list of progress on scheduled imports
Transform History	A historical list of transformations performed during scheduled imports

Warning:
If you have activated an existing integration of the previous version:

1. Activating the 2.0 plugin does not "add to" or "remove" anything from the existing integration.
2. If transitioning from the old integration to this new one, considerations need to be given to customizations already done eg. mapping enhancements or using different coalesce values, these will have to be re-implemented.
3. Both plugins could run at the same time, provided data is coalescing the same way - until there is no need for the older plugin at which time it can be turned off.

Supported Versions

The Altiris integration only supports Altiris version 6.5. The integration does not currently Altiris version 7.0.

Legacy Import set data for Altiris

Learn the module names displayed by the Altiris import set data.

The Altiris import set data section shows a list of Import sets tables used in containing data retrieved from using JDBC to query the Altiris database. The module names (hence the import set tables they point to) match Altiris's table names and structure that it is loading from.

Module names

Note:
Functionality described here requires the Integration - Altiris 2.0 plugin.

- Scheduled Cleanup (Configure a schedule to cleanup/delete import set data that have already been transformed)
- vComputer
- Inv_AeX_OS_Operating_System
- Inv_AeX_HW_CPU

- Inv_AeX_HW_Memory
- Inv_AeX_HW_Serial_Number
- Inv_AeX_HW_Logical_Disk
- Inv_AeX_OS_Add_Remove_Programs
- Inv_AeX_AC_TCPIP

When viewing each of these table lists, at the end of the list you have links to other operational functions of the import set.

Altiris integration application

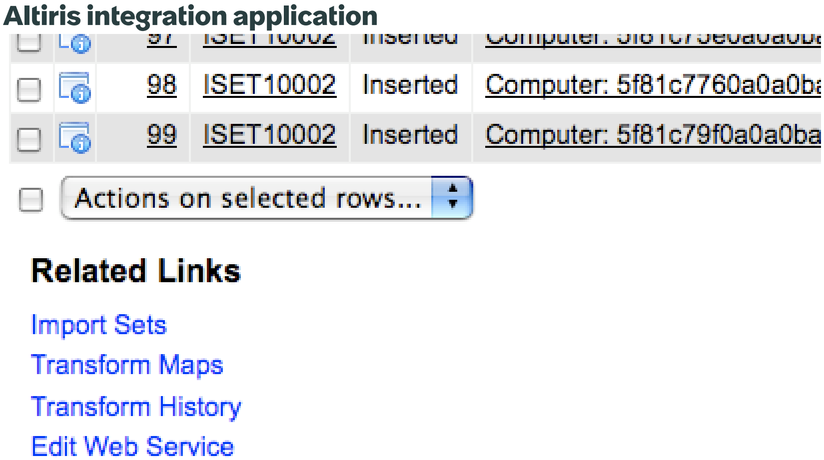

- Import sets
- Transform maps
- **Transform History** - Log of completed import operations (where an import set was transformed into a table).
- Web service import sets

Related topics

Legacy: Web services import set tables for Altiris

Legacy: Altiris integration 2.0

Legacy: Web services import set tables for Altiris

This topic will list the modules that define the web service import set tables - the schema for the import set tables that are receiving the JDBC import.

From each web service, you can add/remove fields as well as access the transform maps to make modifications.

Note:
Functionality described here requires the Integration - Altiris 2.0 plugin.

Import set tables

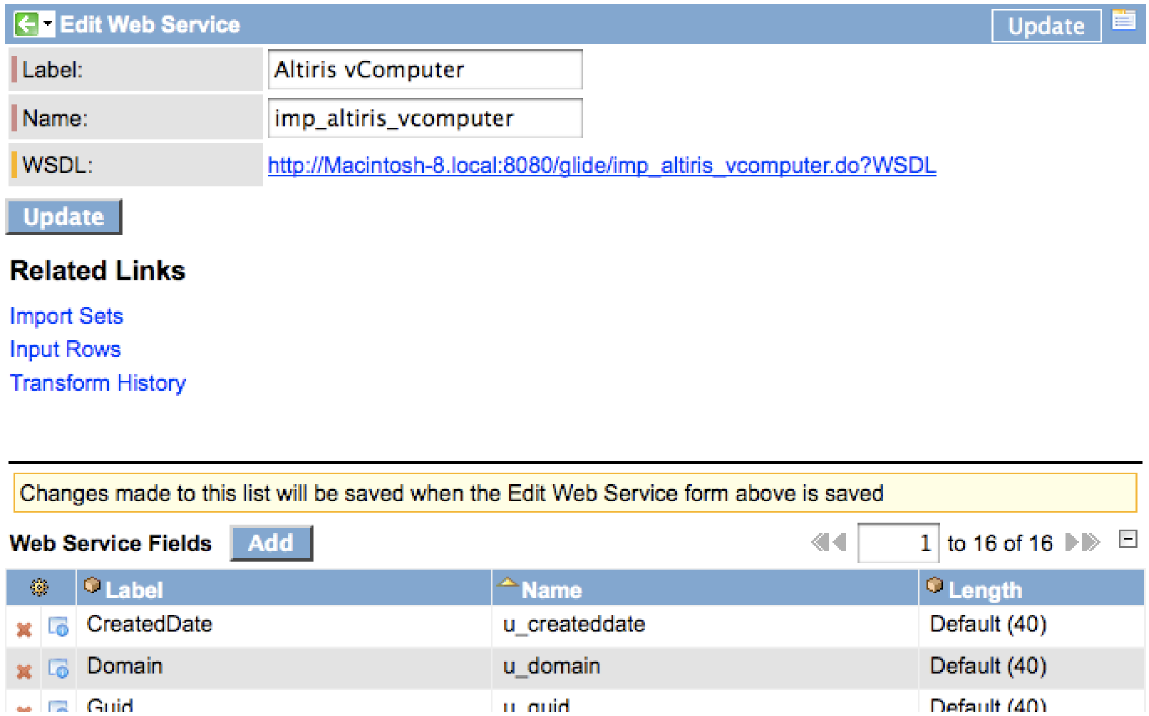

Import set tables

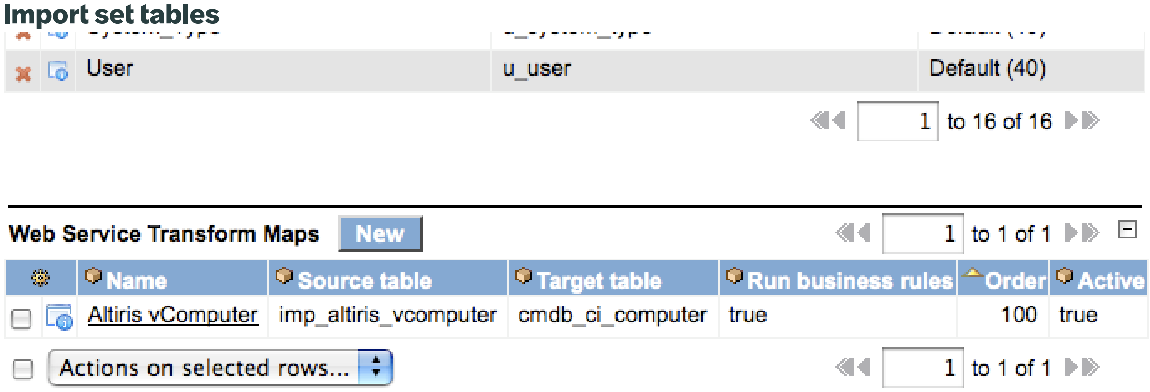

Related topics

Legacy Import set data for Altiris

Legacy: Altiris integration 2.0

Supported integration interfaces

ServiceNow provides a number of interfaces to be able to directly integrate with the platform. These interfaces are considered part of the platform and are provided at no additional charge.

Supported Integration Interfaces

Interface
Email
JDBC
JSONv2 Web Service
LDAP integration

Supported Integration Interfaces (continued)

Interface
SOAP web service
REST API
SAML
Digest token authentication
ODBC driver
Data Export
CTI
Syslog probe

Computer Telephony Integration

Computer Telephony Integration (CTI) is accomplished by the CTI client on the user machine sending a URL to the instance.

The URL must have the following components:

1. The base URL. For example: `https://<instance name>.service-now.com/cti.do?` would get to the instance and ask for CTI processing.
2. Parameters identify what parts of the incident form to display.
   - *sysparm_caller_name*=name where 'name' is the name for a user.
   - *sysparm_caller_phone*=phone where 'phone' is the user's phone number. Either a name or phone should be provided if you want to identify the user on the call. Other parameters may be supplied to identify the user as discussed later.
   - *sysparm_task_id*=taskID where 'taskID' identifies an existing issue that the caller is calling about.
   - *sysparm_view*=view where 'view' is the name of the view to be used to display the data.
   - *sysparm_xxxx*=value where 'xxxx' is the name of a field within the 'incident' record that should be populated with the specified 'value'. For example sysparm_priority=1 would result in the priority field set to value of 1 when the new incident screen is shown.
   - *sysparm_cti_rule*=name where 'name' is the name of a function to be invoked for CTI processing rather than using the default script. The function must be defined in a sys_script entry marked client callable. If the function needs to insert, update, or delete any GlideRecord(s), it must call a separate non-client callable function to perform the update(s).

   Note:
   While the CTI Processing script has been changed to be client callable, the code implementing the task view has been commented out. You must implement a new non-client-callable function for the code that performs the `task.update()`.

   Note:
   To make a script client-callable you must check the client-callable checkbox on the form that displays when the sys_script entry is displayed. The client-callable checkbox might not show by default. To show the client-callable checkbox, you may need to modify the fields that show on the form using the gear icon and slushbucket mechanism.

   Parameters on the URL are available to the business rule as global values. For example:

```
var name = sysparm_caller_name;
```

The business rule you specify must return the URL for the pop-up screen, and set the 'answer' global variable.

An example URL to bring up a screen shot for user Don Goodliffe would look like the following. `https://<instance name>.service-now.com/cti.do?sysparm_caller_name=Don%20Goodliffe` Multiple sysparm parameters can be used, separated by ampersands (&).

CTI Processing script

When the *sysparm_cti_rule* parameter is not specified, the system uses the CTI Processing script to provide the following functionality.

The CTI Processing script does the following:

1. Tries to identify the user by the *sysparm_caller_name* value if it was supplied.
2. If no user has been found, the script tries to identify the user by the *sysparm_caller_phone* value if it was supplied.
3. If a user has been identified then one of the following is done
   a. If the user has open incidents, the popup screen shows information about the current caller and all the user's open incidents.
   b. If the user does not have any open incidents, the popup screen shows a new incident with information provided in the URL shown.
4. If a user was not identified and a taskID is given and the taskID exists, then nothing happens. The code to handle this case is commented out. If you want the popup screen to show the details for the task, you must modify the CTI Processing script to put the functionality in a separate non-client-callable function.

Integrating ServiceNow with your Intranet

You can add a ServiceNow login link to your intranet.

To deliver ServiceNow content from a web page, see Service Portal.

JDBCProbe

A JDBC probe runs on the MID Server to query an external database via [JDBC] and returns results to ServiceNow.

About this task

Probes interact with the MID Server via the ECC Queue, therefore the response of a JDBC probe returns as an XML payload in an "input" ECC Queue record. By default, each response payload will contain up to 200 returned rows, this value can be modified by setting the probe parameter *jdbcprobe_result_set_rows* to the desired number.

Activating the Plugin

Contact Customer Service and Support to activate the **Integration - JDBC** (com.snc.integration.jdbc) plugin.

Procedure

1. Navigate to **System Definition > Plugins**.
2. Right-click the plugin name on the list and select **Activate/Upgrade**.

If the plugin depends on other plugins, these plugins and their activation status are listed.

3. **Optional:** Select the **Load demo data** check box.

   (Optional) Some plugins include demo data—sample records that are designed to illustrate plugin features for common use cases. Loading demo data is a good policy when first activating the plugin on a development or test instance. You can load demo data after the plugin is activated by repeating this process and selecting the check box.

4. Click **Activate**.

Direct JDBC Probe

A direct JDBC probe specifies all the parameters necessary in the outbound ECC Queue XML payload.

It is a standalone probe that is decoupled from a JDBC Data Source and does not insert into an import set. The required field in the ECC Queue record for a direct JDBC probe is **Topic** and it must equal **JDBCProbe**. The **Source** field is reserved for the `sys_id` of the data source record in the case of a JDBC data source.

Note:
If you have ServiceNow Discovery enabled on your platform, add the *`skip_sensor`* parameter to the probe to avoid the discovery sensors from processing the probe and resulting in a "No sensors defined" error.

XML Structure

A direct JDBC probe has the following XML payload structure

```
<parameters><parameter name = "skip_sensor" value =
 "true" /><parameter name = "parameter_name" value =
 "parameter_value" />
  ...
  <parameter name = "work" ><select  ... >
      ...
    </select></parameter></parameters>
```

Related topics

JDBCProbe

JDBC Probes via Data Source

JDBC probes are executed via a JDBC data source when an import is running against the data source.

A JDBC data source JDBC probe is described by the JDBCProbe Topic and the `sys_id` of the data source in the Source field of the ECC Queue output record.

The data source record would look like this

JDBC Data Source

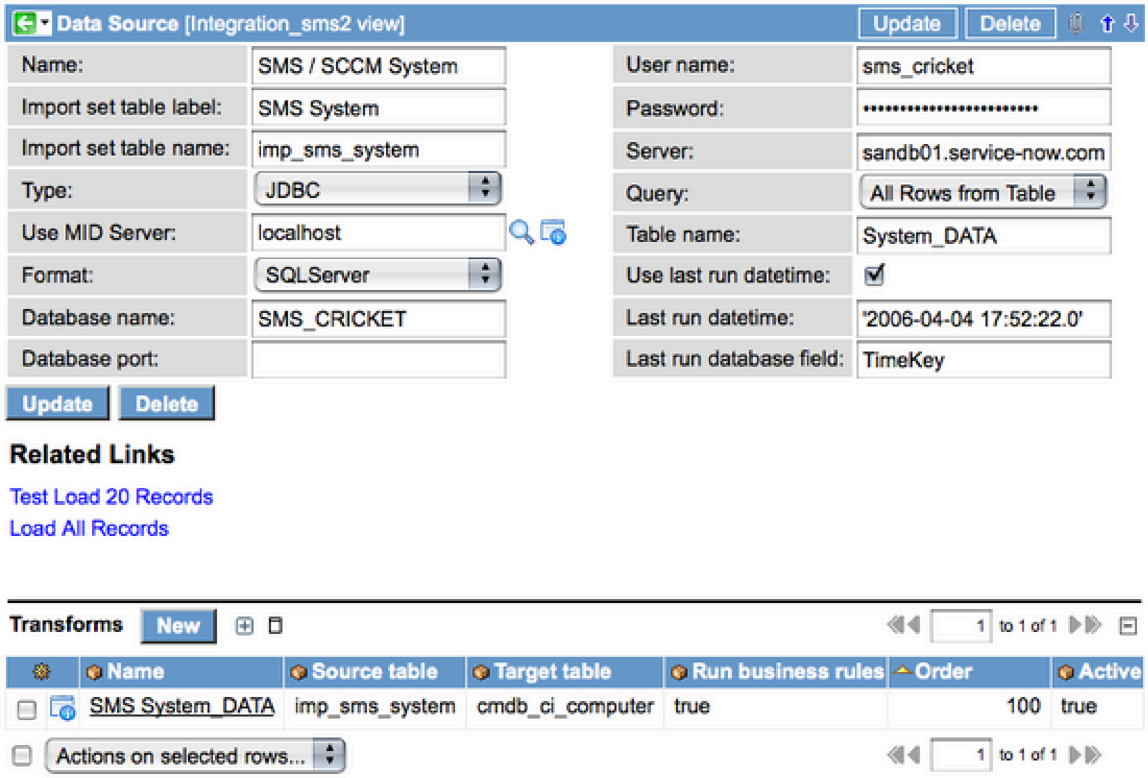

The following ECC Queue output probe will be created when you load from the data source.

JDBC Data Source Probe

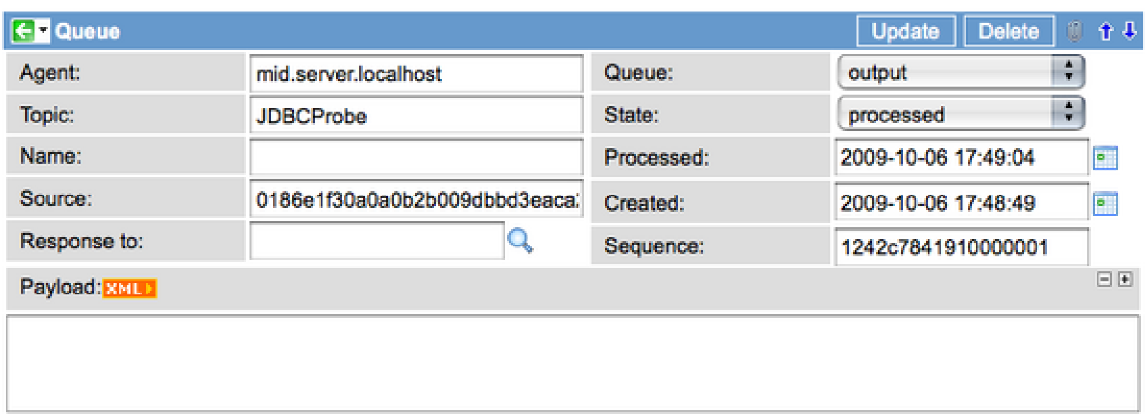

Related topics

JDBCProbe

Select * JDBC Probe short cut

Alternatively, you may specify a `table_name` parameter instead of a work element and the following query could be executed.

```
select * from <table_name>
```

For example, using the following XML payload

```
<?xml version= "1.0" encoding= "UTF-8" ?><parameters><parameter
 name = "jdbc_driver" value =
 "com.microsoft.sqlserver.jdbc.SQLServerDriver" /><parameter
```

```
 name = "connection_string" value =
 "jdbc:sqlserver://
xxx.service-now.com;databaseName=SMS_CRICKET;user=sms;password=s
ms" /><parameter name = "table_name" value =
 "System_DATA" /></parameters>
```

Counting Rows

To count the number of rows in a table, you can indicate a select count(*) type query by including the count_rows parameter with a value of true.

For example

```
<?xml version= "1.0" encoding= "UTF-8" ?><parameters><parameter
 name = "jdbc_driver" value =
 "com.microsoft.sqlserver.jdbc.SQLServerDriver" /><parameter
 name = "connection_string" value =
 "jdbc:sqlserver://
xxx.service-now.com;databaseName=SMS;user=sms;password=sms" /
><parameter name = "count_rows" value = "true" /><parameter name
 = "table_name" value = "System_DATA" /></parameters>
```

The resulting response XML payload

```
<parameters>
  ...
  <result query = "SELECT count(*) as row_count
 FROM System_DATA" ><row id = "1" ><row_count class
 = "java.lang.Integer" length = "11" type = "4"
 >2312</row_count></row></result>
  ...
</parameters>
```

Related topics

JDBCProbe

Parameters

The following parameters are available in a direct JDBC probe.

Parameter	Description
jdbc_driver	Required. Java class name for the JDBC driver to use, the currently supported drivers are: • Oracle: `oracle.jdbc.OracleDriver` • Microsoft SQL Server: `com.microsoft.sqlserver.jdbc.SQLServerDriv` • MySQL: `com.mysql.jdbc.Driver`

Example

```
com.microsoft.sqlserver.jdbc.SQLServerDriver
```

Parameter	Description
connection_string	Required. JDBC connection string/URL for defining the connection, usually contains information about the database server and name, the user ID and password for conenctiong to the database. The syntax is vendor specific, refer to the following links for reference: • Oracle: `[jdbc:oracle:thin:<username/password>@<database>]` • Microsoft SQL Server: `[jdbc:sqlserver://localhost;user=MyUserName;password=*****;]` • MySQL - `[jdbc:mysql://localhost/database?user=username%26password=passwd]`

Example

```
jdbc:sqlserver://xxx.service-now.com;
databaseName=SMS;
user=sms_user;
password=sms_password;
```

Parameter	Description
work	Required if not using the *table_name* parameter short cut. Parent element of an XML fragment describing the SQL command to execute.

Example

```
...
  <parameter name= "work" ><select table = "System_DATA" where =
 "InstanceKey=692" ><MachineID /><SMSIDO /></select></parameter>
  ...
```

Parameter	Description
query_timeout	Optional during SELECT. Number of seconds the JDBC driver will wait for a `query (SELECT)` to complete. Zero means no timeout. If timeout is exceeded, the integration considers the JDBC result inaccessible and places it in an error state.

Parameter	Description
table_name	Optional (except required during `count_rows == true`). Alternate way of executing a `select *` query; instead of using the *work* element. Equivalent to executing `select * from <table_name>`.

Parameter	Description
count_rows	Optional. Indicates whether a `select count(*)` query should return the number of rows (true/false). The result of the count will be returned as a *row_count* element in the *result* element.

Parameter	Description
query	Optional. Type of query. Possible choices are "All Rows from Table" or "Specific SQL". If "Specific SQL", the sql_statement will be required to specify the SQL statement.

Parameter	Description
sql_statement	Optional. Use a specific SQL query. The presence of this element executes a direct query specified in the value attribute.

Example

```
...
  select * from any_table where id = 123
  ...
```

Parameter	Description
skip_sensor	Optional. Determines if Discovery will attempt to process the ECC input from the JDBCProbe. Default = true

For example, to query a table using a direct JDBC probe, requires the following parameters:

- JDBC driver class name
- JDBC connection string
  - Database server
  - Database name

- User name
- User password

- The table name

Queue form



Related topics

JDBCProbe

Using the Work Element

The work element encodes SQL statements to be executed by the probe.

The following are the valid **work** child elements.

- select
- update
- insert
- delete

SELECT

Retrieve rows from the specified table in the database specified by the JDBC connection string. The simplest **select** work just specifies the table name and will retrieve all fields in a row, for example:

```
...
 <parameter name= "work" ><select table =
"System_DATA" /></parameter>
 ...
```

To specify a search criteria, specify the **where** attribute on the **select** element, for example:

```
...
 <parameter name= "work" ><select table = "System_DATA" where =
"InstanceKey=692" /></parameter>
 ...
```

To specify the fields you want returned, embed the fields as child elements of the **select** element, for example:

```
 ...
 <parameter name= "work" ><select table = "System_DATA" where =
"InstanceKey=692" ><MachineID /><SMSIDO /></select></parameter>
 ...
```

To use a direct SQL statement, specify it in the **query** parameter

```
...
 <parameter name= "query" value= "Specific SQL" /><parameter
name = "sql_statement" value = "select * from any_table where
value='test'" />
 ...
```

UPDATE

An **update** SQL can be executed by specifying the **table**, optional **where** clause and required child elements for the fields and their values to update.

```
...
<parameter name= "work" ><update table =
 "alerts.status" where = "ServerName %= 'NCOMS'
 AND ServerSerial=3935" ><Agent>ServiceNow -
 INC10020</Agent><URL>http://
Macintosh-9.local:8080/glide/incident.do?sys_id=17a31f380a0a0bae
0048ca875c8891d0</URL><Severity quoted = "false"
 >3</Severity><Acknowledged quoted = "false"
 >0</Acknowledged></update></parameter>
...
```

INSERT

An **insert** SQL can be executed by specifying the **table**, and required child elements for the fields and their values to insert.

```
...
<parameter name= "work" ><insert table
 = "alerts.status" ><Agent>ServiceNow -
 INC10020</Agent><URL>http://
Macintosh-9.local:8080/glide/incident.do?sys_id=17a31f380a0a0bae
0048ca875c8891d0</URL><Severity quoted = "false"
 >3</Severity><Acknowledged quoted = "false"
 >0</Acknowledged></insert></parameter>
...
```

DELETE

A **delete** SQL can be executed by specifying the **table**, and optional **where** clause.

```
...
<parameter name= "work" ><delete table = "alerts.status" where =
 "ServerName %= 'NCOMS' AND ServerSerial=3935" /></parameter>
...
```

Related topics

JDBCProbe

Build a search provider for your instance

ServiceNow Search Providers allow you search and our Forums from the IE and Firefox search bar.

About this task

In Firefox 3.x you can also assign a keyword to each Search Provider and access them from the address bar. For example, assign w to the wiki search provider and you can search the wiki for Business Rules by typing: `w business rules` in the address bar.

Search Bar

Procedure

1. Create an opensearch description document.
   You can review the OpenSearch standards for details on additional attributes of this document such as including an icon.

```
<?xml version="1.0" encoding="UTF-8"?>
<OpenSearchDescription
 xmlns="http://a9.com/-/spec/opensearch/1.1/">
 <ShortName>Demo Search</ShortName>
 <Description>Demo Search provider</Description>
 <InputEncoding>UTF-8</InputEncoding>
 <Url type="text/html"
 template="https://www.service-now.com/demo/nav_to.do?
 uri=incident_list.do?sysparm_query=active%3Dtrue^123TEXTQUERY3
21%3D{searchTerms}"/>
 </OpenSearchDescription>
```

2. Save the file to a web server with xml extension.
   The method used to install doesn't allow local file calls.
3. Create a simple html page to install the provider.

```
<a
 href="javascript:window.external.AddSearchProvider('http://
yourServer/yourFile.xml');
" title="MySearch" name="ServiceNow Custom Search">Add
 ServiceNow Custom Search Provider</a>
```

   You could try running the JavaScript command from the browser location box instead of using the html file. This works with FF not IE.

Syslog probe

The ServiceNow Syslog probe uses the MID Server to deliver log messages from a ServiceNow instance to another machine, such as a dedicated log server, using the syslog protocol over an IP network.

How the Syslog probe works

The syslog probe is launched by a ServiceNow Script includes (called Syslog) that can be invoked from a Business rules, event, or Orchestration activity and is launched by a MID Server. A syslog server or any server that can receive messages using the syslog protocol must be installed on the recipient (target) machine. Typically, a dedicated log server in the network is configured to receive all internal syslog messages. Some products that accept syslog messages are:

- ArcSight
- Splunk
- LogLogic
- syslog-ng

Example

The Acme Corporation wants to send a log message from their ServiceNow instance to an ArcSight syslog server inside their corporate firewall each time a user login fails. The system administrator uses the login.failed event to trigger a business rule that invokes the Syslog Script Include each time a login fails. Acme's MID Server checks the ECC Queue for work and picks up the syslog probe, which contains the log entry. The MID Server then sends the log message to the ArcSight server, which gathers log messages from all the machine in the internal network.

Syslog probe diagram

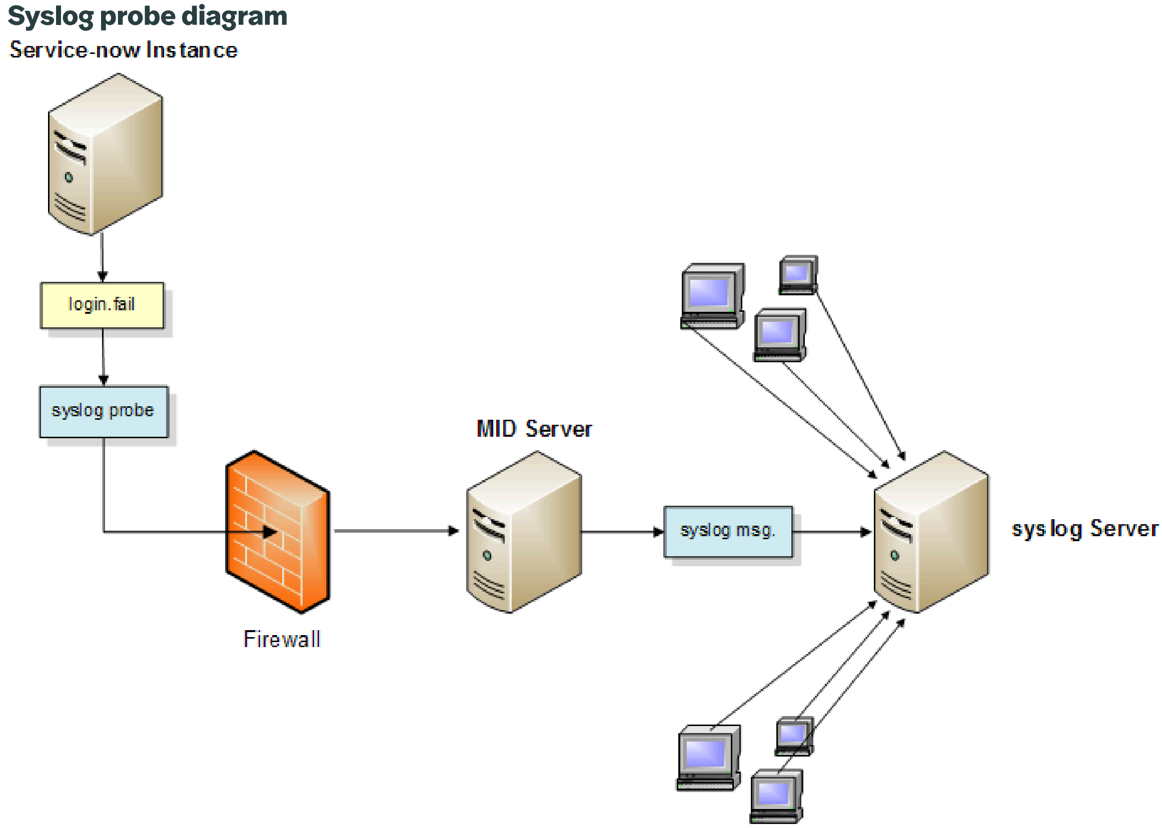

Code sample

The following code sample, included in an event or a business rule, directly calls the Syslog Script Include and sends a syslog message to a designated syslog server:

```
var sl = new Syslog('FQDN of your syslog server',
 'mid.server.Eclipse', 16);
sl.log('This is a sample log message', 6);
```

This code does the following:

- Sends the log message to facility **16**
- Sets the priority to **6** (informational)
- Sends the message to your syslog server
- Launches the probe via the MID Server named "Eclipse"

Managing integration sessions

Manage access to an instance by configuring how long integration sessions last before expiring.

Similar to user sessions, you can secure your instance by configuring when to time out an integration session after a period of inactivity and a maximum active session time. By default, integration session time out only after one minute of inactivity.

Configure a maximum active time for integration sessions

Secure your instance by enforcing a maximum time for sessions regardless of integration activity.

Before you begin
Role required: admin

About this task
By default, sessions expire only after a period of inactivity. Enforcing a maximum active session time ends sessions regardless of if an integration has been active recently. The active session timeout should be greater than the value configured for the inactive session timeout. For example, if sessions are configured to time out after 30 minutes of inactivity, the active session timeout should be greater than 30 minutes.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Select **New**.
3. On the form, fill in the fields.

System Property form

Field	Description
Name	`glide.active.session.timeout.invalidate.se`
Description	`Specifies whether integration sessions can be invalidated by configuring a maximum active session time.`
Type	**true | false**
Value	`true`

4. Select **Submit**.
5. From the System Property [sys_properties] list, open the *glide.integrations.active.session.life_span* property.
6. In the **Value** field, enter the number of minutes before integration sessions time out regardless of user activity.
   The value should be greater than the value of *glide.integration.session_timeout* for inactive integration session timeout. By default, the inactive session timeout for integrations one minute.
7. Select **Update**.

Result

Integrations are logged out of the instance after the time specified and their credentials must be entered again to access the instance. This setting helps secure an instance using integrations, such as the REST API.

Related topics

Configure a maximum active time for user sessions

Modify integration session timeout after inactivity

Specify when to time out integration sessions after a period of inactivity.

Before you begin

Role required: admin

About this task

By default, after one minute of inactivity, the platform ends an integration session and logs the user out automatically.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Select **New**.
3. On the form, fill in the fields.

System Property form

Field	Description
Name	`glide.integration.session_timeout`
Description	`Sets the inactive session timeout for integration sessions, in minutes.`
Type	**string**
Value	Specify the number of minutes.

4. Select **Submit**.

Related topics

Modify user session timeout after inactivity

Domain separation in third-party application and data source integration

This is an overview of domain separation and integration of third-party applications and data sources. 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-Standard

The Basic-Standard level of domain separation is achieved through data visibility as well the integration design.

Basic

- There is business logic to ensure data goes into the proper domain for the application's service provider use cases.
- In the application, the user interface, cache keys, reporting, rollups, aggregations, and so on, all consider domain at run time.
- The owner of the instance needs to be able to set up the application to function normally across multiple tenants.

Standard

- Includes **Basic** level
- Business logic: Processes can be created or modified per customer by the service provider. The use cases reflect proper use of the application by multiple service provider customers in a single instance.
- The owner of the instance needs to be able to configure MVP business logic and data parameters per tenant as expected for the specific application.

Overview

Team should define a few examples with data-driven properties. There should be three different buckets with short description for each one

Related topics

Domain separation for service providers

Notifications

Use ServiceNow® Notifications to manage system email, create system notifications, and configure how your system responds to inbound email.

Email Administration With email notifications use email properties to define your instance and set up email accounts	System Notifications Set up notifications alerts for record changes via email, SMS text message, push notification, or messaging app	Provider Notifications Send notifications directly to recipients via the provider on supported channels
Email Client Send email directly from any record, such as an incident, change request, or problem record	Inbound Email Configure your instance to send replies, create incidents, or update records automatically in response to inbound emails	Notification Preferences With notification preferences you can manage and choose the notifications you want to receive

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Email Administration

All email notifications use the email properties that you define for your instance and the email accounts that you set up. Your email service can also affect the successful transmission of incoming and outgoing email.

Get started

Explore	Configure	Email Accounts
Explore email administration	Configure email administration	Create, filter and service email accounts
Email Encryption	**Reference**	
Secure/Multipurpose Internet Mail Extensions (S/MIME) is an end-end encryption protocol to support data confidentiality, authenticity, and integrity	References for email size limit and email bounce	

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Explore email administration

Learn about email administration.

Email Administration benefits

Benefit	Feature
Define your instance and the email accounts that you set up. Your email service affect the successful transmission of incoming and outgoing email.	Email administration
Email messages can be seen in the System Mailboxes menu, which gives you access to the system Inbox, Outbox, and Sent mail box.	System mailboxes

Benefit	Feature
Email accounts store the connection details and credentials the system uses to access external email servers and services.	Email accounts
Specify which inbound emails to ignore or move to a particular mailbox.	Email filters
The Email Service API provides endpoint to install the Email API on the instance.	Email service
Limit sending emails to IDs that are known to generate bounces as per RFC3463 and allowing the admins to control blocked ids	Email bounce management

Email administration

All email notifications use the email properties that you define for your instance and the email accounts that you set up. Your email service can also affect the successful transmission of incoming and outgoing email.

Email accounts

Email administrators set up email accounts to allow the system to connect to external mail services such as POP3, SMTP, or IMAP servers. You can use the ServiceNow - provided email accounts or create your own accounts for your own email services.

Email properties

Use Email Properties to configure settings for inbound and outbound email. Email properties apply to all email accounts.

Email service size restrictions

Your email service, whether you are using a standard ServiceNow email configuration or an alternate email configuration, determines the successful transmission of incoming or outgoing emails based on a maximum email size. An email configuration consists of your instance, email server, and email client.

- Standard email configurations – For instances that use the standard email configuration, the ServiceNow mail servers cannot send or receive emails larger than approximately 50 MB (before encoding), including the email header, body text, and attachments. Email messages are encoded by an email client, which increases total email message size. The maximum ServiceNow mail server size restriction is 75 MB (after encoding). The maximum email size limit is enforced regardless of any configured attachment size limits.
- Alternate email configurations – If your email service uses a company-owned or third-party server, the approximate size restriction of 50 MB (including the email header, body text, and attachments) may still apply. However, your service might support different maximum total file sizes for inbound and outbound emails. Check with the email administrator of your email service to verify email size limits.

For details on how encoding can affect email message size and email transmission, see KB0521772.

Email service availability

Email service is intentionally shut down on instance clones to prevent resending already delivered email. Upgrades no longer require an extended interruption of email service. Instances

continue to process email during an upgrade. For a possible explanation for undelivered emails, see this blog post by a ServiceNow Technical Support Engineer in the ServiceNow Community.

For instructions on creating and sending custom emails when events on the instance occur, see the tasks in Email and SMS notifications.

Instance-to-instance communication via email

Use Web services to communicate between two instances.

Designate untrusted and trusted email domains

Using Instance Security Center, you can monitor the blocked and allowed incoming email metrics for your instance. For more information, see Designate untrusted and trusted email domains.

Next steps after enabling email

After enabling email on your instance, consider performing several of these important tasks.

- Test the email configuration by routing all email to a single user (set the `glide.email.test.user` property).
- Review the baseline notification categories, email notifications, and templates to determine if they meet your business needs.
- Review the baseline inbound email actions to determine if they meet your business needs.
- Determine if you want to use email layouts to add consistent content elements.
- Determine if you want to use email filters to restrict the email the instance receives.
- Determine if you want to implement a retention policy to archive and destroy email at certain intervals.
- Determine what kind of watermarks outbound email uses to associate records with email messages.
- Determine if you want to create users when the instance receives an email from an unrecognized user.
- Set the precedence of outgoing mail. By default, the instance sends email with a precedence of bulk (set the `glide.smtp.precedence_bulk` property).
- Consider preventing untrusted users from triggering inbound actions to prevent unwanted email from affecting your instance.
- Implement a spam filter to restrict unwanted messages sent to your custom email addresses.

System mailboxes

Email messages can be seen in the System Mailboxes menu, which gives you access to the system Inbox, Outbox, and Sent mail box.

Each provides a filtered view of emails to let you see only what you need to at any given time, including list fields applicable to each. The system mailboxes menu shows your current POP and SMTP status.

The System Mailbox is hosted by ServiceNow, who have sole access to the mailboxes. By default, once the instance pulls an email message, it is deleted from the mail server and stored in the application on the Email [sys_email] table.

For information about the fields that are shown in the System Mailbox, see System email log and mailboxes. You can configure the layout of a system mailbox to show any of the email log fields, not only the ones that are shown by default.

- **Inbound emails**: All inbound mail is placed into the Inbox until it is processed. After it is cleared, the email moves to the Received state. If the email message matches the criteria in an inbound email action, the email is changed to Processed. If not, it is changed to Ready. If the system is restarted for any reason (such as during a system upgrade), all inbound mail waits on the external mail server until the system can request delivery.
- **Outbound emails**: All outbound mail is placed into the Outbox until it is processed. Once cleared, it is moved to Sent (if sent) or Skipped (not sent, as in the case of no valid recipients). If the system is restarted for any reason (such as during a system upgrade), all outbound mail waits in the instance database until the system comes online, and the scheduler looks for mail to deliver.

Configure email administration

Setup basic and advanced email, configure email properties, activate email filter and services and grandt non admin user access to email APIs.

Basic Setup Basic email setup	Advanced Setup Advanced email setup	Email Properties Configure settings for inbound and outbound email
Activate Activate email filters and services	Access Grant non-admin users access to the Email API	

Basic email setup

All production instances can send and receive email using ServiceNow - provided resources. The instance has an email address of instance@service-now.com.

Network layout for standard email configuration

Below is an example of a basic email network with ServiceNow as the domain.

Basic email network

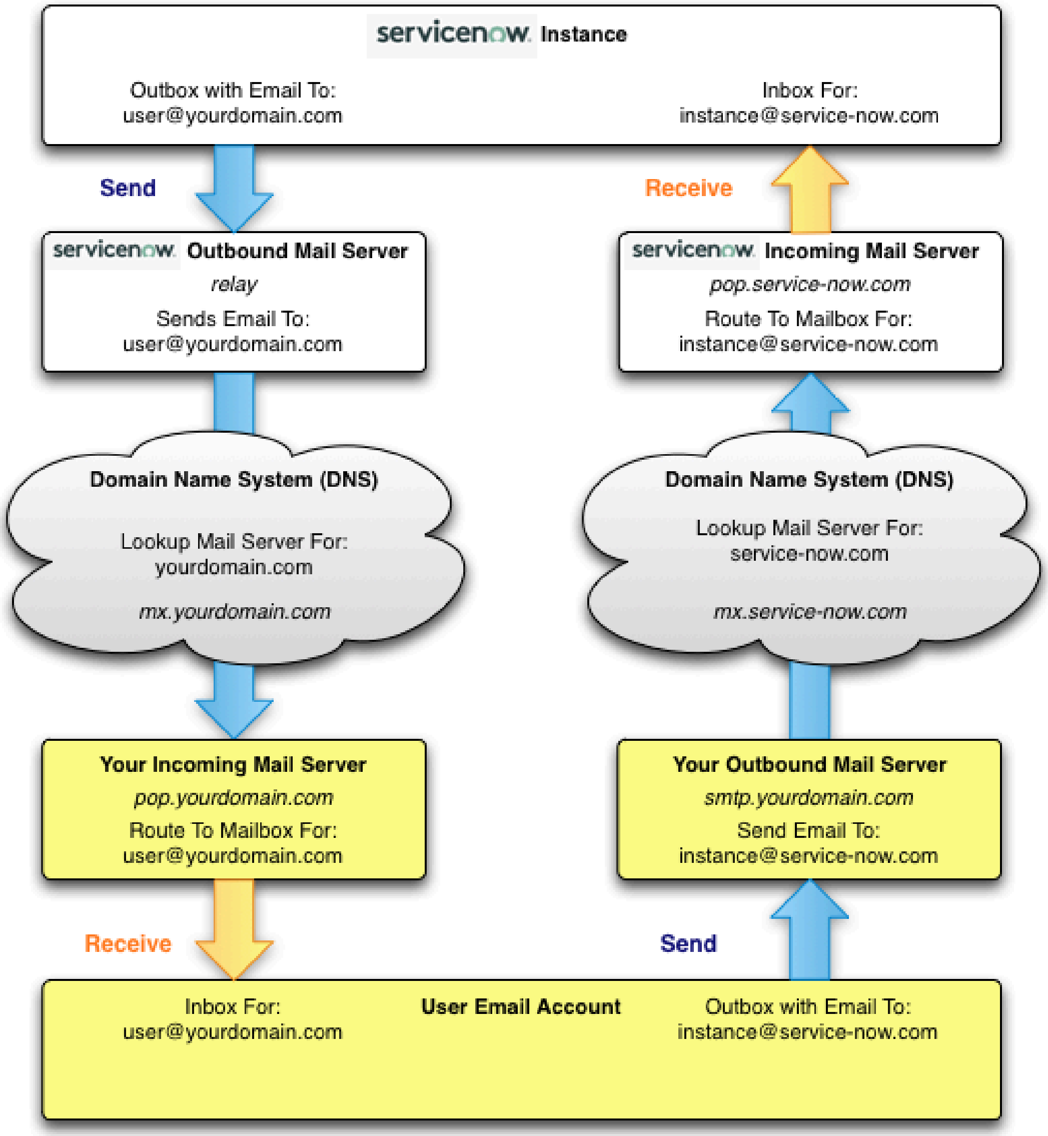

Basic email services and features

- Mail servers maintained by ServiceNow.
  - Encrypt mail with opportunistic TLS (Transport Layer Security) if supported by your mail servers.

    If your internal mail servers send and receive messages via a TLS-encrypted channel, ServiceNow mail servers support that communication.
  - Provide a dedicated mailbox for your instance.
- Pre-configured email accounts to connect to ServiceNow mail servers.
  - An SMTP account sends email to your primary Mail Exchange (MX) server from your instance email address of instance@service-now.com.
  - A POP3 account receives email sent to your instance email address of instance@service-now.com.

- High availability features from ServiceNow datacenters.
- Spam detection for incoming email.

Administrators who want to use basic email services can do so by enabling the email properties for sending and receiving email.

Enable basic email

Enable basic email to use ServiceNow - provided email servers and accounts.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **Email Properties**.
2. Configure these email properties and click **Save**.

Email property descriptions

Property section	Label	System property	Setting required
Outbound Email Configuration	Email sending enabled	glide.email.smtp.active	Yes
Inbound Email Configuration	Email receiving enabled	glide.email.read.active	Yes

Advanced email setup

With an advanced email setup, you can use your own SMTP server, POP3 server, or both.

Setting up your own email environment can be useful if you want to use existing filtering, retention, or compliance aspects of your internal email architecture. You can set up email in several ways:

- Use your own SMTP server to forward email to ServiceNow servers.
- Use your own SMTP server to send email.
- Use your own POP3 server to receive email.
- Use your own SMTP and POP3 servers to send and receive email.
- Use an OAuth 2.0-enabled SMTP server to send email from a third-party service.
- Use an OAuth 2.0-enabled IMAP server to receive email from a third-party service.

The following procedures assume that you enabled basic email properties.

Enable using your own SMTP server

Enable using your own SMTP server so that you can leverage the existing filtering, retention, or compliance aspects of your own SMTP server while also using the ServiceNow POP3 server.

Before you begin

- Role required: admin
- Email server required: SMTP
- Basic email properties: enabled

About this task

You can combine your own internal email architecture with the ServiceNow email architecture to handle email. The following diagram demonstrates how you would use your own SMTP server alongside the ServiceNow POP3 server.

Note:
Supports only one active SMTP account at a time (for outbound emails).

Sending email using your own SMTP server

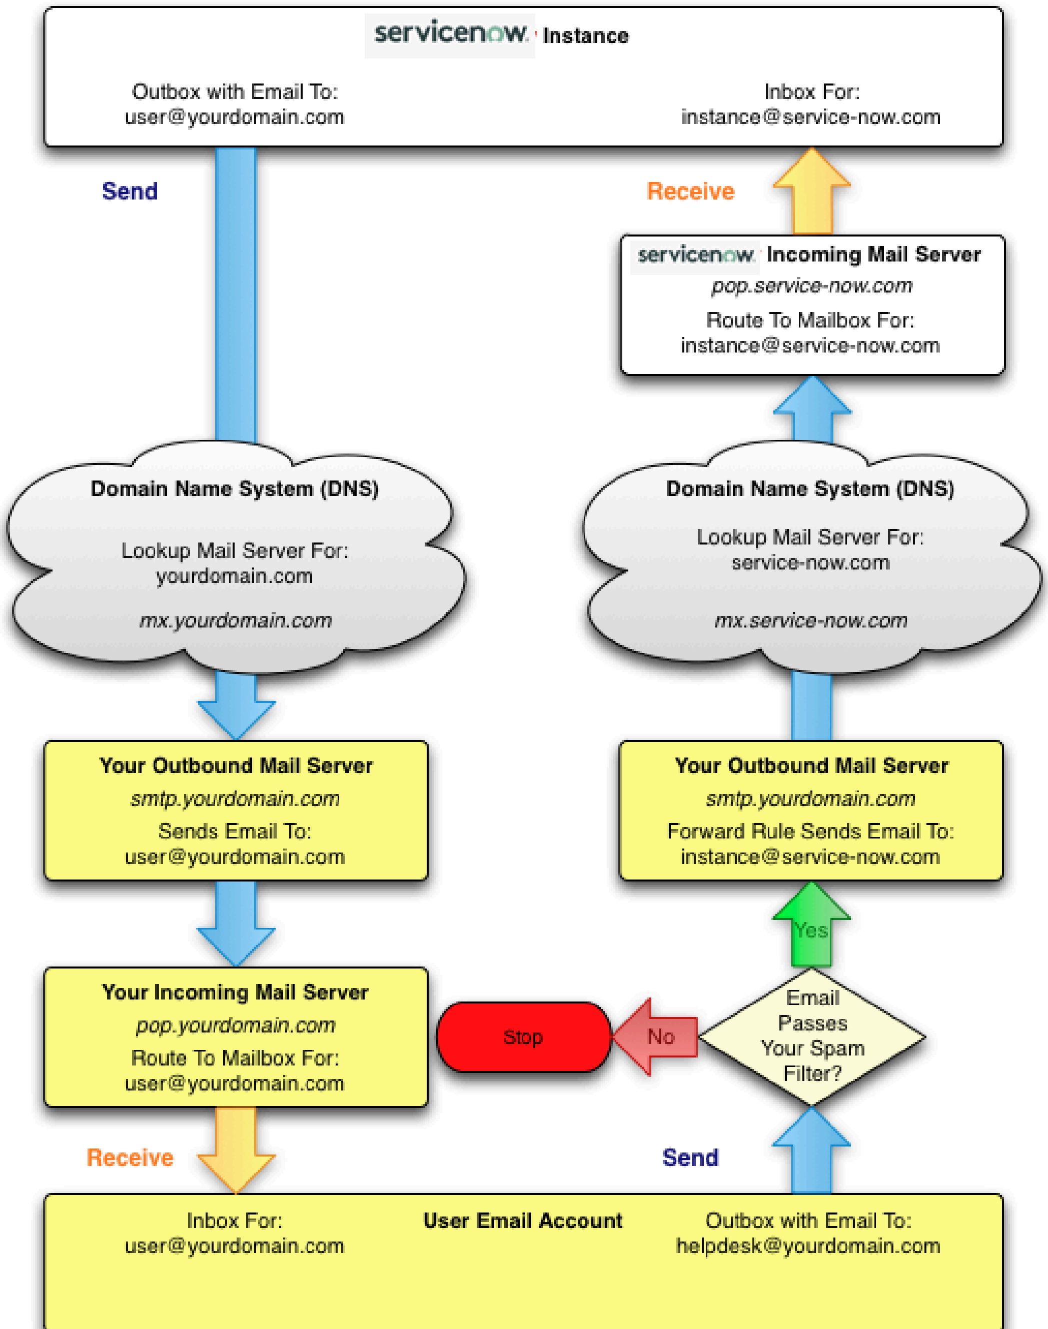

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Accounts**.
   The system displays the list of available email accounts.

2. Locate the record for **ServiceNow SMTP** and change **Active** to **false**.

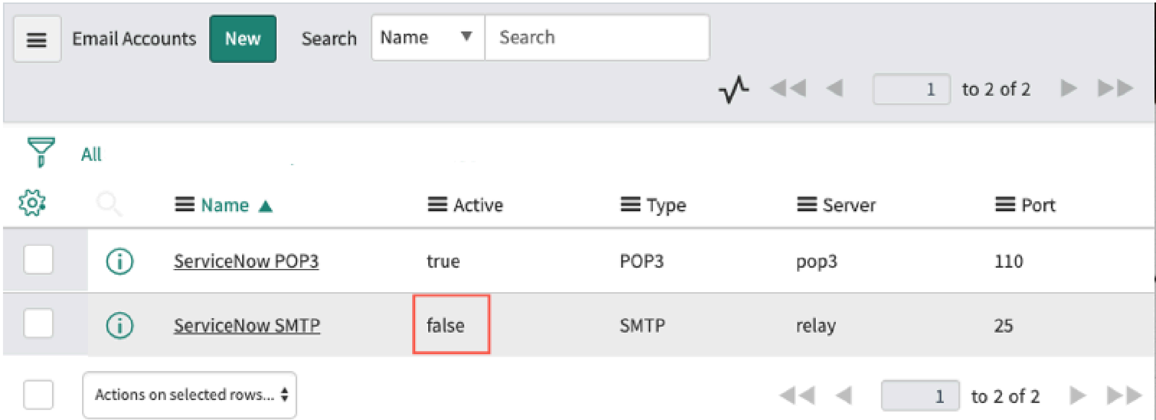

3. Click **New**.

4. Create an email account record for your SMTP server where the **Type** is **SMTP**.

5. On the form, fill in the fields.

Field	Description
Name	Name of the email account.
Type	Type of mail server.
Authentication	Type of authentication used for the email account to connect to the email server. You can select one of the following: ◦ –- None -- ◦ Password ◦ OAUTH ◦ OAUTH 2.0
Server	Remote server to which this account connects. To activate a server for an on-premises installation, enter the full address (FQDN) of the node (for example, `node.customerdomain`).
User Name	Username or ID to authenticate an email address.
Password	Password when the **Authentication** type is **Password**.
Connection Security	Type of secure connection. Choose a setting: **None** No secure protocol is used. **STARTTLS** Upgrades an insecure connection to a secure connection using the SSL/TLS

Field	Description
	encryption protocol, if your email server supports TLS. **SSL/TLS** Connect to an SSL/TLS encrypted port to secure the connection. Email is encrypted between the ServiceNow instance and your mail server. **Warning:** Selecting a less secure protocol like **STARTTLS** or **None** may expose your data. To better ensure the security of data in your email server, select **SSL/TLS**.
Port	Connection TCP port.
System Address Filter	System address filter to apply to the email account. If left blank, the system uses the default system address filter for inbound or outbound email. For more information, see System address filters.
Active	Option to activate the email account.
ServiceNow Configured	Identifier of whether this account is provisioned by ServiceNow. If you create an account, this option is not selected.
Enable Debug Logging	Option to create node logs for the raw data that is exchanged with the email server. You can review the node logs by navigating to **System Logs** > **Utilities** > **Node Log File Browser**. You can enable this field temporarily to diagnose issues related to receiving or sending email.

6. From **Related Links**, click **Test Connection**.

If the email account is valid, the system returns a success message.

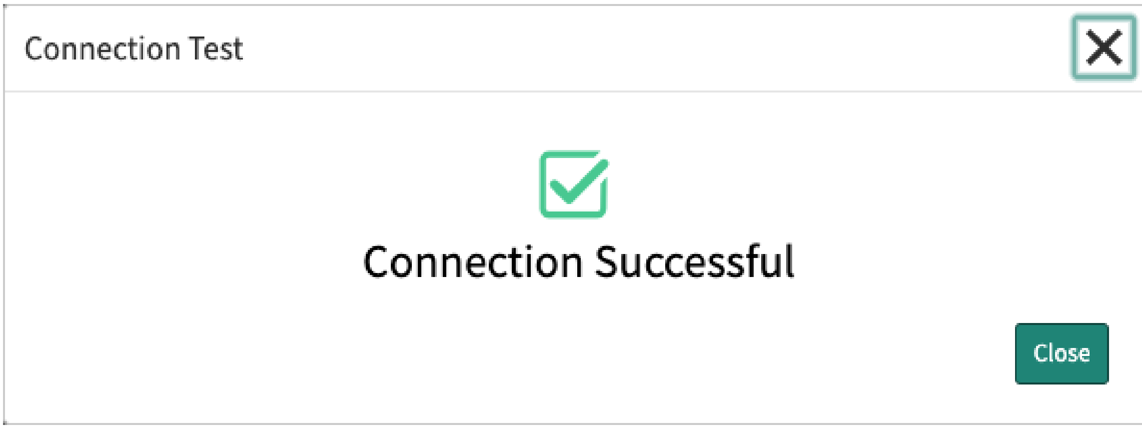

What to do next

Configure the SMTP server in your internal email architecture to forward email from the custom email address to the instance email address. Implement a spam filter on the custom email address.

Related topics

Create an email account

Enable using your own POP3 server

You can use your own POP3 server to store and receive email for the instance.

Before you begin

- Role required: admin
- Email server required: POP3
- Basic email properties: enabled

Procedure

1. On your POP3 server, create a mailbox for your instance that has a custom email address.
   For example, create a mailbox for `service-desk@company.com`.
2. Navigate to **System Mailboxes** > **Administration** > **Email Accounts**.
   The system displays the list of available email accounts.
3. **Optional:** If you do not want to receive, email sent to the instance@service-now.com mailbox, locate the record for **ServiceNow POP3** and change **Active** to **false**.
   An instance can receive email from multiple POP3 accounts at the same time. Leaving the **ServiceNow POP3** account active permits the instance to receive email sent to the default email address.
4. Click **New**.
   The system displays an empty email Account form.
5. Create an email account record for your POP3 server where the **Type** is **POP3**.
6. On the form, fill in the fields.

Field	Description
Name	Name of the email account.

Field	Description
Type	Type of mail server.
Authentication	Type of authentication used for the email account to connect to the email server. You can select one of the following: ◦ –- None -- ◦ Password ◦ OAUTH ◦ OAUTH 2.0
Server	Remote server to which this account connects. To activate a server for an on-premise installation, enter the full address (FQDN) of the node (for example, `node.customerdomain`).
User Name	Username or ID to authenticate an email address.
Password	Password when the **Authentication** type is **Password**.
Connection Security	Type of secure connection. Choose a setting: **None** No secure protocol is used. **STARTTLS** Upgrades an insecure connection to a secure connection using the SSL/TLS encryption protocol, if your email server supports TLS. **SSL/TLS** Connect to an SSL/TLS encrypted port to secure the connection. Email is encrypted between the ServiceNow instance and your mail server. **Warning:** Selecting a less secure protocol like **STARTTLS** or **None** may expose your data. To better ensure the security of data in your email server, select **SSL/TLS**.
Port	Connection TCP port.
System Address Filter	System address filter to apply to the email account. If left blank, the system uses the default system address filter for inbound or outbound email.

Field	Description
	For more information, see System address filters.
Active	Option to activate the email account.
ServiceNow Configured	Identifier of whether this account is provisioned by ServiceNow. If you create an account, this option is not selected.
Enable Debug Logging	Option to create node logs for the raw data that is exchanged with the email server. You can review the node logs by navigating to **System Logs** > **Utilities** > **Node Log File Browser**. You can enable this field temporarily to diagnose issues related to receiving or sending email.

7. From **Related Links**, click **Test Connection**.
   If the email account is valid, the system returns a success message.

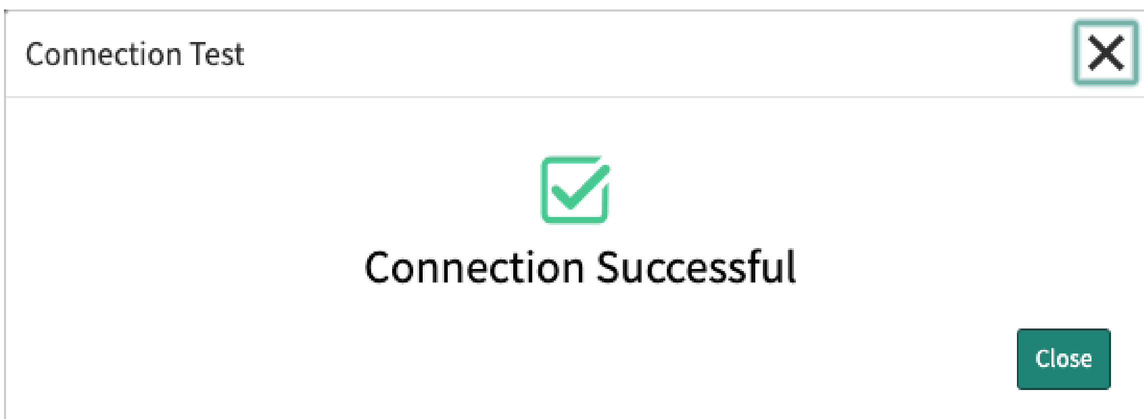

Example:

Receiving email using your own POP3 server

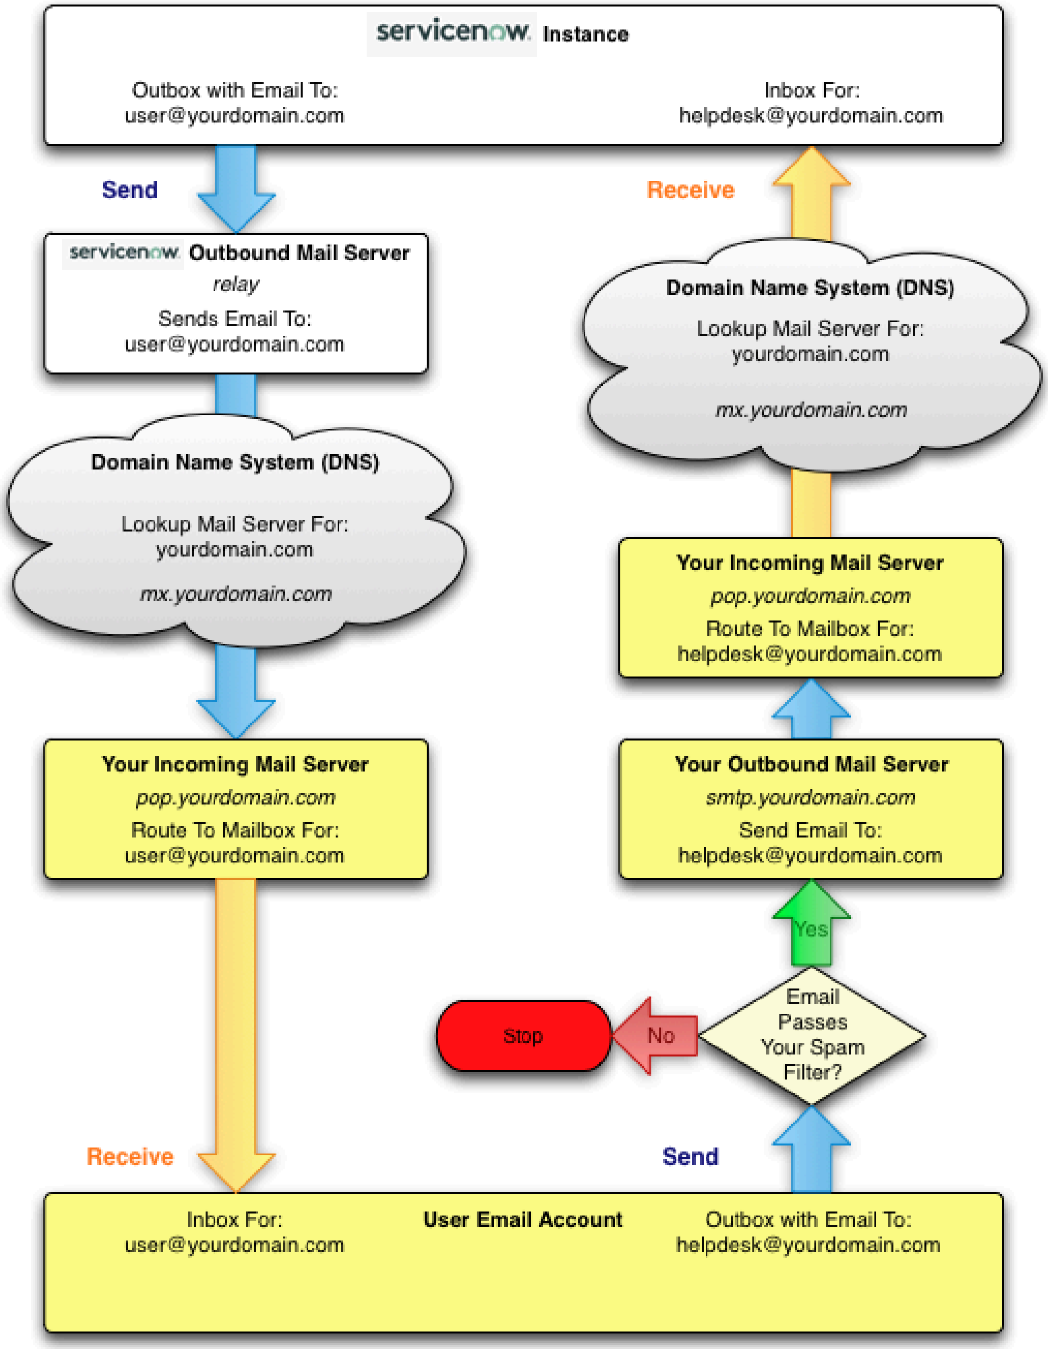

Enable using your own SMTP and POP3 servers

You can use your own SMTP and POP3 servers to send email from the instance and to store and receive email for the instance.

Before you begin

- Role required: admin
- Email servers required:

- SMTP
- POP3

• Basic email properties: enabled

Procedure

1. On your POP3 server, create a mailbox for your instance.
   For example, create a mailbox for `service-desk@company.com`.
2. Navigate to **System Mailboxes** > **Administration** > **Email Accounts**.
   The system displays the list of available email accounts.
3. Locate the record for **ServiceNow SMTP** and change **Active** to **false**.
4. **Optional:** If you do not want to receive email sent to the instance@service-now.com mailbox, locate the record for **ServiceNow POP3** and change **Active** to **false**.
   An instance can receive email from multiple POP3 accounts at the same time. Leaving the**ServiceNow POP3** account active means that the instance receives email sent to its default email address.
5. Click **New**.
   The system displays a blank Email Account form.
6. Create an email account record for your SMTP server where the **Type** is **SMTP**.
7. On the form, fill in the fields.

Field	Description
Name	Name of the email account.
Type	Type of mail server.
Authentication	Type of authentication used for the email account to connect to the email server. You can select one of the following: ○ –- None -- ○ Password ○ OAUTH ○ OAUTH 2.0
Server	Remote server to which this account connects. To activate a server for an on-premise installation, enter the full address (FQDN) of the node (for example, `node.customerdomain`).
User Name	Username or ID to authenticate an email address.
Password	Password when the **Authentication** type is **Password**.
Connection Security	Type of secure connection. Choose a setting: **None** No secure protocol is used. **STARTTLS**

Field	Description
	Upgrades an insecure connection to a secure connection using the SSL/TLS encryption protocol, if your email server supports TLS. **SSL/TLS** Connect to an SSL/TLS encrypted port to secure the connection. Email is encrypted between the ServiceNow instance and your mail server. **Warning:** Selecting a less secure protocol like **STARTTLS** or **None** may expose your data. To better ensure the security of data in your email server, select **SSL/TLS**.
Port	Connection TCP port.
System Address Filter	System address filter to apply to the email account. If left blank, the system uses the default system address filter for inbound or outbound email. For more information, see System address filters.
Active	Option to activate the email account.
ServiceNow Configured	Identifier of whether this account is provisioned by ServiceNow. If you create an account, this option is not selected.
Enable Debug Logging	Option to create node logs for the raw data that is exchanged with the email server. You can review the node logs by navigating to **System Logs** > **Utilities** > **Node Log File Browser**. You can enable this field temporarily to diagnose issues related to receiving or sending email.

8. From **Related Links**, click **Test Connection**.
   If the email account is valid, the system returns a success message.

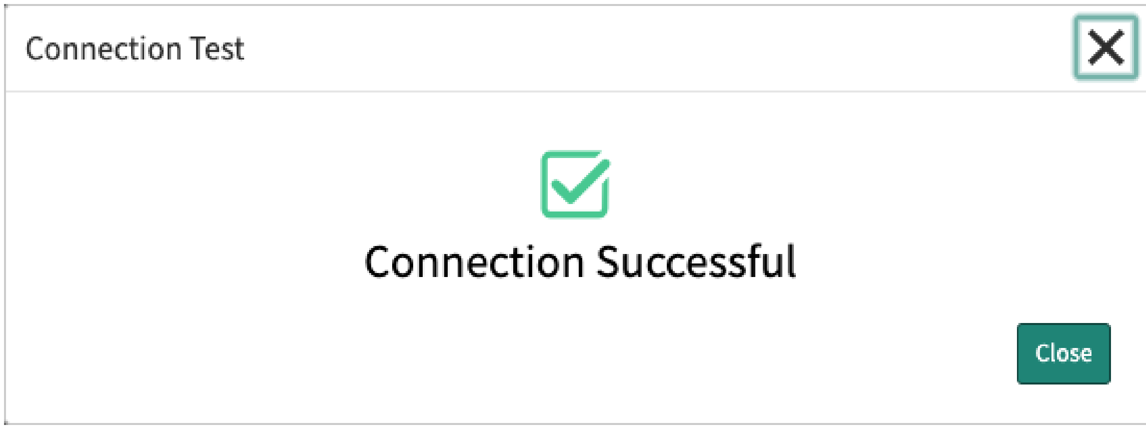

9. Click **New**.
   The system displays a blank Email Account form.
10. Create an email account record for your POP3 server where the **Type** is **POP3**.
11. From **Related Links**, click **Test Connection**.
    If the email account is valid, the system returns a success message.

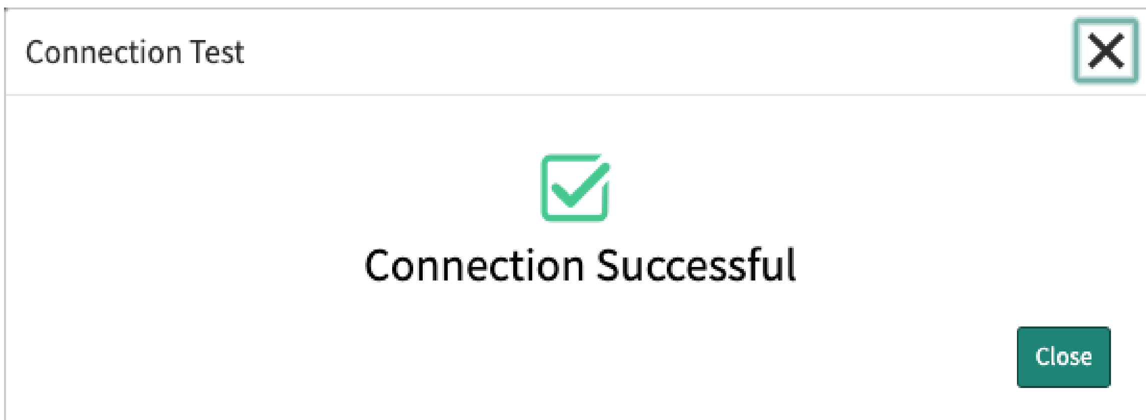

Example:

Using your own SMTP and POP3 servers

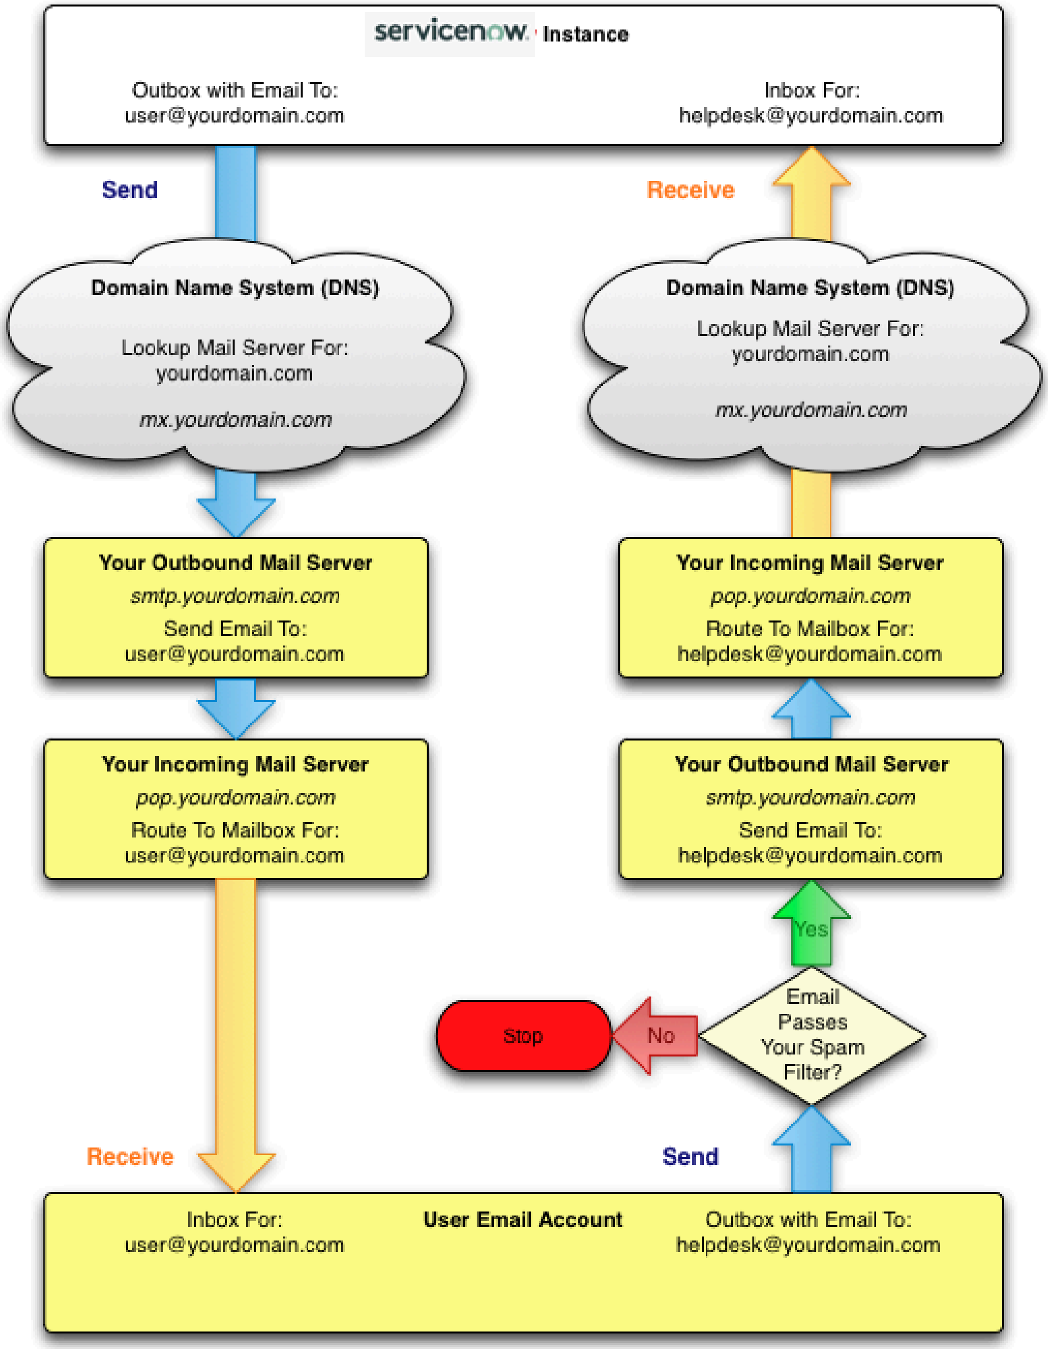

Email bounce management

Help prevent sending emails to addresses that are known to generate bounces by monitoring and filtering them out while sending emails.

Email bounce

An email bounce signifies the non-delivery of your email message according to RFC3464. When an email gets bounced, the sender receives a delivery status email with the reason for the bounce and the associated status code. For more information about these status codes, see Status codes.

An email address is blocked after it reaches a certain threshold of bounced emails. The default count for blocking the bounced email addresses is 10.

Important:
Disabling bounce management is not recommended as it may lead to slow delivery or blocking of emails by email provider.

Status code

The syntax of the status codes is defined as the following:

```
status-code = class "." subject "." detail

      class = "2"/"4"/"5"

      subject = 1*3digit

      detail = 1*3digit
```

The first subcode indicates whether the delivery attempt was successful. The second subcode indicates the probable source of any delivery anomalies. The third subcode indicates a precise error condition.

Example: Status codes

Code 5.1.1

Code	Summary	Description
5.1.1	Bad destination mailbox address	The mailbox specified in the address doesn't exist. For email names, this means the address portion to the left of the "@" sign is invalid.

The default status codes that are used for blocking bounced email addresses are 5.0.*, 5.1.* and 5.7.*, where the asterisks (*) could be any number in the series according to the status code specification.

Note:
There is a possibility that some email providers may use non-standard RFC status codes, potentially leading to functional email addresses being blocked. Administrators action is required to unblock these status codes and email addresses.

To view the blocked status codes go to **All** > **System Logs** > **Bounce Email Addresses**.

Related topics

Configure bounce email address status

Check the status of bounced email addresses

Configure bounce email address status

Configure a bounced email address status according to type, class, subject, and details.

Before you begin
Role required: admin

About this task

For detailed information on status codes and explanations of the components, see RFC Status codes.

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Bounce Email Address Status**.
2. Select **New**.
3. Select the **Type** and enter the **Class**, **Subject**, and **Detail**.

   Note:
   For example: Status code patterns for 5.1.* are blocked by default, if 5, 1, and 2 are entered as the Class, Subject, and Detail, the Type is selected as unblocked then only 5.1.2 is unblocked. All other status codes patterns for 5.1.* are still blocked.

4. Select **Submit**.

Related topics

Check the status of bounced email addresses

OAuth email authentication

OAuth enables your instance to receive and send email through a third-party email account.

Open Authentication (OAuth) is an open standard for authorization that provides administrators with an authorization method when connecting to incoming IMAP and outgoing SMTP servers. OAuth enables an instance to receive and send email from a third-party account, such as Gmail, without having to enter the credentials for that account.

The OAuth 2.0 implementation requires you to obtain an access and refresh token from your third-party email provider for each third-party email account. The tokens are automatically saved to the instance database. They provide authorization for all email communication between the instance and the authorized third-party account. A scheduled job regularly checks to see if email access tokens are valid. If the access token is not valid, but the refresh token is, the instance automatically regenerates a new access token.

OAuth 2.0 support is available starting with the Geneva release. The OAuth implementation supports IMAP and SMTP accounts only. POP3 is not supported. OAuth 1.0 is no longer supported.

Note:
Customer email accounts configured to use OAuth 1.0 authentication with Gmail cease to function as of April 20, 2015. See KB0546976 for more information.

Activating the Email - OAuth support for IMAP and SMTP plugin allows you to use OAuth with email. If you upgrade to Geneva or later instances and are already using OAuth 1.0, activate the plugin again.

Note:
For configuring SMTP and IMAP email accounts with Microsoft Office 365 using OAuth 2.0, see KB0816072.

See Outbound REST with OAuth 2.0 profile tutorial - integrating with Google Contacts API for an example of using an OAuth 2.0 profile to authenticate an outbound REST message with Google to retrieve contact information. Also see OAuth 2.0 for more information on OAuth 2.0 support in the instance.

Activate the OAuth email authentication plugin

You can activate the Email:OAUTH support for IMAP, Microsoft Graph (Receiving), and SMTP plugin (com.glide.email.oauth) for Notifications if you have the admin role.

Before you begin

Role required: admin

About this task

The following items are installed with Email:OAUTH support for IMAP, Microsoft Graph (Receiving), and SMTP:

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Email:OAUTH support for IMAP, Microsoft Graph (Receiving), and SMTP plugin (com.glide.email.oauth) using the filter criteria and search bar.

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

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

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

Enable OAuth 2.0 for email

Setting up OAuth 2.0 for email requires you to obtain access and refresh tokens from your email provider.

Before you begin

- Role required: admin
- Plugin required: Email - OAUTH support for IMAP and SMTP

Procedure

1. Log in to your third-party email account, such as Gmail, and enable OAuth 2.0.
2. Obtain the following from your third-party email account:
   - client ID
   - client secret
   - authorization URL
   - token URL
   - redirect URL
   - token revocation URL
3. Navigate to **System OAuth** > **Application Registry**.

4. Click **New**.
5. Click **Connect to a third party OAuth Provider** to create an application registry record that email uses.
6. Use the information you obtained from your third-party email account to fill in the fields on the form.
   See Use a third-party OAuth provider for instructions.
7. Create the OAuth application registry record and its associated OAuth Entity Profile and OAuth Entity Scope records.
8. Click **Submit**.
9. Navigate to **System Mailboxes** > **Administration** > **Email Accounts**.

   Note:
   Port and Connection Security information is provided by the provider.

   The system displays the list of available email accounts.
10. **Optional:** If you do not want to receive email sent to the default instance email address, locate the record for **ServiceNow POP3** and change **Active** to **false**.
    The system can receive email from multiple POP3 and IMAP email accounts.
11. Locate the records for **ServiceNow SMTP** and change **Active** to **false**.
12. Click **New**.
13. Create an email account record for your OAuth 2.0 SMTP server where the **Type** is **SMTP**.
14. For **Authentication**, select **OAuth 2.0**.
15. For **OAuth Profile**, select the application registry record you created.
16. Click **Authorize Email Account Access** to obtain the access and refresh tokens.

    Another browser window opens asking you to authorize the account access on the third-party email account.
17. Authorize the access.

    After the authorization is successful and the tokens are saved to the instance, the **Authorize Email Account Access** button no longer appears on the Email Account form.
18. Click **New**.
19. Create an email account record for your OAuth 2.0 IMAP server where the **Type** is **IMAP**.
    Use the same **Authentication** and **OAuth profile** settings as the OAuth SMTP email account.

System address filters

Prevent your system from communicating with untrusted domains and email addresses.

System address filters control who can communicate with your instance via email. You can improve email security by filtering domains and email addresses that you find suspicious.

For example, you might identify `example.com` as a suspicious domain. You can stop receiving emails from `example.com` by specifying the domain in a system address filter. You can also use another system address filter to prevent your system from sending emails to `example.com`.

There are two types of system address filters:

Outbound

Controls which domains and email addresses can receive email from your instance.

Inbound

Controls which domains and email addresses can send email to your instance.

When setting up system address filters, you can create one default outbound filter and one default inbound filter. A default outbound filter applies to all active SMTP email accounts automatically. A default inbound filter applies to all active IMAP or POP3 email accounts automatically.

To control senders and recipients for a specific email account, create a non-default filter and then apply it to the account.

Before setup

Before you set up system address filters, consider doing the following:

- Monitor your email to identify suspicious domains and email addresses.

  Using Instance Security Center, you can monitor the blocked and allowed incoming email metrics for your instance. For more information, see Designate untrusted and trusted email domains.

- Designate someone to set up system address filters by assigning them the email_account_admin role.

Setting up system address filters

To set up system address filters for your instance, complete the following tasks:

1. Set email address filters

   Specify which domains and email addresses are allowed or disallowed.

2. Create a system address filter

   Define how email address filters apply to inbound and outbound email.

Next steps

After you set up system address filters, configure email filters for an added layer of security. Email filters enable you to ignore an inbound email or move it to a particular mailbox. For more information on configuring email filters, see Email filters.

Set email address filters

Specify which domains and email addresses are allowed or disallowed.

Before you begin

Role required: email_account_admin

About this task

You can specify two types of email address filters:

- Allow List: A list of allowed domains and email addresses
- Deny List: A list of disallowed domains and email addresses

You can also specify exceptions to the allowed or disallowed domains and email addresses.

Procedure

1. Navigate to **System Mailboxes** > **Administration** > **Email Address Filters**, and then select **New**.
2. Enter the **Name** of the address filter.
3. Select the filter **Type**:
   - Allow List: Allow all specified domains. All other domains are disallowed.
   - Deny List: Disallow all specified domains. All domains that don't match the denied list domains are allowed.
4. In the **Domains** field:

   a. Select the lock icon to unlock it and access the domains controlled by this filter.

   b. Select an existing domain or enter a new domain:
      - If the filter **Type** is Allow List, select the Search icon. Select the domains for which all email addresses are allowed. To add a new domain, select **New**, enter the **Domain**, and select **Submit**.
      - If the filter **Type** is Deny List, select the Search icon. Select the domains for which all email addresses are disallowed. To add a new domain, select **New**, enter the **Domain**, and select **Submit**.

      Note:
      You can specify a wildcard (*) in the domain name, for example `*.com`. If you create an email configuration that has multiple email address filters, all the filters evaluate the given email addresses. The filters determine whether the addresses are valid and the message can be sent as outbound email.

      By default, the maximum number of domains that you can associate with an email address filter is 100. You can reconfigure this limit by setting the *glide.email_address_filter.max_domains* property.

   c. Select the lock icon to lock the **Domains** field.

5. Specify any email addresses that are **Exceptions** to the domains specified in Step 4.

   a. Select the lock icon to open the Email Address Filter Exceptions [sys_email_address_filt_except] table.

   b. Click the search icon to choose an existing email address.

   c. To add a new email address, select **New**, enter the email address exception, and select **Submit**.

      Note:
      You can specify a wildcard (*) in an email address exception.

      By default, the maximum number of exceptions that you can associate with an email address filter is 1000. You can reconfigure this limit by setting the *glide.email_address_filter.max_exceptions* property.

d. Select the lock icon to lock the **Exceptions** field.
The exception is added to the Email Address Filter Exceptions [sys_email_address_filt_except] table.

6. Select **Submit**.
The email address filter is added to the Email Address Filters [sys_email_address_filter] table.

Example: Email address filter

Note:
After the email address filter is added to the Email Address Filters [sys_email_address_filter] table, Create an email client configuration.

Email address filter — allowed exception example

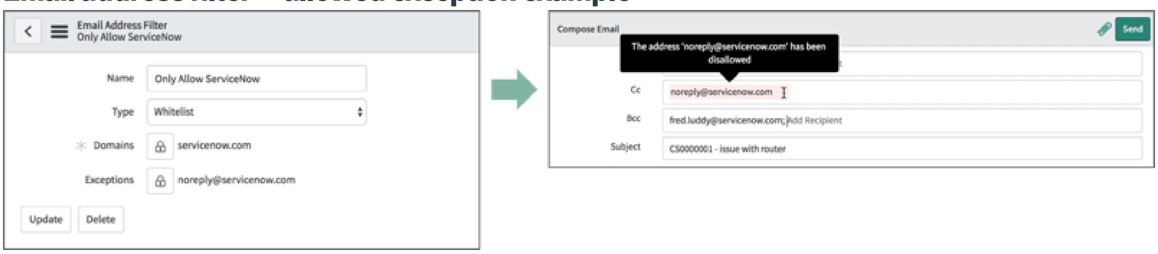

Email address filter — disallowed exception example

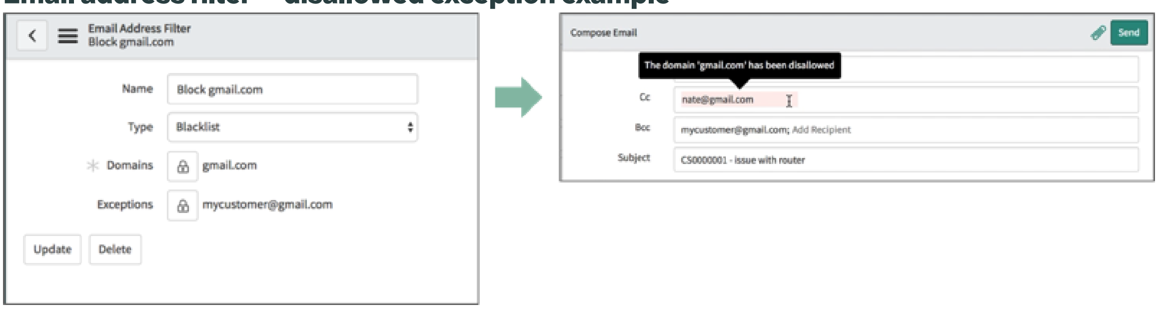

Example:

Email address filter — allowed exception example

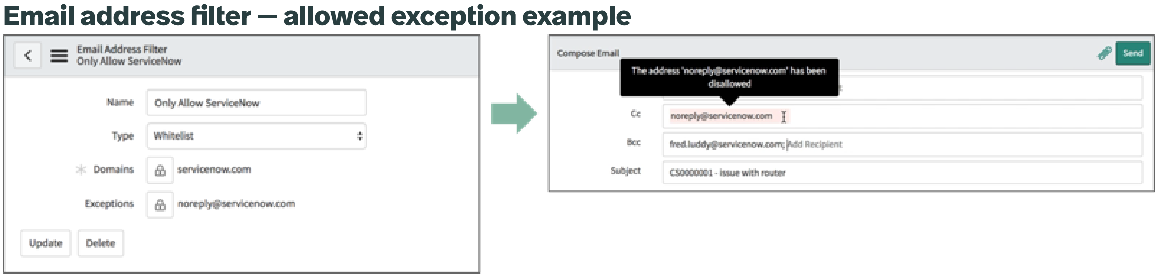

Email address filter — disallowed exception example

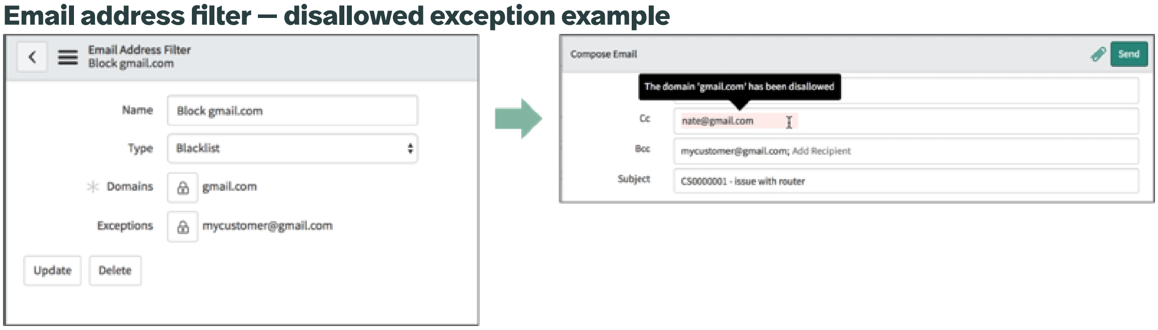

What to do next

Apply the email address filter to a system address filter. For more information, see Create a system address filter.

Reading email using Microsoft Graph

Use Microsoft Graph endpoints to read emails from Microsoft Exchange Online.

Using Microsoft Graph with ServiceNow

In addition to the standard IMAP, SMTP, and POP3, the ServiceNow AI Platform supports Microsoft Graph as an email type.

Email account administrators, which are users who have the email_account_admin role, can configure ServiceNow AI Platform to read emails from an Microsoft Exchange Online tenant using Graph Endpoints by authenticating with a client ID and client secret or with certificates.

> **Note:**
> Antispam and malware scanning are managed by Microsoft Graph account and not ServiceNow.

In the **Email Account** form, select, **Microsoft Graph (Receive)** in the **Type** field.

Activate Email - Support for Email Processing by Microsoft Graph API

You can activate the Email - Support for Email Processing by Microsoft Graph API plugin (glide.email.graph) for Notifications if you have the admin role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Email - Support for Email Processing by Microsoft Graph API plugin (glide.email.graph) using the filter criteria and search bar.

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

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

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

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

Configure an OAuth profile to use a client ID and secret for token generation

Configure an OAuth profile using a client ID and client secret to create an email account for using Microsoft Graph (receive) in your email account type.

Before you begin

Create an Microsoft Azure account to configure OAuth profile.

Complete the set up steps using your Microsoft Azure Developer account. See the Microsoft Azure product documentation for instructions on creating and configuring custom applications.

Role required: admin and Microsoft Azure portal administrator

Procedure

1. Log in to the Microsoft Azure registration portal with your organization credentials.
   For more information, see the Microsoft Azure registration portal documentation.
2. Register a new custom application by filling in the application name, supported account type, and redirect URI, and select **Register**.

   > **Note:**
   > Enter the redirect URI in the following format: `https://<instance>/oauth_redirect.do`.

   An overview of the application's basic information is displayed.
3. Copy the client ID to a text file.
   You will use this ID and the client secret value generated in the next steps to register the app as a third-party OAuth provider on your ServiceNow instance. You use the application ID as the client ID when you connect the application to your instance.
4. Enable the **Mail.ReadWrite** permission.
   a. In Microsoft Azure, navigate to **Manage** > **API permissions**.
   b. Select **Add a permission**.
   c. Select the Microsoft Graph tile.
   d. Select **Application Permissions**.
   e. In the **Select permissions** field, enter `Mail.ReadWrite`.
   f. Select the **Mail.ReadWrite** check box.
   g. Select **Add permissions**.
5. Select **Grant admin consent for <your organization name>**.
6. Select **Yes** to confirm.
7. Add a client secret.
   a. In Microsoft Azure, navigate to **Manage** > **Certificates & secrets**.
   b. Select **New client secret**.
   c. Provide a description and an expiration date and select **Add**.
8. Copy the value (secret value) to a text file.
9. Navigate to **Overview** > **Endpoints** and copy the **OAuth 2.0 token endpoint (v2)** to a text file.

Register an application as an OAuth provider

Use the information generated during Microsoft Azure account configuration to register an application as an OAuth provider.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry**.
2. Select **New**.
3. On the **What kind of OAuth application** screen, select **Connect to a third-party OAuth Provider**.
4. On the form, fill in the fields.

Application Registries form

Field	Value required
Name	Name to uniquely identify the record.
Client ID	Application ID of the application you created in Microsoft Azure.
Client Secret	The client secret you generated when you created the application in Microsoft Azure. **Note:** When using certificates, the **Client Secret** is a dummy value.
OAuth API Script	OAuth API script name. For more information see, OAuth API Script. **Note:** This is required only while using certificates.
Default Grant type	Client Credentials.
Token URL	Token URL copied after configuring the Microsoft Azure account.
Redirect URL	https://<instance>./oauth_redirect.do **Note:** This URL should be the same as the URL in Microsoft Azure.

5. Right-click the form header and select **Save**.
   An OAuth Entity Profile record is created.
6. In the OAuth Entity Scopes related list, add scopes to match the permissions you defined when you configured the application.
7. Select **Insert a new row**.
8. Enter `Default` as the name and `.default` as OAuth scope.
9. Right-click the form header and select **Save**.
10. In the OAuth Entity Profiles embedded list, select the profile created by default.
11. In the OAuth Entity Scopes embedded list, add a new row and select the **Default** scope.
12. Select **Update**.

What to do next

Create an email account for Microsoft Graph (receive) using the OAuth profile.

Configure an OAuth profile to use certificates for authentication with Microsoft Azure

Configure an OAuth application profile to authenticate using certificates.

Before you begin

Create an Microsoft Azure account to configure OAuth profile.

Complete the set up steps using your Microsoft Azure Developer account. See the Microsoft Azure product documentation for instructions on creating and configuring custom applications.

Role required: admin and Microsoft Azure portal administrator

Procedure

1. Log in to the Microsoft Azure registration profile with your organization credentials.
   For more information, see the Microsoft Azure registration portal documentation.
2. Register a new custom application by filling in the application name, supported account type, and redirect URL, and then select **Register**.

   Note:
   Enter the redirect URL in the following format: `https://<instance>/oauth_redirect.do`.

   An overview of the application's basic information is displayed.
3. Copy the client ID to a text file.
   You will use this ID and the secret value generated in the next steps to register the app as a third-party OAuth provider on your ServiceNow instance. You use the application ID as the client ID when you connect the application to ServiceNow.
4. Enable the **Mail.ReadWrite** permission.
   a. In Microsoft Azure, navigate to **Manage** > **API permissions**.
   b. Select **Add a permission**.
   c. Click the Microsoft Graph tile.
   d. Select **Application Permissions**.
   e. In the **Select permissions** field, enter `Mail.ReadWrite`.
   f. Select the **Mail.ReadWrite** check box.
   g. Select **Add permissions**.
5. Select **Grant admin consent for <your organization name>**.
6. Select **Yes** to confirm.
7. Upload the certificate
   a. In Microsoft Azure, navigate to **Manage** > **Certificates & secrets**.
   b. Select **upload certificate**.
   c. Provide a description, a start date, and an expiration date and upload the certificate.
8. Copy the client ID and directory ID to a text file.
9. Navigate to **Overview** > **Endpoints** and copy the **OAuth 2.0 token endpoint (v2)** to a text file.

What to do next

Configure a JWT provider

Configure a JWT provider

Configure a JWT provider on the ServiceNow AI Platform to configure an OAuth application profile to authenticate using certificates.

Before you begin

Upload a Java Key Store certificate. Attach a JKS certificate containing the certificate uploaded on Microsoft Azure to your instance to use when enabling the JWT client authentication. For more information, see Upload Java Key Store certificate.

Create a JWT signing key to assign to your Java KeyStore (JKS) certificate. For more information, see Configure a JWT signing key.

Role required: admin

Procedure

1. Add a JWT provider to your ServiceNow instance and fill in the claim values.

Claim Name	Claim Value
aud	https://login.microsoftonline.com/{tenantId}/oauth2/v2.0/token **Note:** Tenant ID is the directory ID copied from Microsoft Azure.
iss	Client ID copied from Microsoft Azure
sub	Client ID copied from Microsoft Azure

   For more information, see Create a JWT provider with a JWT signing key.
2. From the JWT Provider record menu, select **Copy sys_id** and copy the JWT provider sys_id.

What to do next

Generate a SHA-1 thumbprint

Generate a SHA-1 thumbprint

Generate a SHA-1 thumbprint using the JWT provider's sys_id and JKS certificate's sys_id and certificate's alias to be added to the GraphCertificateOAuthTemplate script.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Scripts-Background**.
2. Paste the following script to generate the SHA-1 thumbprint value.

```
var certId = "<sys_id of the certificate record>";

var certAlias = "<alias name for the certificate>";

var gce = new GlideCertificateEncryption();

```

```
var thumbprint = gce.getThumbPrintFromKeyStore(certId,
 certAlias,"SHA-1");

gs.log(thumbprint);
```

3. Select **Run script**.

What to do next

Create OAuth API script

Create OAuth API script

Create and duplicate OAuth API script for application registry.

Before you begin

Generate a SHA -1 thumbprint. For more information see Generate a SHA-1 thumbprint.

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Script Includes**.
2. Select and open **GraphCertificateOAuthTemplate**.
3. Duplicate the script by changing the name in the **Name** field,select **Insert and Stay** from the record menu.

   Note:
   The name should begin with **OAuth**.

   The script is copied and created with a different name.
4. In the script, enter the JWT provider sys_id and SHA -1 thumbprint.
5. Select **Update**.

What to do next

Register an application as an OAuth provider.

Create an email account for Microsoft Graph (receive)

Create an email account for reading emails from Microsoft Exchange Online using Microsoft Graph Endpoints.

Before you begin

Role required: email_account_admin

About this task

Antispam and malware scanning are managed by Microsoft Graph account and not ServiceNow

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Accounts**.
2. Select **New**.
3. On the form, fill in the fields.

Email Accounts form

Field	Description
Name	Name of the email account.
Type	Microsoft Graph (Receive).
Authentication	OAUTH 2.0 - Type of authentication used for the email account to connect to the email server.
Inbox email address	Address of the inbox for reading emails.
OAuth Profile	OAuth profile configured during application registry.
Active	Option to activate the email account.
Enable debug output	Option to create node logs for the raw data that is exchanged with the email server. You can enable this field temporarily to diagnose issues related to receiving or sending email.

4. Select **Submit**.
5. Select the test connection link at the bottom to test the account.

Send email using client credential flow

Use the client credential flow to send emails from Microsoft Exchange Online.

Email account administrators, who are users that have the email_account_admin role, can configure ServiceNow AI Platform to send emails from an Microsoft Exchange Online tenant. The email account administrators can use SMTP by authenticating with a client ID and client secret or with certificates.

Activate OAuth support plugin

You can activate the Email:OAUTH support for IMAP, Microsoft Graph (Receiving), and SMTP plugin (com.glide.email.oauth) for Notifications if you have the admin role.

Before you begin

Role required: admin

About this task

The following items are installed with Email:OAUTH support for IMAP, Microsoft Graph (Receiving), and SMTP:

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Email:OAUTH support for IMAP, Microsoft Graph (Receiving), and SMTP plugin (com.glide.email.oauth) using the filter criteria and search bar.

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

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

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

Configure client credential flow for SMTP OAuth2 using a client id and secret

Configure outbound email accounts in a ServiceNow instance using a client id and secret.

Before you begin

Create an Microsoft Azure account to configure an OAuth profile.

Complete the setup steps using your Microsoft Azure Developer account. See Microsoft Azure product documentation for instructions on creating and configuring custom applications.

Role required: Microsoft Azure portal administrator and admin

Procedure

1. Log in to the Microsoft Azure registration portal with your organization credentials.
   For more information, see the Microsoft Azure registration portal documentation.
2. Register a new custom application by filling in the application name, supported account type, and redirect URL, and select **Register**.

   Note:
   Enter the redirect URL in the following format: `https://<instance>/oauth_redirect.do`.

   An overview of the application's basic information is displayed.
3. Copy the client ID to a text file.
   You'll use this client ID and the client secret value generated in the next steps to register the app as a third-party OAuth provider on your ServiceNow instance. You use the application ID as the client ID when you connect the application to your instance.
4. Enable the **SMTP.SendAsApp** permission.
   a. In Microsoft Azure, navigate to **Manage** > **API permissions**.
   b. Select **Add a permission**.
   c. Select the Microsoft Office 365 tile.
   d. Select **Application Permissions**.
   e. In the **Select permissions** field, enter `SMTP SendAsApp`.
   f. Select the **SMTP SendAsApp** check box.
   g. Select **Add permissions**.
5. Select **Grant admin consent for <your organization name>**, then select **Yes** to confirm.
6. Add a client secret.
   a. In Microsoft Azure, navigate to **Manage** > **Certificates & secrets**.
   b. Select **New client secret**.
   c. Provide a description and an expiration date, then select **Add**.

7. Copy the value (secret value) to a text file.
8. Register service principals in Microsoft Exchange using PowerShell.

   a. Run the following commands.

   ```
   Install-Module -Name ExchangeOnlineManagement -allowprerelease
   Import-module ExchangeOnlineManagement
   Connect-ExchangeOnline -Organization <tenantId>
   ```

   For more information, see Authenticate SMTP connection using OAuth

   b. Register an Microsoft Azure AD application's service principal by running the following command.

   ```
   New-ServicePrincipal -AppId <APPLICATION_ID> -ObjectId <OBJECT_ID>
   ```

   Important:
   For the OBJECT_ID value, go to **Overview** > **Managed applications in local directory** and copy the Object ID.

   c. Add permission to user email ID using the following command.

   ```
   Add-MailboxPermission -Identity "<email_id>" -User <object_id>
    -AccessRights FullAccess
   ```

9. Navigate to **Overview** > **Endpoints** and copy the **OAuth 2.0 token endpoint (v2)** to a text file. You use the **OAuth 2.0 token endpoint (v2)** when registering an OAuth provider. For more information, see Register an OAuth Provider.

What to do next

Register an OAuth Provider

Register an OAuth Provider

Use the information generated during the Microsoft Azure account configuration to register an application as an OAuth provider.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry**.
2. Select **New**.
3. On the **What kind of OAuth application** screen, select **Connect to a third-party OAuth Provider**.
4. On the form, fill in the fields.

Application Registries form

Field	Description
Name	Name to uniquely identify the record.
Client ID	Application ID of the application that you created in Microsoft Azure.

Field	Description
Client Secret	The client secret that you generated when you created the application in Microsoft Azure. **Note:** When using certificates, the **Client Secret** is a dummy value.
OAuth API Script	OAuth API script name. For more information see, OAuth API Script. **Note:** This field is required only while using certificates.
Default Grant type	Client Credentials
Token URL	Token URL that you copied after configuring the Microsoft Azure account.
Redirect URL	`https://<instance>./oauth_redirect.do` **Note:** This URL should be the same as the URL in Microsoft Azure.

5. Select and hold (or right-click) the form header and select **Save**.
   An OAuth Entity Profile record is created.
6. In the OAuth Entity Scopes related list, add scopes to match the permissions that you defined when you configured the application.
7. Select **Insert a new row**.
8. Enter `smtp` as the name and `https://outlook.office365.com/.default` as the OAuth scope.
9. Select and hold (or right-click) the form header and select **Save**.
10. In the OAuth Entity Profiles embedded list, select the profile created by default.
11. In the OAuth Entity Scopes embedded list, add a new row and select the **smtp** scope.
12. Select **Update**.

What to do next

Use the OAuth profile to create an email account with SMTP OAuth2.0. For more information, see Create an email account with SMTP OAuth2.0.

Configure client credential flow for SMTP OAuth2 using certificate-based authentication

Configure an OAuth application profile to authenticate using certificates for outbound email.

Before you begin

Create an Microsoft Azure account to configure an OAuth profile.

Complete the setup steps using your Microsoft Azure Developer account. See Microsoft Azure product documentation for instructions on creating and configuring custom applications.

Role required: Microsoft Azure portal administrator and admin

Procedure

1. Log in to the Microsoft Azure registration profile with your organization credentials.
   For more information, see the Microsoft Azure registration portal documentation.
2. Register a new custom application by filling in the application name, supported account type, and redirect URL, and then select **Register**.

   Note:
   Enter the redirect URL in the following format: `https://<instance>/oauth_redirect.do`.

   An overview of the application's basic information is displayed.
3. Copy the client ID to a text file.
   You'll use this client ID and the secret value generated in the next steps to register the app as a third-party OAuth provider on your ServiceNow instance. You use the application ID as the client ID when you connect the application to ServiceNow.
4. Enable the **SMTP.SendAsApp** permission.
   a. In Microsoft Azure, navigate to **Manage** > **API permissions**.
   b. Select **Add a permission**.
   c. Select the Microsoft Office 365 tile.
   d. Select **Application Permissions**.
   e. In the **Select permissions** field, enter `SMTP.SendAsApp`.
   f. Select the **SMTP.SendAsApp** check box.
   g. Select **Add permissions**.
5. Select **Grant admin consent for <organization name>**, then select **Yes** to confirm.
6. Upload the certificate.
   a. In Microsoft Azure, navigate to **Manage** > **Certificates & secrets**.
   b. Select **upload certificate**.
   c. Provide a description, a start date, and an expiration date, then upload the certificate.
7. Copy the client ID and directory ID to a text file.
8. Register service principals in Microsoft Exchange using PowerShell.
   a. Run the following commands.

   ```
   Install-Module -Name ExchangeOnlineManagement -allowprerelease
   Import-module ExchangeOnlineManagement
   Connect-ExchangeOnline -Organization <tenantId>
   ```

   For more information, see Authenticate SMTP connection using OAuth

   b. Register an Microsoft Azure AD application's service principal by running the following command.

   ```
   New-ServicePrincipal -AppId <APPLICATION_ID> -ObjectId <OBJECT_ID>
   ```

   Important:
   For the OBJECT_ID value, go to **Overview** > **Managed applications in local directory** and copy the Object ID.

c. Add permission to user email ID using the following command.

```
Add-MailboxPermission -Identity "<email_id>" -User <object_id>
 -AccessRights FullAccess
```

9. Navigate to **Overview** > **Endpoints** and copy the **OAuth 2.0 token endpoint (v2)** to a text file. You use the **OAuth 2.0 token endpoint (v2)** when configuring a JWT provider. For more information, see Configure a JWT provider.

What to do next

Configure a JWT provider

Configure a JWT provider

Configure a JWT provider on the ServiceNow AI Platform to configure an OAuth application profile to authenticate using certificates.

Before you begin

Upload a Java Key Store (JKS) certificate. Attach a JKS certificate containing the certificate uploaded on Microsoft Azure to your instance to use when enabling the JWT client authentication. For more information, see Upload Java Key Store certificate.

Create a JWT signing key to assign to your JKS certificate. For more information, see Configure a JWT signing key.

Role required: admin

Procedure

1. Add a JWT provider to your ServiceNow instance and fill in the claim values.

Claim Name	Claim Value
aud	https://login.microsoftonline.com/{tenantId}/oauth2/v2.0/token **Note:** The tenant ID should be replaced with the directory ID copied from Microsoft Azure.
iss	Client ID copied from Microsoft Azure
sub	Client ID copied from Microsoft Azure

   For more information, see Create a JWT provider with a JWT signing key.

2. From the JWT Provider record menu, select **Copy sys_id** and copy the JWT provider sys_id.

What to do next

Generate SHA-1 thumbprint

Generate SHA-1 thumbprint

Generate an SHA-1 thumbprint using the JWT provider's sys_id, the Java Key Store (JKS) certificate's sys_id, and the JKS certificate's alias to the GraphCertificateOAuthTemplate script.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Scripts-Background**.
2. Paste the following script to generate the SHA-1 thumbprint value.

```
var certId = "<sys_id of the certificate record>";

var certAlias = "<alias name for the certificate>";

var gce = new GlideCertificateEncryption();

var thumbprint = gce.getThumbPrintFromKeyStore(certId,
 certAlias,"SHA-1");

gs.log(thumbprint);
```

3. Select **Run script**.

What to do next

Create an OAuth API script

Create an OAuth API script

Create and duplicate an OAuth API script for application registry.

Before you begin

Generate an SHA-1 thumbprint. For more information, see Generate a SHA-1 thumbprint.

Generate a sys_id when you configure a JWT provider. For more information, see Configure a JWT provider.

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Script Includes**.
2. Select and open **EmailCertificateOAuthTemplate**.
3. Duplicate the script by changing the name in the **Name** field, select **Insert and Stay** from the record menu.

   Note:
   The name must begin with **OAuth** to create an OAuth script.

   The script is copied and created with a different name.
4. In the copied script, enter the JWT provider sys_id and SHA-1 thumbprint.
5. Select **Update**.

What to do next

Register an OAuth Provider.

Create an email account with SMTP OAuth2.0

Create an email account from Microsoft Exchange Online with SMTP OAuth2.0.

Before you begin
Role required: email_account_admin

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Accounts**.
2. Select **New**.
3. On the form, fill in the fields.

Email Accounts form

Field	Description
Name	Name of the email account.
Type	SMTP
Authentication	Select **OAUTH 2.0**. Type of authentication used for the email account to connect to the email server.
Server	Server used for creating the email accounts.
Email user label	Display value used for outgoing messages. This field is for SMTP type accounts only.
User name	User name or ID to authenticate an email address. The value in this field is also the **From** address when the instance sends an email. If you're using SMTP, the value must be a full email address. The value in the **From** field can override this field for SMTP accounts.
From	The From email address used for notifications sent with this account. This address takes precedence over the **User name** field. If nothing is present in this field, the User name is used as the return address for notifications sent from the instance. **Note:** This value is for SMTP only.
OAuth Profile	OAuth profile configured during application registry.
Connection Security	If your email server supports TLS, an insecure connection is upgraded to a secure connection using the SSL/TLS encryption protocol.
Port	Connection TCP port.
Active	Option to activate the email account.
ServiceNow Configured	Identifier of whether this account is provisioned by ServiceNow. If you create an account, this option isn't selected.
Enable debug output	Option to create node logs for the raw data that is exchanged with the email server. You can enable this field temporarily to diagnose issues related to receiving or sending email.

servicenow

4. Select **Submit**.
5. Test the connection by selecting the **Test Connection** link.

Email properties

The Email Properties page is where you can configure settings for inbound and outbound email.

Email properties are available from either of these modules:

- **System Mailboxes** > **Email Properties**
- **System Properties** > **Email Properties**

Troubleshooting steps to take when a ServiceNow instance does not receive email

Email Properties page

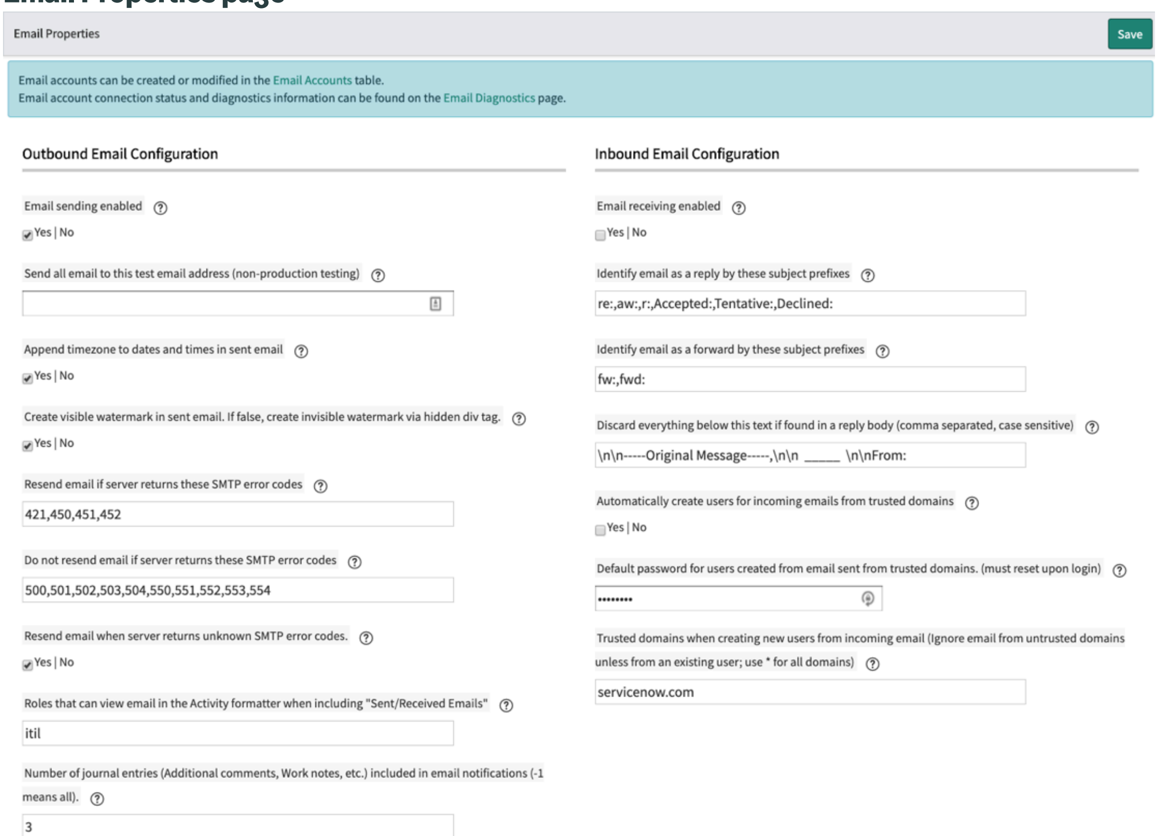

Email accounts are configured in the **System Mailboxes** > **Administration** > **Email Accounts** module. For more information and instructions, see Create an email account.

Email diagnostics are available from the **System Mailboxes** > **Email Diagnostics** module.

Note:
To learn more about the security properties that affect email processing, see SMTP authentication (deprecated) in Instance Security Hardening Settings.

Outbound mail configuration

The **Outbound Mail Configuration** section of the Email Properties page contains properties for sending email.

Outbound email configuration

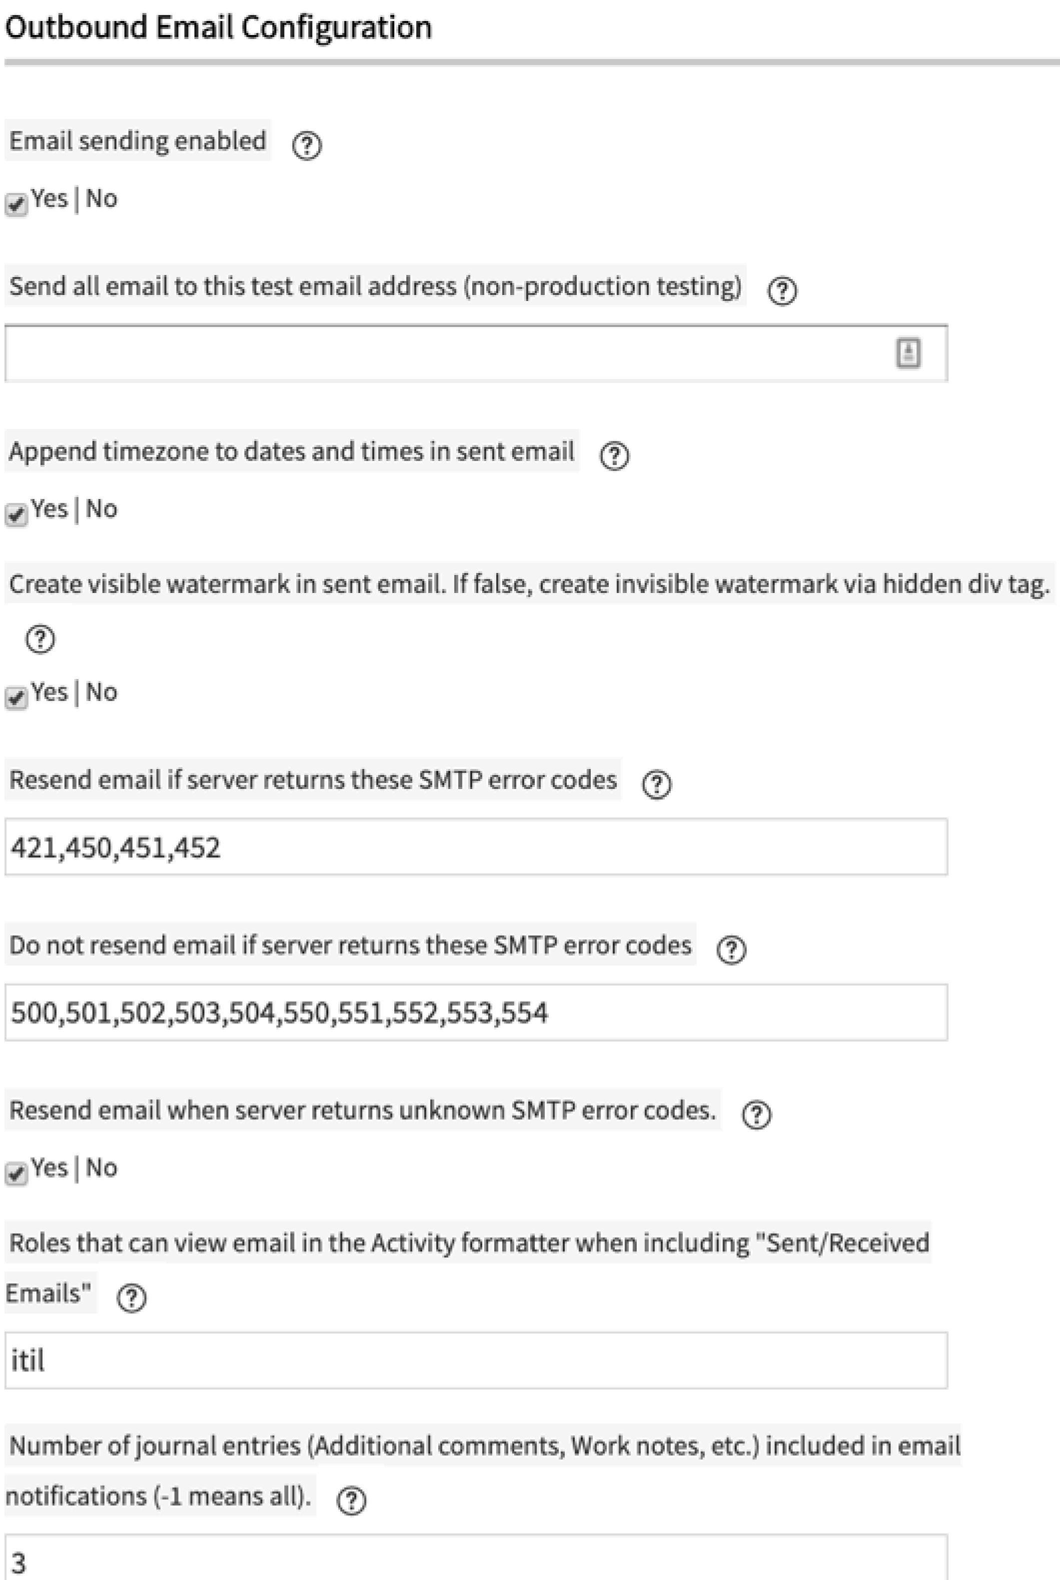

Outbound email properties

Property	Label	Description
glide.email.smtp.active	Email sending enabled	Specifies whether to enable or disable the outgoing mail server. • Type: true | false • Default value: true
glide.email.test.user	Send all email to this test email address (non-production testing)	Specifies the comma-separated list of email addresses to which the instance sends all email messages. Typically used in non-production instances for testing purposes. • Type: string • Default value: none
glide.email.append.timezone	Append time zone to dates and times in sent mail	Specifies whether to append the system time zone to date and date/time values in outbound emails. For example, 2010-07-02 04:01:14 PST. • Type: true | false • Default value: true
glide.email.watermark.visible	Create visible watermark in sent mail. If false, create invisible watermark via hidden div tag.	Indicates whether the watermark in email notifications is visible (true) or wrapped in a hidden div tag (false). • Type: true | false • Default value: false
glide.smtp.defer_retry_ids	Resend email if server returns these SMTP error codes	Specifies the comma-separated list of SMTP error codes that force the instance to resend email. • Type: string • Default value: 421,450,451,452
glide.smtp.fail_message_ids	Do not resend email if server returns these SMTP error codes	Specifies the comma-separated list of SMTP error codes that prevent the instance from resending email. • Type: string • Default value: 500,501,502,503,504,550,551,552,553,554
glide.smtp.default_retry	Resend email when	Enables (true) or disables (false) resending email when an unknown SMTP error

Outbound email properties (continued)

Property	Label	Description
	server returns unknown SMTP error codes.	code is encountered. The instance only recognizes the SMTP error codes defined in the *glide.smtp.defer_retry_ids* property. • Type: true | false • Default value: true
glide.ui.activity.email_roles	Roles that can view email in the Activity formatter when including "Sent/ Received Emails"	Specifies the comma-separated list of roles that can view email in the activity formatter. • Type: string • Default value: itil
glide.email.journal.lines	Number of journal entries (Additional comments, Work notes, etc.) included in email notifications (-1 means all).	Specifies the number of entries from a journal field, such as Additional comments or Work notes, included in email notifications. A value of -1 includes all journal entries. • Type: integer • Default value: 3 • Learn More:Restrict the Number of Entries Sent in a Notification

Inbound mail configuration

The **Inbound Mail Configuration** section of the Email Properties page contains properties to control inbound email.

Inbound email properties

Property	Label	Description
glide.email.read.active	Email receiving enabled	Specifies whether to enable or disable the inbound mail server. • Type: true | false • Default value: true
glide.email.reply_subject_prefix	Identify email as a reply by these	Specifies the comma-separated list of prefixes in the subject line that identify an email reply.

Inbound email properties (continued)

Property	Label	Description
	subject prefixes	• Type: string • Default value: re:,aw:,r:,Accepted:,Tentative:,Declined: **Note:** The case of the reply prefix in the email, for example RE:, must exactly match the case of the prefixes defined in this property. If, for example, an email contains the Re: prefix and only RE: is defined in the property, the email will not be recognized as a reply. Therefore, it is a best practice to define multiple versions of the prefix, including mixed-case versions, such as RE:, Re:, and so on.
glide.email.forward_subject_prefix	Identify email as a forward by these subject prefixes	Specifies the comma-separated list of prefixes in the subject line that identify a forwarded email. • Type: string • Default value: fw:,fwd: **Note:** The case of the forward prefix in the email, for example fw:, must exactly match the case of the prefixes defined in this property. If, for example, an email contains the Fwd: prefix and only fwd: is defined in the property, the email will not be recognized as a forward. Therefore, it is a best practice to define multiple versions of the prefix, including mixed-case versions, such as FWD:, Fwd:, and so on.
glide.pop3readerjob.create_caller	Automatically create users for incoming emails from trusted domains	Controls the behavior when an instance receives an email from an email address not associated with a user record. If this property is set to **true**, the instance creates a new user record for the email address and places that new user in the **Caller** field of any tickets created. If the property is set to **false**, the instance places **Guest** in the **Caller** field of any tickets created.

Inbound email properties (continued)

Property	Label	Description
		• Type: true | false • Default value: false • Learn More:Enabling Automatic User Creation
glide.user.default_password	Default password for users created from email sent from trusted domains. (must reset upon login)	Specifies the password for new users created from incoming email. Users must reset the password at first login. • Type: string • Default value: password • Learn More:Enabling Automatic User Creation
glide.user.trusted_domain	Trusted domains when creating users from incoming email.	Comma-separated list of trusted domains for which the instance automatically creates a user based on incoming emails. Use an asterisk (*) to trust all domains. If an email is not from a trusted domain, the instance processes the inbound email as a "guest user" but it does not create a guest user in the instance. • Type: string • Default value: servicenow.com • Learn More:Enabling Automatic User Creation

Email image filtering properties

Use email image filtering properties to control how inbound email images attach to a target record.

When a user sends an email to the system, email images (such as logos or email signatures) attach to the target record by default. These images are also visible in the activity stream of the target record. The system then replicates these images, which can cause duplicate image attachments to the target record and also duplicate images in the activity stream.

To filter images from emails and reduce duplicate image attachments to target records, configure the following properties in the System Properties [sys_properties] table. Use these properties to:

- Specify the email image sizes eligible for filtering.
- Choose an action that controls image attachment behavior and image visibility in the activity stream.

Note:
These properties do not delete attachments. You can verify attachment records in the Attachments [sys_attachment] table.

Email image filtering properties

Name	Description
glide.email.inbound.image_sys_attachment.filter.minimum_bytes	Sets the minimum image size in bytes for inbound email images to attach to the target record. Any image that is below this size limit is subject to the image filtering behavior that you set in the *glide.email.inbound.image_sys_attachment.filter.action* property. To disable image filtering entirely and attach all email images to the target record, set the minimum byte size to `0`. • Type: Integer • Default value: 0
glide.email.inbound.image_sys_attachment.filter.action	Specifies the image filtering behavior. **Note:** Regardless of the action you select, the image is visible when you choose to **Preview HTML Email** or **Show Email Details** in the activity formatter. • Type: Choice list ○ AttachTarget - Associate the image to the target record. The image is visible in the activity formatter and in the attachment to the target record. ○ AttachEmail - Attach the image to the email record. The image is not visible in the activity formatter of the target record nor in the attachment to the target record. ○ AttachNone - Do not attach the image to a record. You can select this option to attach the image manually at a later time. • Default value: AttachEmail

Related topics

Add a system property

Email digest properties

Several properties are available to manage digest intervals for email digests.

The following properties are available for the email digest feature.

Note:
To open the System Properties [sys_properties] table, enter `sys_properties.list` in the navigation filter.

Properties for email digest intervals

Property	Description
glide.email.digest.default_interval	The sys_id of the default email digest interval available to users. • **Type**: string • **Default value**: 28d157e07f1332007f005212bdfa9116 • **Location**: System Property [sys_properties] table
glide.email.digest.max_intervals	The maximum number of email digest intervals that can be defined. • **Type**: integer • **Default value**: 100 • **Location**: System Property [sys_properties] table • **Learn more**: For details on digest intervals, see Create or modify email digest intervals.

Advanced email properties

Use advanced email properties to fine-tune the way your instance sends or receives email.

Add these properties to the System Property [sys_properties] table before they can be used to overwrite the defaults.

com.glide.email.max_read

Specifies the maximum number of emails a POP3 reader should process concurrently.

- Type: integer
- Default value: 20

com.snc.on_call_rotation.reminders.showtz

Specifies whether to show a user's time zone.

- Type: true | false
- Default value: false

glide.cms.use_email_override_url

Forces the system to use the *glide.email.override.url* property, rather than the *glide.servlet.uri* property, when a notification has a link to a CMS page in an instance.

- Type: true | false
- Default value: false
- Location: System Property [sys_properties] table

glide.email_address_filter.max_domains

Sets the maximum number of domains that can be associated with a single email address filter.

- Type: integer
- Default value: 100

glide.email_address_filter.max_exceptions

Sets the maximum number of exceptions that can be associated with a single email address filter.

- Type: integer
- Default value: 1000

glide.email.allow_duplicate_message_ids

Allows storing of emails with duplicate message IDs and adds error message "Duplicate message-id encountered, prevented loop."

- Type: true | false
- Default value: false

glide.email.append.timezone

Appends the user time zone to all date and time values in outbound email messages (for example, 2018-07-02 04:01:14 PST). The time zone does not show if the system property *glide.sys.time_format* does not have a "z" value. For more information on configuring the time zone in email notifications, see Time zone for email notifications.

- Type: true | false
- Default: false

glide.email.email_client.check_attachment_availability

Enables a system scan of each file that a user attaches to an email draft in the email client. If a file is found to be infected with a virus, the user can't send the email with the attached file. The user must remove the attachment to send the email.

- Type: true | false
- Default value: true

glide.email.email_with_no_target_visible_to_all

Email that is missing a target record or whose target record is the email are made visible to any user regardless of their roles. By turning this property to false, the user is restricted unless they sent in the record or have an admin role.

- Type: true | false
- Value: false

glide.email.inbound.calendar_behavior

Specifies how the system stores calendar data, such as an invitation or an invitation response. Enter one of these options (not case sensitive):

- Attach: Store the calendar data as an attachment on the associated record, such as the incident or change that triggers an invitation.
- Ignore: Discard the calendar data.
- Inline: Store the calendar data as text in the email Body field.
- Type: string
- Default value: Attach

glide.email.inbound.check_attachment_availability

Prevents users from downloading and using virus-infected files that are attached to inbound emails.

The system detects infected files in an email by checking the ServiceNow virus header. If there are multiple email attachments, the system scans the files via Antivirus Scanning. All attachments that are found to be infected are made unavailable for use.

- Type: true | false
- Default value: true

glide.email.inbound.convert_html_inline_attachment_references

Specifies whether to convert inbound email HTML so email images appear in the email HTML body preview. The system displays broken cid (content ID) links in place of images received when this property is disabled. The format in which the system displays an email image depends on the property setting at the time the email is received, not the current property setting.

- Type: true | false
- Default value: true

glide.email.inbound.generate.missing.html.part

Enables the system to generate HTML text for inbound emails that contain only a plain text message.

- Type: true | false
- Default value: true

glide.email.inbound.generate.missing.text.part

Enables the system to generate plain text for inbound emails that contain only an HTML message.

- Type: true | false
- Default value: true

glide.email.inbound.max_attachment_count

Sets the maximum number of attachments allowed per inbound email.

- Type: integer
- Default value: 30
- Location: System Property [sys_properties] table

glide.email.inbound.max_body_chars

Sets the maximum body size in characters allowed per inbound email.

- Type: integer
- Default value: 524288

glide.email.inbound.max_total_attachment_size_bytes

Sets the maximum total attachment size in bytes allowed per inbound email.

- Type: integer
- Default value: 18874368
- Location: System Property [sys_properties] table

glide.email.inbound_action_set_email_ignored_when_user_disabled

Sets email to received ignored when met with one of two conditions:

- User account (guest or any other account used to run the inbound email action) is inactive.
- The user account (guest or any other account used to run the inbound email action) is locked out and the glide.pop3.process_locked_out property is set to false.

- Type: true | false
- Default value: true

glide.email.mail_to

Specifies the email address for sending notifications that use the ${mailto:} variable.

- Type: string
- Default value: SMTP email address that is active by default

glide.email.name_split

Specifies the delimiter used between first and last names in an email address. For example, a delimiter of "." (period) in the email address john.smith@company.com tells the system to look for a user record for John Smith.

- Type: string
- Default value: period (.)

glide.email.notification.save_when_no_recipients

Controls whether a notification-generated `sys_mail` record is saved even if there are no recipients. Used along with other notification recipient logging properties, this property enables investigating problems with notifications.

- Type: true | false
- Default value: false

glide.email.outbound.check_attachment_availability

Prevents users from sending emails that contain virus-infected files.

If an outbound email contains an attachment that was previously marked as infected via Antivirus Scanning, then the system doesn't send the email. An event is created to indicate that the email attachment is not available for use.

- Type: true | false
- Default value: true

glide.email.outbound.header.auto_submitted

Stores the value used in the "Auto-submitted" outbound email header. Clear the property value to remove the "Auto-submitted" header from all outbound emails. Some spam filters flag auto-generated email as spam.

- Type: string
- Default value: auto-generated

glide.email.outbound.max_attachment_count

Sets the maximum number of attachments allowed per outbound email.

- Type: integer
- Default value: 30
- Location: System Property [sys_properties] table

glide.email.outbound.max_body_chars

Sets the maximum body size in characters allowed per outbound email.

- Type: integer
- Default value: 524288

glide.email.outbound.max_total_attachment_size_bytes

Sets the maximum total attachment size in bytes allowed per outbound email.

- Type: integer
- Default value: 18874368
- Location: System Property [sys_properties] table

glide.email.override.url

Sets the URL to use in emailed links in place of the instance URL. The URL should end with nav_to.do. An example value is `https://servicenow.customerdomain.com/production.nav_to.do`. This property is suitable for customers who use custom redirect URLs for their instances.

- Type: string
- Default value: instance URL

glide.email.remove_illegal_address_quotes

Removes invalid quotation marks from an inbound email address, thus allowing the inbound email address to be accepted as valid. For example, when you set the property to true, the address "john.doe@example.com" is changed to john.doe@example.com. However, only addresses that are formatted as John Doe <'john.doe@example.com'> are removed entirely.

- Type: true | false
- Default: false

glide.email.smtp.max_recipients

Specifies the maximum number of recipients the instance can list in the To:, CC:, BCC: lines combined for a single email notification. Notifications that would exceed this limit instead create duplicate email notifications addressed to a subset of the recipient list. Each email notification has the same maximum number of recipients.

- Type: integer
- Default value: 100

glide.email.smtp.max_send

Specifies how many emails to send through each new SMTP connection. If there are more emails to send than the specified value, the instance establishes a new SMTP connection.

- Type: integer
- Default value: 100

glide.email.smtp.max_recipients_overall

Specifies the overall maximum number of recipients an email can be sent to. Specifies the maximum number of recipients the instance can list in the To:, CC:, BCC: lines for a single email notification.

- Type: integer
- Default value: 100000

glide.email_system_address_filter.max_address_filters

Sets the maximum number of email filters that can be associated with a single system address filter.

- Type: integer
- Default value: 100

glide.email.text_plain.strip_xhtml

Converts the XML of both outbound and inbound emails that are shown in comments to plain text. If false, the system preserves the XML.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging

Enables or disables logging all reasons a recipient was excluded. If set to true, the subsequent properties dealing with the exclusion of logging are enabled. If false, the properties are disabled. This property cannot suppress log messages generated by the `glide.email.test.user` property.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.device_inactive

Logs recipients who are excluded because their chosen notification device record is marked as inactive.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.device_schedule

Logs recipients who are excluded based on the **Schedule** field on the Notification Device [cmn_notif_device] record for their chosen notification device.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.event_creator

Logs recipients who are excluded because they initiated the notification event, such as updating an incident record. The **Send to Event Creator** check box is then cleared on the notification record.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.invalid_email

Logs recipients who are excluded because the email address for that user is invalid, for example the @ is missing, or empty.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.user_calendar_integration_disabled

Logs recipients of calendar invitations who are excluded. Recipients are excluded because the **Calendar Integration** field is set to **None** on the user record.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_logging.user_inactive

Logs recipients who are excluded because the **Active** check box is cleared on the user record.

- Type: true | false
- Default value: true

glide.notification.recipient.exclude_ logging.user_notification_disabled

Logs recipients who are excluded because the **Notification** field is set to **Disabled** on the user record.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging

Enables or disables logging all reasons a recipient was included. If set to true, the subsequent properties dealing with the inclusion of logging are enabled. If false, the properties are disabled.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.delegate

Logs recipients who are included because they are delegates of another user.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.event_parm

Logs recipients who are included because they are in the **parm1** or **parm2** fields of the event record.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_ fields

Logs recipients who are specified in the **Users/Groups in Field** field for the notification record. The recipient_fields are fields in the target record that contain a recipient to add. For example, to include an incident assignee as a recipient, set the **Users/Groups in Field** field to **Assigned to**.

- Type: true | false
- Default value: true

glide.notification.recipient.include_ logging.recipient_groups.group_email

Logs users who are included in a notification that is sent to a group. The property also logs users in the event **parm1** or **parm2** field.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_groups.manager

Logs recipients who manage a recipient group. The property also logs users in the event **parm1** or **parm2** field.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_groups.membership

Logs recipients who are part of a notification recipient group. The property also logs users in the event **parm1** or **parm2** field.

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.recipient_users

Logs recipients who are included via notification record's **Users** field (recipient_users).

- Type: true | false
- Default value: true

glide.notification.recipient.include_logging.subscription

Logs recipients because they are subscribed via the user notification preferences.

- Type: true | false
- Default value: true

glide.pop3.parse_start

Specifies the text that indicates the beginning of the email body section. The instance parses name:value pairs within this section to set or update field values when processing inbound email actions.

- Type: string
- Default value: none

glide.pop3.parse_end

Text indicating the end of the email body section where the instance should parse name:value pairs to update field values when processing inbound email actions.

- Type: string
- Default value: none

glide.smtp.dateformat

Specify the date format to use for outgoing email notifications.

- Type: string
- Default value: date format listed in email sender's user record `[sys_user.date_format]`

glide.smtp.precedence_bulk

Specifies whether outbound email includes the header "Precedence: bulk". Some spam filters flag bulk email as spam. Set the value to **false** to remove this header from outbound email.

- Type: true | false
- Default value: true

glide.smtp.send_partial

Splits outgoing email between valid and invalid recipients. The email is sent only to the valid recipients. A new email containing the invalid recipients is created in the **send-ready** state and attempts to send again.

- Type: true | false
- Default value: false

glide.smtp.timeformat

Specify the time format to use for outgoing email notifications.

- Type: string
- Default value: time format listed in email sender's user record `[sys_user.time_format]`

glide.ui.activity.email.use_display

Specifies whether to display email addresses or user IDs (display value from the User table) in email headers. If true, the instance searches for a user record with a matching email address. If it cannot find a matching user record, it displays the email address.

- Type: true | false
- Default value: false

glide.activity.rule.sys_email.expand_max_body_size

Specifies the limit of the volume of data which can be processed via a single click of Expand All or getEmailBodies operation. If limit value is hit the user is notified that The maximum allotted email text threshold has been reached and some mails are not expanded.

glide.ui.email_client.email_address.disambiguator

Sets the columns from the User `[sys_user]` table that the auto-complete list displays. Separate each column name with a semicolon character (;).

- Type: string
- Default value: name

glide.ui.incident_activity.max_addresses

Specifies the maximum number of addresses to list in an email audit record. If the number of addresses exceeds this limit, the instance truncates the list after the maximum value and displays an ellipsis character (...).

- Type: string
- Default value: 5

NotifyAffectedCI.max_rel_level

Sets a value that is used by the **Affected ci notifications** business rule. The business rule notifies subscribers when tasks affect configuration items (CIs). The business rule generates notifications for parent CIs up to the level defined by this property. You can adjust the property value according to the complexity and depth of your CI relationships.

- Type: integer
- Default value: 5

glide.email.inbound.attachment_extensions_blacklist

Restrict attachments with questionable file extensions for inbound email.

glide.email.smtp.claim.lookback.hrs

Sets the default value of 24 hours. Any email older than 24 hours will not be sent by the system. This protects the system from having a huge number of sent emails at a time or if the outbox has a large number of emails waiting to be sent.

Set to a value greater than 24 (hours) and increase it in intervals to not send a large number of emails at one time. Depending on how big the backlog of emails is and how spaced out they are created in time, you might want to start with a value of 36 (1.5 days worth), let those go out, then set the value again to maybe 48 (two days), and keep increasing the value until all of the older emails are sent.

- Type: integer
- Value: 24

glide.email.outbound.static_translation.enabled

Enables the system property for static translation of notifications for the global application.

- Type: true | false
- Default value: true

glide.notification.translation.dynamic

Enables the system property for dynamic translation of notifications for the global application.

- Type: true | false
- Default value: true

Activate email administration

Activate email filter and email service.

Activate email filters

Administrators can activate the Email Filters plugin (com.glide.email_filter).

Before you begin
Role required: admin

Procedure

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

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

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

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

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

What to do next
Create an email filter

Activate Email Service

Users with the admin role can activate the Email Service plugin (com.glide.email.service) to enable the Email API.

Before you begin
Role required: admin

Procedure

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

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

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

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

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

What to do next
Create a system address filter

Related topics

List of plugins (Washington DC)

Grant access to the Email API

To grant non-admin users access to the Email API, administrators must create an access control and assign a special role.

Before you begin

Role required: admin

About this task

For additional details, see Email API .

Procedure

1. Elevate privileges to security_admin.
2. Navigate to **System Security** > **Access Control (ACL)**.
3. Click **New**.
4. Create an access control for the Email [sys_email] table.

Access control values

Field	Value
Type	record
Operation	create
Admin overrides	Selected
Name	sys_email
Requires role	email_api_send

5. Assign the email_api_send role to the non-admin user.
6. Verify the non-admin user either has read access to all records in the notification target table or has read access to a specific record in the table.
   For example, to work with incident notifications the non-admin user needs access to the Incident table. You can provide record access by:
   - Granting the user an appropriate role, such as the itil role.
   - Selecting a specific incident that was opened on behalf of the user.

Creating an email account

Use email administrator to create email accounts, email filter, and create an email account for Microsoft Graph.

Email accounts

Email accounts store the connection details and credentials the system uses to access external email servers and services.

By default, instances come with email accounts for ServiceNow - managed SMTP and POP3 servers to send and receive email. These accounts cannot be modified, but can be deactivated if you do not want to use them. To connect to other email servers and services you must create

email accounts for them. You can create email accounts for servers and services that use these protocols:

- SMTP
- POP3
- IMAP

Note:
The system only allows one SMTP email account to be active at a time and sends all email through this account. You can however receive email from multiple POP3 or IMAP accounts.

The system stores individual email accounts in the `sys_email_account` table. Create separate email accounts to send and receive email. Use email properties to define how the system processes email for all email accounts.

By default, the ServiceNow POP3 server provides each instance with its own mailbox at the address instance@service-now.com. Likewise, the ServiceNow SMTP server sends email from the address instance@service-now.com. To create another mailbox, provision your own POP3 or IMAP server and create an email account to the server on the instance. To change the sent from email address, update the address in the email account you use for sending email.

Emails received via a POP3 or IMAP server account are stored in the Email [sys_email] table and then deleted from the account.

Note:
ServiceNow has exclusive access to an account you configure. After a mail is read from a POP or IMAP account that you have configured, ServiceNow deletes the email from the server.

Create an email account

You can create email accounts in addition to the accounts provisioned for you.

Before you begin

- Role required: admin
- Email server: a compatible email server
  - SMTP
  - POP3
  - IMAP

Warning:
An SMTP server that uses a localhost or loopback (127.0.0.1) IP address is not accepted as a compatible email server. For more information, see KB0724199.

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Accounts**, and then click **New**.
2. On the form, fill in the fields.

Field	Description
Name	Name of the email account.
Type	Type of mail server. You can select one of the following:

Field	Description
	○ POP3 ○ IMAP ○ SMTP **Note:** The system can use only one active SMTP account at a time.
Authentication	Type of authentication used for the email account to connect to the email server. You can select one of the following: ○ –- None -- ○ Password ○ OAUTH ○ OAUTH 2.0 The Email - OAUTH support for IMAP and SMTP (com.glide.email.oauth) plugin must be active for the OAuth options to be visible.
OAuth Provider	OAuth application registry record for this account. This field appears if you selected **OAuth 2.0**.
Server	Remote server to which this account connects. To activate a server for an on-premise installation, enter the full address (FQDN) of the node (for example, `node.customerdomain`).
Active	Option to activate the email account.
ServiceNow Configured	Identifier of whether this account is provisioned by ServiceNow. If you create an account, this option is not selected.
Enable debug output	Option to create node logs for the raw data that is exchanged with the email server. You can review the node logs by navigating to **System Logs** > **Utilities** > **Node Log File Browser**. You can enable this field temporarily to diagnose issues related to receiving or sending email.
Email user label	Display value used for outgoing messages. This field is for SMTP type accounts only.
User name	User name or ID to authenticate an email address. The value in this field is also the **From** address when the instance sends email. If you are using SMTP, this must be a full email address. The value in the **From** field can override this (for SMTP accounts). **Note:** The address in the **From** field on the Notification [sysevent_email_action] form takes precedence over this field.
Password	Password when the **Authentication** type is **Password**.

Field	Description
	Note: You may need to increase the size of this field to accommodate longer passwords. By default, this field has a size of 40 characters.
From	(SMTP only) The From address used for notifications sent with this account. This address takes precedence over the **User name** field. If nothing is present in this field, the User name is used as the return address for notifications sent from the instance.
Connection Security	Type of secure connection. Choose a setting: **None** No secure protocol is used. **STARTTLS** Upgrades an insecure connection to a secure connection using the SSL/TLS encryption protocol, if your email server supports TLS. **SSL/TLS** Connect to an SSL/TLS encrypted port to secure the connection. Email is encrypted between the ServiceNow instance and your mail server. **Warning:** Selecting a less secure protocol like **STARTTLS** or **None** may expose your data. To better ensure the security of data in your email server, select **SSL/TLS**.
Port	Connection TCP port.
System Address Filter	System address filter to apply to the email account. If left blank, the system uses the default system address filter for inbound or outbound email. For more information, see System address filters.

3. Click **Submit**.
4. Click the test connection link at the bottom to test the SMTP, POP3, or IMAP account.

   The system displays a pop-up window with the results of the connection test.

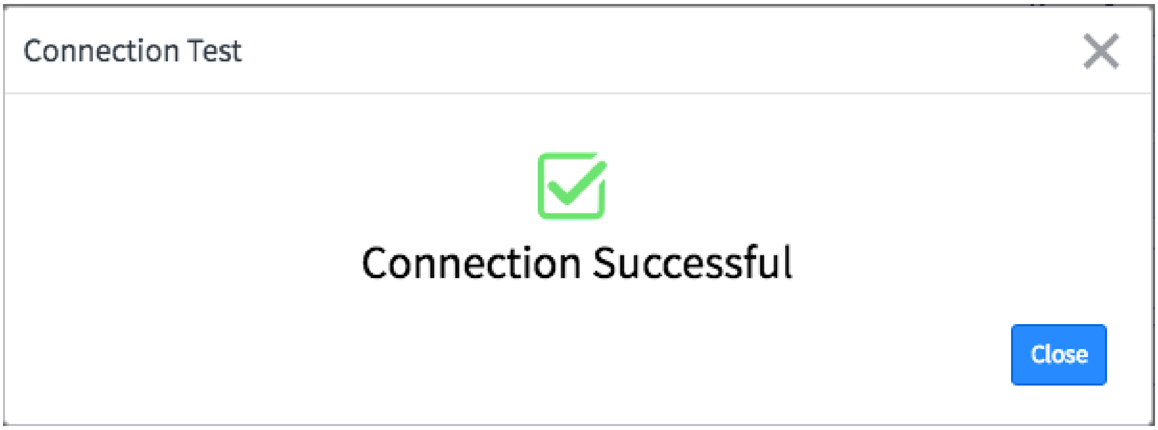

If the test succeeds, click **Close** to return to the email account record. If the test fails, click **View Logs** to display more information about the test results.

Result
The system uses the email account to send or receive email.

What to do next
Enable the system to send or receive email.

Multiple email readers

Organize POP3/IMAP email accounts into email account groups (subsets of email accounts) that can be processed by multiple email reader jobs. To improve the performance of inbound email account processing, you can incrementally add email account groups and if needed, email reader jobs to process email account groups concurrently.

How email account group processing works

By default, the email reader job processes POP3/IMAP email accounts serially, which can result in longer processing times when you have many email accounts. It can also be difficult to identify any problematic email accounts that might contribute to slower processing times.

To reduce email account processing time, you can create one or more email account groups that contain subsets of your email accounts. Organizing your email accounts into groups enables the default email reader job to process each email account group separately. After creating one or more email account groups, you can then add another email reader job to process account groups concurrently.

Use the `glide.email.inbound.account_group_processing` system property to activate email account group processing.

During email account group processing:

- The system stores the email accounts in the Email Account Groups [sys_email_account_group] table.
- The default email reader job runs every two minutes. The email reader chooses the account group to be processed based on the account group processing state and the date and time that the group was last processed.
  - Claimed state: If an account group has a claimed state, the email reader is processing the account group. If you defined other email reader jobs, no other email reader job can claim the group for processing.
  - Unclaimed state: If an account group is unclaimed, the account group is available for processing by the email reader.

- The email reader job processes the email accounts contained in the selected account group.
- The email reader job chooses the next available account group for processing and continues processing that account group.

You can monitor the processing of each email account group by using the Email Account Groups [[sys_email_account_group]] table. Each account group record includes the processing duration (how long it takes for the email reader to process all the accounts in the group).

After you review the processing times for your email account groups, determine whether to make further adjustments in email account group processing. For example, you can create another email account group or create another email reader job to process your email account groups in parallel. You define additional email reader jobs using the Schedule [sys_trigger] table.

Set up email account group processing

To set up email account group processing, you must have the email_account_admin or admin role.

1. **Create email account groups.**

   Define one or more email account groups that contain a subset of your POP3/IMAP email accounts. Your base system includes a default email account group. Any POP3/IMAP accounts not contained in an email account group are processed as part of the default email account group.
2. **Enable email account group processing**

   Navigate to sys_properties.list and locate the `glide.email.inbound.account_group_processing` system property. Set the **Value** to **true**. The email reader job starts processing the email account groups that you created.
3. **Monitor email account group processing.**

   Check the status of your email account groups to see if processing time has been reduced. Determine whether the additional groups sufficiently reduce account processing time or consider adding another email reader job to further reduce processing time.
4. **Create an email reader job.**

   As part of fine-tuning email account group processing, create an email reader job to process email account groups concurrently, in addition to the default email reader job.

Create email account groups

Define an email account group that contains a subset of your POP3/IMAP email accounts. The email reader job automatically processes each email account group as scheduled.

Before you begin

Role required: admin or email_account_admin

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Account Groups** and click **New**.
2. Fill in the form.

Field	Description
Name	The name of the account group.

Field	Description
Email Accounts	A subset of POP3/IMAP email accounts. Select the email accounts that make up the group. **Note:** When you add an email account to a group, it cannot be reused in other account groups.
Active	Check box that enables the email account group for processing. New account groups are active by default.
Default	Read only. Check box that indicates the account is the default email account.
Status	Read only. Processing state of the email account group: Unclaimed or Claimed.
Claimed by	Read only. ID of the email reader job currently in progress.
Last Processed	Read only. The last time that the email reader job processed this group.
Processing Duration	Read only. The length of time taken by an email reader job to process the account group.

3. Click **Create**.
   The system updates the Email Account Group [sys_email_account_group] table with the new account group. Use the this table to monitor email account group processing.

What to do next

- Review email account group processing.
- Determine if you want to continue fine-tuning email account processing. You could add another email account group or another email reader job to process email account groups concurrently.

Enable email account group processing

Enable the email reader job to start processing your email account groups so that you can reduce the time it takes to process inbound email accounts.

Before you begin

Create email account groups.

Role required: admin

Procedure

Add a system property with the following settings:

- Name: *glide.email.inbound.account_group_processing*
- Type: true | false
- Value: **true**

What to do next

Check the status of your email account groups to see if processing time has been reduced. For more information, see Monitor email account groups.

To further reduce processing time, consider creating another email reader job to process inbound email accounts. For more information, see Create an email reader job.

Monitor email account groups

Use the Email Account Groups [sys_email_account] table to check the status of email account groups processed by email reader jobs.

Before you begin

1. Create email account groups.
2. Enable email account group processing.

Role required: email_account_admin or admin

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Account Groups**.
2. Review the processing details for each account group:

Field	Description
Name	The name of the email account group.
Active	The operational state of the email account group. The value true indicates that the account is enabled for email account processing.
Claimed by	The ID of the email reader job if the job is running on the account group.
Default	The default indicator. The value true indicates that the email account group is the default.
Email Accounts	The names of the email accounts contained in the email account group.
Last Processed	The day and time that the email reader job last processed this account group.
Processing Duration	The length of time (in hours, minutes, seconds) to process the account group.
Status	Processing state of the email account group: ◦ Unclaimed: The account group is available for processing by the next email reader job. ◦ Claimed: An email reader job is processing the current account group.

Example:

Email account group — unclaimed processing state

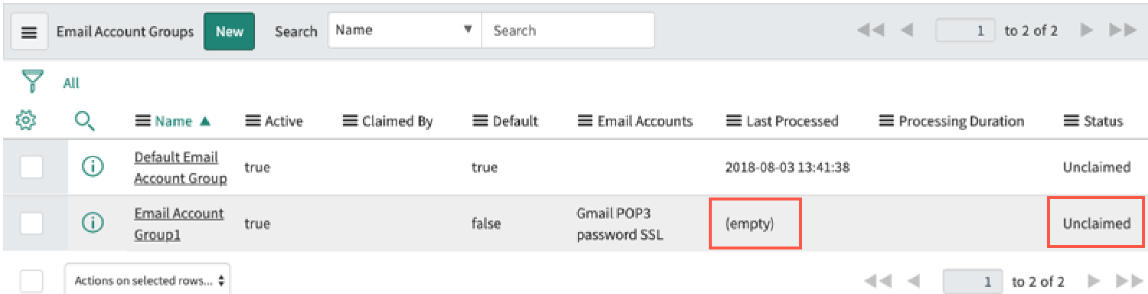

Email account group — claimed processing state

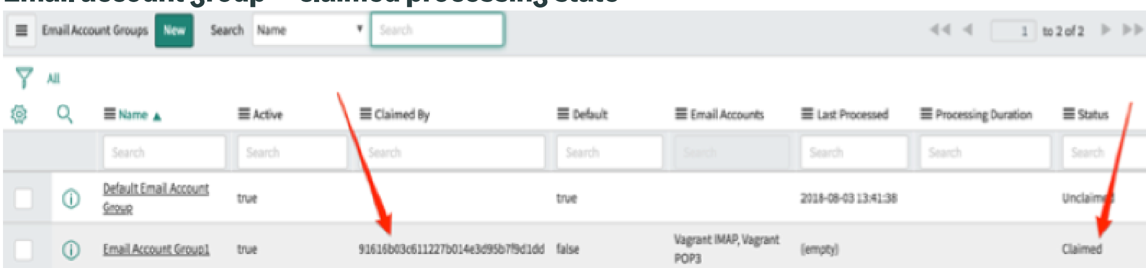

What to do next

If fine-tuning email accounts, consider doing one of the following:

- Create another email account group.
- Create another email reader job to process email account groups concurrently.

Create an email reader job

If you created email account groups to fine-tune inbound email account processing, you can create an email reader job to process those account groups concurrently, in addition to the default email reader job.

Before you begin

1. Create email account groups.
2. Enable email account group processing

Role required: email_account_admin or admin

Procedure

1. In the Navigation filter, type `sys_trigger.list`.
2. In the sys_trigger table, select the Email Reader job record.
3. In the Email Reader job form:
   a. Enter a unique name for the email reader job to be added so that you can differentiate it from the default email reader job.
   b. If needed, change the time interval for this email reader job.
4. From the form context menu, click **Insert**.
   Another email reader job is created. This job runs concurrently with the default email reader job to process email account groups.

What to do next

Monitor email account group processing and determine if the additional reader job reduces processing time. If needed, you can continue fine-tuning email account processing. For example, you might consider adding another email account group.

Email filters

Specify which inbound emails to ignore or move to a particular mailbox.

When an email is ignored, the email is saved to your instance but is not processed. You can access an ignored email by viewing its Email [sys_email] record.

The Email Filters (com.glide.email_filter) plugin is active by default for releases after Kingston, for releases before Kingston, it needs to be activated, for more information see Activate email filters.

Default email filters

By default, the following filters are available from the **System Mailboxes** > **Administration** > **Filters** module:

Default email filters

Filter	Description
Ignore VCAL	Ignores all email containing vCalendar requests. This filter prevents inbound email actions from creating unnecessary incident records when the instance receives a response to sent email. vCalendar requests in email responses are identified by the EmailUtils script include.
Ignore header	Ignores email that contains specific headers. This filter overrides the *`glide.pop3.ignore_headers`* property.
Junk email - sender equals recipient	Filters out emails that are from the same user as the SMTP user to the Junk folder.
Ignore subject	Ignores email with specific terms or phrases in the subject line. This filter overrides the *`glide.pop3.ignore_subjects`* property. This filter might not apply to emails arriving from unknown users. Unknown users can be locked out.
Move spam to junk folder	Moves email identified as spam to the Junk folder. This filter checks for the value of the ServiceNow spam header. If the header is `X-ServiceNow-Spam-Status:Yes`, the filter moves the email to the Junk folder

Note:
Filter action should be configured to indicate how the system should react when the conditions of this filter evaluate to true. The two choices in the base system are: Mark as Ignored and Move to Junk. You can use the Action script to perform additional email tasks.

To learn more about enabling spam scoring and filtering, see Email spam scoring and filtering (instance security hardening) in Instance Security Hardening Settings.

Email filter script include

Email filters use a script include called `EmailUtils` that contains a simple utility function to determine if `vCalendar` is in the body of the response email. The results of this query are used in a condition script in the *Ignore VCAL* email filter.

Spam scoring and virus scanning

Every message sent through email servers is assessed for the likelihood of being spam. Based on this assessment, the instance adds headers to each message that can be used for filtering within the customer instance using the Email Filters plugin.

The system also adds the `X-ServiceNow-Virus:INFECTED` header to an email that contains one or more virus-infected attachments. The system ignores the email.

Spam scoring and virus scanning are available only for instances that use the ServiceNow email infrastructure. For more information on spam scoring and filtering, see KB0549426.

Related topics

Create an email filter

Create an email filter

You can create email filters to apply a custom action script or filter actions when email matches your filter's conditions.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Filters**.
2. Click **New**.
3. On the form, fill in the fields.

Field	Description
Name	Name of the email filter.
Application	Application scope of the email filter.
Order	Priority that the email filter takes over other filters for the same table. To give the filter higher priority over other filters, enter a lower number.
Active	Option to activate the email filter.
Short description	Description of what the filter does.
Conditions	Conditions to define which emails to filter.
Condition script	Script to define which emails to filter. Use this script to access other tables, variables, or methods. For example, the condition script for the Ignore VCAL filter looks for the results of the EmailUtils script include, which detects vCalendar attachments.
Action script	Script to define additional behavior to be performed when the filter condition evaluates to true. This might include an email action not included in the default filter actions. In the base system, the action script is used to produce the error string that is published to the email log (**System Logs** > **Email**).

4. Right-click in the header bar and select **Save** from the context menu.

   The *Filter Actions* related list appears.

5. Click **New** in the related list.
6. Select how the system should react when the conditions of this filter evaluate to true.

   The two choices in the base system are: *Mark as Ignored* and *Move to Junk*. You can use the *Action script* to perform additional email tasks.

   Note:
   Before creating additional email actions with a script, see Inbound Email Actions. Like business rules, inbound email actions use both conditions and scripts and can provide a number of useful actions on emails the instance receives.

Email service

The `Email Service` API provides endpoint to install the Email API on the instance.

By default, `Email Service` only allows admin users to access the Email API. Administrators can grant access to the Email API by:

1. Adding an access control for creating records in the Email [sys_email] table linked to the email_api_send role.
2. Assigning the email_api_send role to the non-admin user.
3. Verifying the non-admin user has read access to the notification target table. For example, an email notification for an incident, requires read access to the Incident table.

Non-admin users with the email_api_send role can access the `Email` API from the REST API Explorer.

Create a system address filter

Define how email address filters apply to inbound and outbound email.

Before you begin

Complete the steps in Set email address filters.

Role required: email_account_admin

About this task

You can use a combination of allowed list and denied list email address filters in a system address filter.

A system address filter blocks a domain or address that appears only on denied lists and never on allowed lists. To ensure that a system address filter effectively blocks a certain domain or address, check for the following:

- The domain or address is included in at least one denied list
- The domain or address is not included in any allowed list

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **System Address Filters**, and then click **New**.
2. On the form, fill in the fields.

Field	Description
Name	Name of the system address filter.

Field	Description
Active	Option to activate the system address filter.
Type	Type of system address filter. Select one of the following: **Outbound** Control which domains and email addresses can receive email from your instance. Select this option if you plan to use the system address filter for an SMTP server. **Inbound** Control which domains and email addresses can send email to your instance. Select this option if you plan to use the system address filter for an IMAP or POP3 server.
Application	Application scope of the system address filter.
Default	Option to use this system address filter as the default inbound or outbound filter. To control senders and recipients for a specific email account, leave this field cleared.
Short description	Description of the system address filter.
Email Address Filters	Email address filters to apply to the system address filter. By default, the maximum number of email address filters that you can associate with a single system address filter is 100. You can reconfigure this limit by setting the `glide.email_system_address_filter.max_address_filters` property.

3. Click **Submit**.

Result
A default outbound filter applies to all active SMTP email accounts automatically. A default inbound filter applies to all active IMAP or POP3 email accounts automatically.

For a non-default filter, the next step is to apply the filter to an email account manually. For more information on applying a filter to an email account, see Create an email account.

Email encryption - S/MIME protocol

Secure/Multipurpose Internet Mail Extensions (S/MIME) is an end-end encryption protocol for sending digitally signed and encrypted emails that support data confidentiality, authenticity, and integrity.

Introduction to S/MIME

An administrator with privileges can enable and configure S/MIME. Understanding of the following is required when using the full capabilities of S/MIME:

- Digital signatures and signature verification
- Message encryption and decryption

- Public key
- Digital certificates

Digital signatures and verification

With digital signature, S/MIME verifies the identity of the sender of the email. This verification ensures the following:

- Message in the email is the exact message sent by the sender.
- Message is received from the right sender and not someone pretending to be the sender.

Message encryption and decryption

S/MIME uses encryption to protect the content of the email, which ensures that only the receiver can decrypt the content. Encryption creates coded information so that it cannot be read or understood until it is decoded and readable. Message encryption helps with the two key security factors of confidentiality and data integrity.

Public key

S/MIME uses key pairs and asymmetric cryptography. A private key in a key pair belongs only to the sender. If the private key has been used, the owner of that key has used it.

Public key cryptography ensures secure communication between the sender and the receiver. Both have a key-pair, with one being private and the other public.

Public keys are shared between the sender and the receiver. A public key is paired to only one private key. The corresponding public key is used to identify its paired private key and only its paired private key. A public key can be used by multiple recipients.

A key pair can be used to

- Sign and verify a signature
- Encrypt and decrypt the content of an email

S/MIME digital signatures and encryption require each sender and recipient to have it enabled. They also need to send or exchange public keys though digital certificates to identify each other.

For more information about key management and cryprographic module, see Key Management Framework Reference.

Digital certificates

Digital certificates help in delivering the public key in the key pair. A digital certificate is a digital credential that provides information about the identity, validity, and any other required information. Digital certificates are issued by a certification authority (CA) and are valid for only a specific period of time.

Note:
ServiceNow® does not provide S/MIME certificates for ServiceNow mail infra users. Users should get their S/MIME certificates issued from the third party S/MIME certificate solution providers.

S/MIME outbound emails

Signing outbound mails

The ServiceNow AI Platform uses the private key of the sender (instance email account) and the receiver uses the public key to verify signatures.

Encrypting outbound mails

The ServiceNow AI Platform uses public keys of the recipients to encrypt the emails and every recipient uses their private key to decrypt the email.

S/MIME for inbound email

Sign verification for inbound mails

The sender uses a private key to sign the email and the ServiceNow AI Platform uses the public key of the sender to verify the signature.

Decrypting inbound mails

The sender uses the public key to encrypt the email and the ServiceNow AI Platform uses the private key to decrypt the email.

Activate S/MIME Email

You can activate the S/MIME Email plugin (com.glide.email.smime) for Notifications if you have the admin role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the S/MIME Email plugin (com.glide.email.smime) using the filter criteria and search bar.

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

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

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

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

What to do next

Import an S/MIME key pair

Setting up S/MIME for email

S/MIME is a protocol for sending digitally signed and encrypted emails to ensure the confidentiality, authenticity and integrity.

Set up S/MIME for outbound mails (signing and encryption)

You can use S/MIME for outbound mails for the following purposes:

- Digital signature
- Encryption
- Digital signature and encryption

To set up S/MIME, the admin must have the following:

- email_account_admin and sn_kmf.cryptographic_manager roles
- Key pair (private key and public key)
- CA certificate
- Email certificate

Upload the instance email account key pair and email certificates, and enable outbound S/MIME properties. For more information, see Enable S/MIME.

If there are multiple recipients and some of the recipients do not have valid certificates, the email will be sent only to recipients with a valid certificate.

Set up S/MIME for inbound mails (sign verification and decryption)

S/MIME for inbound mails can be used for the following:

- Signature verification
- Decryption
- Signature verification and decryption

For information about enabling inbound S/MIME properties, see Enable S/MIME.

If the system fails to decrypt a message, no inbound actions will run on the email and it is moved to the received-ignored status.

To set up the system to ignore inbound signed emails if the signature cannot be verified or is invalid, admins can create the `email.inbound.smime.ignore_unverified_emails` and set it to true.

For more information about key management and cryprographic module, see Key Management Framework Reference.

Import an S/MIME key pair

Import an S/MIME key pair consisting of the private key and certificate to sign outbound emails or decrypt emails.

Before you begin
Role required: email_account_admin and sn_kmf.cryptographic_manager

Upload the private-public key pair corresponding to the instance email account.

You must have the key alias.

Note:
When you create a PKCS12 file, ensure that it's created with a key alias.

Make sure the file format is PKCS #12.

About this task

This video shows you how to perform the following procedure. This video shows how to import an S/MIME key pair

Procedure

1. Navigate to **All** > **System Mailboxes** > **Administration** > **Email Accounts**.
2. Select the email account.

   Note:
   The keypair is associated with the email address and not the email account.

3. Select the **Import SMIME Key Pair** related link.
4. Select the **Crypto Specifications** tab and select the key alias.
   The algorithm definition screen is displayed.
5. Complete the Algorithm Definition form.

   Algorithm Definition form

Field	Description
Crypto module	Read only. Name of the selected cryptographic module displays.
Crypto purpose	SMIME Crypto Purpose
Algorithm	Type of algorithm used to accomplish the crypto purpose. The algorithm also controls the key origin. Adjusts automatically based on the selected crypto purpose.
Enroll module for resource exchange	Allows clones or backups to get their own unique keys. **Note:** Do not check the box.

6. Select **Next**.
7. In the **Lifecycle Definition** screen, select a Key Lifecycle from the Applies to column.

   Key Lifecycle fields

Applies to	Selected key that the lifecycle applies to.
For field	Control for the key that the lifecycle applies to.
Type	Select if the valuation for the key lifecycle is a relative value or an absolute value or none. Relative value depends on other data entries in the system, such as key generation, activation, and deactivation.

	Absolute value is an exact value, such as a date.
Lifecycle default	Read only. Displays a value if set.
Order	Enter the sequence in which to process the key lifecycle state for the crypto specification.
Relative duration	Number of years, months, or days the key is valid.
Relative duration type	Duration of the lifecycle: Years, Months, or Days.
Relative operation	**Before** or **After**.
Relative to	Field the duration is relative to. Displays if a relative duration or operation is selected.

8. Select **Update**.
9. Select **Next**.
10. In the Key Origin screen, in the **Origin** field select **Import from PKCS12** and enter the key alias in the **Key Alias** field.
11. Select **Next**.
12. In the Key Creation screen, select **Import Key**.
13. In the Import Keystore/Certificate dialog box, provide the keystore password and import the keystore/certificate.
    a. Provide the password in the **Enter Keystore Password** field.
    b. In the **Import Keystore/Certificate** field, select **Browse** and select the certificate to be imported.
    c. Select **OK**.
14. Select the **Module Keys** tab to view the keys.

    Secure information for the keys will be stored on the Module Keys tab along with the private key and public key. The **Key lifecycle state** is set to **Active**.

Upload a CA certificate

Upload a digital CA (Certificate Authority) certificate to validate email certificates for secure communication.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **SMIME** > **CA Certificate**.
2. Select **New**.
3. Select the attachment icon to upload the CA certificate.
4. On the form, fill in the fields.

Fields	Descriptions
Name	Unique name for the CA certificate.
Type	The type of CA certificate.
Expiration notification	Option to enable expiration.
Warn in days to expire	Number of days to send a warning prior to the expiration date.
Active	Option to enable the CA certificate.
Short description	Description for the CA certificate.

5. Select **Submit**.

What to do next

Upload an email certificate.

Upload an email certificate

Upload an email certificate to validate a signature for inbound email or encrypt an outbound email or both for secure communication.

Before you begin

- To enable signing in the instance, upload the private-public key pair corresponding to the instance email account.
- To enable encryption in the instance and send encrypted email to a recipient's email address, upload the public key/certificate corresponding to the recipient’s email address.

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **SMIME** > **Email Certificate**.
2. Select **New**.
3. Select the attachment icon to upload the email certificate.
4. On the form, fill in the fields.

S/MIME email certificate form

Fields	Descriptions
Email	Email address for sending the email.
Expiration notification	Select the check box to enable expiration.
Notify on expiration	Add yourself to be notified before expiration.
Warn in days to expire	Number of days to send the warning prior to the expiration date.
Active	Check box to enable the email certificate.
Key store password	Key password to protect PKCS #12 file.

Fields	Descriptions
Short description	The description for this email certificate.

5. Select **Submit**.

What to do next
EnableEnable S/MIME.

Enable S/MIME

Configure S/MIME settings for inbound and outbound email.

Enable S/MIME for Inbound and Outbound emails

You can configure email properties from the Email Properties page through the **System Mailboxes** or **System Properties** module.

Email properties are available from either of these modules:

- **System Mailboxes** > **Email Properties**
- **System Properties** > **Email Properties**

Outbound S/MIME configuration

For encryption, you must upload the email certificate for the recipients in PEM format and the CA certificates for the issuing authority. For more information, see Upload an email certificate and Upload a CA certificate.

For signing, you must upload the email account key pair in P12 format. For more information, see Import an S/MIME key pair.

Outbound S/MIME configuration properties

Configuration	Related property
Digitally sign your outbound emails.	email.outbound.smime.signing.enabled
Send a public certificate in the outbound email.	email.outbound.smime.signing.send_public_cert
Encrypt message contents and attachments.	email.outbound.smime.encryption.enabled
Encryption algorithm. Select **AES-CBC** or **AES-GCM**. **Note:** Before selecting AES-GCM, check whether GCM is supported by the client for S/MIME.	email.outbound.smime.encryption.algo

Inbound S/MIME configuration

For decryption, you must upload the email certificate for the recipients in PEM format and the CA certificates for the issuing authority. For more information, see Upload an email certificate and Upload a CA certificate.

For signature verification, you must upload the email account key pair in P12 format. For more information, see Import an S/MIME key pair.

Inbound S/MIME configuration properties

Configuration	Related property
Verify the signature for inbound emails.	email.inbound.smime.verify_sign
Decrypt inbound emails.	email.inbound.smime.decrypt

Enable S/MIME for email notification form

To digitally sign or encrypt your emails, go to **All** > **Email** > **Notifications**, click **New** and select **Digitally sign your emails** checkbox for digitally signing emails and **Encrypt emails** checkbox for email encryptions.

Enable S/MIME for email notification form

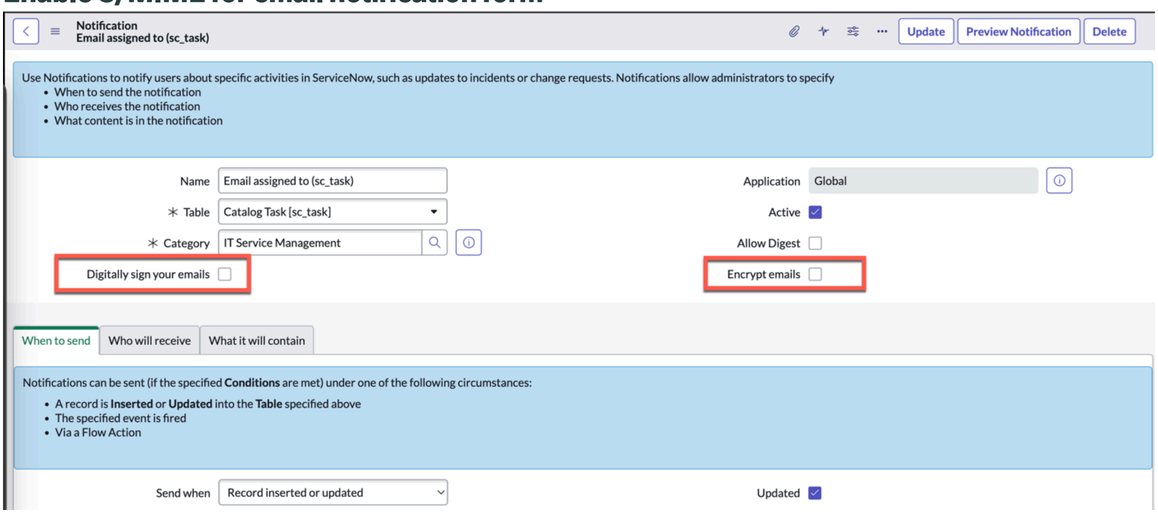

For more information, see Create an email notification.

Enable S/MIME for email client

In the compose email form, select **Digitally sign your emails** checkbox for digitally signing emails and **Encrypt emails** checkbox for email encryptions.

Enable S/MIME for email client

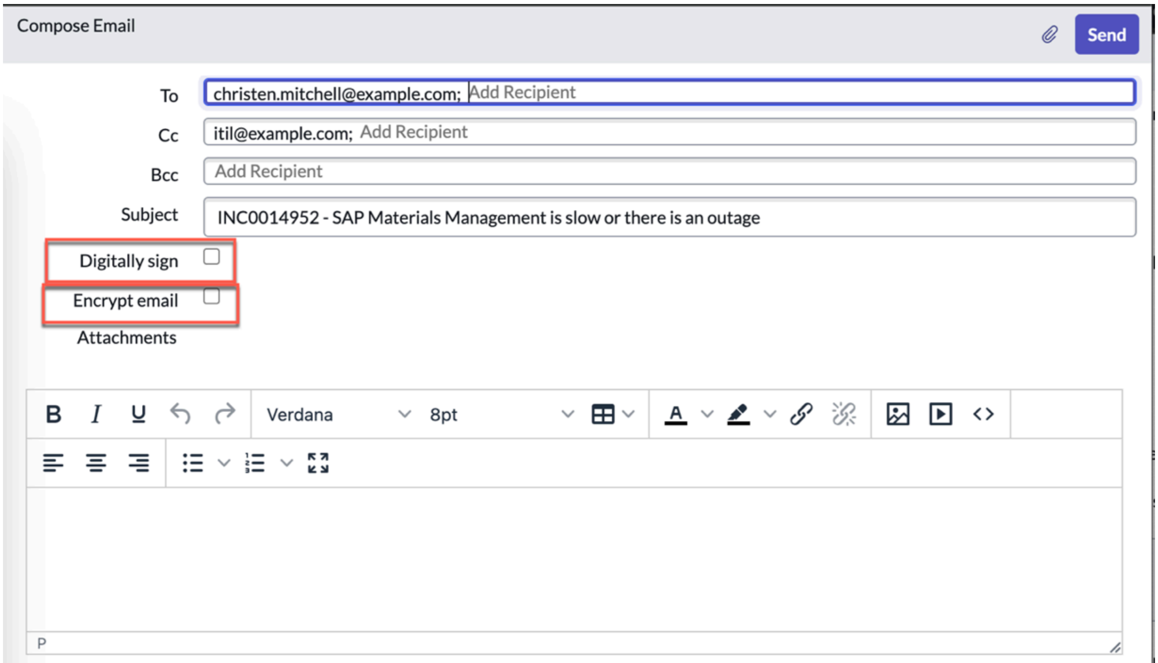

Reference

Reference section provides additional information about email administration.

Email size limits

To prevent issues with large email messages, the system enforces configured limits on the maximum allowed email body size, total attachment file size, and number of attachments per email.

An email or attachment that exceeds the system capacity may not be processed as expected. To ensure that your emails and attachments are processed as expected, you can configure properties that limit the size of emails and attachments.

Message body size limit properties

Add properties to control the maximum email body size allowed for inbound and outbound email messages.

Properties

By default, the system processes up to 524,000 characters in the email **body** and **body_text** fields.

You can add and configure system properties `glide.email.inbound.max_body_chars` and `glide.email.outbound.max_body_chars` to increase or decrease the email body size limits. The system truncates the email body when it exceeds the configured limit. For example, if you configure the body size limit of `glide.email.inbound.max_body_chars` to 1,000 characters but receive an inbound email that is 2,000 characters, then the system truncates the email body at 1,000 characters.

Message body size limit properties

Name	Description
glide.email.inbound.max_body_chars	Sets the maximum body size in characters allowed per inbound email. • Type: integer • Default value: 524288
glide.email.outbound.max_body_chars	Sets the maximum body size in characters allowed per outbound email. • Type: integer • Default value: 524288

Note:
A different property, *com.glide.attachment.max_size*, sets the maximum file size allowed for any attachment in the system and overrides any larger values of *glide.email.inbound.max_total_attachment_size_bytes* and *glide.email.outbound.max_total_attachment_size_bytes*.

Inbound email body processing

For inbound emails, the system enforces the maximum body size as set by the *glide.email.inbound.max_body_chars* property. When the body size exceeds the configured value:

- The system does not run inbound email actions that would otherwise be triggered by the email.
- The system truncates the Body text on the Email form.
- The system logs a warning and sets the Error string field on the Email form. The log message for such an email might look like this: `Email set to receive-ignored because its size exceeds the value set in` *glide.email.inbound.max_body_chars*.

Outbound email body processing

For outbound emails, the system enforces the maximum body size as set by the *glide.email.outbound.max_body_chars* property. When the body size exceeds the configured value:

- The system does not send the email.
- The system truncates the Body text on the Email form.
- The system logs a warning and sets the Error string field on the Email form. The log message for such an email might look like this: `Email set to send-ignored because its size exceeds the value set in` *glide.email.outbound.max_body_chars*. `1337 character(s) were truncated from the body field.`

Related topics

Advanced email properties

Attachment limit properties

Several properties control email attachment limits.

To see the status of all email attachments, go to the Email Attachments [sys_email_attachment] table.

Properties

All the properties are located in the System Property [sys_properties] table. Setting any of the following properties to an excessively large value may cause performance issues.

Attachment limit properties

Name	Description
glide.email.inbound.max_attachment_count	Sets the maximum number of attachments allowed per inbound email. • Type: integer • Default value: 30 • Learn more: Inbound Email Attachment Processing
glide.email.inbound.max_total_attachment_size_bytes	Sets the maximum total attachment size in bytes allowed per inbound email. • Type: integer • Default value: 18874368 • Learn more: Inbound Email Attachment Processing
glide.email.outbound.max_attachment_count	Sets the maximum number of attachments allowed per outbound email. • Type: integer • Default value: 30 • Learn more: Outbound

Attachment limit properties (continued)

Name	Description
	Email Attachment Processing
glide.email.outbound.max_total_attachment_size_bytes	Sets the maximum total attachment size in bytes allowed per outbound email. To send an email, the system must encode the contents of the email. This process may significantly increase the size of the email, including any attachments. It is best to set this property to a value well below the maximum email size. • Type: integer • Default value: 18874368 • Learn more: Outbound Email Attachment Processing

Note:
A different property, *com.glide.attachment.max_size*, sets the maximum file size allowed for any attachment in the system and overrides any larger values of *glide.email.inbound.max_total_attachment_size_bytes* and *glide.email.outbound.max_total_attachment_size_bytes*.

Inbound email attachment processing

For inbound emails, the system enforces the maximum number and size of attachments as set by the *glide.email.inbound.max_attachment_count* and *glide.email.inbound.max_total_attachment_size_bytes* properties. When an attachment for an inbound email exceeds either value, the attachment is discarded. A record is created in the Email Attachments [sys_email_attachment] table containing the discarded file name and the reason it was discarded. The discarded file remains in the Email [sys_email] record. The order in which the system processes the attachments determines which attachments are discarded. This order may not be consistent from email to email.

When an inbound email attachment is discarded, the system also fires the *inbound.email_attachments.discarded* event. You can use the event to trigger a notification that alerts the email sender. For an example of a notification that can be triggered by the event, see the inactive "Demo Email Attachment(s) Discarded" notification.

Inbound email attachments are prevented from attaching to the target record if an identical attachment exists.

Outbound email attachment processing

For outbound emails, the system enforces the maximum number and size of attachments as set by the *glide.email.outbound.max_attachment_count* and *glide.email.outbound.max_total_attachment_size_bytes* properties. Email records are created from various sources and may exceed the configured attachment limits.

Emails that are ready to be sent from the Email [sys_email] table are subject to the outbound attachment limits. Emails that exceed either limit trigger a warning in the email system log and are sent with attachments up to the maximum number or total file size.

The log message for such an email might look like this: `Maximum combined attachment size exceeded. (max:15728640 bytes). One or more attachment records ignored.`

Emails for notifications, scheduled reports, and exported tables

Notifications can be set to include all the attachments from the record that triggers the notification. If the attachments exceed either of the outbound email attachment limits, the system excludes the excess attachments from the email and logs a warning message.

Reports can be scheduled for email distribution as attachments. Large reports may exceed the outbound attachment size limit. In this case, the system sends the scheduled report email without the report attached and logs a warning message. To avoid the issue, send links to large reports instead of sending the reports as attachments.

If a user attempts to export numerous records from a list that exceeds a configured warning threshold, a dialog box offers the option to email the exported records as an attachment. If the attachment exceeds the outbound attachment size limit, the system sends the email without the exported record list attached and logs a warning message.

Check the status of bounced email addresses

Choose the email address record that you want to modify, and update the state field as needed.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Logs** > **Bounce Email Addresses**.
2. From the state of the mail address, select the state and choose the desired state to block or unblock the email address

   > **Important:**
   > Before modifying the status, address the reason for the failure of that particular email address.

3. Select the confirm icon (✓) to confirm the status.

System notifications

Set up notifications to alert users of record changes that may concern them. You can notify users via email, SMS text message, push notification, or messaging app.

Get started

Email and SMS Notifications	Push Notifications	Notifications in Messaging Applications
		
Use email notifications to send selected users email or SMS notifications in the system	Push notifications to mobile devices	Use the Now Actions app to configure the messaging apps for your instance

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Email and SMS notifications

Use email notifications to send selected users email or SMS notifications about specific activities in the system, such as updates to incidents or change requests.

Email notifications allow administrators to specify:

- When to send the notification
- Who receives the notification
- What content is in the notification
- Whether the notification can be delivered in an email digest and if so, the digest content

Additional email notification options are available. Users can subscribe to notifications, and administrators can make some notifications mandatory.

Administrators also have the option of converting existing email notifications to a rich HTML format. This format provides several advantages, including:

- Raw HTML content is converted into a WYSIWYG format.
- The content can be edited in a feature-rich HTML editor.

- Mail scripts are condensed into a single, easy-to-read line that can be reused in multiple email notifications.
- To prevent broken links, items like images and incidents, that are linked with URLs relative to an instance are converted to absolute links. For example, if an incident is linked using a relative URL, the link is converted to an absolute link.

Note:
The rich HTML format is the default for all new email notifications.

Note:
Instances cannot send or receive encrypted email messages. The system strips out the body of the encrypted email because it cannot process the encrypted content in plain text or HTML.

An email notification can also send as an SMS notification if the recipient has subscribed to the notification on an SMS channel. The system uses the subject line of the email notification and converts it to an SMS message. If the administrator doesn't want to use the email notification subject for the SMS notification, they can define an alternate SMS message in the email template form or email notification form. For more information, see Create an email notification.

For more information on creating an SMS channel, see Create notification channels.

For more information on setting up SMS notification see KB0712569.

If you want to change how the instance processes incoming email, see Inbound email actions. See System email log and mailboxes for examples of messages the system displays when notifications or inbound email actions are not processed.

Create notification categories

You can create notification categories to identify and group related notifications that are listed in the notification settings for your users.

Before you begin

Role required: admin

About this task

Before adding new categories, review the base system notification categories in the Notification Categories [sys_notification_category] table to determine if you need a new category.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notification Categories**.
2. Click **New**.
3. Enter the **Name** of the category to identify the family of notifications.
   Provide a category name that is meaningful to your users, so that they can easily find their notifications under the appropriate category in their notification settings.
4. Enter a **Short description** to identify the category.
5. Click **Submit**.

Result

The system adds the new category to the Notification Categories [sys_notification_category] table. After you create at least one, active notification that uses the new category, the category and notification are listed in the notification settings for users (**Notifications** tab in the System Settings window).

Create an email notification

Create an email notification specifying when to send it, who receives it, what it contains, and if it can be delivered in an email digest.

Before you begin

Role required: admin

Consider the following items when you create or update a notification:

- Your notification recipients must be active users and have a valid email address defined.

  ServiceNow users or members of groups must be defined as active users in the User [sys_user] table. They must also have a valid email address defined for their primary channel (device) in the Notification Device [cmn_notif_device] table. If users don't have an active profile and a valid email address, the notifications won't be received.

- Your notification recipients must have the appropriate notification preferences enabled.

  If the notification is subscribable, each ServiceNow user or group member must have the notification and channels (devices) for the notification enabled in their notification preferences. Admins can impersonate users to review and configure their notification preferences.

- To have your email notification also send as an SMS notification, recipients must subscribe to the notification on an SMS channel. For more information on creating an SMS channel, see Create notification channels.
- After you create or update the notification, use the **Preview Notification** option to examine it. For example, you can test links that you may have added and verify the notification recipients. For details, see Preview email notifications.
- For information on operators available for filters and queries, see Operators available for filters and queries.

About this task

Watch the video to learn about the actions that the instance can take in response to messages from users and shows how to create or modify email notifications to users.

https://player.vimeo.com/video/995190091?badge=0&autopause=0&player_id=0&app_id=58479

Use the following tabs in the Notification form to configure an email notification:

- **When to send** — Conditions required to send the notification.
- **Who will receive** — Recipients of the notification.
- **What it will contain** — Contents of the notification.
- **What Digest will contain** — Contents of the email digest if the notification can be delivered in a digest.

Note:
If you don't see all the fields on the form, switch to the **Advanced** view.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Select **New**.
3. Fill in the fields at the top of the Notification form, as appropriate.

Field	Description
Name	Enter a unique name for the email notification. Descriptive names help identify the purpose of the email notification. For example, Incident Opened & Unassigned.
Table	Select the database table to link the notification to, for example Incident [incident]. **Important:** Don't select the Task [task] table. This table is for extending other tables. Notifications that run on the Task table directly aren't supported. **Note:** Only tables and database views that are in the same application scope appear in the list.
Type	Select the type of notification that you're creating: EMAIL or Meeting Invitation.
Active	Select the check box to enable the email notification.
Category	Select the category to which this notification belongs. A category identifies and groups related notifications. This notification, if active and subscribable, is listed in the selected **Category** in the notification preferences for each user (**Notifications** tab in the System Settings window). **Note:** Don't leave the category as **Uncategorized**, as users may not be able to find the notification in their list of notifications. If you need a new category, see Create notification categories.
Allow Digest	Select the check box if an email digest is to be created for the notification. If selected, the **What Digest will contain** tab is displayed for creating the digest content.
Description	Type a description for this notification.

4. Fill in the fields on the **When to send** tab.

Note:
If the same trigger generates multiple notifications, the system only sends one notification. The system considers all other notifications, even if they have a different subject and body, as duplicates. The Ignore Duplicates business rule controls this functionality.

Field	Description
Send when	Select under what condition that the notification is sent: ◦ When a record is inserted or updated ◦ When a particular event is fired ◦ When Notification step in Flow Designer
Weight	[Required] Set a numerical value for the notification priority relative to duplicate notifications.

Field	Description
	Notifications that have the same target table and recipients are considered duplicates if the weights are different. If weights are same, an additional evaluation is done to check if the subject and the body (excluding watermark) are same to qualify as duplicate notification. When there are duplicate notifications, the system only sends the notification with the highest weight. All other notifications are moved from the Outbox to the Skipped mailbox. The default value 0 causes the system to always send the notification (assuming the conditions are met). For example, suppose that a service desk agent adds a comment to an incident and shortly thereafter closes it. By default, these actions trigger both the Incident commented and Incident Closed notifications. However, both notifications are from the Incident table and also notify the incident caller. The system only sends the notification with the highest weight, which in this case is the Incident Closed notification. **Note:** The SMTP Sender scheduled job determines how often to send email. By default, this job runs every minute.
Conditions	Use the condition builder to select the conditions under which this notification is sent. For example, select **Priority** > **greater than** > **3 - Moderate** to send the notification only for High and Critical priority incidents.
Inserted	Select the check box to enable email notification when a record is inserted. This field appears when you set the **Send when** field to **Record inserted or updated**.
Updated	Select the check box to enable email notification when a record is updated. This field appears when you set the **Send when** field to **Record inserted or updated**.
Event name	Select the event that triggers this notification. This field appears when you set the **Send when** field to **Event is fired**.
Advanced condition	Create a script to perform certain actions, like sending a notification based on the current email record, changing field values, or changing system properties. The advanced condition script must return **true** or set a global answer variable to **true** to send the notification. The advanced condition script uses the following global variables: ◦ *current*: Contains the current record from the table to which the notification is linked. ◦ *event*: Contains the event that triggered the notification. **Note:** The **Advanced condition** field is evaluated in addition to other conditions that you set on the notification. Both the condition and advanced condition must evaluate to true to send the notification.

5. Fill in the fields on the **Who will receive** tab.

The following example shows the default view of the tab. The advanced view contains additional fields (see table).

Tip:
Consider limiting the recipient list of any notification to 1000 users. By default, if a notification has more than 100 intended recipients, the system creates multiple notification messages with up to 100 recipients each. If you want to change the recipient limit, set the system property `glide.email.smtp.max_recipients`.

Field	Description
Users	Select the users that you want to receive the email notification. You can search for users with the reference lookup icon or manually add their email addresses. This list of users is static.
Users/ groups in fields	Select fields from the target record that reference the User [sys_user] or Group [sys_user_group] tables. The system sends the notification to users or groups in each field that you select. For example, if a notification uses the Incident [incident] table, you can select the **Opened by** field to send the notification to users or groups who opened the incident. This list of users or groups is variable and depends upon the values of the associated task record. You can also select a field that includes an email address string to send a notification to that address. **Note:** You can dot-walk to values in reference fields by selecting the plus sign in the field selector and then selecting the related field. If you address the notification to a user with an inactive record in the User [sys_user] table, the system doesn't send the notification to that user.
Groups	Select the groups that you want to receive the email notification from. You can search for groups with the reference lookup icon or by manually entering the group name. This list of groups is static. **Note:** Group members receive individual notifications only if **Include members** is selected in the group record.
Exclude delegates	Select this option to help prevent the instance from sending email notifications to delegates of the users and members of the groups you selected.
Send to event creator	Select this check box to send the notification to the person who performed the action that started the notification process if the person is also a recipient. If the event creator isn't specified in one of the recipient fields, the event creator doesn't receive a notification regardless of the setting in this field. For new notifications, this option is selected by default. If you want to know why you may not be receiving certain email notifications, see the blog post Troubleshooting email notifications- Send to the Event Creator by a ServiceNow Technical Support Engineer in the ServiceNow Community.
Event parm 1 contains recipient	Select this check box if the event parameter 1 contains one or more notification recipients (in a comma separated-list). This field is visible only when the **Send when** field is set to **Event is fired**.

Field	Description
Event parm 2 contains recipient	Select this check box if the event parameter 2 contains one or more notification recipients (in a comma-separated list). This field is visible only when the **Send when** field is set to **Event is fired**.
Subscribable	Select this check box to enable all users to subscribe to this notification. See Subscription-based notifications for more information. **Note:** If the record contains sensitive or protected data, consider restricting the recipient list only to those users and groups who normally have access to it, and don't enable the **Subscribable** option. You can also configure your notification content so that private or sensitive data isn't exposed. For example, you could insert a link back to the associated record, so that details aren't revealed in the notification.

The system doesn't exclude recipients based on access controls. Recipients can receive email about records that they can't normally access from the user interface. For example, requesters can receive email about incidents and catalog requests opened on their behalf even though they normally don't have access to these records. If a notification includes the record details, verify that all recipients need these details.

Note:
By default, the system doesn't send email notifications to itself. For example, an email notification from instanceABC@service-now.com doesn't send to instanceABC@service-now.com. The system helps prevent this behavior to avoid looping.

6. Fill in the fields on the **What it will contain** tab.
The following example shows the default view of the tab. The advanced view contains additional fields (see table).

Field	Description
Email template	If you want to reuse existing content, select an email template to add content to the email notification. You can only select an email template that meets one of the following conditions: ◦ shares the same scope and table as the notification ◦ shares the same scope but has no specified table ◦ shares the same table and is in the global scope
Subject	Enter the subject line for the email message. The subject can include variables from the **Select variables** column. If empty, the system uses the **Subject** value from the **Email template**. If you enter a value in this field, it overrides the template value. If your recipients subscribe to the email notification on an SMS channel, then the system sends the email notification subject as an SMS message. To send a different SMS message, fill in the **SMS alternate** field on the email template form or the email notification form.
Message HTML	Enter the content of the email notification message. The message can include variables from the **Select variables** column.

Field	Description
	Variables map to column names available from the notification table, its parent tables, and reference tables. Use variables to include values from a record in the table such as an incident short description or comments and work notes. The **Message HTML** field is visible only if you set the content type to **HTML and plain text** or **HTML only**. If empty, the system uses the **Message HTML** value from the **Email template**. If you enter a value in this field, it overrides the template value. To avoid adding extra `<p>` and `<div>` elements to your email notifications, see the blog post Extra line spacing with paragraph tags in email client by a ServiceNow employee in the ServiceNow Community.
SMS alternate	Enter the notification message to send to an SMS device. The SMS alternate message is limited to 140 characters. If empty, the system uses the **SMS alternate** value from the **Email template**. If you enter a value in this field, it overrides the template value. If you don't fill in the **SMS alternate** field on this form or the email template, then the system uses the email notification subject as the SMS message.
Importance	Set the importance of the email message to low or high.
Content type	Select the content type for the email notification: ◦ HTML and plain text ◦ HTML only ◦ Plain text only By default, **HTML only** is enabled.
Include attachments	Select this check box to send all attachments from the triggering record as email attachments.
Omit watermark	Use this check box to apply or remove the watermark attached to each email. If the email doesn't contain a watermark, the system reviews the conditions of the inbound actions to create or update task records. For more information, and an alternative way to hide watermarks, see Watermarks on notification emails.
Message Text	Enter the notification message to send in plain text. This field appears when you set the content type to **HTML and plain text** or **Plain text only**. If empty, the system uses the **Message Text** value from the **Email template**. If you enter a value in this field, it overrides the template value.
From	Enter the email address that you want the email notification to use in the From field. For example, helpdesk@yourcompany.com. The email must be in a valid format, otherwise a notification message appears near the field. Changing this address may require an advanced email setup such as enabling email forwarding, for example when using Sender Policy Framework (SPF) records for spam detection.

Field	Description
Reply to	Enter the email address that you want people to use when replying to the email notification. For example, helpdesk@yourcompany.com. The email must be in a valid format, otherwise a notification message appears near the field. You can add no more than one email address in this field. Changing this address requires an advanced email setup such as enabling email forwarding.
Push message only	Select this option to send this notification only as a push notification to a mobile device. The Push Notification feature must be active.
Push messages	Associate one or more push messages with this notification. The Push Notification feature must be active. **Note:** The push message and notification must be for the same table.

7. If you selected the **Allow Digest**check box, fill in the fields on the **What Digest will contain** tab to create the email digest content for the notification.
The following example shows the default view of the tab. The advanced view contains additional fields (see table).

Field	Description
Digest Template	If you want to reuse existing content, such as headers or footers, select an email template to add content to the email digest.
Digest Subject	Enter the subject line for the email digest. The subject can include variables from the **Select variables** column. If empty, the system uses the **Subject** value from the **Email template**. If you enter a value in this field, it overrides the template value.
Digest HTML	Enter the recurring content for the email digest. The digest content can include variables from the **Select variables** column. Variables map to column names available from the notification table, its parent tables, and reference tables. Use variables to include values from a record in the table such as an incident short description or comments and work notes. To avoid adding extra `<p>` and `<div>` elements to your email digest, see the blog post Extra line spacing with paragraph tags in email client by a ServiceNow employee in the ServiceNow Community.
Digest Separator (HTML)	Use the line to separate each item summarized in the digest.
Digest From	Enter the email address to be used in the From field of the email digest. For example, helpdesk@yourcompany.com. The email address must be in a valid format, otherwise a notification message appears near the field.

Field	Description
	Changing this address requires an advanced email setup such as enabling email forwarding.
Digest Reply To	Enter the email address that you want people to use when replying to the email digest. For example, helpdesk@yourcompany.com. The email must be in a valid format, otherwise a notification message appears near the field. Changing this address requires an advanced email setup such as enabling email forwarding.
Digest Text	Enter the recurring content of the email digest to send in plain text. This field appears when you set the content type to **HTML and plain text** or **Plain text only**.
Digest Separator (text)	[Optional] Use the dash character as a line to separate each item summarized in the digest. This field appears when you set the content type to **HTML and plain text** or **Plain text only**.

When you save or update the notification, the email **Digest**option is available for the notification in the notification settings of your users.

8. When you finish creating the notification, select **Submit** or, if you're done modifying the notification, select **Update**.

What to do next

Use the **Preview Notification** option to check your notification, for example, you can see:

- How the **Subject** and **Message** fields are displayed.
- Which users will or won't receive the notification, including the reasons why the users won't receive it.

Convert legacy email notifications to rich HTML

By default, new email notifications are created in the rich HTML format. But you can also convert legacy notifications to rich HTML.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. On the **Email Notifications** list screen, click the name of the email notification you want to convert.
3. Click the **What it will contain** tab.
4. Click **Switch to Rich HTML Editor**.

   The system copies any raw HTML from the **Message** field and converts it to rich HTML in the **Message HTML** field. Additionally, any mail scripts in the body are automatically saved to the

Email Script `[sys_script_email]` table and are replaced in the notification body with an embedded script tag. This makes the notification body easier to read.

Note:
The string "div" at the bottom of the screen shows the location of your cursor within the **Message HTML** field. In this case, the cursor is in a line containing an HTML `<div>` tag.

What to do next

When you convert an email notification that was created in a version prior to Eureka to rich HTML, mail scripts are automatically moved to the Email Script [sys_script_email] table and an embedded script tag with the name of the script is automatically inserted into the body of the notification.

When creating new email notifications, write mail scripts using **System Notification** > **Email** > **Notification Email Scripts**. When the scripts are completed, add a `${mail_script:script name}` embedded script tag to the email notification body. This makes it easy to use the same scripts in multiple email notifications. All you need to copy and paste from one notification to the next is the embedded script tag.

If you manually enter a mail script, any text bounded by `<mail_script> </mail_script>` in the body of a new or converted email notification or template which is saved to the record, a message asks whether the mail script should be converted.

Invalid mail script message

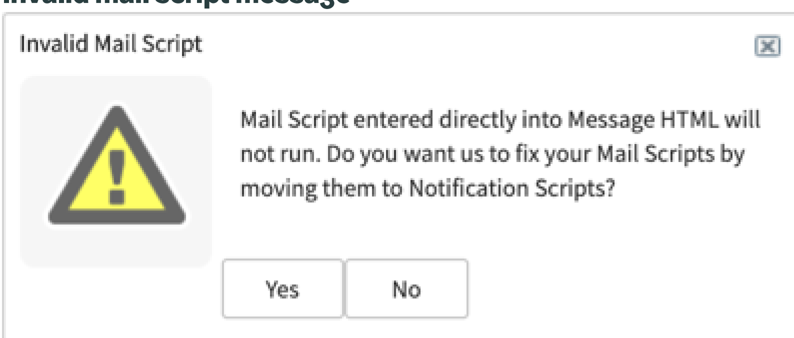

In many cases, an unconverted mail script fails to run from inside the HTML editor. If you select **Yes**, the script is added to the Email Script [sys_script_email] table and is automatically replaced in the body with an embedded script tag. You can view the mail scripts in their original form by opening the email notification and clicking the **Show Notification Scripts** related link.

Control visibility to email records generated by notifications

Define conditions that restrict read access to target email records containing sensitive information.

Before you begin

Role required: admin

About this task

Use the Email Access Restriction [email_access_restriction] table to define conditions that control read access to an email record generated by a notification containing sensitive information. These conditions are processed by the base system ACL for the Email [sys_email] table.
When you specify conditions to restrict access for a notification, the email records are visible only to users that match the conditions specified for the notification.

Procedure

1. Navigate to **All** > **System Notification** > **Email Access Restriction**.
2. Click **New**.
3. Complete the form.

Field	Description
Notification	Choose the notification that generates the email record that should be restricted.
Application	Displays the application scope.
Table	Displays the target table for the notification that you selected.
Condition	Use the condition builder to determine who can access the email record. The conditions must evaluate to true to enable read access to the email record.
Description	Enter a short description of the read access restriction.

4. Click **Submit**.

Advanced conditions for email notifications

Use an advanced condition to send a notification based on the current email record, changing field values, or system properties.

To send a notification using an advanced condition, you can:

- Call a function that returns a value, or
- Set the global variable `answer` using a script

For example, in the following code, you call a function to prevent the system from sending an email notification if the sender of a self-service request is a member of the XYZ group:

```
(function() {
  var groupMember = gs.getUser();
  return !groupMember.isMemberOf('XYZ');
})();
```

Alternatively, you can script the same advanced condition by using the `answer` variable:

```

var groupMember = gs.getUser();
  if(groupMember.isMemberOf('XYZ')){
    answer = false;
  } else {
    answer = true;
  };
```

Note that the script must set the `answer` variable to **true** to send the notification. If you script no conditionals, the value of `answer` is equal to the last value that you set for the variable.

You can add a script-based condition in the **Advanced condition** field by configuring the Email Notification form and adding the field. You can access the field in the **Advanced** view without configuring the form.

The advanced condition script uses the following business rule global variables:

- current: contains the current record from the table to which the notification is linked.
- event: contains the event that triggered the notification.

Note:
The **Advanced condition** field is evaluated in addition to other conditions you set on the notification. Both the **Condition** and **Advanced condition** must evaluate to true in order to send the notification.

Related topics

Create an email notification

Edit HTML content in an email notification

For added control over the content of an email notification, you can edit the underlying HTML.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Open an email notification record.
3. On the form, open the **What it will contain** tab.
4. In the **Message HTML** field, click the source code icon (< >) to open the HTML source code editor.

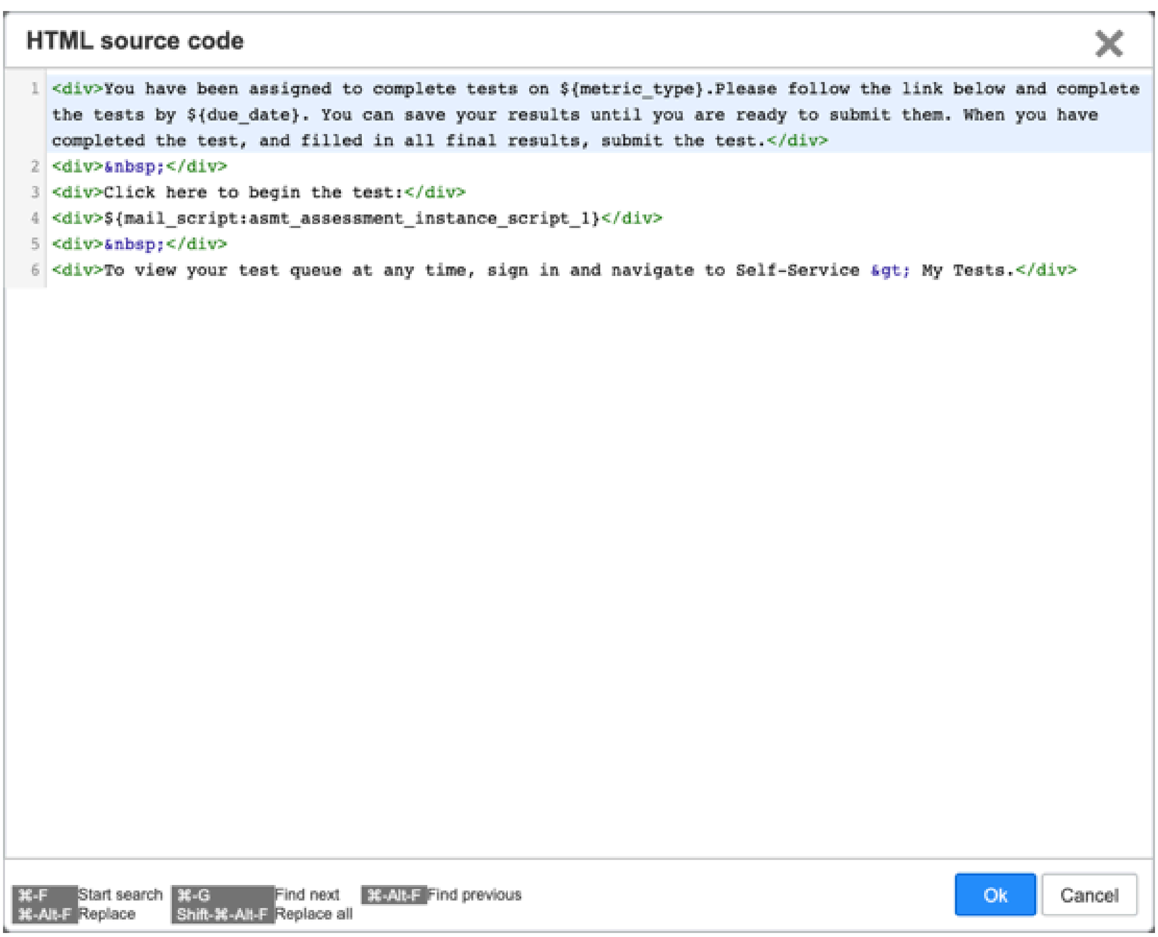

5. Make the needed changes to the HTML.
6. On the window, click **Ok**.
7. On the form, click **Update**.

Document attachments on an email notification

You can attach documents and reports to email notifications by scripting or linking to the sys ID of the record.

You can include all attachments from the source record with the notification. For example, if an incident update generates a notification, you can include all attachments from the incident record with the notification. To include all attachments from the source record, go to **Advance view** and select the check box for the **Include attachments** under **What will it contain** tab. The email messages, including attachments, cannot exceed the maximum email size. This size includes MIME encoding. For details on MIME encoding, see Email service size restrictions.

Attaching documents with scripting

Using scripting, you can attach documents by linking to them, or you can attach various types of reports by specifying their IDs in the system.

Linking to an attachment

You can add an attachment to a notification by linking to the attachment record in the message of the notification. Upon clicking the link, email recipients log in to the instance to view the attachment record.

Linking to attachment records in this fashion requires using email notification scripting. For example:

```
template.print ( 'Attachment: <a
 href="/sys_attachment.do?sys_id=' + now_GR. sys_id + '">' +
 now_GR. file_name + '</a>\n ' ) ;
```

Attaching reports using the Sys ID

You can also attach various types of reports, including gauges, dashboards, and charts, to a notification. The scripts to attach these reports take the following syntax:

```
${report:X:Y}
```

where:

X is the type of report you want to attach (reportID, gaugeID, dashboardID, or chartID).

Y is the sys ID of the report, gauge, dashboard, or chart to be attached.

For example:

- ${report:reportID:<abc123>}
- ${report:gaugeID:<abc123>}
- ${report:dashboardID:<abc123>}
- ${report:chartID:<abc123>}

Note:
Multilevel pivot reports can’t attach to email notifications.

Line breaks in email notifications and rich HTML

Rich HTML provides additional control over line breaks in your email notifications and templates.

To provide control over line breaks, a **Newlines to HTML** check box is available in the **Email Script** form.

Selecting the **Newlines to HTML** check box indicates that the method for handling line breaks in earlier versions carries forward for email notifications and templates. When an email notification or template is converted to rich HTML, the **Newlines to HTML** check box is automatically selected.

For new mail scripts, add correct HTML line breaks to template.print() statements.

If an email notification or template is not converted to rich HTML, newlines are automatically wrapped with <div> tags, the same as previous versions. The old mail scripts still work. However, the administrator does not enjoy the benefits of working in the rich HTML format, and does not have as much control over exact HTML formatting.

HTML line breaks in new scripts

When writing new scripts, insert explicit HTML line breaks and clear the **Newlines to HTML** check box so that no HTML tags are injected when email notifications are generated. For existing notifications and templates, replace template.print("\n") JavaScript function calls with template.print("
"). This replacement gives you better control over the HTML formatting of your email notifications.

Preview email notifications

You can preview what notifications look like before you actually enable the instance to send them.

Before you begin
Role required: admin

About this task
You can preview both types of notifications as specified by the **Send when** field on the Notification form:

- **Record inserted or updated**: A change to record in the instance triggers the notification.
- **Event is fired**: An event, such as the expiration of a certificate or an inbound email action, triggers the notification.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Open the notification or create one.
   You must save the record before you can view the preview accurately.
3. Click **Preview Notification** on the form header.
4. On the Notification Preview window, verify that the notification works as expected.

Field	Description
Preview records for this breakdown source	The type of event that triggers the notification. This list appears if you preview an event-triggered notification. Select one of the following: ◦ **Generated Event**: Preview the notification with a generic event that the previewer creates. This does not actually generate an event record. ◦ **Existing Event**: Preview with an existing event record in the instance. If you select this option, select the event in the **Event Record** field.
Event Record	An existing event to preview an event-created notification. This option appears if you select **Existing Event** as the event type (for event-triggered notifications only).
Event Creator	The user triggering the notification for the preview. The event creator defaults to the user who clicked **Preview Notification**. You can change the creator as needed. You can change the preview record as needed to see the changes in the notification content.
Preview Record	The record triggering the notification for preview. The preview record defaults to one of the records in the table specified in the **Table** field on the Email Notification form. You can change the preview record as needed to see the changes in the notification content.
Users	The users who will receive the notification, as specified in the **Who will receive** section of the Email Notification form: ◦ All users that you specify on the form appear, but only the users that will actually receive the notification with the current preview settings appear in black text. ◦ Users that are specified but for whatever reason will not receive the notification appear in red, strikethrough text. Place the cursor over any of these names to see the reason the user will not receive the notification. For example, one reason could be that the user's notification settings are turned off.
Subject and Body	The content of the notification as defined by the template. The **Subject** and **Body** sections on the preview display the content in the corresponding **Subject** and **Message** fields on the template. If the template includes a link to the record that triggered the notification, the **Preview Record** link is used. Click the link to go to that record.

5. After you have reviewed the notification, exit the preview window.

6. Make the necessary changes to the notification or template, if necessary.

Time zone for email notifications

A system property controls the time zone that the instance uses for the date and time stamp of a message.

The date and time stamp of a notification uses the system time zone, not the time zone of any recipient. The email property `glide.email.append.timezone` controls whether to append the time zone. If true, the system time zone of the instance is appended to any Date/Time fields in outbound email messages (for example, 2018-07-02 04:01:14 PST).

Specify alternative outbound email addresses for notifications

By default, the system sends all outbound email notifications from the default email address of the instance, but you can specify an alternative address.

Before you begin

Role required: admin

About this task

For organizations that need to send email messages from specific email addresses, such as from multiple service desks, or they want to send notifications in different languages, the platform supports configuring multiple outbound addresses.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Select an existing notification record for the desired event, such as **Incident Closed**.
3. Create a copy of this notification for each outbound email address.
4. Open one of the notification copies, and click the **Advanced view** related link.
5. In the **What it will contain** section, add an email address to the From field that is different from the default instance address.
   For more information on how to construct the From address, refer to section 3.6.2 of RFC 2822.
6. Add a different email address than the**From** address to the **Reply to** field if you want replies to this notification to go to a different address.

   The system checks the **From** field for an address. If this field is empty, then the system uses the default address for the instance. If the **Reply to** field is empty, then all replies are sent to the address from which the notification was sent. If the **Reply to** field contains an email address, then the system sends all replies to the notification to this address.
7. Create mutually-exclusive conditions for notifications of the same type, so only the desired notification is sent when the event is fired.

   For example, if the **Company** is a certain value, then the notification comes from a unique email address entered in the **From** field.
8. Click **Update**.

Specify an outbound email address for a particular language

You can specify a different email address for each language your instance supports.

Before you begin

Role required: admin

Procedure

1. Create or copy a notification record for the desired event.
2. In the **What will it contain** section, enter a new email address in the **From** field.
3. Create the **Subject** and **Message** content in the desired language.
4. In the **When to send** section, create a condition as follows:
   a. In the list of **Condition** fields, select **Show Related Fields** from the bottom of the choice list.
   b. From the choice list of **Related Fields**, select the field that identifies the recipient.

      For example, select **Caller** > **User** fields to send the notification to the user who called in an incident, or **Assigned to** > **User** fields to send the notification to the user to whom an incident is assigned.
   c. From the choice list of user fields, select **Language**.
   d. Select the **is** operator.
   e. Complete the condition by selecting the language of the desired user.
5. Click **Update**.
   All notifications for that event originate from the specified email address and go out in the language of the recipient.

Notification variables

Use notification variables to display dynamic information in the body of a notification such as a field value, a link to a record, or a link to system preferences.

Syntax

Specify a notification variable using this syntax:

```
${variable-name+variable-parameters}
```

The `variable-name` portion is always required. Not all notification variables support the `variable-parameters` portion. When available, most variable parameters are optional. See the list of available notification variables for variable names and available parameters.

Available variables

The system provides these notification variables.

Available notification variables

Variable	Description	Available parameters
${`field-name`}	Display the value of the specified field.	None
${`image-field-name`}	Display an image associated with a record. This variable is typically used with HTML to specify the source of an image element.	None

Available notification variables (continued)

Variable	Description	Available parameters
• ${URI} • ${URI_REF}	Display a link to the current record. • *URI*: The link text is the word `LINK`. • *URI_REF*: The link text is the display value of the record. **Note:** These variables don't apply to records in Workspace. To link to a record in Workspace, create a mail script that prints a URL to a notification. For more information, see Linking to a record in Workspace.	Any valid *`sysparm`* URL parameter. For example: • sysparm_scriptlet • sysparm_view
• ${`reference-field`.URI} • ${`reference-field`.URI_REF}	Display a link to the record listed in a reference field. • *URI*: The link text is the word `LINK` • *URI_REF*: The link text is the display value of the record.	Any valid *`sysparm`* URL parameter. For example: • sysparm_scriptlet • sysparm_view
${CMS_URI}	Display a link to the specified record within a CMS page.	*<CMS-`site`>/<CMS-`page`>*: The required relative path to the CMS page.
${notification:body}	Display the body contents of an email template or email notification. Use this notification variable to specify where to display body content in an email layout.	None
${mail_script:`script-name`}	Run the specified mail script.	None
${NOTIF_UNSUB}	Display a link unsubscribe from this notification.	*`link_text`*: specify the text to display as a link within quotation marks.
${NOTIF_PREFS}	Display a link to set notification preferences.	*`link_text`*: specify the text to display as a link within quotation marks.
${comments:n}	Display the most recent comments that were made on the target record. The number of comments to display is n. For example, `${comments:3}` displays the last three comments that were made to the record. To display all comments, use the variable `${comments}`.	Any number greater than 0.

Available notification variables (continued)

Variable	Description	Available parameters
${comments_and_work_notes:n}	Display the most recent comments and work notes that were made on the target record. The number of comments and work notes to display is n. For example, `${comments_and_work_notes:3}` displays the last three comments and work notes that were made to the record. To display all comments and work notes, use the variable `${comments_and_work_notes}`.	Any number greater than 0.

Examples

Refer to the following examples to see how each notification variable is rendered in the output:

Variable	Example
`${field-name}`	Source: `Incident ${number} - comments added` Output: `Incident INC1000001 - comments added`
`${image-field-name}`	Source: `<img src='${picture}?t=medium'/>`  Output:
• ${URI} • ${URI_REF}	${URI} Source: `Click here to view incident: ${URI}` ${URI} Output: Click here to view incident: LINK

Variable	Example
	${URI_REF} Source: `Click here to view incident: ${URI_REF}` ${URI_REF} Output: 
• `${reference-field.URI}` • `${reference-field.URI_REF}`	Source: `Click here to view Incident: ${URI_REF}` `Click here to view Related Problem: ${problem_id.URI_REF}` Output: 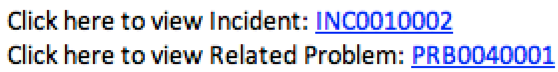
${CMS_URI}	Source: `${CMS_URI+ess/incident_detail}` Output: a link to a target CMS page such as `https://<instance name>.service-now.com/ess/incident_detail.do?sysparm_document_key=incident,46e18c0fa9fe19810066a0083f76bd56`
${notification:body} ${mail_script:`script-name`} ${NOTIF_UNSUB} ${NOTIF_PREFS}	Source: 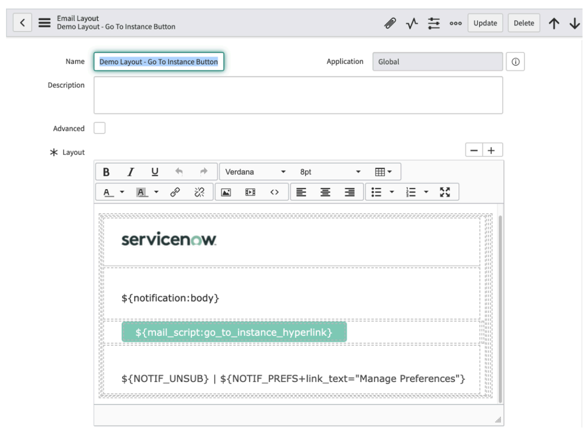 Output: 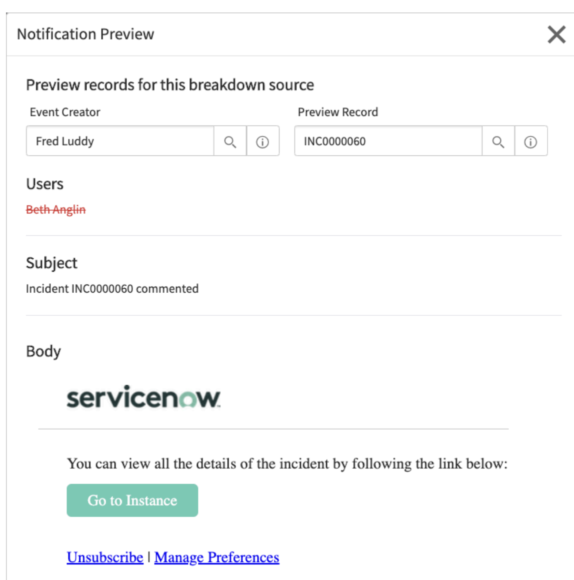

Links to records in email notifications

Adding the *${URI}* parameter to an outbound email body or template creates a link to a specific record.

When a user clicks the word LINK, the instance prompts the user to log in if not already logged in, and then redirects the user to the record specified in the URI.

Link displayed by ${URI} parameter

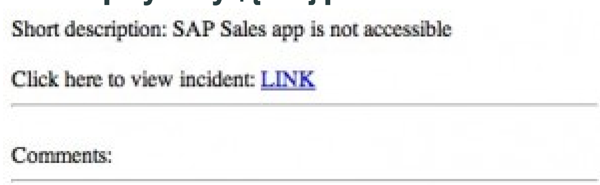

The ${URI} parameter has an extension called the ${URI+} format to specify additional arguments in the email link, such as sysparm terms, in addition to the automatically created URI. For example (whitespace added for improved readability):

```
${URI+&sysparm_scriptlet=current.assigned_to=gs.getUserID()
&sysparm_scriptlet_condition=current.assigned_to.nil()
&sysparm_view=incident_active}
```

This example executes the JavaScript:

```
current.assigned_to=gs.getUserID()
```

when the condition of

```
current.assigned_to.nil()
```

is satisfied. Additionally, the script sets the view to `incident_active`.

Linking to a record in Workspace

The `${URI}` and `${URI_REF}` variables don't apply to records in Workspace. To link to a record in Workspace, create a mail script and reference it in your notification. For more information on using mail scripts, see JavaScript in emails.

The mail script that you create should print a URL to the notification. The URL must have the following format:

```
https://<instance_name>/now/workspace/<workspace_name>/record/<t
able_name>/<sys_id>
```

The following example script shows the logic that a mail script must include to create a link to a record in Workspace.

```
// Dynamically construct an Agent Workspace URL and insert a
 link in the notification
  var agentURL = '<a href="' +
 gs.getProperty('glide.servlet.uri') +
 '/now/workspace/agent/record/'+ current.getTableName() + '/' +
 current.sys_id + '">' + current.number + '</a>';
  template.print(agentURL + "<br />");
```

Enable links to records

Adding the special *${URI}* parameter to an outbound email body or template creates a link to a specific record.

When a user clicks the word **LINK**, the instance prompts the user to log in if not already logged in, and then redirects the user to the record specified in the URI.

URI email notification


Short description: SAP Sales app is not accessible

Click here to view incident: LINK

Comments:

The *${URI}* parameter has an extension called the *${URI+}* format to specify additional arguments in the email link, such as sysparm terms, in addition to the automatically created URI. For example (whitespace added for improved readability):

```
${URI+&sysparm_scriptlet=current.assigned_to=gs.getUserID()
&sysparm_scriptlet_condition=current.assigned_to.nil()
&sysparm_view=incident_active}
```

This example executes the JavaScript:

```
current.assigned_to=gs.getUserID()
```

when the condition of

```
current.assigned_to.nil()
```

is satisfied. Additionally, the script sets the view to *incident_active*.

Change the link text

To show the display value of the record as the link text instead of the word LINK, use the ${URI_REF} parameter instead of the ${URI} parameter.

URI_REF email notification

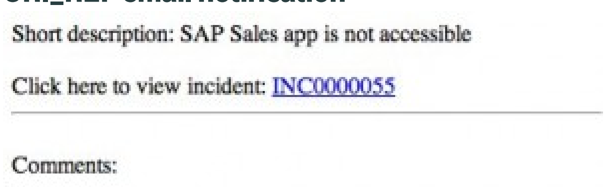
Short description: SAP Sales app is not accessible

Click here to view incident: INC0000055

Comments:

For example, if the URL displays an incident record, the link text is the incident number, which is the display value for incidents. If the URL displays a user record, then the link text is the user name.

Link to related records

A notification can link to a related record by specifying a reference field in front of the *${URI}* or *${URI_REF}* parameters.

Format the related record link as follows:

- `${<reference field that contains the related record you want to display>.URI}`
- `${<reference field that contains the related record you want to display>.URI_REF}`

For example:

Related records

Related record to provide link to	Notification record table	Reference field	Samples
Related task record to be approved from an approval notification	Approval [sysapproval_approver]	Approval for [sysapproval]	• ${sysapproval.URI} • ${sysapproval.URI_REF}
Related problem record in an incident notification	Incident	Problem [problem_id]	• ${problem_id.URI} • ${problem_id.URI_REF}

For example, the following notification template produces the email links in the picture below:

```
Click here to view Incident: ${URI_REF}
Click here to view Related Problem: ${problem_id.URI_REF}
```

Related record link

Click here to view Incident: INC0010002
Click here to view Related Problem: PRB0040001

Content page links in email notifications

Links to CMS pages can be put in notifications to make it easy for the reader to access the pages.

The link takes the following format: *${CMS_URI+<site>/<page>}*.

For example, to link the email recipient to a page called Incident in the content site ESS, with the current incident as the target document, use the following format: *${CMS_URI+ess/incident_detail}*

The resulting email URL has this format: https://<instance name>.service-now.com/ess/incident_detail.do?sysparm_document_key=incident,46e18c0fa9fe19810066a0083f76bd56

Email unsubscribe

Administrators can add unsubscribe links to notifications so that users can stop receiving particular email messages.

Note:
The base system notifications include unsubscribe and notification preferences links.

The system offers two types of macros to create unsubscribe links.

- An unsubscribe link that creates an email message to the instance.
- An unsubscribe link that opens the notification preferences for the user on the instance.

Available unsubscribe macros

Unsubscribe type	Macro used	Description	Available parameters
Unsubscribe by email	${NOTIF_UNSUB}	The system generates an HTML *mailto* hyperlink. When users select the link, their browser or	*link_text*: Specify the text to display as

Available unsubscribe macros (continued)

Unsubscribe type	Macro used	Description	Available parameters
(Unauthenticated)		email client creates a pre-formatted unsubscribe email message to the instance.	a link within quotation marks.
Unsubscribe by notification preferences (Authenticated)	${NOTIF_PREFS}	The system generates an instance link directly to the notification preferences for this notification type.	`link_text`: Specify the text to display as a link within quotation marks.

Administrators can add unsubscribe macros to any notification record type such as:

- Email layouts
- Email templates
- Email notifications

Unsubscribe by email

Unsubscribe using list unsubscribe header: Unsubscribe from email notification using list unsubscribe header. You can directly unsubscribe form notification emails by selecting **Unsubscribe** next to the email address. An unsubscribe request goes through and a pre-formatted email message is automatically sent.

Note:
The unsubscribe option or the placement of the button can vary based on the email client being used.

Important:
Unsubscribe with one click is supported from Washington Patch 1 onwards.

Unsubscribe with pre-formatted email message: Unsubscribe by email requires the user email client or browser to create a pre-formatted email message containing these elements.

- The **To** field has the email address of the instance.
- The **Subject** starts with the string `Unsubscribe from`.
- The **Body** has a JSON string with a name-value pair of *`Unsubscribe`* and an array value that contains two more name-value pairs.
  - The *`notification_id`* parameter specifies the Sys ID of the notification the user wants to unsubscribe from.
  - The *`unsub_token`* parameter specifies an instance ID the system uses to verify that the email came from a ServiceNow instance.

Note:
Some email clients and web browsers, such as GMail on Chrome, require extra client configuration to support mailto hyperlinks. Administrators can provide an alternative unsubscribe method for users whose email client or browser does not support mailto links. See RFC6068 for information about the mailto URI scheme.

The **Unsubscribe from Notification** inbound action processes the email and unsubscribes the sender from the listed notification. Unsubscribing by email message does not require users to authenticate with the instance first.

Note:
For users with multiple email address subscribed, when you unsubscribed from the primary email address it automatically unsubscribes you from all other email addresses. Unsubscribing through secondary email addresses is currently not supported and the request won't be fulfilled.

Unsubscribe by notification preferences

Unsubscribe by notification preferences requires the user's browser to navigate to the notification preferences page on the instance. After logging in, the system displays the notification preferences for this particular notification.

Users can set preferences for this notification such as disabling notifications for a particular device. Users must save their notification preferences for changes to take effect.

Example: Unsubscribe links

This email layout adds several unsubscribe links to the bottom of each email notification.

```
${NOTIF_UNSUB} from this notification by email or
 ${NOTIF_UNSUB+link_text="click
      here"}.
Manage your ${NOTIF_PREFS} or ${NOTIF_PREFS+link_text="click
 here"}.
```

When rendered in an email notification, the unsubscribe links only display the link text.

Sample email with unsubscribe links

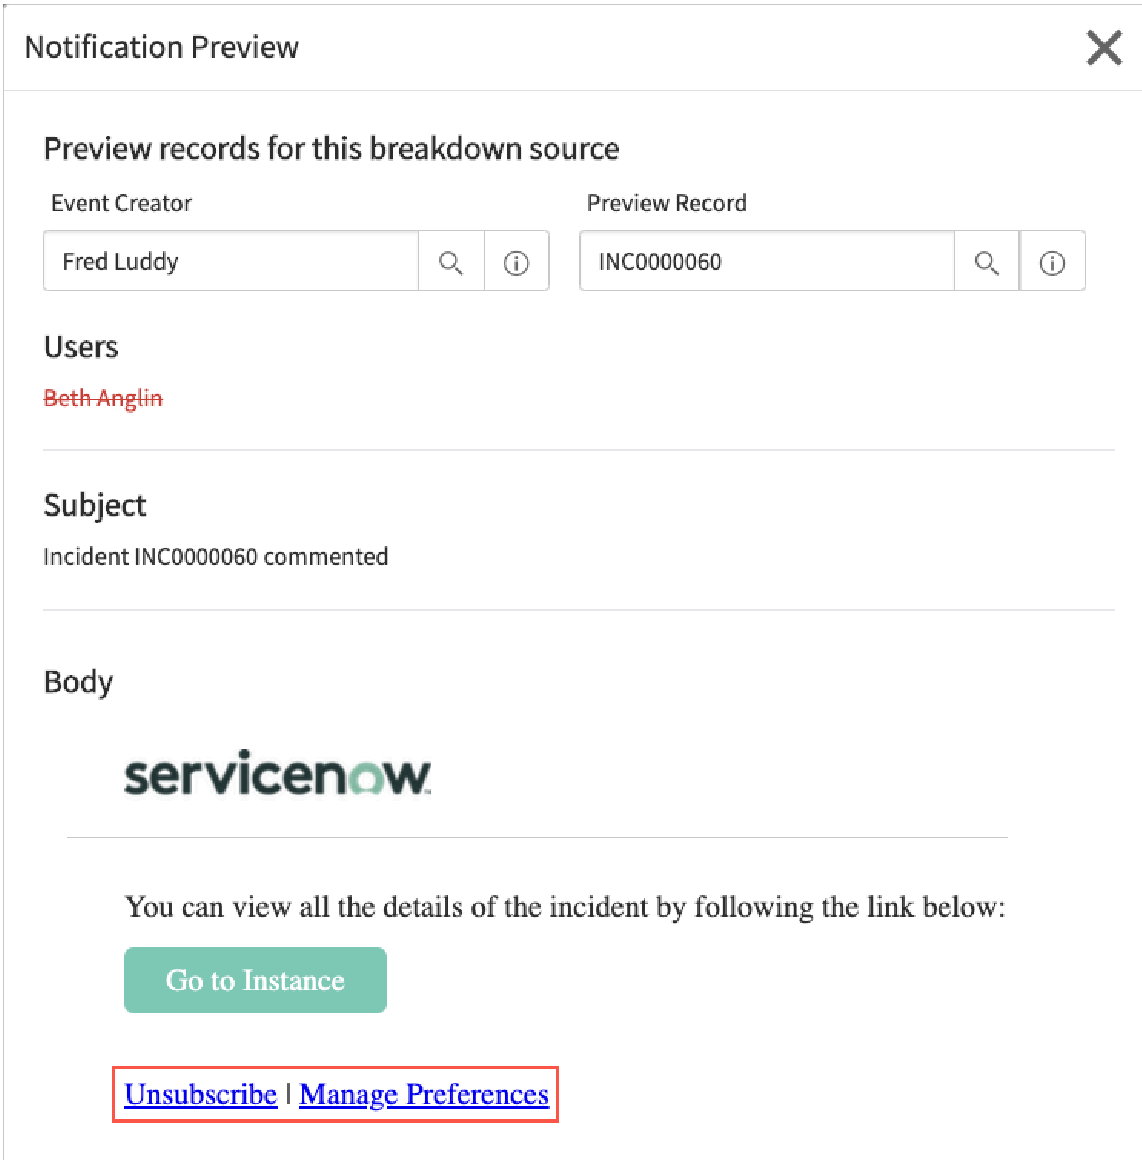

If a user clicks the **Unsubscribe** link, the email client creates a message such as this:

Sample unsubscribe by email message

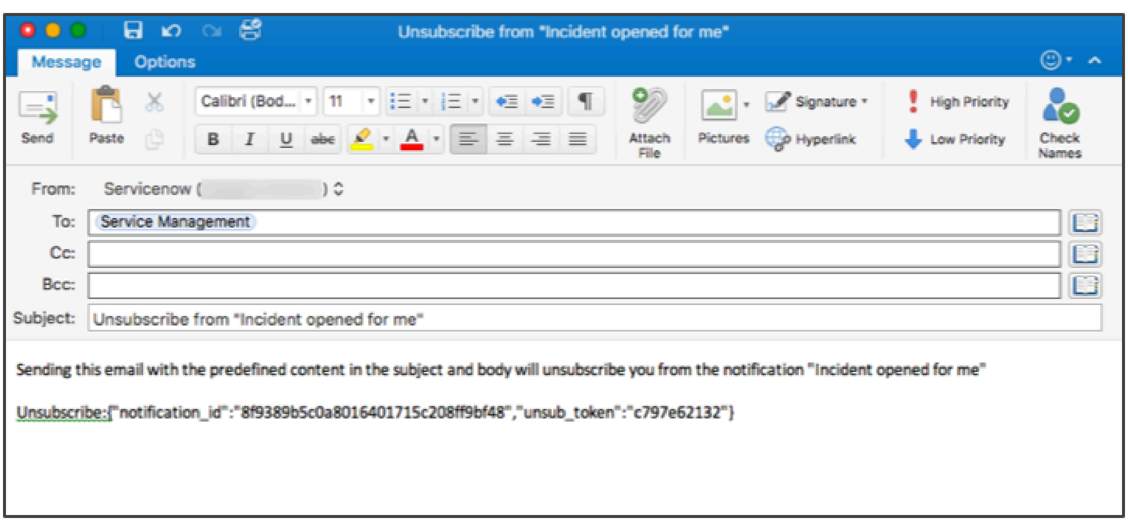

Scripting for email notifications

Email scripts allow for business rule-like scripting within an outbound email message.

With mail scripts, you can dynamically change the email output of your system based on different criteria. Mail scripts allow you to perform simple tasks, such as displaying incident data, and complex ones, such as making advanced database queries.

You can add a `${mail_script:script name}` embedded script tag to the body of the email notification or template, replacing *`script name`* with the name of the script you created. This makes it easy to use the same scripts in multiple email notifications or templates.

If you manually enter a mail script bounded by `<mail_script>` and `</mail_script>` in the body of a new or converted email notification or template, and then attempt to save the record, a message asks whether the mail script should be converted. In many cases, an unconverted mail script fails to run from inside the HTML editor. If you select **Yes**, the script is added to the Email Script [sys_script_email] table and is automatically replaced in the body with an embedded script tag (`${mail_script:script_name}`).

JavaScript in emails

Create mail scripts in **System Notifications** > **Email** > **Notification Email Script**, and refer to them by using `${mail_script:script name}` in the script field.

To print text into the body of the message, use the *`template.print("a string")`* function.

JavaScript in templates

Script

```
template.print("<p></p>Requested items:<br />");

var gr = new GlideRecord("sc_req_item");
gr.addQuery("request", current.sys_id);
gr.query();
while(gr.next()) {
    template.print(gr.number + ":  " + gr.cat_item.getDisplayValue() + ", Stage: " + gr.stage.getDisplayValue() + "<br />");
}

```

The *`event.parm1`* and *`event.parm2`* parameters that come from the originating event can also be used.

For additional information on creating mail scripts, see Notification Email Scripts on the Developer Site.

Mail script variables

Certain variables are available when processing mail_script scripts.

For examples of mail scripts, see Example scripting for email notifications.

Mail script variables

Variable	Object Description
template	Handles printing from the mail script to the email message. `template.print("message"); //outputs message to the email body.`

Mail script variables (continued)

Variable	Object Description
	`template.space("number of spaces"); //outputs spaces to the email body.`
email_action	GlideRecord object for the email notification (sysevent_email_action).
event	GlideRecord object for the event that fired the notification (sysevent).
email	EmailOutbound object Available methods: • `addAddress(String type, String address, String displayname)`: type can be cc or bcc. • `setFrom(String address)`: override the sender address. • `setReplyTo(String address)`: override the reply to address. • `setSubject(String subject)`: override the subject of the message. • `setBody(String message)`: override the body of the message. The email address that is passed by `setFrom` and `setReplyTo` needs to be in a valid form such as `helpdesk@sn.com` or `Display Name <helpdesk@sn.com>`. If the email address includes a 'Display Name', then that value overrides the instance's display name. For more information, see GlideEmailOutbound - Scoped.

Example scripting for email notifications

Examples of scripting for email notifications.

Scripting examples for email notifications

A simple text string is the most basic example of the way a mail script works. This script prints out "Incident number - INC00001".

```
template.print("Incident number - "+ current.number);
```

More advanced scripts, like this one, can be found by browsing through the base system email templates.

```
template.print("Summary of Requested items:<br />");
var now_GR = new GlideRecord("sc_req_item");
now_GR.addQuery("request", current.sysapproval);
now_GR.query();
while(now_GR.next()) {
  template.print(now_GR.number + ":  " + now_GR.quantity + " X "
 + now_GR.cat_item.getDisplayValue()
                   + " at " +
 now_GR.cat_item.price.getDisplayValue() + " each <br />");
}
```

To dynamically change field values within an email, use the following functions within <mail_script> syntax:

```
...
email.setFrom(current.caller_id.email);
email.setReplyTo("joe.employee@yourcompany.com");
email.setSubject("This is the new subject line");
email.setBody("This is the new body");
...
```

Using the *instance_name* property ensures that the notification still works when migrated between instances.

```
dothis();

function dothis(){
    var now_GR =new GlideRecord('sys_attachment');
    now_GR.addQuery('table_sys_id',current.sys_id);
    now_GR.query();while(now_GR.next()){
        template.print('Attachment: <a href="https://'+
 gs.getProperty('instance_name')+'
            .service-now.com/sys_attachment.do?sys_id='+
 now_GR.sys_id+'">'+ now_GR.file_name+'</a>');}}
```

You can specify copied and blind copied recipients by using the email object within a mail script.

```
//email.addAddress(type, address, displayname);
    email.addAddress("cc","john.copy@example.com","John
 Roberts");
    email.addAddress("bcc","john.secret@example.com","John
 Roberts");
```

The following is an example script to add users from *watch_list* as copied recipients.

```
if(!current.watch_list.nil()){
   //get watch list addresses and add to cc
   var watcherIds = current.watch_list.split(",");

   //get user records
   var user = new GlideRecord("sys_user");
   user.addQuery("sys_id", watcherIds);
   user.addQuery("notification",2);
   //email
   user.addQuery("email","!=","");
   user.query();

   while(user.next()){
      //add to cc list
      email.addAddress("cc", user.email,
 user.getDisplayValue());}}
```

Related topics

TemplatePrinter API

GlideEmailOutbound API

Useful attachment scripts

This is a searchable version of the Useful Attachment Scripts.

Warning:
The customization described here was developed for use in specific instances, and is not supported by Now Support. This method is provided as-is and should be tested thoroughly before implementation. Post all questions and comments regarding this customization to our community forum.

Copy attachments from record to record

Use the following script to copy an attachment from one record to another record:

```
GlideSysAttachment.copy('sourcetable','sys_id','destinationtable
','sys_id');
```

Note:
GlideSysAttachment.copy copies all attachments; it cannot select specific attachments.

Delete duplicate attachments

Use the following script in a business rule, scheduled job, or background script to delete duplicate attachments located in the Attachments [sys_attachment] table:

```
function fixDuplicateImages(){
  var now_GR = new GlideRecord('sys_attachment');
  now_GR.addQuery('table_name','LIKE','ZZ_YY%');
  now_GR.orderBy('table_sys_id');
  now_GR.orderByDesc('sys_created_on');
  now_GR.query();
  var lastID ='not_a_match';
  var lastFile ='not_a_match';
  while(now_GR.next()){
    var isDuplicate = (lastID == now_GR.table_sys_id)&&(lastFile
 == now_GR.file_name);
    lastID = now_GR.table_sys_id;
    lastFile = now_GR.file_name;
    gs.print(now_GR.table_sys_id + ' ' + now_GR.table_name +
 ' ' + now_GR.file_name + ' ' + now_GR.sys_created_on + ' ' +
 isDuplicate);
    if(isDuplicate)
      now_GR.deleteRecord();
  }
}
```

Display whether tasks have attachments in list view

Use the following script in a business rule to display whether tasks have attachments when viewed in the record list view. Note that the script needs a custom field on the Has Attachments [u_has_attachments] table.

```
checkAttachment();

function checkAttachment(){
  // if inserting then the task has an attachment
  if(current.operation()=='insert') {
    hasAttachment('true'); // if deleting attachment check for
 other attachments
    if(current.operation()=='delete') {
      var timeNow3 =new GlideDateTime();
```

```
      gs.log('has_attachment br: gliderecord query start date
 time is: '+ timeNow3.getNumericValue(),'jwtest');
      var attachCount = new GlideAggregate('sys_attachment');
      attachCount.addQuery('table_sys_id',current.sys_id);
      attachCount.addAggregate('COUNT');
      attachCount.query();
      var numAttachments ='0';

      // if no other attachments task does not have attachment
      if(attachCount.next()){
        numAttachments = attachCount.getAggregate("COUNT");
        if(numAttachments >0){
          hasAttachment ='true';
        } else {
          hasAttachment('false');
        }
        var timeNow4=new GlideDateTime();
        gs.log('has_attachment br: gliderecord query start date
 time is: '+ timeNow4.getNumericValue(),'jwtest');
      }

function hasAttachment(answer){
  var task = new GlideRecord('task');
  task.addQuery('sys_id',current.table_sys_id);
  task.query();

  if(task.next()){
    task.u_has_attachment= answer;
    task.autoSysFields(false);  //Don't set the lastUpdatedTime
 or the Simultaneous Update Alert will likely get triggered
    task.setWorkflow(false);  //Don't allow other business rules
 to run, otherwise multiple notifications will likely be sent
    task.update();
  }
}
```

Note:
Schedule the business rule to run after insert/delete.

Link to attachments in an email notification

Use the following script in an email notification or template to include links to attachments:

```

printattachments();

function printattachments(){
  var now_GR =new GlideRecord('sys_attachment');
  now_GR.addQuery('table_sys_id',current.sys_id);
  now_GR.query();
  while(now_GR.next()){
    template.print('Attachment: <a
 href="http://'+gs.getProperty("instance_name")+'.service-now.c
om/sys_attachment.do?sys_id='+ now_GR.sys_id+'">'+
 now_GR.file_name+'</a>');
  }
}
```

Note:
Replace "instance_name" with your instance name.

Attachment Logging

Whenever a user downloads an attachment, the action writes an *attachment.read* event record to the event log. If desired, you can process these events with a Script Action or an Email Notification. This can be useful if you want to do something when an attachment is read. For example, you can record when and by whom certain attachments are downloaded. For this functionality, the *current* variable must point to a sys_attachment record, and the event record must use the following parameters:

- parm1: File name
- parm2: Table name

Baseline email notifications

The baseline system provides several email notifications.

Some events listed in this table do not appear in a business rule and are fired by other conditions in the platform. Some events are hardcoded and are not user configurable.

Note:
The below table list is not an exhaustive list. Products can create their own notifications.

Baseline email notification descriptions

Email notification	Description	Triggering event	Business Rule Controlling Event
Appointment Invite	Type: Meeting Invitation Adds a meeting invitation to the recipient's calendar by sending an iCalendar formatted email	itil_appointment.inserted	Global business rule
Appointment Update	Type: Meeting Invitation Updates an existing meeting in the recipient's calendar by sending an iCalendar formatted email	itil_appointment.updated	Global business rule
Approval Rejected	Type: EMAIL A task-based approval has been rejected – includes the approver's name.	approval.rejected	approver events
Approval Rejected by Other	Type: EMAIL	approval.rejected.by.other	approval events

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
	A task-based approval has been rejected – includes the approver's name.		
Approval Request	Type: EMAIL Sends an email for the recipient to reply with an approval decision. In the approval email, the recipient selects a link that builds the appropriate reply email. These emails come only from task-based approvals.	approval.inserted	approver changes
Catalog Approval Rejected	Type: EMAIL A catalog request has been rejected – includes the approver's name.	request.approval.rejected	approval events
Catalog Approval Request	Type: EMAIL A catalog request for which you were an approver has been cancelled.	request.approval.cancelled	approver changes
Catalog Approval Request	Type: EMAIL A catalog request for which you are an approver has been made.	request.approval.inserted	approval events
Certificate Expired	Type: EMAIL Notification that the X.509 certificate has expired.	certificate.expired	certificate events
Certificate Expiring	Type: EMAIL Notification that the X.509 certificate is expiring in N days.	certificate.expiring	certificate events

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
Change approved	Type: EMAIL A change request has been approved.	change.approved	Change events and task events
Change assigned to me	Type: EMAIL A change request has been assigned to you.	change.assigned	Change events
Change assigned to my group	Type: EMAIL A change request has been assigned to your group.	change.assigned.to.group	Change events
Change commented (to assignee)	Type: EMAIL A comment has been added to a change request. The person assigned to the change request receives an email notification.	change.commented	Change events
Change commented (unassigned)	Type: EMAIL A comment has been added to a change request. The assignment group assigned to the change request receives an email notification.	change.commented	Change events
Change Notification	Type: EMAIL Notification of a change in the fields label in a form.	label.notify	
Change rejected	Type: EMAIL A change request has been rejected.	change.rejected	Change events and task events
Change Task worknoted (unassigned)	Type: EMAIL A work note has been added to a change task.	change_task.worknoted	Change task events

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
	The assignment group assigned to the change task receives an email notification.		
Change Task worknoted (to assignee)	Type: EMAIL A work note has been added to a change task. The person assigned to the change task receives an email notification.	change_task.worknoted	Change task events
Change Task assigned to my group	Type: EMAIL A change task has been assigned to your group.	change_task.assigned.to.group	Change task events
Change Task assigned to me	Type: EMAIL A change task has been assigned to you.	change_task.assigned	Change task events
Change worknoted (to assignee)	Type: EMAIL A work note has been added to a change request. The person assigned to the change request receives an email notification.	change.worknoted	Change events
Change worknoted (unassigned)	Type: EMAIL A work note has been added to a change request. The assignment group assigned to the change request receives an email notification.	change.worknoted	Change events
Email assigned to	Type: EMAIL An incident has been assigned to you (the recipient of the email).	incident.assigned	incident.events

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
Email assigned to (sc_task)	Type: EMAIL A task has been assigned to you (the recipient of the email). **Note:** This email notification was used for the legacy Delivery Plan system of email notifications. In order to use this email notification with a workflow, the **work_start** field needs to be set by the workflow using a Set Value activity.	sc_task.assigned.to.user	sc_task_events
Email assigned To Group	Type: EMAIL An incident has been assigned to an assignment group of which you are a member (the recipient of the email).	incident.assigned.to.group	incident events
Email assigned to group (sc_task)	Type: EMAIL A Service Catalog task has been assigned to an assignment group of which you are a member (the recipient of the email).	sc_task.assigned.to.group	sc_task_events
Incident Closed	Type: EMAIL An incident opened by you (the recipient of the email), has been closed.	incident.updated	incident events
Incident Commented	Type: EMAIL An incident opened by you (the recipient of the email) has had comments added.	incident.commented	incident events

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
	This notification uses a template for an employee self-service (ESS) user.		
Incident Commented	Type: EMAIL An incident assigned to you (the recipient of the email) has had comments added. This notification uses a template for an ITIL user.	incident.commented	incident events
Incident Opened	Type: EMAIL An incident has been opened for you (the recipient of the email) by someone else. This notification uses a template for an employee self-service (ESS) user.	incident.inserted	incident events
Incident Opened & Unassigned	Type: EMAIL An incident has been opened and is unassigned. This notification uses a template for an ITIL user.	incident.inserted	incident events
Incident Resolved	Type: EMAIL An incident opened by you has been resolved, and feedback is required to determine if the incident should be closed.	incident.updated	incident events
Knowledge Closed Created	Type: EMAIL A contributor's submission to the Knowledge Base was accepted and an article was created.	kb.submission.closed_created	KB event
Knowledge Closed Duplicate	Type: EMAIL	kb.submission.closed_duplicate	KB event

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
	A contributor’s submission to the Knowledge Base was determined to be a duplicate, and no article was created.		
Knowledge Closed Invalid	Type: EMAIL A contributor’s submission to the Knowledge Base was determined to be invalid (unusable).	kb.submission.closed_invalid	KB event
Notify Change Calendar	Type: Meeting Invitation Notifies the recipients of the schedule for a change request and exports the schedule to the Microsoft Outlook calendar. The email is in the format of iCalendar formatted email.	change.calendar.notify	change events
Notify Change Calendar Remove	Type: Meeting Invitation Notifies the recipients that a scheduled change has been closed or assigned to someone else and removes the entry from the Microsoft Outlook calendar. The email is in the format of iCalendar formatted email.	change.calendar.notify.remove	change events
Problem Task assigned to me	Type: EMAIL A problem task has been assigned to you.	problem_task.assigned	Problem task events
Problem Task assigned to my group	Type: EMAIL A problem task has been assigned to your group.	problem_task.assigned.to.group	Problem task events
Problem Task worknoted (to assignee)	Type: EMAIL	problem_task.worknoted	Problem task events

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
	A work note has been added to a problem task. The person assigned to the problem task receives an email notification.		
Problem Task worknoted (unassigned)	Type: EMAIL A work note has been added to a problem task. The assignment group assigned to the problem task receives an email notification.	problem_task.worknoted	Problem task events
Problem worknoted (to assignee)	Type: EMAIL A work note has been added to a problem. The person assigned to the problem receives an email notification.	problem.worknoted	Problem events
Problem worknoted (unassigned)	Type: EMAIL A work note has been added to a problem. The assignment group assigned to the problem receives an email notification.	problem.worknoted	Problem events
Reminder Insert	Type: Meeting Invitation Creates a calendar reminder regarding an open task. The email is in the format of iCalendar formatted email.	reminder.notify	
Reminder Insert	Type: Meeting Invitation A task has been closed or deleted. This notification removes the task reminder from Outlook. The email is in the format of iCalendar formatted email.	reminder.notify.delete	

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
Reminder Insert Email	Type: EMAIL Sends the recipient a reminder email about a specific task.	reminder.notify.email	
Request Approved	Type: EMAIL A Service Catalog request, opened by the recipient, has been approved.	sc_request.approved	sc_request events
Request Assigned	Type: EMAIL A Service Catalog request has been assigned to the recipient.	sc_request.assigned	sc_request events
Request Completed	Type: EMAIL A Service Catalog request, opened by the recipient, has been completed.	sc_request.updated	sc_request events
Request Item Assigned	Type: EMAIL An item requested from the Service Catalog has been assigned to you.	sc_req_item.assigned	sc_request events
Request Item Delivery	Type: EMAIL An item requested from the Service Catalog by the recipient is being delivered.	sc_req_item.delivery	sc_request events
Request Opened on Behalf	Type: EMAIL A Service Catalog request has been opened on behalf of the recipient.	sc_request.requested_for	sc_request events
Reset Password	Type: EMAIL	reset.password	

Baseline email notification descriptions (continued)

Email notification	Description	Triggering event	Business Rule Controlling Event
	The recipient's password has been reset as requested.		
Scheduled Import Completed	Type: EMAIL A scheduled import set has completed.	scheduled_import_set.completed	
System Upgraded	Type: EMAIL The recipient's system has been upgraded.	system.upgraded	
Task approved	Type: EMAIL An ITIL task has been approved.	task.approved	Change events and task events
Text Index Completed	Type: EMAIL A scheduled system index has completed.	text_index.complete	
Unscheduled Change	Type: EMAIL A named configuration item has changed, and no active change request exists.	cmdb.unscheduled.change	

Note:

Notification example: notify an assignment group of updates to Priority 1 Incidents

Notify users by email when there are updates to high priority incidents.

Before you begin

Role required: admin

About this task

Send emails to an assignment group whenever there are updates to an incident in which the **Priority** is **1 - Critical**. Include information that is of interest to the recipients, such as the incident number, category, assignees, and any comments that were added to the incident.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**, and then click **New**.
2. On the email notification form, enter the following values:

Field	Value
Name	Priority 1 Incident Updated
Table	Incident [incident]
Active	Selected
Category	Incident Alert
Send when	Record inserted or updated
Inserted	Selected
Updated	Selected
Conditions	[Priority][is][1 - Critical] AND [Updated][changes]
Users/Groups in fields	Assignment group
Subject	`Priority ${priority} Incident updated` **Note:** In this notification, the variable ${priority} returns the value **1 - Critical**.

3. In the **Message HTML** field, enter the following message and script:

```
Short Description: ${short_description}
Click here to view incident: ${URI}
Incident number: ${number}
Category: ${category}
Assigned to: ${assigned_to}
Assignment group: ${assignment_group}
<hr/>
Comments:
${comments}
```

4. From the form context menu, click **Save**.
5. Preview the email notification to ensure it includes all the needed information.

   a. On the notifications form, click **Preview Notification**.

   b. Note that the email includes the following information:
      - Short description
      - A link to the incident record
      - Incident number
      - Category
      - The name of the user to whom the incident is assigned
      - The group assigned to the incident
      - Comments from the incident record

6. Test that the email notification sends to an assignment group when its Priority 1 Incident is updated.
   a. Create a user who has an email address that you can monitor, and then create a group that includes the user that you created.
   b. Navigate to **Incident** > **Open**, and then open an incident in which the **Priority** is **1 - Critical**.
   c. In the **Assignment group** field, enter the group that you created.
   d. From the form context menu, click **Save**.
   e. Add comments to the form to update the incident, and then click **Update**.
   f. Check the email account of the user member in the assignment group.

Notification example: notify task assignees

Notify users who are assigned a Task [task] record.

Before you begin

Role required: admin
Set up your email as a test email address. Navigate to **System Properties** > **Email Properties**, and then enter your email address under **Send all email to this test email address**.

Procedure

1. Navigate to **System Notification** > **Email** > **Notifications**, and then click **New**.
2. On the form, enter the following values:

Field	Value
Name	Task Assigned
Table	Task [task]
Active	Selected
Category	Uncategorized
Send when	Record inserted or updated
Inserted	Selected
Updated	Selected
Conditions	[Assigned to][changes]
Users/Groups in fields	Assigned to
Subject	Task Assigned

3. In the **Message HTML** field, add a message to send to whomever the task is assigned to.
4. From the form context menu, click **Save**.
5. To see a mock version of the system email that you created, click **Preview Notification** on the notification form.
6. Test the notification sends to a task assignee.
   a. Assign some task records.
   b. Check your email for assignment notifications.

Email templates

Email templates enable administrators to create reusable content for the subject line and message body of email notifications.

Templates deliver consistent information on specific system activities and improve the efficiency of creating multiple email notifications for similar actions. If necessary, you can make minor changes on the fly to an email that uses a template by overriding the subject line and message body content in the notification form. A common practice is to define the message body in a template and create new subject lines for different types of recipients. There is no limit to the number of templates that you can create.

Email templates are created in rich HTML format, and administrators have the option of converting existing email templates to rich HTML (starting with the Eureka release). This format provides several advantages, including:

- Raw HTML content is converted into a WYSIWYG format.
- The content can be edited in a feature-rich HTML editor.
- Mail scripts are condensed into a single, easy-to-read line that can be reused in multiple email notifications.
- To prevent broken links, images linked using URLs relative to a particular instance are converted to absolute links.

Create an email template

You can create an email template with rich HTML formatting, rather than plain text.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notifications** > **Email** > **Templates**.
2. Click **New**.
3. Fill in the form fields (see table).

Field	Description
Name	Enter a unique name for the template. For example, `change.update.risk`.
Application	The type of scoped application.
Email layout	If the message body for the template is to be formatted using a predefined email layout, select the layout.
Table	Select the name of the table involved. For example, **Change Request** `[change_request]`.
Subject	Enter a subject line that explains the purpose of the email. Select the appropriate variables for the subject line from the fields available on the selected table. Place the cursor where you want the variable to appear, and click the field name in the **Select variables** column. For example: `${sys_class_name} ${number} with ${risk} risk has been assigned to you.`
Message HTML	Enter the content of the email template message. You can use the HTML editor toolbar to format the HTML, and you can include variables from the

Field	Description
	Select variables column. Variables map to column names available from the notification table, its parent tables, and reference tables. Use variables to include values from a record in the table such as an incident short description or comments and work notes. Notes: ◦ When a template is applied to a notification, the contents of this field are used when you select **HTML and plain text** or **HTML** in the **Content Type** field in the Email Notification form. ◦ You cannot use HTML code to control the appearance of the contents in the `${comments}` variable. ◦ If you want to include a link to the record that triggered the notification, see Enable links to records.
Message Text	Enter the notification message to send in plain text. **Note:** When a template is applied to a notification, the contents of this field are used when you select **HTML and plain text** or **Text** in the **Content Type** field in the Email Notification form.
SMS alternate	Create a different message to be delivered to an SMS device. Enter a brief message, showing the most important information only. If this message field is blank, the contents of the **Message Text** field are used for the SMS message. **Note:** The message in this field is used when the message is sent to a device configured as SMS.

Related topics

Create an email layout

Apply a template to an email notification

After you create an email template, you can apply it to a notification.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Open the email notification record that should use the template.
3. Navigate to the **What it will contain** section.
4. From the **Email template** field, select the appropriate template.
5. Click **Update**.

Convert an email template to rich HTML

If you have an email template created prior to the Eureka release, you can convert it to rich HTML.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Policy** > **Email** > **Templates**.
2. Open the email template you want to convert.

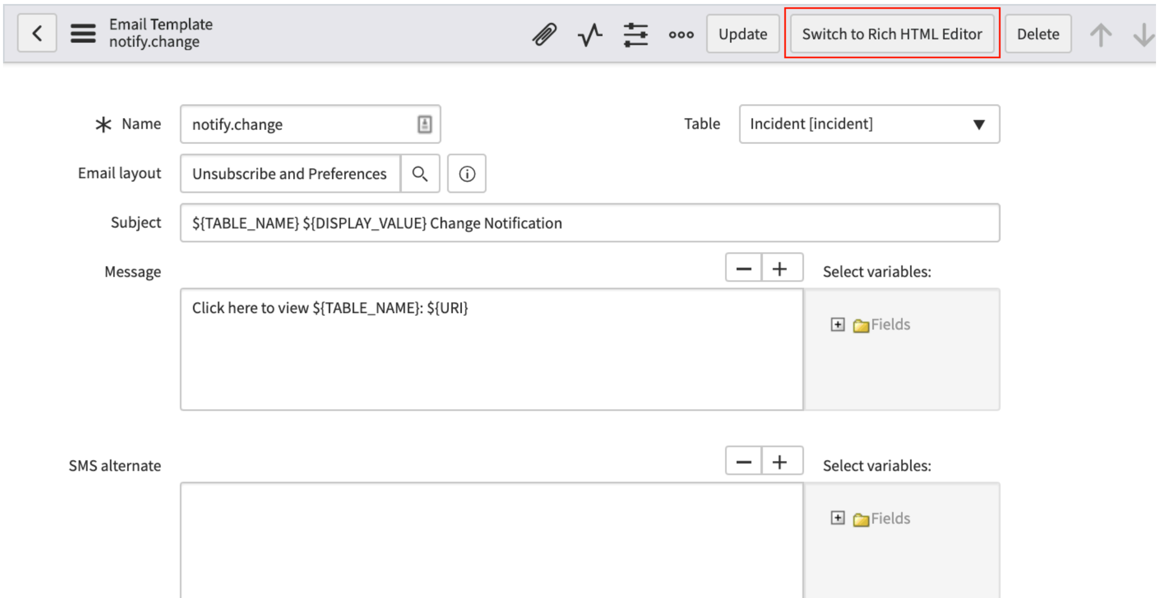

3. Click **Switch to Rich HTML Editor**.

   When content is converted, these changes are made to the content:

   - HTML Editor: Any raw HTML in the **Message** field is rendered as WYSIWYG text in the **Message HTML** field. The content can be edited in a feature-rich HTML editor.
   - Simplified Mail Scripts: Any mail scripts in the body are automatically saved to the Email Script [sys_script_email] table and are replaced in the notification body with an embedded script tag. This makes the template body easier to read and makes it easier to reuse scripts in multiple email templates.

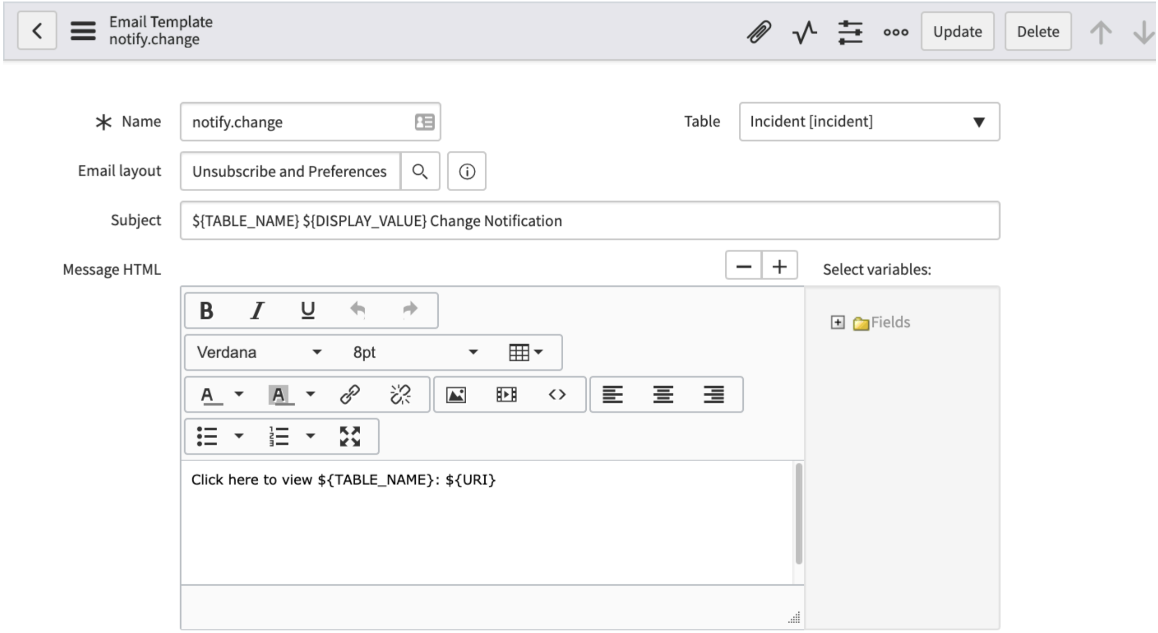

4. Modify the message text, as needed, using the WYSIWYG editor or by editing the underlying HTML.

   Note:
   The letter "P" at the bottom of the screen shows the location of your cursor within the **Message** field. In this case, the cursor is in a line containing an <HTML> tag.

   Note:
   Email templates that are already formatted with rich HTML do not show the **Switch to Rich HTML Editor** button.

Example:
Whether you are working with templates converted from earlier versions or creating new templates in the rich HTML format, it is a good practice to write mail scripts in **System Notification** > **Email** > **Notification Email Scripts**. When the scripts are completed, a `${mail_script:script name}` embedded script tag should be added to the email template body. This makes it easy to use the same scripts in multiple email templates. All that must be copied and pasted from one template to the next is the embedded script tag.

If you manually enter a mail script in the body of a new or converted email notification or template, and then attempt to save the record, a message asks whether the mail script should be converted. Unconverted mail scripts often fail to run from inside the HTML editor. If you select **Yes**, the script is added to the Email Script [sys_script_email] table and is automatically replaced in the body with an embedded script tag.

You can view the mail scripts in their original form by opening the email template and clicking the **Show Notification Scripts** related link.

Construct an email message with a template

Email templates provide a list of the fields from database tables that are available for constructing an email message.

An email template can include a **mailto** automatic response link, which enables the email recipient to simply choose a link that sends a preformatted response back to the instance. An example is the email template used for notification that an approval is required. The following example shows the base email:

Email notification with the template specified

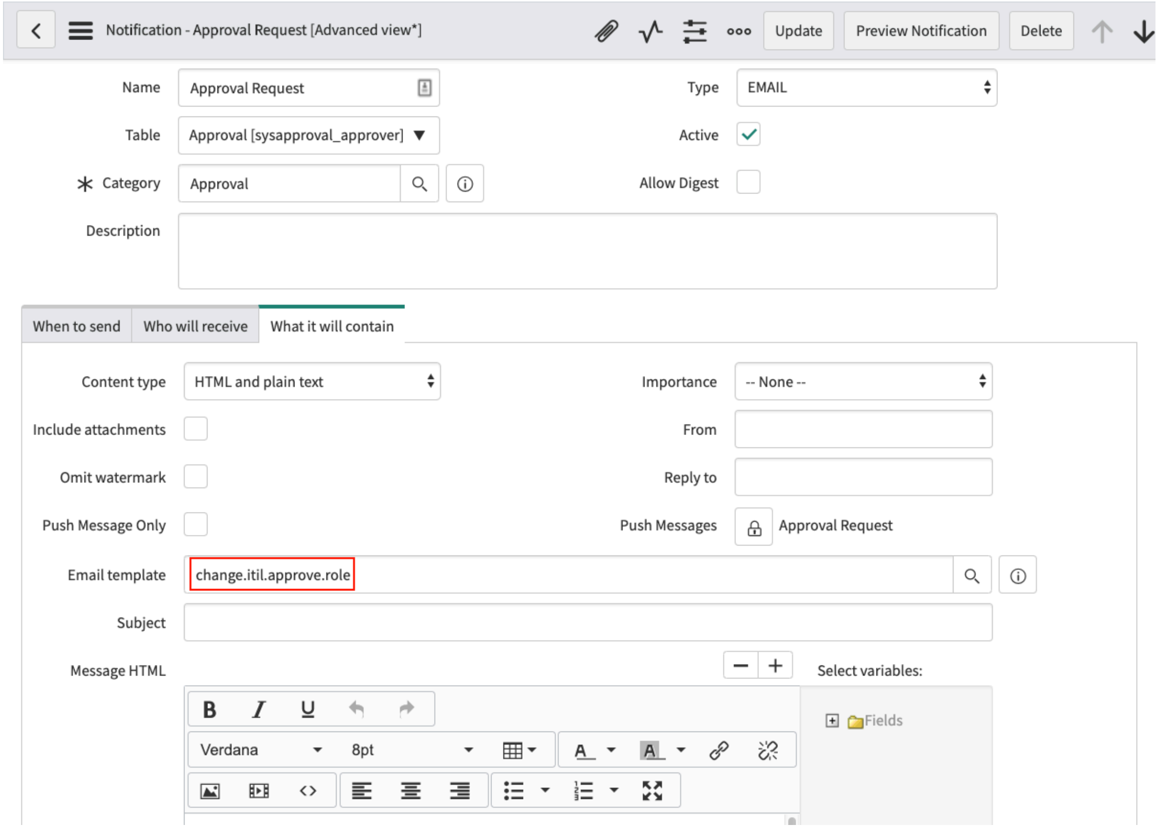

In this case, when a change request approval is requested, an update is made to the `sysapproval_approval` table. The "approval events" business rule is then executed, which creates the "approval.inserted" event. The "Approval Request" email event is defined to process the approval.inserted event, and this email event is defined to use the following change.itil.approve.role email template. Notice that this template has two mailto: items specified. One for "mailto.approval" and one for "mailto.rejection." This email template builds an automatic response that lets the email receiver simply click a link in the email to either approve or reject the change request.

Here is what the **change.itil.approve.role** template looks like using the rich HTML editor:

change.itil.approve.role template

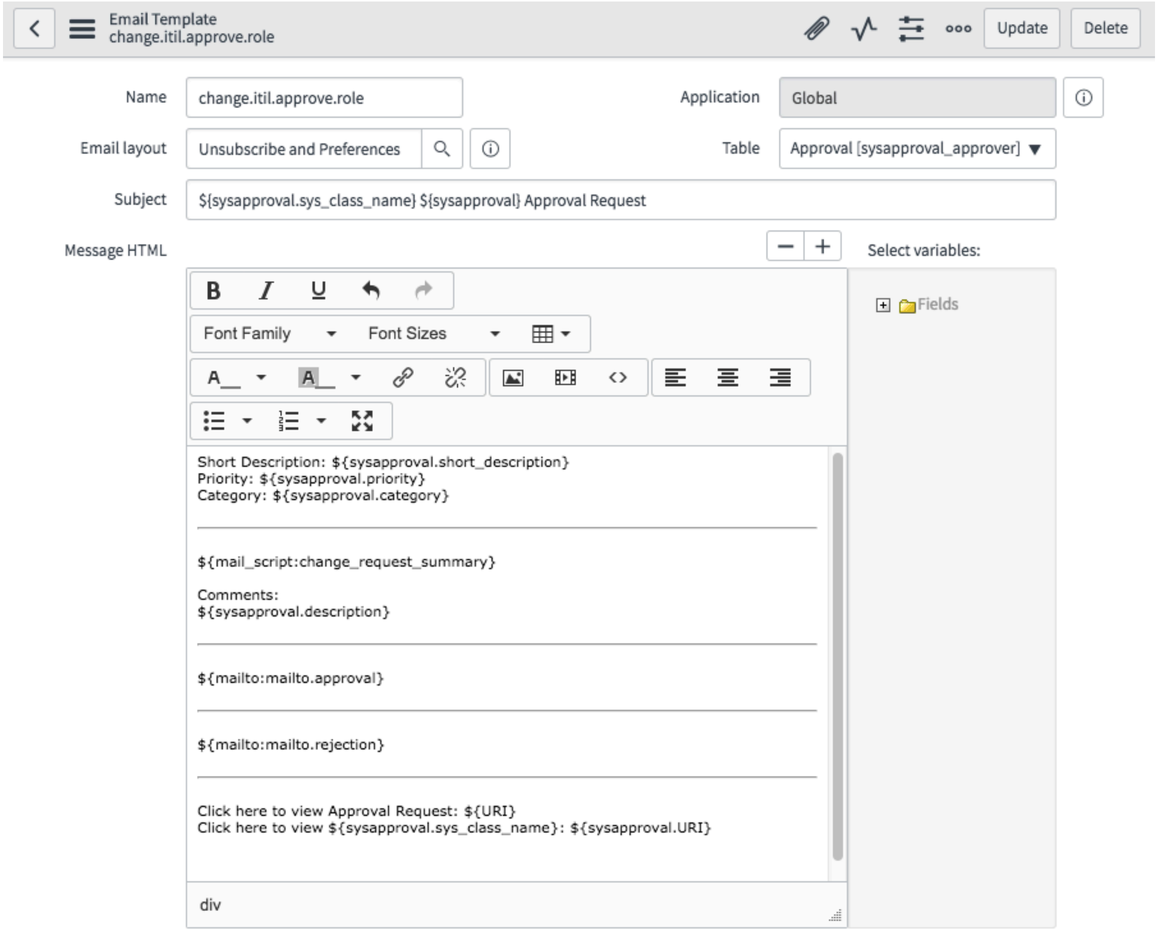

The selection list on the right shows the available fields from the `sysapproval_approver` table.

Notice the **change.itil.approve.role** template also uses a predefined email layout (Unsubscribe and Preferences), which provides links for the recipient to unsubscribe from approval notifications and set notification preferences.

And here is the **mailto.approval** template:

mailto approval template

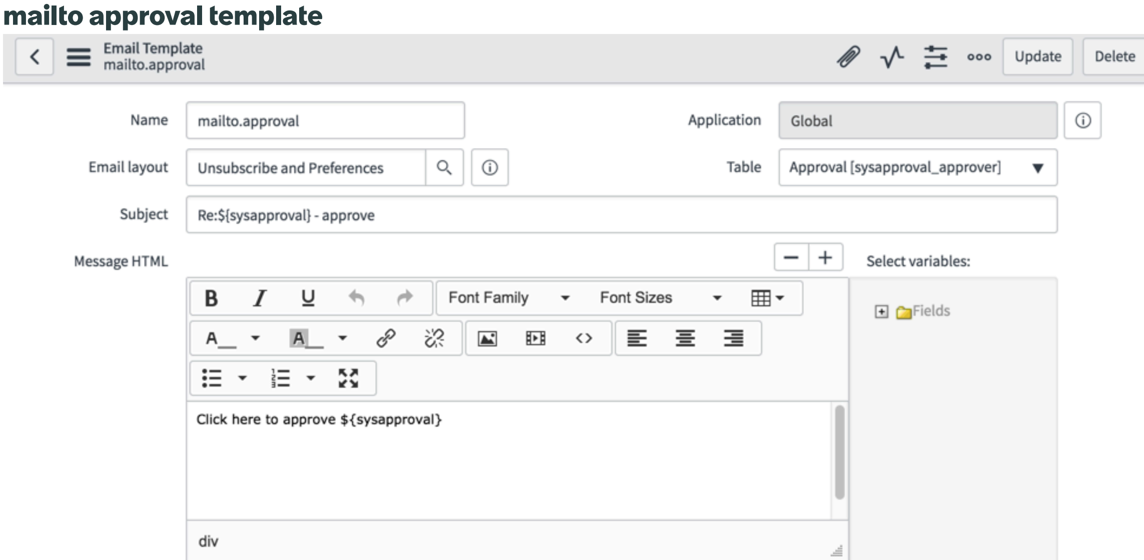

Note:
If you put text on the lines following the **Click here to approve ${sysapproval}** line, this text forms the actual body of the email.

The combination of templates used would generate an email similar to the following example:

Template generated email message

Short Description: Upgrade to a new version of Windows Office
Click here to view Approval Request: LINK
Click here to view Change Request: LINK

Priority: 4 – Low
Category: Software
Comments:
All matches for the sales department must be updated to the latest version of Windows Office

Click here to approve CHG30003
Click here to reject CHG30003

Unsubscribe | Notification Preferences

Notice that the email recipient can use the following links in the mail:

- A link to view all the details of the change request
- A link to view the approval record
- A link that generates an automatic email response to approve the change
- A link that generates an automatic email response to reject the change
- A link to unsubscribe from approval notifications and another link to set notification preferences

Related topics

Email layouts

Add blank lines in an email template

You can easily add blank lines in an email template using HTML tags.

Judicious use of these HTML tags can make your email easier to read.

- You can use **
** to insert a line break
- You can wrap paragraphs in **<p> ... </p>** tags to format your email nicely.

Calendar integration

With email notifications, you can use import export maps to leverage information about records and integrate with Outlook or another calendar.

For example, an email notification can create a calendar event based on the planned start and end dates of a change request. To enable integration with a calendar, the following iCalendar variables are available to be added to an email template message and reference the email template from the notification. The variables must be added to the **Message Text** field.

Variable	Description
${dtstart}	Start Date
${dtend}	End Date
${location}	Location
${alarm_time}	Alert or reminder time

By default, the instance computes the value of iCalendar variables using import and export maps for the Appointment [itil_appointment] and Change Request [change_request] tables. The instance also uses import and export maps for other tables, depending on which plugins you activate. For example, if you activate the Incident (com.snc.incident) plugin, the instance uses the import and export maps of the Incident [incident] table.

Each import export map can specify a different set of iCalendar fields. For example, the icalendar.change_request import export map only maps two iCalendar fields.

The instance uses the **External Name** value as the variable name in the email template. For example, the icalendar.change_request import export map defines the *dtstart* and *dtend* variables.

External name	Associated variable name	Table	Field mapped	Field label
dtstart	*${dtstart}*	change_request	start_date	Planned start date
dtend	*${dtend}*	change_request	end_date	Planned end date

Map date fields to iCalendar variables

You can specify what fields provide the date information in calendar invitation notifications by changing the field mappings of the *dtstart* and *dtend* variables in the import export map for the iCalendar invitation.

Before you begin

Role required: admin

Procedure

1. In the navigation filter, enter `sys_impex_map.list`.
2. Open a map to edit.
3. In the Field Maps related list, click either the **end_date** or **start_date** mapped field to change the mapping for *dtstart* or *dtend*, as needed.
4. Change the **Database** field to the field you want to use to set the start date or end date.
5. Click **Update**.

Create iCalendar invitations for custom tables

To generate iCalendar invitations that use field values from custom tables, create an import export map that computes the values of the iCalendar fields.

Before you begin

Role required: admin

Procedure

1. Create the custom table and fields using the **Date/Time** field type for the fields that map to the *${dtstart}* and *${dtend}* variables.
2. In the navigation filter, enter `sys_impex_map.list`.
3. Click **New**.
4. Set the following fields:

   - **Name**: Use the following naming convention: `icalendar.<table name>`. For example, **icalendar.u_my_custom_table**.
   - **Table**: Select the custom table you created.
   - **Type**: Select **icalendar**.
5. Right-click the form header and select **Save**.
6. In the Field Maps related list, click **New**.
7. In the Mapping Entry Wizard, select **Mapping to a Database Field**.
8. Create field mappings for *dtstart* and *dtend*.

   Example
   For example, map the **External Name** *dtstart* to the**u_meeting_start_time** field in u_my_custom_table.

   Note:
   These variables are required.
9. Click **Submit**.
10. Create field mappings for any of the following iCalendar fields as necessary.
    - dstart
    - dtend
    - location
    - alarm_time
11. Click **Update**.

    Here are sample field mappings between iCalendar variables and custom fields in a custom table, u_my_custom_table:

    Sample field mappings

External name	Database field	Type	Map
dtstart	u_meeting_start_time	field	icalendar.u_my_custom_table
dtend	u_meeting_end_time	field	icalendar.u_my_custom_table

12. Create an email template that defines what to include in the iCalendar invitation.

    a. Set the **Table** field to the custom table you created.

    b. In the **Message** text field, use the following format to define the iCalendar invitation.

```
BEGIN:VCALENDAR
PRODID:-//Service-now.com//Outlook 11.0 MIMEDIR//EN
VERSION:2.0
METHOD:REQUEST
BEGIN:VEVENT
ATTENDEE;ROLE=REQ-PARTICIPANT;RSVP=TRUE:MAILTO:${to}
DTSTART:${dtstart}
DTEND:${dtend}
UID:${sys_id}
DTSTAMP:${dtstamp}
SUMMARY:${u_meeting_summary}
END:VEVENT
END:VCALENDAR
```

Note:
Mail script is not allowed or processed in meeting invitation email templates.

iCalendar invitation template details

iCalendar template line	Required?	Notes
BEGIN:VCALENDAR	Yes	
PRODID:-//Service-now.com// Outlook 11.0 MIMEDIR//EN	Yes	
VERSION:2.0	Yes	
METHOD:REQUEST	Yes	
BEGIN:VEVENT	Yes	
ATTENDEE;ROLE=REQ-PARTICIPANT;RSVP=TRUE:MAILTO: ${to}	Yes	The ${to} contains the recipients as defined in the notification.
DTSTART:${dtstart}	Yes	You must use the import export map to map dtstart to a start time field on the custom table.
DTEND:${dtend}	Yes	You must use the import export map to map dtend to an end time field on the custom table.
UID:${sys_id}	Yes	You must provide the name of a field that uniquely identifies the record, such as the sys_id or the record number field.
DTSTAMP:${dtstamp}	Yes	
SUMMARY:${u_meeting_summary}	No	To include a summary from a text field on the custom table, provide the name of the field, such as u_meeting_summary. The summary field value must not contain line breaks. You do not need to create a field mapping.
END:VEVENT	Yes	
END:VCALENDAR	Yes	

13. Create an email notification to trigger the iCalendar invitation and ensure the following fields are set accordingly:

    - **Table**: Select the custom table.
    - **Type**: Select **Meeting Invitation**.
    - **Content type**: Select **Plain text** only.
    - **Email template**: Select the template you created.

Email layouts

Create reusable content for the message body of email templates.

Email layout overview

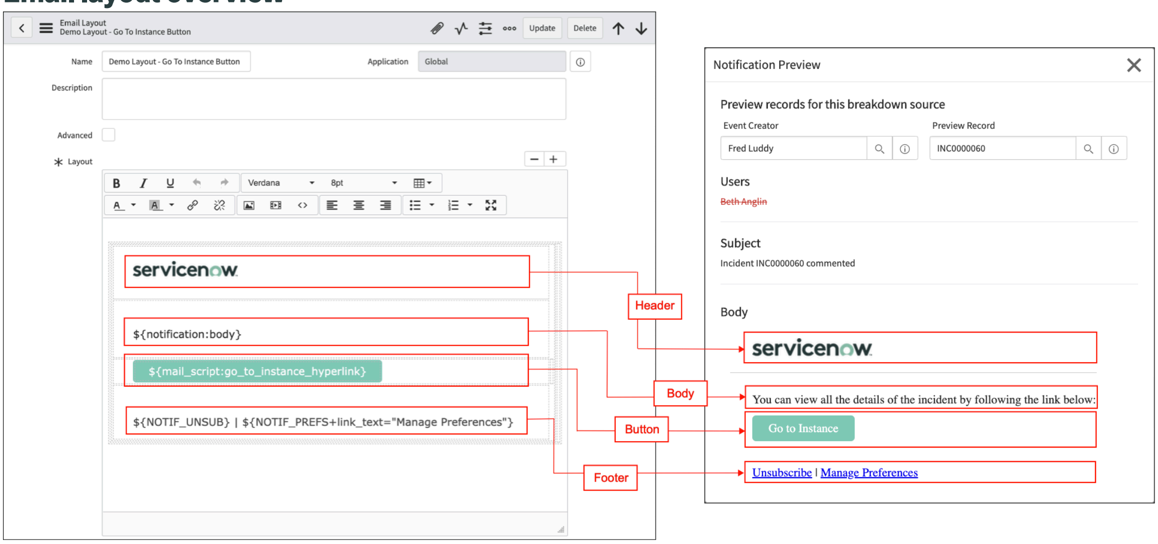

Administrators can use email layouts to:

- Ensure all email notifications have a consistent layout such as always displaying a header, body, and footer.
- Display static content on all email notifications such as a company logo or a background.
- Declare inline styles available for use in the message body of an email template such as setting a text font, size, and color.
- Provide users with links to common response actions such as unsubscribe from a notification or manage notification preferences.

Note:
To display dynamic content such as mail scripts, use email templates. For more information, see Create an email template.

Email layouts insert HTML elements into the message body of email templates. Any style elements you define in the email layout are available to the email template. Email layouts support style sheets in these formats.

- Internal style sheets defined within a `<style>` element.
- Inline styles within a `style` attribute.

By default, the system includes several sample layouts administrators can use to create their own layouts. Administrators can create email layouts using an inline HTML editor or

manually entering HTML code. The system stores email layout records in the Email Layout [sys_email_layout] table.

Note:
If there is no content in the notification from either the notification record or the template, the layout won't be applied.

Create an email layout

Create an email layout to specify the HTML content you want to appear in the body of one or more email templates.

Before you begin

- Role required: admin
- Record required: email template record

Procedure

1. Navigate to **All** > **System Policy** > **Email** > **Layouts**.
   The system displays the list of existing email layouts.
2. Click **New**.
   The system displays a blank email layout form.
3. Fill in the form.

   Email layout fields

Field	Description
Name	Specify a unique name for the record.
Application	Lists the parent application to which this record belongs.
Description	Specify an optional description of the layout.
Advanced	Select whether to display the **Advanced Layout** field.
Layout	Use the inline editor to add HTML elements. The system displays HTML from this field in the body of any email template that uses the layout. Include notification variables to show content from the related record. For more information notification variables see Notification variables. **Note:** The editor automatically formats any HTML code you enter from the **Source code** view.
Advanced Layout	Use this field to manually enter HTML code. The system displays HTML from this field in the body of any email template that uses the layout. **Note:** To enter text in this field, select **Click here to disable syntax highlighting and script formatting**.

4. Click **Submit**.
   The system creates the email layout record.

5. Navigate to **System Notification** > **Email** > **Templates**.
   The system displays the list of existing email templates.
6. Select the email template which you want to use an email layout.
   The system displays the email template record.
7. In **Email layout**, select the email layout you want to use to format the body of email messages.
8. Click **Update**.
   The email template uses the selected email layout to format the body of email messages.

Email retention

You can archive and eventually destroy email messages that you no longer need or if your Email table is excessively large.

Email retention is available starting with the Helsinki release.

Email archive and destruction plugins

The email archiving and destruction feature uses the Data archiving and Email Retention plugins. The Data Archiving plugin must be active to archive and destroy email records. The Email Retention plugin provides a set of rules that specify when the system archives and destroys email records.

Note:
The Email Retention plugin also prevents the system from deleting watermarks, which are required for inbound email actions to continue to function.

The Email Retention plugin and associated archive and destroy rules are active by default on new instances. On upgraded instances, you must manually activate both the plugin and the archive and destroy rules. ServiceNow recommends that you review and approve these rules before activating them.

If your instance already has a process to manage email records, you do not need to activate the Email Retention plugin. If you want to replace your current process with Email Retention, be sure to deactivate the current process before activating the archive and destroy rules.

Archiving and destroying email records

Archiving means moving records from the Email [sys_email] table to the Archive Email [ar_sys_email] table when they exceed the archive rule time limit. Destroying means deleting records in the Archive Email table when they exceed the destroy rule time limit.

Note:
When a destroy rule deletes email records, associated watermarks are not deleted. They are preserved to ensure that your inbound email actions continue to function.

Default archive and destroy rules

Email Retention provides these email archive rules:

- **Emails - Ignored and over 90 days old:** archives email message records that were created more than 90 days prior to the current date and are of type **received-ignored** or **sent-ignored**.
- **Emails - Over a year old:** archives email message records that were created more than 365 days prior to the current date.

Email Retention also provides this email destroy rule:

Email Archive - Over a year old: destroys email records that have been archived for more than 365 days prior to the current date.

With these default settings, your email messages are kept on the instance for a total of two years: one year in the Email table, and one year in the Email archive table. At the end of this period, the system deletes the expired email records from the Email archive table.

Note:
By default these rules are active on new instances and inactive on upgrades. The system runs archive and destroy rules when you activate them.

Compatibility with other record management implementations

If you are already using another method to manage email records, such as table cleaners, you do not have to use the Email Retention feature. To prevent unexpected record deletion, ServiceNow recommends that you avoid using multiple email management processes on the same instance at the same time.

Note:
For assistance replacing your existing record management implementation with Email Retention, contact your professional services or sales representative.

Effects of archiving and deleting email records

Inbound email actions copy the body of an email to the work notes of the related record. If the inbound email record is later deleted, Email Received section is no longer visible in the work notes though it still contain a text copy of the email.

When the system sends an email message about a record, the activity formatter displays a **Sent Email** section with a link to the email message. If the system archives the email message, the activity formatter removes **Sent Email** section. When the system deletes the email message, it is no longer visible in the activity formatter nor the work notes.

Note:
Set the archive time length long enough so your users can access sent emails though the activity formatter.

Archiving email records changes the methods available to the system to identify inbound email as a reply. After archiving an email record, the system can no longer use the **In-Reply-To** field to match an incoming email to an email record. However, the system can still match incoming email to an existing record from a record number or watermark.

Activate the Email Retention plugin

The Email Retention plugin provides archive and destruction rules for email messages. It is active by default for new instances, but must be activated for upgrades.

Before you begin
Role required: admin

About this task
The Email Retention plugin requires these plugins:

- Data archiving
- System Mailboxes

Procedure

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

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

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

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

Archive email manually

You can archive email messages manually on demand instead of waiting for the instance to archive them based on a scheduled job.

Before you begin

Role required: admin

About this task

Email Retention provides these email archive rules:

- **Emails - Ignored and over 90 days old:** archives email message records that were created more than 90 days prior to the current date and are of type **received-ignored** or **sent-ignored**.
- **Emails - Over a year old:** archives email message records that were created more than 365 days prior to the current date.

You can manually archive email messages that meet these archive rules or any additional archive rules that you create.

Note:
An archive record must be active for the instance to be able to process records with it.

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Rules**.
2. Open the email archiving rule that you want to modify.
3. Click **Recalculate Estimate** to see how many records in the Email [sys_email] table are going to be archived.

   Note:
   The estimate appears in the **Record estimate** field.
4. Click **Run Archive Now**.

   Note:
   You can also Create an archive rule.

Watermarks on notification emails

By default, the system generates a watermark label at the bottom of each notification email to allow matching incoming email to existing records. Each watermark includes a random 20-character string that makes it unique.

Starting with the Jakarta release, the system automatically generates randomized watermarks for notification emails in base systems. The random 20-character string reduces the possibility of a watermark being guessed or coincidentally matching the watermark of an email from another instance.

Note:
If you are upgrading from a release before Jakarta, random watermark support is optional and requires the Random Watermark Support plugin to be activated.

Watermark format

The email watermark always begins with "**Ref:**" to identify the label as a watermark. After this identifier, the default label is 31 characters in length and consists of:

- Customizable prefix — The default prefix is **MSG**.
- Auto-numbered identifier — The numeric string identifying the source record, such as incident, problem, or change request.
- An underscore character followed by a random 20-character string

Randomized watermark example

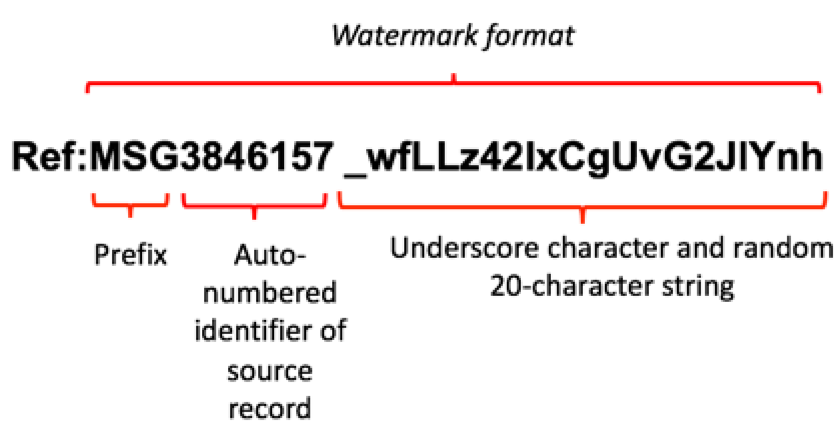

When inbound emails are processed, the system matches random watermarks to the appropriate source records.

Watermark configuration

Watermarks are always generated, but you can configure them to:

- Create a custom watermark prefix for each instance to prevent accidentally triggering events in the wrong instance.
- Have custom prefix characters after MSG
- Be hidden globally
- Be omitted from individual email messages

If watermarks are omitted from email notifications, inbound email actions might not work properly. Without a watermark, the system processes inbound email messages as described in Criteria for matching email to inbound actions.

Note:
Email clients that use the plain text version of the email still show the watermark.

Create a custom watermark prefix for email notifications

By default, email notifications use the watermark prefix **MSG**, but you can create a custom watermark prefix.

Before you begin

Role required: admin

About this task

Any email notifications that are forwarded from one instance to another might be more easily distinguished with different prefixes for each instance. To avoid unintentionally triggering events in the wrong instance, create a unique watermark prefix for each instance.

Note:
Do not use colons (:) in custom watermark prefixes. Colons are a reserved character and may cause the watermark to be ignored.

Procedure

1. Navigate to **All** > **System Definition** > **Number Maintenance**.
2. Open the **MSG** record in the Email Watermark [sys_watermark] table.
3. Enter the unique **Prefix** for this instance.
4. Click **Update**.

Result

Your custom watermark applies to all new email notifications. Email notifications that existed before you created a custom watermark keep the same watermarks as before.

Omit an email notification watermark

You can omit watermarks on email notifications if you do not want the instance to match the notification to an existing record.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Select the email notification to update.
3. Click the **Advanced View** related link.
4. In the **What it will contain** section, select the **Omit watermark** check box.

What to do next

When incoming email does not contain a watermark, the system searches the subject line and message body for a record number. The system attempts to match any record number that it finds to an existing record. If there is a matching record number, the system updates the record with the values in the incoming email. To ensure that response email messages don't update records, remove the record number variable `${number}` from the **Subject** and **Message HTML** fields.

Hide email watermarks globally

Rather than omitting watermarks, it is possible to hide watermarks for global application using HTML markup.

Before you begin

Role required: admin

About this task

Watermarks can only be hidden in the HTML message for the global application. The text version of the message, because it does not have markup allowing show/hide semantics, will always have the watermark.

Procedure

1. Navigate to `sys_properties.list` in the **Application Navigator**.
2. Create a new property named `glide.email.watermark.visible` and set it to **false**.

   This ensures that all watermarks are hidden on all email messages. This cannot be done on a per-email basis.

Parse an email thread

When an email is received, parse the most recent message in an email thread by creating **Email Reply Separators** using specified string or regular expressions.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Policy** > **Email** > **Email Reply Separators**.
   A list of email reply separators shows up.
2. Click **New** to create a new email reply separator.
3. On the **Email Reply Separator** form, fill in the fields.

Fields	Description
Name	Name of the separator.
Application	Scope of the separator.
Active	Option to activate the separator.
Type	Type of separator. ◦ String: Searches for the specified string to parse out an email thread. ◦ Regex: Searches for the specified regular expression to parse out an email thread.

Fields	Description
	Note: Complex regular expressions may take time to evaluate the email thread. In such cases, the regular expression evaluation is timed out after 10 seconds and skipped for that regular expression and the email processing is continued.
Order	Priority of the separator.
Separator	Separator value based on the selection of the **Type**.
Short description	Description about the separator.
Language	Language of the reply separator. **Note:** All email clients don't indicate a language, so all active reply separators are applied to locate a separator in the email message.

4. Click **Submit**.

Email digests

An email digest is a single email that summarizes the activity for a selected notification and its target record during a specified time interval. You can enable an email digest to reduce the number of notifications received when frequent updates to the associated record occur within a short time period.

Example email digest

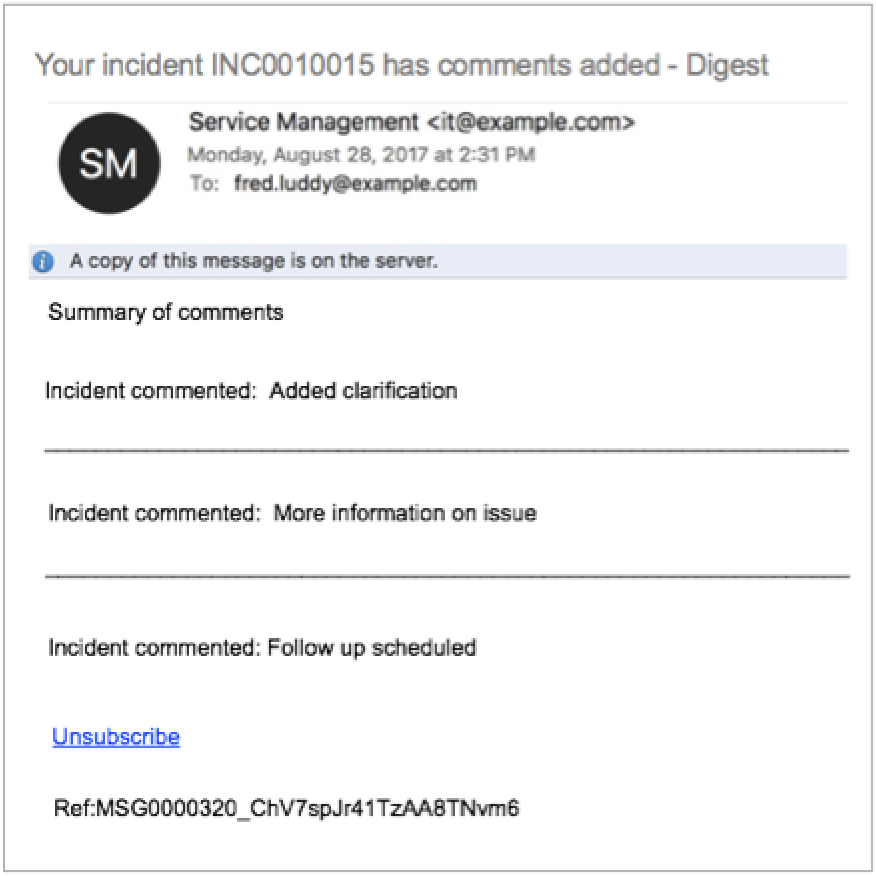

Note:
Email digests apply to email notifications only and are not supported for SMS messages, push notifications, and activity streams.

How email digests work

In new and upgraded instances, the Email digest (com.glide.email_digest) plugin is activated by default. Your instance must use the Core UI interface, since the email digest feature involves setting user notification preferences in the System Settings window.

Admins determine which notifications can be delivered in an email digest and configure the digest content for those notifications. Admins can also control the intervals for digests. An interval is the length of time that notifications are collected for the digest, such as daily or hourly.

Users enable the digest and select the digest interval for a specific notification in their notification preferences. The system accumulates the notifications that normally would be sent during the specified interval and summarizes them in the email digest.

When processing an email digest, the system:

- Stores the digest configuration for the notification in the Notifications [sysevent_email_action] table.
- Temporarily stores the notification content accumulated for a user in the Email Digest Parts [sys_email_digest_part] table and the Email Digest Part Users [sys_email_digest_part_user] table.
- Runs an email digest job every 15 minutes to check when a digest is ready to be sent to a user.

The system uses the digest interval to determine when the digest is ready to be sent. The digest interval begins when the first notification is triggered for the user and stops at the end of the interval time.

- Sends the email digest to the user soon after the selected digest interval ends.

  For example, if a user selects an hourly digest interval and the first notification is triggered at 08:15, the interval starts at 08:15. When the digest interval ends, the system generates the email digest approximately one hour later, at about 09:15 or shortly after, depending on when the email digest job ran.

Set up email digests

What to do — admins

1. Review the base system digest intervals and if needed, create or modify intervals.

   The base system digest intervals are one day (24 hours), one hour, every four hours, and seven days (weekly).

2. Determine the email notifications that can be delivered in a digest and configure the email digest content for those notifications.

   When determining which notifications are appropriate for an email digest, consider the notification content and intervals that your users can select. For example, digests that can be generated weekly might be better suited for notifications that are non-urgent.

What to do — all users

For a notification that can be delivered in an email digest, enable the email digest in your notification preferences (**Notifications** tab in the System Settings window). You also specify the interval time that the notifications are accumulated.

Next step

If you are an admin, begin the digest setup process by reviewing the base system email digest intervals and create or modify the intervals.

Create or modify email digest intervals

Admins can create or modify the email digest intervals (length of time) during which activity for a selected notification is accumulated in an email digest. Users select a digest interval when they enable an email digest in their notification preferences.

Before you begin

Role required: admin

About this task

Before adding an email digest interval, review the base system digest intervals in the Digest Intervals [sys_email_digest_interval] table. You can modify intervals, including the base system intervals, which are one day (24 hours), one hour, every four hours, and seven days (weekly).

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Digest Intervals**.

Action	Description
To add an interval	Click **New**.
To modify an interval	Select the interval to be changed.

2. Enter the new or changed interval information:

Field	Description
Name	Interval name that describes the interval length of time, for example: Every 2 hours.
Interval	Length of the interval, either number of **Days** or **Hours** (hours, minutes, and seconds). The minimum time length allowed is one hour, and the maximum time length allowed is seven days (one week).

3. If you created an interval, click **Submit**, or if you changed an interval, click **Update**.
The system updates the Digest Intervals [sys_email_digest_interval] table with the new or modified interval. The digest interval is listed as an option in the notification preferences (**Notifications** tab in the System Settings window) when a user enables a digest for the notification. For example:

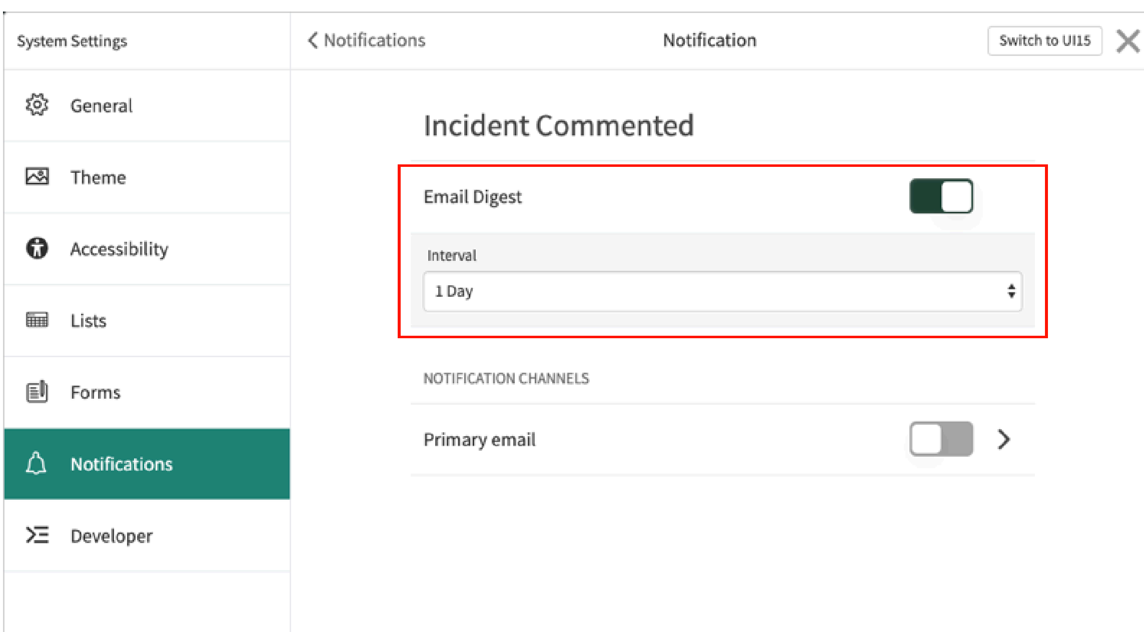

4. To delete an interval, select the interval to be deleted in the Digest Intervals [sys_email_digest_interval] table and click **Delete**.

Note:
If the interval is in use, the system does not remove the interval.

What to do next

Determine the email notifications that can be delivered in a digest and configure the email digest content for those notifications.

Related topics

Email digest properties

Configure email digests

Use the Notification form to create or modify the content of an email digest for a notification. You can also disable the email digest for a notification so that it isn't available in the notification preferences of your users.

Before you begin

Role required: admin

About this task

When you create or update a notification, use the **What Digest will contain** tab of the Notification form to configure the email digest for the notification. After you define the digest content, the digest option for the notification is available in the notification settings of your users.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.

Action	Description
To create a notification	Select **New** to open the Notification form.
To modify a notification	Select the notification to be changed.

2. In the Notification form, select the **Allow Digest** check box to display the **What Digest will contain** tab and define the digest content.
3. Complete the fields in the **What Digest will contain** tab.
   The image shows the default view of the tab. The advanced view contains additional fields (see table).

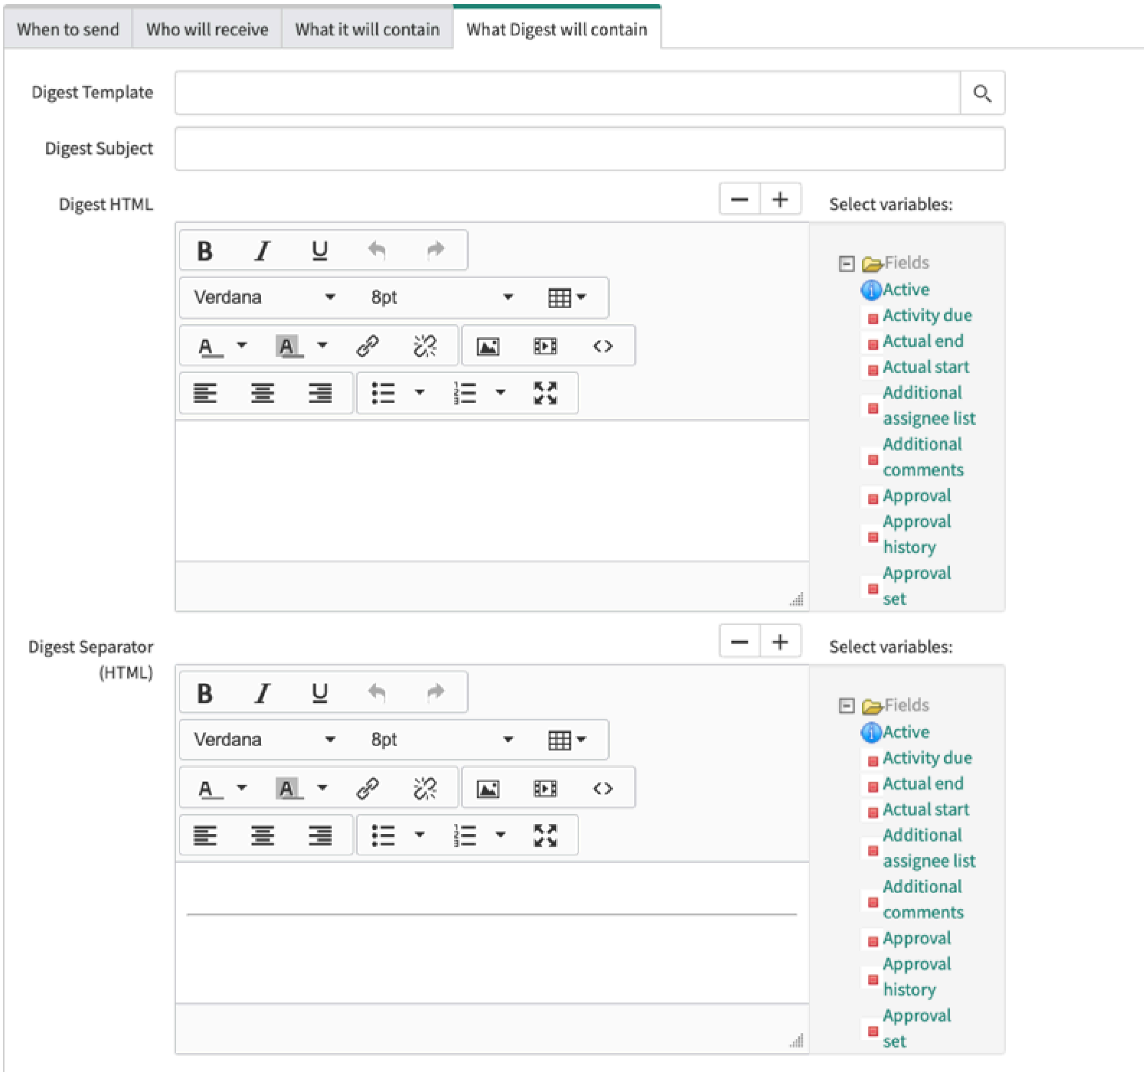

Field	Description
Digest Template	If you want to reuse existing content, such as headers or footers, select an email template to add content to the email digest.
Digest Subject	Enter the subject line for the email digest. The subject can include variables from the **Select variables** column. If empty, the system uses the **Subject** value from the **Email template**. If you enter a value in this field, it overrides the template value.
Digest HTML	Enter the recurring content for the email digest. The digest content can include variables from the **Select variables** column. Variables map to column names available from the notification table, its parent tables, and reference tables. Use variables to include values from a record in the table such as an incident short description or comments and work notes. To prevent adding extra `<p>` and `<div>` elements to your email digest, see the blog post Extra line spacing with paragraph tags in email client by a ServiceNow employee in the ServiceNow Community.

Field	Description
Digest Separator (HTML)	Use the line to separate each item summarized in the digest.
Digest From	Enter the email address to be used in the From field of the email digest. For example, helpdesk@yourcompany.com. The email address must be in a valid format, otherwise a notification message appears near the field. Changing this address requires an advanced email setup such as enabling email forwarding.
Digest Reply To	Enter the email address that you want people to use when replying to the email digest. For example, helpdesk@yourcompany.com. The email must be in a valid format, otherwise a notification message appears near the field. Changing this address requires an advanced email setup such as enabling email forwarding.
Digest Text	Enter the recurring content of the email digest to send in plain text. This field appears when you set the content type to **HTML and plain text** or **Plain text only**.
Digest Separator (text)	[Optional] Use the dash character as a line to separate each item summarized in the digest. This field appears when you set the content type to **HTML and plain text** or **Plain text only**.

4. When you finish creating the notification, select **Submit**.
   The email **Digest** option is displayed in the notification settings of your users.

   Note:
   The email recipients identified in the **Who will receive** tab will receive the digest after they enable the digest in their notification preferences. If the user isn't a designated recipient for the notification, a digest isn't generated even though the user may have the digest enabled for a notification. In this case, the user receives the actual notifications when they're generated.

5. **Optional:** To disable the email digest for a notification, unselect the **Allow Digest** check box and select **Update** when you finish your changes to the notification.
   When you disable the digest content, the digest option is also no longer available for the notification and isn't listed in the notification preferences for your users. If a user previously enabled the digest option for the notification, the system stops collecting those notifications for a digest and doesn't generate the digest.

   Note:
   If you turned off the digest and want to make the digest available again, select the **Allow Digest** check box. The system retains the previously saved digest content and displays it in the **What Digest will contain** tab.

Related topics

[Email digests]

[Create or modify email digest intervals]

[Enable an email digest in Core UI]

Enable email digest in Next Experience

Reduce the number of email notifications you receive by enabling email digest notifications. An email digest is a single email that summarizes the activity for a selected notification and its target record during a specified time interval.

Before you begin

Role required: none

About this task

Notifications do not include the option to be received as an email digest by default. Only admins can create an option to enable email digest for a selected notification.

Procedure

1. From the user profile, navigate to **Preferences** > **Notifications** and then select **Advanced Preferences**.
2. Select **System Notifications**.
3. Expand the notification.
4. Enable **Email Digest** and select the digest interval in which you want to receive the notification communication.

Enable an email digest in Core UI

Reduce the number of notifications you receive for a selected target record during a specified time interval by enabling an email digest. The digest is a single email that collects individual notifications for a targeted record generated during the specified interval. The email digest of two different target records will not be clubbed in a single email.

Before you begin

Role required: none

About this task

The email digest option is available only if your admin configured email digest content for that notification.

Procedure

1. Click the gear icon () in the banner frame to open the System Settings window, and click the **Notifications** tab.
2. Select the notification:
   a. In the Notifications By Category section, click the row or right arrow (**>**) of the category that you want to view.
   b. In the list of notifications for the category, click the row or right arrow (**>**) next to the notification that you want to edit.
3. To enable or disable an email digest for the notification, if an email digest is available:

a. Click the **Email Digest** switch.
   The email digest is enabled when the switch is green and disabled when the switch is grey.

b. If you enabled the digest, select the **Interval** (length of time) during which the notifications are accumulated.

   The digest interval begins with the first occurrence of notification activity. For example, if you select the daily (1 Day) interval and the first notification occurs at 07:00, the system begins accumulating notifications at 07:00 and stops at 07:00 the next day.

Soon after the interval ends, the system sends the email digest to you instead of sending the individual notifications generated during the digest interval. If you disabled the digest, the system immediately stops accumulating the notifications for the digest and does not send the email digest. The system resumes sending the notifications as they are generated.

Related topics

Apply notification conditions

Email digests

Multilingual email notifications

Using multilingual email notifications, you can send email notifications in the recipient's preferred language and can easily manage notification content in different languages.

Email notification translation

You can customize email notifications for users across multiple regions based on their preferred language. For example, a recipient whose language is set to French will receive notification emails in French, while a German requester would be notified in German.

Two types of translations are available for email notifications:

- Static translation

  Customize email notifications for recipients across multiple regions based on their preferred language. With static translations, the translation request goes to a translator.

- Dynamic translation

  Customize email notifications automatically for users across multiple regions based on their preferred language.

You can also send emails using a combination of both static and dynamic translation. For example, if an email has to be sent in English, French, Spanish, and German and the translated templates for French and Spanish already exist, then the dynamic translation will only happen for German.

Groups

If there is no group email ID, a notification triggered for a group with multilingual recipients can send a group email to different people with different language preferences.

If there is a group email ID, an email is sent in a single language to all the recipients in the group. The recipient's preferred language is not considered.

Email notification translation plugins

Activate the glide notification translation plugin to send multi-lingual emails based on the recipient's preferred language. The localization framework plugin is activated by default along with the Glide notification translation plugin.

Activate the translated email retention plugin to archive emails.

Activate Glide Notification Translation

You can activate the Glide Notification Translation plugin (com.glide.notification.translation) for Notifications if you have the admin role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Glide Notification Translation plugin (com.glide.notification.translation) using the filter criteria and search bar.

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

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

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

Activate Translated Email Retention

You can activate the Translated Email Retention plugin (com.glide.email_retention.translation) for Notifications if you have the admin role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Translated Email Retention plugin (com.glide.email_retention.translation) using the filter criteria and search bar.

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

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

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

Static translation

Using static translation, you can customize email notifications for recipients across multiple regions based on their preferred language.

About static translation for email notifications

Author localized content for the various content fields on notifications, email templates, and email layout forms using static translation after activating the translation plugin.

For static translation, the translation request goes to a translator. The translator provides translations for notification content in different languages.

Not all of the fields on the notification forms are translatable.

Translatable fields

Table name	Translatable fields
Email Notification	Subject, message, message_html, and message_text
Email Template	Subject, message, message_html, and message_text
Email Layout	Layout

Note:

- The static content that needs to be translated in mail scripts must be present inside **gs.getMessage()** and translations for the strings must be available in the Message [sys_ui_message] table.
- You can stop the translation of tokens by creating the *`glide.email.outbound.static_translation.session.language.change.enabled`* system property and setting it to false.

Localization Framework

Use the Localization Framework settings to request translations for the following artifact configurations:

- Email Layout Configuration
- Email Notification Configuration
- Email Template Configuration

Related topics

Request a translation for an email notification, template or layout

Enable static translation

Enable static translation of notifications for the global application.

Before you begin

Role required: admin

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. Select **New**.
3. On the form, fill in the fields.

Field	Description
Name	Name for the system property, which is `glide.email.outbound.static_translation.ena`
Application	Application for the system property. **Global** applies the system property to all applications when the system property is set to true.
Description	Description displayed in the System Properties list.
Choices	Description displayed in the System Properties list.
Type	Type of system property, either true or false.
Value	Value for the **Type** selected, which should be set to true.
Ignore cache	Option to ignore the cache.
Private	Option to keep the system property private.
Read roles	Roles that can read the system property.
Write roles	Roles that can modify the system property.

4. Select **Submit**.

Request a translation for an email notification, template or layout

Request a translation to send an email in the recipient's preferred language.

Before you begin

An email notification must exist. For more information, see Create an email notification.

The Localization Framework settings should be configured to request the translation for the artifact configurations and add languages.

Role required: admin

About this task

Notification recipients must be defined as active users in the User [sys_user] table and have a valid email address defined for their primary channel (device) in the Notification Device [cmn_notif_device] table.

Procedure

1. Navigate to the email item for which you want to request a translation.
   - For notifications, navigate to **All** > **System Notification** > **Email** > **Notifications**
   - For templates, navigate to **All** > **System Notification** > **Email** > **Templates**.
   - For layouts, navigate to **All** > **System Policy** > **Email** > **Layouts**.
2. Select the item you want to translate from the list.
3. Select **Request Translation** and select the languages to translate into.
4. Select **Submit**.

Dynamic translation

Using dynamic translation, you can customize email notifications automatically for users across multiple regions based on their preferred language.

About dynamic translation for email notifications

Dynamic content notifications are user-specific. You can use dynamic translation to translate notifications automatically based on the language selected by the recipient.

For example, a requester whose language is set to French will automatically receive notification emails in French, while a Spanish requester would be notified in Spanish.

Using dynamic translation involves the following tasks:

1. Activate Dynamic Translation
2. Set system property for dynamic translation
3. Enable dynamic translation

After dynamic translation is enabled, you can send emails and view or preview the translated email in the preferred languages of the recipients by accessing the **Translated Email Contents** tab.

Note:
The language preference of users mentioned in **cc** and **bcc** fields is not considered.

Enable dynamic translation

Enable dynamic translation of notifications for the global application.

Before you begin

Role required: admin

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. Select **New**.
3. On the form, fill in the fields.

Enable dynamic translation form

Field	Description
Name	Name for the system property, which should be `glide.notifications.translation.dynamic`.
Application	Application for the system property. **Global** applies the system property to all applications when the system property is set to **true**.
Description	Description displayed in the System Properties list.
Choices	Description displayed in the System Properties list.
Type	Type of system property, which should be **true | false**.
Value	Value for the **Type** selected, which should be. `true`.
Ignore cache	Option to ignore the cache.
Private	Option to keep the system property private.
Read roles	Roles that can read the system property.
Write roles	Roles that can write the system property.

4. Select **Submit**.

Configure the note banner for translated emails

Configure the content for a translated email note banner displayed at the top of the page to display a customized message to the recipient.

Before you begin

Role required: admin

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. Search for and select the properties to configure and customize the note banner.

System property	Description
glide.email.translation.dynamic.note.html_body	Displays the message as a header for the translated email html body.
glide.email.translation.dynamic.note.text_body	Adds the message to the text body of the translated email body text.

System property	Description
glide.email.translation.dynamic.note	Applies CSS styles to the note.

Note:
If you want to disable the note banner, set the *glide.email.translation.dynamic.note.enabled* system property to false.

Enable dynamic translation for email notifications

Enable dynamic translation for a notification to send dynamically translated emails to recipients in their preferred language.

Before you begin

An email notification must exist. For more information, see Create an email notification.

Role required: admin

About this task

Notification recipients must be defined as active users in the User [sys_user] table. They must also have a valid email address defined for their primary channel (device) in the Notification Device [cmn_notif_device] table.

Procedure

1. Navigate to **All** > **System Notifications** > **Email** > **Notifications**.
2. Select the name of the email notification you want to send.
3. Select the **What it will contain** tab.
4. If **Advanced view** is now already selected, select it.
5. Select the **Enable Dynamic Translation** check box.

What to do next

Preview email notifications

Domain separation and Notifications

Domain separation is supported in the Notifications 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

Domain separation is specifically supported in Notifications but not in email accounts. Notifications are not data-separated but they are process-separated. Notifications are also triggered by specific actions.

Note:
Subscription-based notifications are not domain aware and cannot support domain-specific settings.

How domain separation works in Notifications

There are two basic components of domain separation and Notifications.

- Notifications are process-separated (not data-separated).
- Notifications are triggered in two main ways:

1. When a record is Inserted or Updated

   Notifications with matching conditions AND in the same domain and global domain as the inserted/updated record are processed.

2. When an event defined in the notification is triggered
   a. Events typically have a target record. For example, [incident.inserted] event references the incident record being inserted.
   b. When an event is fired, notifications configured for that event in the same domain and global domain as the event's target record are processed.

Domains and email accounts

Domain separation is not supported in email accounts for these reasons:

1. Sending mail: There is only one SMTP sender per account. This prohibits providing domains for each account, and they are not configurable.
2. For receiving Inbound mail: You can set up multiple email accounts but cannot meaningfully set the domain of an inbound email action. Inbound Actions are processed in the domain of the user who sent the email. For example: User_A in Domain A sends an email to a ServiceNow email account which executes the "Create an incident" inbound email action. The resulting new incident created by the inbound action is in Domain A.

To learn more see Inbound email actions.

Note:
If the number of email accounts exceeds 20, reception of email slows down.

Use case

If an instance is using the Domain separation plugin and a new email notification is defined for a domain that has the same event as the notification on the global domain, the user receives two emails for the same event.

Solution: Set the [sys_overrides] field on the notification that belongs to the domain so it overrides the setting on global. For more information, see Delegated administration.

Related topics

Domain separation for service providers

Email FAQs and troubleshooting notification emails

Logs and diagnostics are provided to help determine whether notification emails are being sent and received successfully, what any issues are, and who receives the email.

Diagnostics help monitor the overall health of the system and troubleshoot general problems, such as not receiving any incoming mail. Logs help identify problems with individual emails, and different logs are useful for diagnosing different types of problems.

Log checking scenarios

Problem	Log
Need to check whether an individual notification email was successfully sent	Check the **Sent** System Mailbox for that email. Also check the **Failed** System Mailbox for failure notifications.
Individual email failed	• Check the message log of the individual email. • Check the **Error string** field of the email record. **Warning:** Some email servers do not return error strings
Email not received by end user	• Check the **Junk** System Mailbox for notifications about returned emails. These emails also appear in the Emails log with a **Type** of received-ignored. • Check the **Error string** field of the email record. **Warning:** Some email servers do not return error strings • Check the message log of the individual email for the reasons different recipients were included or excluded.

Email FAQs

Email FAQs

Troubleshooting resources

See the following resources on troubleshooting inbound and outbound email problems.

Error or symptom	Solution
Inbound Email FAQ	See KB0957760.
Instance not receiving inbound email	See Inbound email troubleshooting, ServiceNow KB0524472 and KB0520595. These KB articles also provide links to a video series on troubleshooting inbound emails.
Instance not sending outbound email	See Troubleshooting Outbound email, ServiceNow KB0521382. This KB article also provides links to a video series on troubleshooting outbound emails.
Email from Outlook produces an empty Incident record containing an attachment named winmail.dat.	Configure the local Outlook client or Exchange server not to send Rich Text formatted (RTF) data to the instance.
Emails are getting set to type send-ignored for no obvious reason	See KB0790932.
Duplicate emails generated by the instance	See Troubleshooting duplicate emails generated by the instance, KB0529413.
SMTP Sender Job Stuck	See SMTP Sender Job Stuck > 30 Mins alert, KB0755061.
Connection to smtp.office365.com failed	See KB0825391.
IMAP with OAuth error AUTHENTICATE failed	See KB0963959.
Error connecting the POP email account on secure port	See KB0829289.
Inbound email processing is stuck alerts	See KB0855277.
Email Reader / SMTP Sender job stuck	See KB0679998.
Email Reader Job stuck/running for more than 2 hours alerts	See KB0755063.

Email diagnostics

The Diagnostics and Connection page provides information on the current state of your email configuration. This page includes status on email properties, scheduled jobs, and email account connections that affect how your instance receives and sends email.

Email diagnostic information can help you identify problems with inbound or outbound email. The following video shows different steps to determine why your instance is not receiving inbound emails. One of the steps uses email diagnostics to check the email connection status and email reader scheduled job.

Troubleshooting steps to take when a ServiceNow instance does not receive email

Navigate to the Diagnostics and Connection page from either of these modules:

- **System Mailboxes** > **Email Diagnostics**
- **System Diagnostics** > **Email Diagnostics**

Email Diagnostics and Connection page

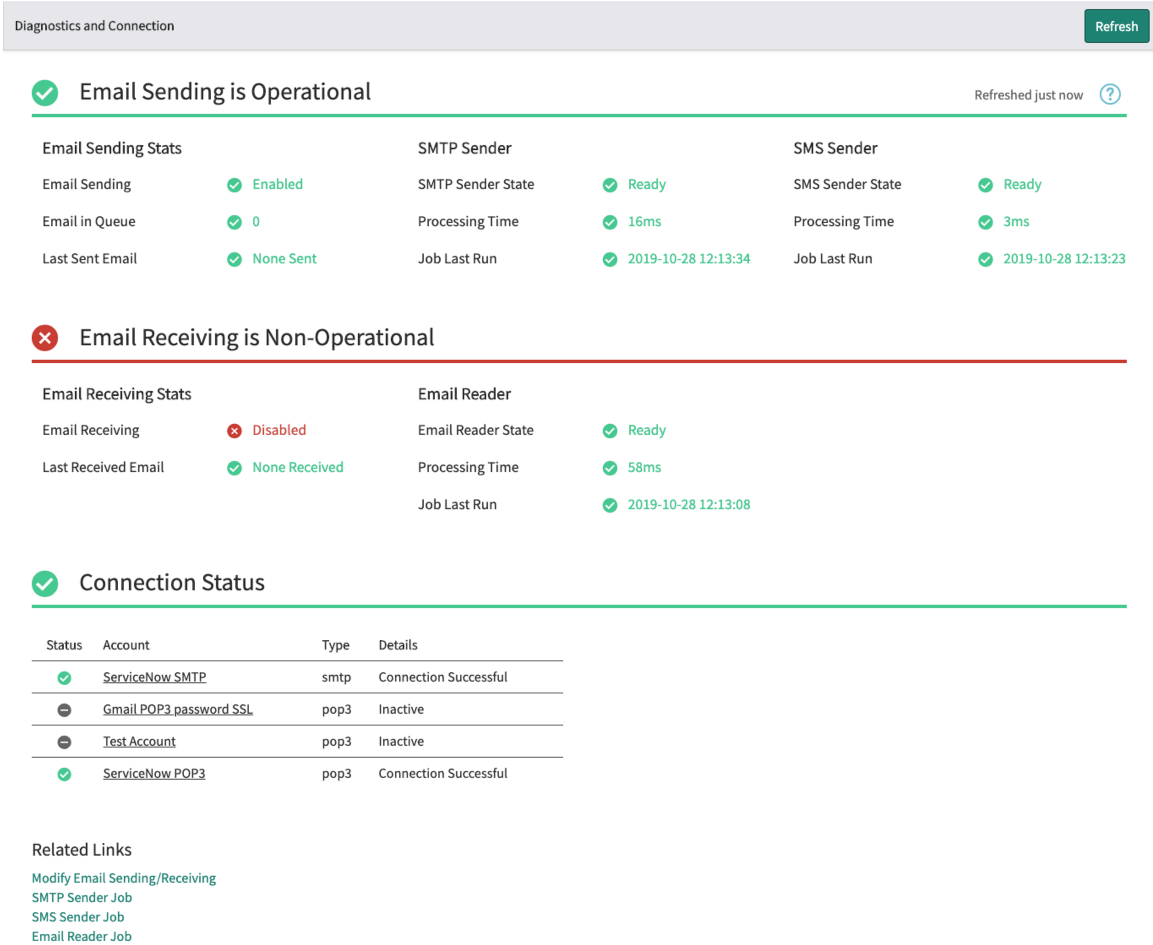

For quick status on a field, point to the green check mark or red X icon. The check mark icon indicates that the item is operational or healthy, while the X icon indicates that the item is not operational or within the expected range.

Mail diagnostics

Diagnostic	Description
Email Sending is [Status]	
Email Sending	Status of outbound email as either **Enabled** or **Disabled**. To change the status, click the **Modify Email Sending/Receiving** related link, and in the Email Properties form, update the `Email sending enabled` property.
Email in Queue	Number of email messages that are ready to be sent.
Last Sent Mail	Date and time the last email message was sent.
SMTP Sender State	Current state of the SMTP Sender job, which sends email on a recurring schedule. By default, this job runs every minute. To update the SMTP Sender state, click the **SMTP Sender Job** related link, and in the Schedule table, open the SMTP Sender record to be changed.

Mail diagnostics (continued)

Diagnostic	Description
SMTP Processing Time	Duration of the last SMTP Sender job run. This value should be shorter than the SMTP Sender interval.
SMTP Job Last Run	Date and time when the SMTP Sender job last ran.
Default SMTP Status	Indication of whether the SMTP connection was successful, shown only if the email accounts feature is active. Click **Default SMTP** to change your SMTP account settings.
SMS Sender State	Current state of the **SMS Sender** job, which sends SMS notifications on a recurring schedule. By default, this job runs every minute. To update the SMS Sender job, click the **SMS Sender Job** related link, and in the Schedule table, open the SMS sender record to be changed.
SMS Sender Processing Time	Duration of the last **SMS Sender** job run. This value should be shorter than the **SMS Sender** interval.
SMS Sender Job Last Run	Date and time when the **SMS Sender** job last ran.
Email Receiving is [Status]	
Email Receiving	Status of inbound email as either **Enabled** or **Disabled**. To change the status for email receiving, click the **Modify Email Sending/ Receiving** related link, and in the Email Properties form, update the *`Email receiving enabled`* property.
Last Received Email	Date and time the last email message was received.
Email Reader Status	Current state of the email reader job, which downloads any email waiting on the mail server and creates email.read events. To update the Email Reader Status job, click the **Email Reader Job** related link, and in the Schedule table, open the email reader record to be changed.
Email Reader Processing Time	Duration of the last reader job run. This value should be shorter than the reader interval.
Job Last Run	Date and time when the reader job last ran.
Connection Status	
{Accounts}	The result of the test connection to the accounts. The connection is tested every time you load the page.

Related topics

Scheduled jobs

System email log and mailboxes

The system email log records all emails that the instance creates or receives. System mailboxes are filtered views of this log.

Every notification email that the instance creates or receives is recorded in an Email [sys_email] record. You can navigate to a log of these records at **System Logs** > **Emails**.

The System Mailboxes are filtered views of the Emails [sys_email] table. The instance assigns an email record to a system mailbox depending on the values of the **Type** and **State** fields. For more information, see System mailboxes.

The following fields can be included in the layout of the system log and any of the system mailboxes:

Email log

Field	Description
Mailbox	The system mailbox that lists this email record. The instance sets the value of this field according to the values of the **Type** and **State** fields.
State	The current state of the email (Error, Ignored, Processed, or Ready).
Receive type	The type of inbound email (None, Forward, New, or Reply).
Type	The status of the email. Choices are: • received: The server received this email. • received - ignored: The server received this email, but it was ignored by the instance for inbound email action purposes. Typically, these emails are either spam or auto-replies. See the Error String field for details. • send - failed: The server has attempted to send the email and failed. See the Error String field for details. • send - ignored: The server skipped sending this email. Typically, this is for an email which was generated but lacked a recipient email address or is a duplicate email. See the Error String field for details. • send - ready: The email is ready to be sent, but has not been sent out by the mail server. Typically, an email remains in this state for only a short time. • send - translation - ready: The email is generated during email translation and sent out. Typically, an email remains in this state for only a short time. • sent: The email was sent by the instance without any errors or issues.
Target	A Document ID reference to the record if the email is generated by an insert, update, or delete of a particular record.
User	The name of the user, from the user record, of the instance from which the email notification was sent. **Note:** This is a string field.

Email log (continued)

Field	Description
Notification Type	The type of notification. Choices are: • None • SMS • SMTP
UID	The unique ID of the email stored on the server.
Created	The date and time of the email activity for the locale of the machine running the instance.
Deleted	For inbound email, indicates whether the email was deleted from the email server.
Weight	The weight of the email, which determines the sending priority relative to other notifications on the same table.
Importance	An indication that the email was sent with a changed level of importance, such as Urgent.
Originating Event and Notification	For emails generated by notifications, an embedded list that stores the event and notification that created the email.
Subject	The email subject. For notifications, you create the subject text in **System Notification** > **Email** > **Notifications**.
Error String	The error string captured from the email server to determine why the email was not sent. This is logged only if the email is send-failed.
Recipients	The email addresses of the recipients.
Body	The body of the email, displayed in raw HTML markup. Use the related link Preview HTML Body to see the body text as rendered HTML.
Content type	The email content type.
Headers	Any headers embedded in the email.

Message logs for individual emails

The email log entries for an individual notification email are accessible as a related list in the email record. The **Error string** field in the email record can provide additional information.

Every email record contains an **Email Log** related list. This list shows the Email Log Entry [syslog_email] records. For troubleshooting purposes, the most useful fields are probably **Level** and **Message**.

To diagnose problems with outbound emails, also examine the **Error string** field in the email record. However, not every receiving email server sends back an error string.

Note:
The Email Log Entry record is available for only 7 days after the email record is created.

Inbound actions on received emails

Check these logs as the first step to diagnose issues with inbound actions. The logs show which inbound action did or did not apply and for what reason.

Note:
If an inbound action did not update any field on the target record, the log indicates that the inbound action was skipped.

Reasons for including or excluding recipients

For outbound notifications, the email system log provides reasons that recipients were included or excluded.

Each log entry corresponds to a reason for inclusion or exclusion. For example, all users who were excluded because they are inactive appear in a single log entry.

A series of system properties can be used to fine-tune the information to be logged. Two properties, *`glide.notification.recipient.include_logging`* and *`glide.notification.recipient.exclude_logging`*, control all recipient inclusion and exclusion logging. Several other properties allow you to tailor the information reported in the logs to meet your needs. All of the properties are enabled by default.

Push notifications

In addition to sending email and SMS notifications, an instance can send push notifications to mobile devices.

A push notification is a text message that appears on a user's mobile device to alert them about something important or to ask them to perform an action. Your instance supports push notifications.

An example push notification

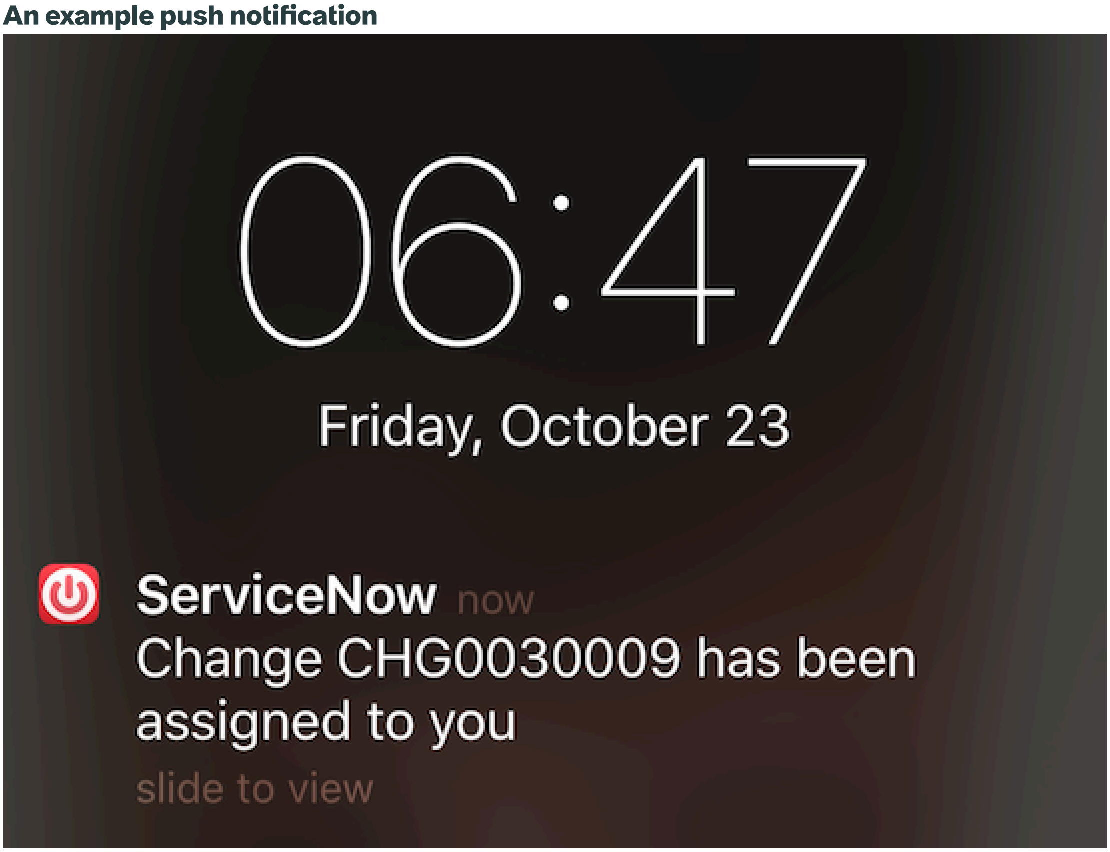

Use push notifications to send messages to users when certain conditions are triggered on your instance, such as the assignment of an incident to the user. A push notification can even ask for a reply, and the instance can process the reply by acting on the related records. For example, you can have the instance send an approval request for a Change to a user. You can let the user approve or deny the Change by clicking a response button on the push notification. The user's response can then update the status of the Change record.

You can set up push notifications in a similar manner to email and SMS notifications. Determine:

- who to send the notification to
- when it should be sent
- what it should contain

The ServiceNow mobile app is available as ServiceNow Classic and as Mobile Agent. By default, ServiceNow Classic supports push notifications, but you can also develop your own push application and configure your instance to send push notifications to it. For more information on setting up push notifications in Mobile Agent, go to Mobile push notifications.

Push notification setup

Push notification setup differs depending on the mobile or push application that you want your users to use. For an overview, see: Push notification setup with a custom push application for ServiceNow mobile apps. If you create your own application, you must understand how push notifications and the Apple Push Notification Service system work. For more information, see the APNs Overview in theLocal and Remote Notification Programming Guide for Apple developers.

> **Important:**
> Apple does not guarantee delivery of all push notifications. Review the Quality of Service (QoS) information in the Local and Remote Notification Programming Guide for Apple developers.

Push notification system

The push notification system involves several key elements that manage the delivery of push messages and push notification responses.

Customer instance

Your ServiceNow instance.

Push proxy

An instance that collects all push notifications that go to the ServiceNow Classic mobile application and forwards them to the ServiceNow applications for Apple iOS or Android. If you create a custom push application, you do not use the push proxy.

Push provider

The provider of push messages, which is the Apple Push Notification service (APNs) or Firebase Cloud Messaging (FCM) service, for the ServiceNow Classic mobile app or custom mobile apps.

Push application on a mobile device

The application, such as the ServiceNow Classic mobile application. You can also build and customize your own push application.

Push notification process

The push notification process is as follows:

1. Activity on the instance triggers a push notification.
2. The instance looks for who to send the notification to and checks the user notification preferences to find their push device settings.
3. The instance sends the push notification to the push notification service, either the Apple Push Notification service (APNs) or Firebase Cloud Messaging (FCM) service. If you are using the ServiceNow Classic mobile app, the instance sends the notification through a push proxy instance, which then forwards the notification to the APNs or FCM service.
4. The push notification service sends the notification to the user's push application. When the instance sends the notification to the push notification service, it includes message content along with the instance name and ID. This information ensures that a response to the notification is sent back to the correct instance. If the user can reply to the notification, such as sending an approval to a change request, the response is sent via REST message back to the instance.
5. The instance identifies a script to run to handle the response.
6. The script performs an action on the instance, such as marking a change request approved. If there is invalid JSON or if the script fails, an error response is sent back to the device.

Push notification system with the ServiceNow Classic mobile app

Push notification responses

With custom push applications, you can have your users act on records in the instance by responding to a push notification.

The actions that they can take are defined in a script that you associate with the notification message. The response is sent to the instance in this format:

```
https://{instance_name}/api/now/v1/push/{application
 Name}/action/{action}
```

The application name is the push application the user is using. The application name must match the application name in the Push Application [sys_push_application] table. The action is the Sys ID of the action in the Push Notifications Actions [sys_push_notif_act_script] table.

Activate push notifications

Several plugins must be activated to use push notifications. If you have the Mobile UI (com.glide.ui.m) plugin active, push notification plugins are automatically activated.

Before you begin
Role required: admin

About this task
Ensure the following plugins are active on your instance:

- Push Notification: Provides the necessary components and REST APIs to send push notifications to mobile devices.
- Notification System Push Addon: Adds support for push notifications to the existing notification system.
- Push Feedback: Handles feedback from Apple on which devices are no longer valid so they do not keep receiving push notifications.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. On the All Applications page, select **Request Plugin** to open the request form on Now Support.
3. On Now Support, select the link to access the Now Support Service Portal Service Catalog.

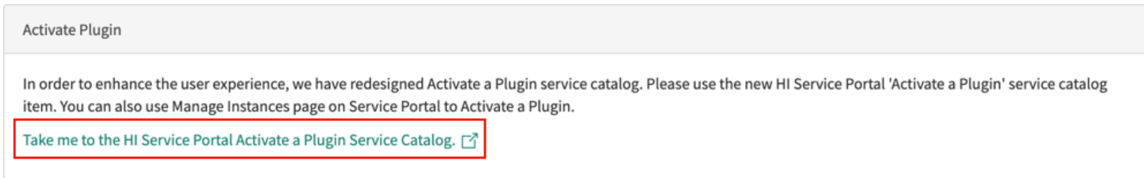

4. Select your instance.
5. Select **Actions > Activate Plugin**.
6. On the form, fill in the fields.

Activate Plugin request form

Field	Description
Target Instance	Instance on which to activate the plugin.
Plugin Name	Name of the plugin to activate.
Specify the date and time you would like this plugin to be enabled	The date and time must be at least two business days from the current time. **Note:** Plugins are activated in two batches, once in the morning and once in the evening, on every business day in the US Pacific time zone. If the plugin must be activated at a specific time, enter the request in the **Reason/Comments** field.
Reason/Comments	Information that would be helpful for the ServiceNow personnel who are activating the plugin. For example, if you need the plugin activated at a specific time instead of during one of the default activation windows, specify that in the comments.

7. Select **Submit**.

Installed with push notifications

Several types of components are installed with the push notifications plugins.

Tables installed with push notifications

These tables are installed with push notifications.

Table name	Description
Push Applications [sys_push_application]	Push applications registered to the instance to receive push messages.
Push Default Registrations [sys_push_notif_default_reg]	Contains all the notifications users are automatically subscribed to for a given push application. You can add notifications to this table on the Push Default Registrations related list on the Push Application form.
Push Feedbacks [sys_push_feedback]	Feedback from the APNs that tells the instance which push devices can no longer receive push messages. The feedback is handled either on the same instance where your notifications are triggered, or on a central instance that uses a REST call to tell your instance which push applications to deactivate.
Push Message Attribute Definitions [sys_push_notif_msg_attr_def]	The attribute definitions used for push message content specification.
Push Message Attribute Values [sys_push_notif_msg_attr_val]	The values associated with push messages.
Push Notifications [sys_push_notification]	The push notifications that the instance attempted to send to users.
Push Notification Actions [sys_push_notif_act_script]	The action scripts that the instance uses in response to an actionable push message.
Push Notification Installations [sys_push_notif_app_install]	The devices with push apps where users agreed to receive push notifications. This table lists the records by the token for the push app. The instance uses this information to know which push device and app to send notifications to. Records are created in this table when a user logs on an instance with the push app. The Mobile Devices [sys_mobile_devices] table, which is installed with the Mobile UI plugin, contains all the user devices that logged in to the instance with the ServiceNow mobile application.
Push Notification Messages [sys_push_notif_msg]	Messages customized for push notifications. These messages can be associated with a notification.
Push Notification Message Contents [sys_push_notif_msg_content]	The entire content, including JSON, for push messages.
Push Platforms [sys_push_platform]	The platforms that are supported for push notifications, and the maximum payload size.

Business rules installed with push notifications

These business rules are installed with push notifications.

Business rule	Table	Description
Build Device and Subscriptions	Push Installation [sys_push_notification_installation]	Automatically creates a notification device on a user's notification preferences when the user registers a mobile application with the instance.

Outbound REST messages installed with push notifications

These outbound REST messages are installed with push notifications.

REST message	Description
ServiceNowMobileApp Push	The REST message that you can use with your custom iOS mobile app.

Roles installed with push notifications

These roles are installed with push notifications.

Role	Description
push_admin	Can create and modify push notifications.

Push components installed with push notifications

These components are installed with push notifications.

Push applications installed with push notifications

Push application	Description
ServiceNowPushApp	The push application record for the ServiceNow ServiceNow mobile application. **Note:** The ServiceNow mobile app automatically subscribes users to several push notifications. You can see the list of these notifications in the Push Default Registrations related list on the ServiceNowPushApp application record.

Message content installed with push notifications

Message content	Description
ConnectMessageContent	The payload for ServiceNow connect messages.
Generic Approval (Background)	The payload that provides an approval and rejection option to the user receiving the push message.
Generic Approval (Foreground)	The payload that provides an approval and rejection option to the user receiving the push message.
Generic Record Payload	A generic payload that you can use to send push messages to users.

Push notification actions installed with push notifications

Push notification action	Description
Approval - Approve	Allows an administrator to approve a request.
Approval - Reject	Allows an administrator to reject a request.

Push notifications installed with push notifications

Push notifications	Table	Description
ConnectMessagePushNotification	Live Feed Message [live_message]	Sends a push notification when a new live feed message

Push notification properties

Push notifications provides several properties to customize the setup.

Add these properties to the System Properties [sys_properties] table.

`glide.push.apns.version`

Controls which version of the Apple Push Notification service (APNs) to use in your instance. You can use APNs version 1 or 2. The system uses APNs version 2 by default. To use APNs version 1 instead, set the value to **1**.

- Type: integer
- Default value: 2

`glide.push.debug`

Creates entries in the system log for push notification errors.

- Type: true | false
- Default value: false

`glide.push.enabled`

Enables or disables push notifications.

- Type: true | false
- Default value: true

`glide.push.feedback.debug`

Creates entries in the system log for feedback sent by the APNs for custom iOS push applications.

- Type: true | false
- Default value: false

`glide.push.notification.ttl_seconds`

Specifies the number of seconds after which a queued push notification expires. After a push notification is triggered, it is queued and finally processed by a scheduled job. If the time the notification sits in the queue exceeds this value, the notification is not sent. Check the Push Notifications log for more information.

- Type: integer
- Default value: 21600

Push notification setup with the ServiceNow mobile app

The ServiceNow mobile application automatically enables push notifications and configures mobile devices for push notifications when users initially log in to the instance.

Note:
Push notifications with the ServiceNow mobile application are not supported in on-premise instances.

Setting up push notifications involves both the system administrator and users.

- Admin creates or updates push notifications for the ServiceNow mobile app. The admin does not need to set up the mobile devices for users nor update their user preferences for receiving push notifications.
- Users install the ServiceNow mobile mobile on their iOS or Android mobile devices and set their preferences for receiving notifications on the ServiceNow mobile app.

Push notification setup tasks for the ServiceNow mobile app

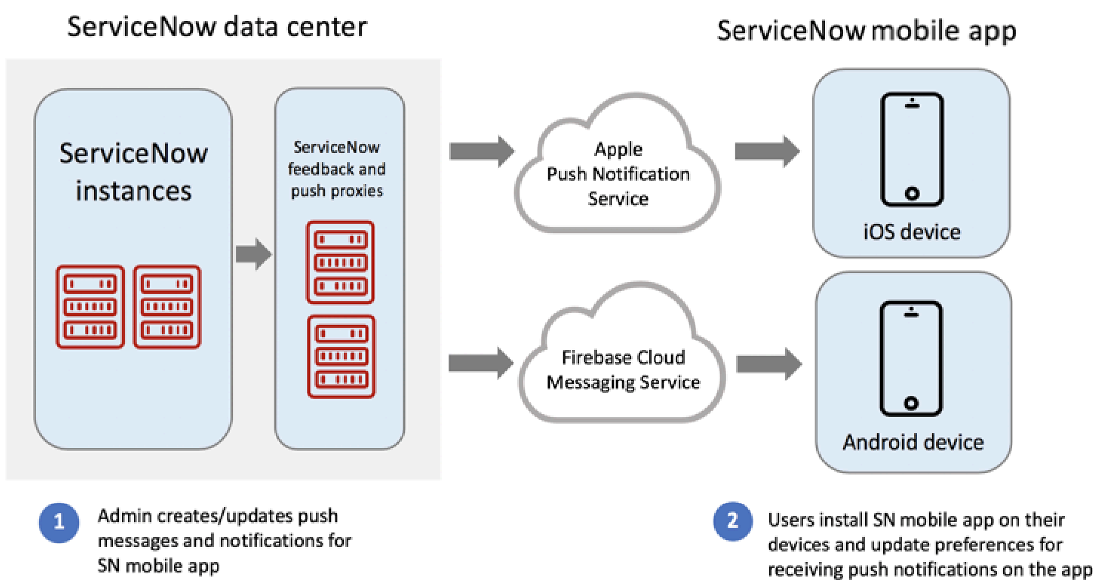

What to do — admins

Admins create push notifications, similar to setting up email notifications. A push notification has two main parts: the push message and the notification, which includes the push message.

Note:
The push message and notification must be for the same table.

Create the push message

Before creating a push notification, create the push message with the actual message content (JSON payload) to be included in the push notification.

Create the push notification that includes the push message

Create the push notification that includes the push message and message content. A push notification specifically sends the push message. You can update a

standard platform notification and use it as a push notification that includes the push message.

Add the new push notification to the Push Default Registration table

To ensure that a new push notification is included and enabled in the notification preferences of your users, register a new push notification in the Push Default Registration table. The notification is listed in the notification preferences of your users the next time that they log in to the ServiceNow mobile app.

What to do — ServiceNow mobile app users

Download the ServiceNow mobile app

Install the ServiceNow mobile app on an iOS or Android mobile device.

Agree to accept push notifications

After installing the ServiceNow mobile app, users are asked to accept push notifications. When the instance receives the push notification acceptance message from a user, it creates a record in the Push Notification Installations [sys_push_notif_app_install] table. This record is what the instance uses to identify the device + mobile application combination necessary to identify a push notification recipient.

The device then performs the following actions:

- Obtains a token that identifies the device.
- Triggers the creation of the device in the user notification preferences.

Log in to the company instance from the ServiceNow mobile app

Logging in automatically subscribes users to the push-specific notifications registered in the Push Default Registration table.

Update preferences for receiving notifications on the ServiceNow mobile app

Users can enable and disable through **Notification Settings** in the ServiceNow mobile app. They can set additional notification preferences, such as conditions or filters that affect notification delivery, by using the System Settings window on a desktop or mobile browser. For details, see Setting notification preferences in Core UI.

Next steps

If you are an admin, review the base system notifications and determine if new push message content and push notifications are needed.

Create a push message

Before you create a push notification, create the push message with the actual message content for the notification.

Before you begin

The Push notification plugin must be active. The plugin is active by default.

Role required: admin

About this task

The push message and notification must be for the same table.

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Messages** and click **New**.
2. Fill out the fields on the form (see table).
3. Click **Submit**.

Push Notification Message form fields

Field	Description
Name	Enter a descriptive name for the message.
Push App	Select the ServiceNow mobile application.
Push Message Content	Select the JSON content to be included in the push notification payload. **Note:** For details on defining payloads that control push notification behavior, see KB0622333. For example, you can specify predefined button pairs (Yes/No, Approve/Reject, Accept/Decline) as part of the push message content.
Message	Enter the message. You can add variables just as you would for other notifications. Any message you enter here overrides the message in the notification.
Related list	
Push Message Attribute Values	Optional. Select the attributes that apply to this notification. For details, see Create an attribute value or action for a push message.

What to do next

Set up the push notification that contains the message created or update an existing push notification to use the push message.

Create a notification using a push message

Email administrators can create a notification that specifically sends a push notification.

Before you begin

Configure the push message before performing these steps.

Role required: admin

About this task

You can associate a push message with a standard notification. A push message specifies the text the system sends as part of the push notification to the mobile device.

Procedure

1. Navigate to **All** > **System Notification** > **Create Push Notification**.
2. Fill out the notification form as necessary (see Create an email notification for descriptions of the form fields.
3. Click the **What it will contain** tab.
4. Next to **Push Messages**, click the lock icon and select a push message.

> **Note:**
> The push message and notification must be for the same table.

5. If you want this notification to be sent only as a push notification and not as any other type of notification, select **Push Message Only**.
6. Click **Submit**.

   If the notification fails, the user is not notified. If the message fails to send because it exceeds the maximum payload, the instance logs the failure in the System Log.

What to do next

Add the push notification to the Push Defaults Registrations table so that the push notification is listed in the notification preferences for users. Users can then select which notifications they want to receive for the ServiceNow mobile app.

Add a push notification to the Push Default Registrations table

After you create a new push notification, add it to the Push Default Registration table. Push notifications registered in this table are listed and automatically enabled in the notification settings of your ServiceNow mobile app users.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Application**.
2. Select the **ServiceNow Mobile Application**.
3. In the Push Application form for the ServiceNow mobile app, select the **Push Default Registrations** tab and click **New**.
4. In the Push Default Registration New record, select the notification to be registered.
5. Click **Submit**.
   The next time that users log in to their instance, the new push notification is included in their list of notifications.

Push notification setup with a custom push application for ServiceNow mobile apps

If you are using your own custom mobile or push application, you must configure your app for use and set up the push contents.

Setting up a push notification infrastructure that uses a custom push app involves a push admin (also called push app developer or mobile app developer) and system administrator. When push admins create a customized push app, they also configure the app, its push message content (payload generators), and optional attributes, such as push action scripts. The admin creates and updates the push notifications for the custom push app. After users install the custom push app and initially log in to their instance from their mobile device, the system automatically creates a device (channel) for the custom app.

> **Note:**
> These instructions are intended for users who develop their own customized push application. You do not need to configure the ServiceNow mobile push application.
>
> This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

Push notification setup tasks for a custom push app

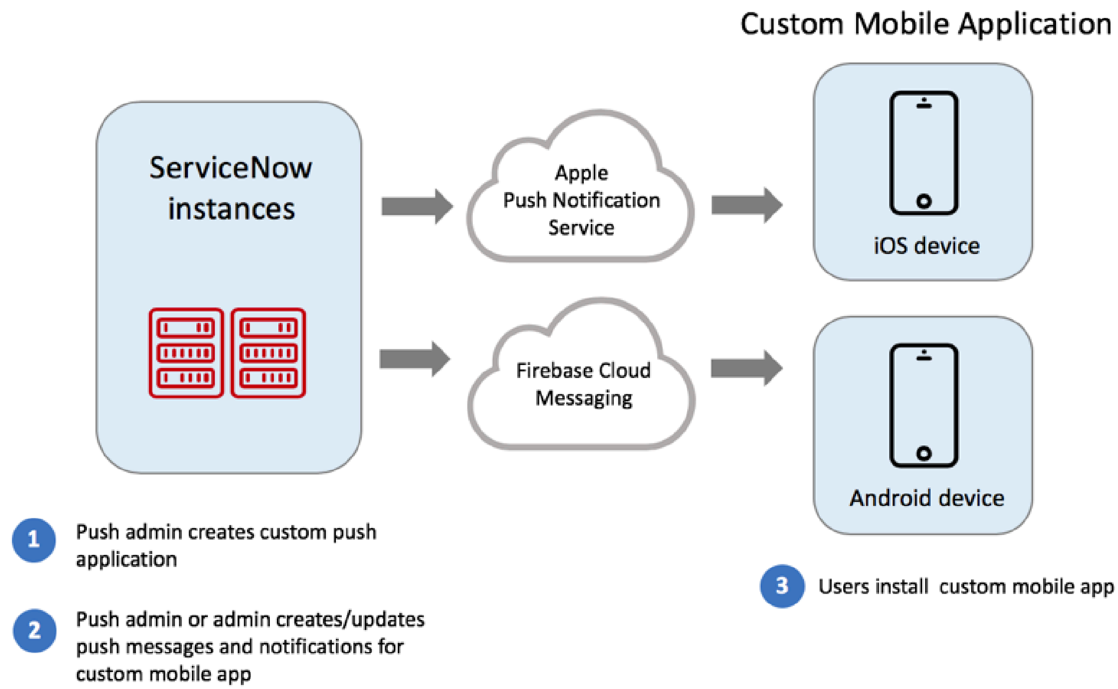

Before you begin

Complete the steps in Activate push notifications.

Assign the push_admin role to your organization's mobile app developer.

What to do — push admin

Configure push notifications for your custom push app. This process differs for iOS and Android devices.

1. (iOS only) Upload a push certificate to your instance

   Upload a push certificate to your instance so that you can use it later to connect your iOS device to the push notification system.

2. Create a push application record for your custom app

   Register your customized mobile application with your instance to receive push notifications for the application. The instance uses this push application record to identify the device + push application combination necessary to determine a push notification recipient.

3. Create push message content

   Create a JSON content payload for different types of push notifications. The content determines how a push notification appears on the push application, and whether the user can send a message in response to the push notification. The push admin can create attribute definitions that specify a default push action script or string, for use in the push message content.

4. (Optional) Create a push message attribute definition

   Push message attribute definitions allow you to create reusable properties for push message content specification.

5. (Optional) Create an attribute value or action for a push message

   You can create attribute values that override the default attribute definitions used in the push message content.

6. (Optional) Create a push action

   A push action is a server-side script that runs when the instance receives a response to an actionable push message.

What to do — admin

Create the push messages and push notifications: The administrator creates the push messages, sets up push notifications for the custom push app, and if desired, the content.

Next steps

Use the Push Installation API to opt users in to receive push notifications. For more information, see Push Installation API.

Upload a push certificate to your instance

Upload a push certificate to your instance so that you can use it later to connect your iOS device to the push notification system.

Before you begin

Using your Apple developer environment, create a push certificate and convert it to a PKCS (.p12) file. For details on generating the .p12 file with the required certificate and private key, see Enable push notifications and Communicate with APNs using a TLS certificate in the Apple documentation.

> **Note:**
> Ensure that your certificate is a push certificate and not a developer certificate. Also, if you're developing an app for testing purposes, ensure that you create a Sandbox certificate. If you're developing an app that you plan to launch in the Apple App Store, ensure that you create a Production certificate.

Role required: push_admin or admin

> **Note:**
> This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

About this task

After you create a PKCS (.p12) file that contains the iPhone developer certificate and the private key used to sign the certificate, you need to upload the file to your instance. The system uses the information in the PKCS file to communicate with the APNs.

Procedure

1. Navigate to **All** > **System Definition** > **Certificates**, and then click **New**.
2. On the form, fill in the following fields:
   - Name: Enter a certificate name.
   - Type: Select **PKCS12 Key Store**.

3. Attach the PKCS (.p12) file to the record.
4. Click **Submit**.

What to do next
Create a push application record for your custom app

Create a push application record for your custom app

Register your customized mobile application with your instance to receive push notifications for the application.

Before you begin
Complete the following:

1. Activate push notifications
2. (iOS only) Upload a push certificate to your instance

Role required: admin or push_admin

Note:
This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

About this task
Push notifications are application-specific: They are sent to one type of mobile application, regardless of how many users have this application installed. If you develop your own mobile application, you must configure it for use by creating a record for it in the Push Application [sys_push_application] table. By default, ServiceNow mobile applications are automatically set up and ready to use.

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Application**, and then click **New**.
2. On the form, fill in the fields.

Field	Description
Display Name	Name of the push application.
Name	Database name of the push application. Enter a name that doesn't include spaces.
Push	Type of push service to enable for the application. Because you're setting up a custom push app, select **Direct**.
Feedback	Type of push feedback to use for the application. Because custom push apps don't receive push feedback, select **None**.
Apple	
Certificate	The X.509 certificate that you created on the Apple notifications portal. The certificate enables a device to talk to the APNs.
Sandbox Certificate	Sandbox certificate that you created on the Apple notifications portal. The certificate enable a device to talk to the APNs.

Field	Description
APNS Bundle Id	Bundle ID included in the certificate that you created on the Apple notifications portal. The system populates this field automatically after you specify a certificate.
Google	
API key	FCM API key you obtained from Google for the Android push notification.

3. Click **Submit**.

What to do next

Create push message content that contains the JSON content in the push notification payload.

In the Push Default Registrations related list, you can select the notifications that you want automatically subscribed to users who use this application. Users are subscribed to only active notifications.

Create push message content

Push message content specifies additional JSON content in the push notification payload that is sent to the push provider.

Before you begin

Complete the following:

1. Activate push notifications
2. (iOS only) Upload a push certificate to your instance
3. Create a push application record for your custom app

You must know how to use JSON with push messages.

Role required: admin or push_admin

> **Note:**
> This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

About this task

Push message content defines the style of push notification that can be sent out for your custom app. You can add custom content, such as a picture, icons, or action buttons for the user to respond to the notification. Use the following variables in the script:

- `current`: properties of the current record.
- `message`: push message sent as the body of the entire push content.
- `attributes`: object of the push message attributes that you define.

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Message Content**.
2. In the Push Notification Message Contents table, click **New**.

3. Fill out the fields on the Push Notifications Message Content form (see table).
4. Click **Submit**.

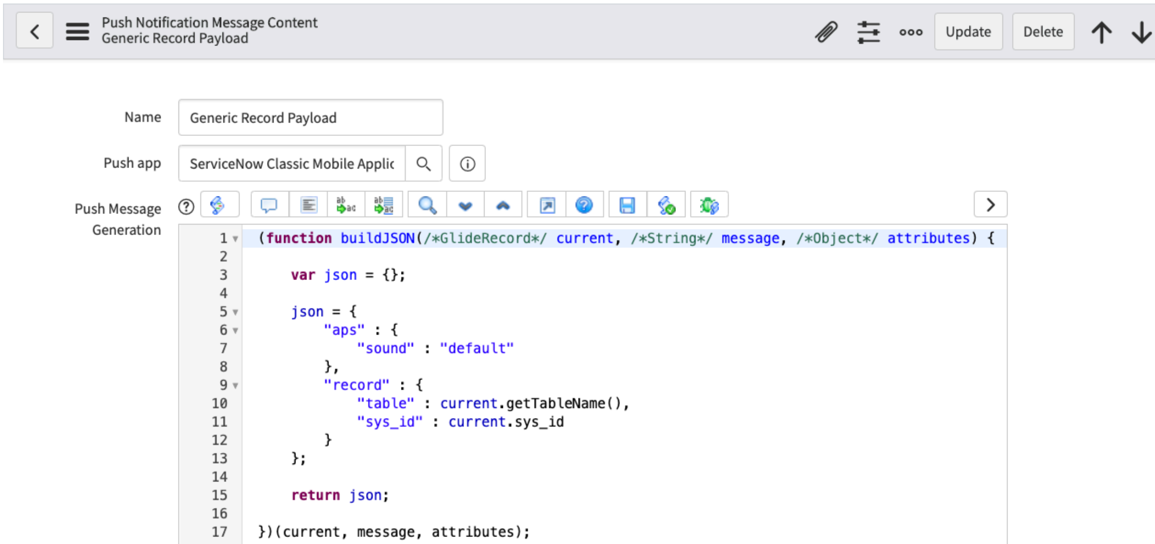

Notification Message form fields

Field	Description
Name	Enter a descriptive name for the message content.
Push app	The push application the content can be used with.
Push Message Generation	Enter a script that determines the message content. See the example scripts.
Related list	
Push Message Attribute Definitions	Select the attributes that apply to this notification. Attributes can be a value or an action. These attributes are used as default values for the content items you create in the content script. However, any attributes you create with the push message can override these attributes. See Create a push message attribute definition for information on creating attributes.

Example:
The following is an example of a content record that creates a two-button layout, one to approve something, such as a change request, and one to decline it.

```
var json = {
  "table" : current.getTableName(),
  "sys_id" : current.sys_id,
  "template" : {
    "type": "2 button",
    "button1" : {
      "title" : "Approve",
      "action" : attributes.button_action,
      "parameters" : {
        "response" : "approve"
      }
    },
```

```
    "button2" : {
      "title" : "Decline",
      "action" : attributes.button_action,
      "parameters" : { "response" : "decline"
      }
    }
  }
};
json;
```

What to do next

Define a push message attribute to specify a default push action script or string that you can use in the push message content.

Create a push message attribute definition

Push message attribute definitions allow you to create reusable properties for push message content specification.

Before you begin

1. Activate push notifications
2. (iOS only) Upload a push certificate to your instance
3. Create a push application record for your custom app
4. Create push message content

Role required: admin or push_admin

> **Note:**
> This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

About this task

Use an attribute definition to specify a default push action script or string that you can then use in the push message content.

You can override these attributes by creating any push message attribute values.

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Message Content**.
2. Select a content record or script record.
3. In the Push Message Attribute Definitions related list, click **New**.
4. Fill out the fields on the form (see table).
5. Click **Submit**.

Push Message Attribute Definition form fields

Field	Description
Attribute Name	Enter a descriptive name for the layout.
Attribute Type	Select the type of attribute:

Field	Description
	○ **Action**: An action to take on the instance, as defined by a script. ○ **String**: An arbitrary string to send as part of the message content. The string can specify items like a button label in the message.
Default Script	Select the mobile action script that tells the instance what to do when it receives a response from the push notification. This option appears if you select **Action** for the **Type**.
Default Value	Specify an arbitrary string value to be placed in the attribute that is used by the message content. For example, the string could specify a button label. This option appears if you select **String** for the **Type**.

What to do next

If you want to override default attribute definitions, see Create an attribute value or action for a push message.

Create an attribute value or action for a push message

Attribute values are associated with push messages to provide a way for the message to include certain types of information in the push notification.

Before you begin

Complete the following:

1. Activate push notifications
2. (iOS only) Upload a push certificate to your instance
3. Create a push application record for your custom app
4. Create push message content
5. Create a push message attribute definition

Role required: admin or push_admin

Note:
This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

About this task

These push message attribute values override values that you set in an attribute definition on the Push Content form.

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Messages**.
2. Select a message.
3. In the Push Message Attribute Values related list, click **New**.
4. Fill out the fields on the form (see table).
5. Click **Submit**.

Push Message Attribute Values form fields

Field	Description
Attribute	Select a push message attribute definition.
Value/ Action	Enter the value for the attribute or select the push action. This field changes to Value or Action depending on the type of attribute you select in the Attribute field.

What to do next

(Optional) Create a push action to perform an action on the instance.

Create a push action

A push action is a server-side script that runs when the instance receives a response to an actionable push message.

Before you begin

Complete the following:

1. Activate push notifications
2. (iOS only) Upload a push certificate to your instance
3. Create a push application record for your custom app
4. Create push message content
5. Create a push message attribute definition
6. Create an attribute value or action for a push message

Role required: admin or push_admin

Note:
This process describes configuration used in the ServiceNow mobile app. Push Notification configuration for the current ServiceNow mobile UI can be found at Mobile push notifications

About this task

Create a push action to perform an action on the instance. For example, you might have an actionable push message that lets the user approve a change request. The action that handles the response should update the **Approval** field on the relevant Change Request record.

You can use global variables or, optionally, current variables and parameters passed through the JSON content.

Procedure

1. Navigate to **All** > **System Notification** > **Push** > **Push Action**.
2. Fill in the form fields (see table).
3. Click **Submit**.

Push Notification Action form fields

Field	Description
Name	Enter a descriptive name for the action.
Script	Enter the script.

Requeue failed push notification messages

Push notification delivery might fail for various reasons. You can view which messages failed and requeue them to be sent out if necessary.

Before you begin
Role required: admin

About this task

Note:
There is no way for the instance to guarantee or confirm push message delivery. For more information on how Apple handles push notifications, see the Local and Remote Notification Programming Guide for Apple developers.

Messages can also fail to send for several reasons, such as a large message queue or other issue with the instance. You can use a system property to control how long the instance queues a push notification after it is triggered.

Push notification message payload size is limited by provider:

- Apple iOS: 2,048 bytes
- Google: 4,096 bytes

Your instance does not send push messages that exceed this limitation. System logs save any failed messages. For details on Apple payload limitations, see the Local and Remote Notification Programming Guide for Apple developers.

Procedure

1. Navigate to **All** > **System Logs** > **Push Notifications**.
   By default, the messages that were created today appear. You can change the filter if necessary. The Type column can have these values:
   - **failure**: The message could not be sent.
   - **pending**: The message is queued for processing.
   - **success**: The message was successfully sent, although not necessarily received by the mobile device.
2. Put any failed messages back into the queue to be resent by selecting the check boxes next to the failed messages, and then selecting **Re-queue failed push notifications** from the choice list.
3. If some push notifications continue to fail, consider increasing the value in the *`glide.push.notification.ttl_seconds`* property.
   See Push notification properties for more information.

Push notification retention

You can archive and eventually destroy push notifications that you no longer need or if your Push Notification table is excessively large.

Push notification retention is available starting with the Jakarta release.

Push notification archive and destruction plugins

The push notification archiving and destruction feature uses the Data Archiving and Push Retention plugins. The Data Archiving plugin must be active to archive and destroy push

notification records. The Push Retention plugin provides a set of rules that specify when the system archives and destroys push notification records.

In new instances, the Push Retention plugin and associated archive and destroy rules are active by default. On upgraded instances, you must manually activate both the plugin and the archive and destroy rules. Be sure to review and approve the archive and destroy rules before activating them.

If your instance already has a process for managing push notification records, you do not need to activate the Push Retention plugin. If you want to replace your current process with Push Retention, be sure to deactivate the current process before activating the archive and destroy rules.

Archiving and destroying push notification records

Archiving means moving records from the Push Notification [sys_push_notification] table to the Push Notification Archive [ar_sys_push notification] table when they exceed the archive rule time limit. Destroying means deleting records in the Push Notification Archive table when they exceed the destroy rule time limit.

Default archive and destroy rules

Push retention provides the following push archive rules:

- **Push Notification - Over a year old:** archives push notification records that were created more than 365 days prior to the current date.
- **Push Notification Archive - Over a year old:**destroys push notification records that have been archived for more than 365 days prior to the current date.

With these default settings, your messages are kept on the instance for a total of two years: one year in the Push Notification table, and one year in the Push Notification Archive table. At the end of the period, the system deletes the expired notification records from the Push Notification Archive table.

Note:
By default these rules are active on new instances and inactive on upgrades. The system runs archive and destroy rules when you activate them.

Compatibility with other record management implementations

If you are already using another method for managing push notification records, such as table cleaners, you do not have to use the Push Retention feature. To prevent unexpected record deletion, avoid using multiple push notification management processes on the same instance at the same time.

Note:
For assistance replacing your existing record management implementation with push notification retention, contact your professional services or sales representative.

Activate the Push Retention plugin

The Push Retention plugin (com.glide.push_retention) provides the retention policy for push notifications, so that you can specify when the system archives and destroys push notification records.

Before you begin
Role required: admin

Procedure

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

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

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

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

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

Related topics

List of plugins (Washington DC)

Archive push notifications manually

You can manually archive push notifications on demand instead of waiting for the instance to archive them based on a scheduled job.

Before you begin

Role required: admin

About this task

Push retention provides these push notification archive rules:

- **Push Notification - Over a year old:** archives push notification records that were created more than 365 days prior to the current date.
- **Push Notification Archive - Over a year old:**destroys push notification records that have been archived for more than 365 days prior to the current date.

You can manually archive push notification messages that meet the default archive and destroy rules or any additional archive rules that you create.

Procedure

1. Navigate to **All** > **System Archiving** > **Archive Rules**.
2. Open the push notification archiving rule, for example Push Notification - Over a year old
3. Click **Recalculate Estimate** to see how many records in the Push Notification [sys_push_notification] table are going to be archived.
   The estimate appears in the **Record estimate** field.
4. Click **Run Archive Now**.

Notifications in messaging applications

Enable users to receive their ServiceNow platform notifications in Slack or Microsoft Teams messaging applications. Use the Now Actions app to configure the messaging apps for your instance.

Admin setup

To enable ServiceNow notifications within Slack or Microsoft Teams, you must be an administrator for both ServiceNow and the third-party application.

1. **Request Integration Hub**

   The Slack and Microsoft Teams spokes require a Starter subscription to Integration Hub. To activate Integration Hub and view available subscription types, see Request an Integration Hub plugin.

2. **Request the Messaging Notification plugin**

   To activate this feature, request the Messaging Notification plugin (com.glide.notification.messaging) through the HI Customer Service system. This plugin activates related plugins if they are not already active.

3. **Install the Now Actions messaging app for Slack or Microsoft Teams**

   As an administrator for both ServiceNow and the third-party application, find and install the Now Actions app. Install the application in ServiceNow to associate the app with your instance.

   Note:
   The primary owner of a Slack workspace or organization owns and manages the application. To avoid losing admin access to the workspace or organization, transfer primary ownership before any administration changes occur. If you transfer primary ownership, update the JSON payload for the Slack installation in the Notification Integration page. For more information about editing the Slack installation, see Install the Now Actions messaging app.

   If the Now Actions app is uninstalled, the workspace or team displays as inactive in the Notification Integration page. To view the Notification Integration page, navigate to **System Notification** > **Messaging** > **Messaging Integration Configuration**.

4. **Configure message content**

   Create a message in the Message Contents [messaging_content] table that you can use in multiple notifications. Messages can be informative, or they can request action from an individual Slack or Teams user. For example, the Approval Request message includes Approve and Reject buttons to take the associated action in ServiceNow. Only users that have linked their Slack or Teams account with their ServiceNow account can take action in ServiceNow from the third-party application.

5. **Create a messaging notification**

   Notifications define when a message is sent, who it is sent to, and what it contains. To send a notification to Slack or Teams, add a messaging content record to the notification and configure the notification messaging channel. Notifications that require action must be sent to individual users instead of groups.

Note:
Message Channel Sync Job can sync more than 1000 channels using pagination.

User setup

To enable individual ServiceNow notifications in Slack or Microsoft Teams, and to take action on notifications, configure your system settings and link your user accounts.

Link your ServiceNow user account to your Slack or Microsoft Teams account

Link your ServiceNow account to your Slack or Teams account to take action on ServiceNow notifications from within Slack or Teams. If you do not link accounts, you can still receive notifications within group channels. However, you cannot to approve ServiceNow requests or take action on notifications from within Slack or Teams.

Configure system settings

Each user can enable ServiceNow notifications in a third-party application by configuring the communication channel in their system settings. Opt in to receive approval requests and other actionable notifications in third-party channels.

Notification channels

In ServiceNow, a notification channel is a method or device for receiving notifications. For example, in the base system you can elect to receive notifications through email or SMS channels. If messaging applications are enabled on your instance, you can enable notifications to Slack and Teams using Notification Channels in your system settings.

Actions and events

The Message Actions [messaging_observer_handler] table processes communication channel events, such as adding or deleting a Slack channel, to keep the instance and the third-party application in sync. Only developers creating Slack or Teams applications should add actions to this table. Register new actions with the instance through the Message Event Registry [messaging_observable] table. To remove a base system action, clear the active flag instead of deleting the record.

Logging

The Inbound and Outbound message logs contain information on actions, events, and activities between ServiceNow and the third-party application. Use the logs to identify the content, state, and origin of communications to and from your instance.

Request the Messaging Notification plugin

To activate this feature, request the Messaging Notification plugin (com.glide.notification.messaging) through the HI Customer Service system. This plugin activates related plugins if they are not already active.

Before you begin

Role required: admin

About this task

There are two ways to request a plugin:

- Access the Now Support Service Catalog directly by selecting **All** > **Service Catalog** > **Activate Plugin** on Now Support.
- Access the Now Support Service Catalog through the All Applications page on your instance by following these steps.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. On the All Applications page, select **Request Plugin** to open the **Activate Plugin** form on Now Support.

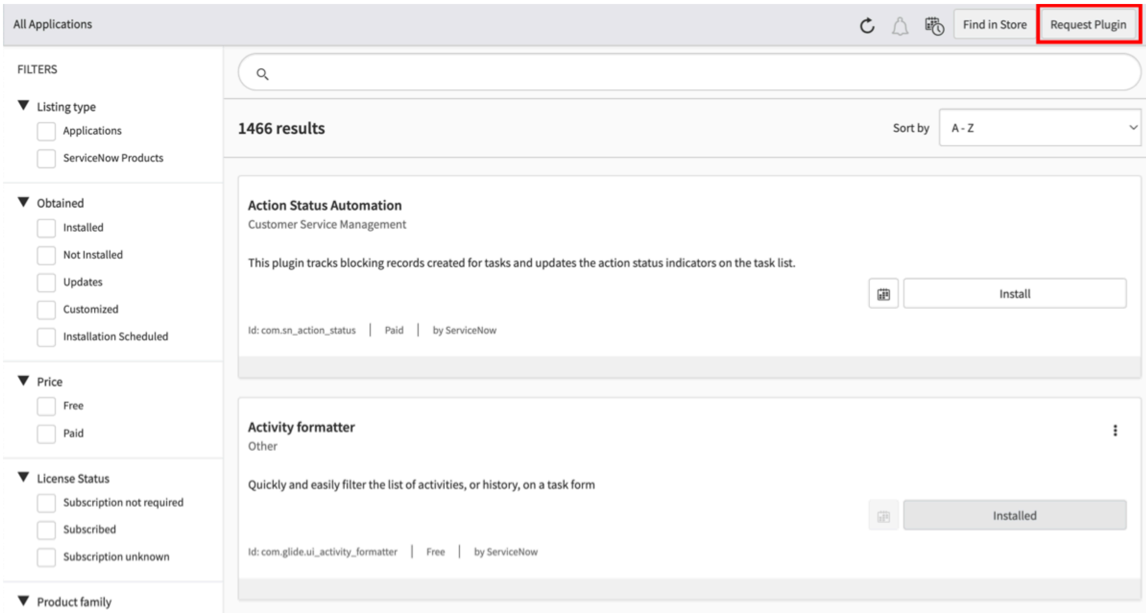

3. On Now Support, select the link to access the Now Support Service Portal Service Catalog.

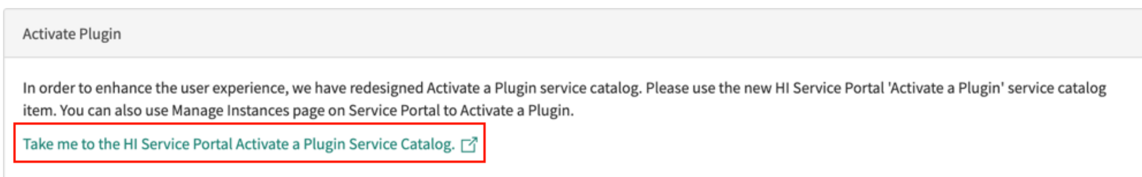

4. Select your instance.
5. Select **Actions > Activate Plugin**.
6. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Instance on which to activate the plugin.
Which plugin would you like to activate	Name of the plugin to activate. **Note:** If the system does not list the plugin you want or if you are activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	The date and time to activate the plugin.

Field	Description
	Note: Plugins are activated in two batches, once in the morning and once in the evening, on every business day in the US Pacific time zone. If the plugin must be activated at a specific time, enter the request in the **Reason/Comments** field.

Example

For example, see the following form to activate the CSM Workspace plugin on an instance named My Instance.

Activate Plugin form

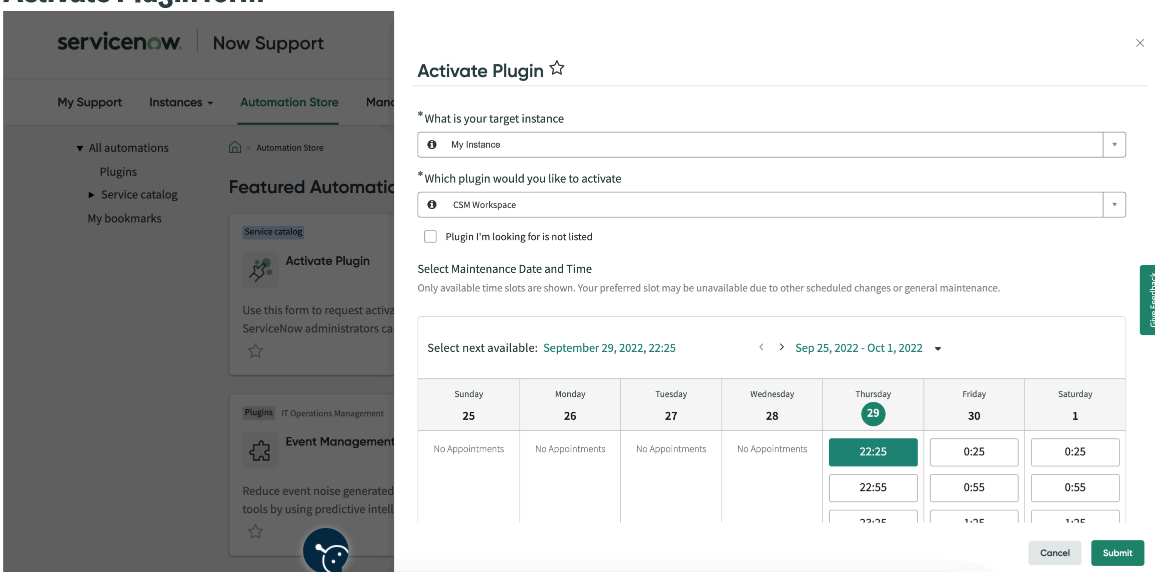

7. Select **Submit**.
   For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

Related topics

List of plugins (Washington DC)

Install the Now Actions messaging app

As an administrator for both ServiceNow and the third-party application, find and install the Now Actions app. Install the application in ServiceNow to associate the app with your instance.

Before you begin

To activate this feature, request the Messaging Notification plugin (com.glide.notification.messaging) through the HI Customer Service system. This plugin activates related plugins if they are not already active. For more information, see Request the Messaging Notification plugin.

Roles required:

- Admin or messaging_admin in ServiceNow
- Administrator for Slack or Teams

Install the Now Actions messaging app for Slack

Install the Now Actions messaging app from the Slack App Directory and associate the app with your instance.

Before you begin

Role required: admin.

Procedure

1. Install the Now Actions app from the Slack App Directory.

   a. Find and install the Now Actions app in the Slack App Directory.

   b. Click Authorize when prompted.

   c. **Optional:** Enter a customized name, app description, and icon.

   d. Configure your Now Actions web address.
   This is the base URL for your instance.

   Example
   For example, `https://<instance_name>.service-now.com/`.

   e. Copy the JSON configuration found in the **Add your Slack credentials to your app** section for use in later steps.

2. Associate your Slack app with your instance.

   a. Navigate to **System Notification** → **Messaging Integration Configuration**.

   b. Click the **Install** button next to Slack.

   c. In the JSON Configuration pop-up, paste in the JSON configuration you copied in earlier steps.

3. Check for installation verification.
   The Slack workspace should appear below the Slack heading on the Notification Integration page.

   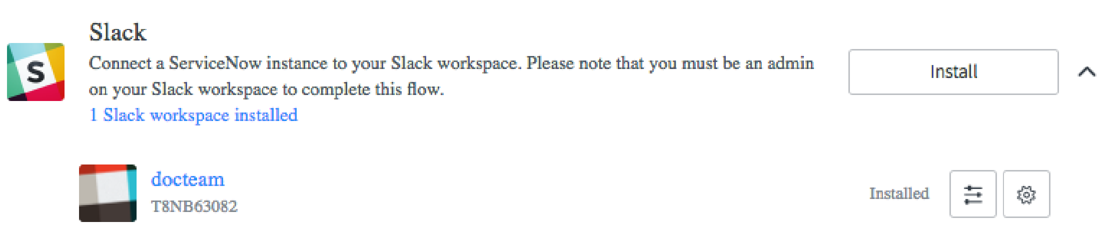

   During installation, the instance synchronizes with all public Slack channels in the workspace. New public channels added after installation will automatically synchronize with the instance. You can browse to **System Notification** > **Messaging Channels** on your instance to verify that the new channel appears on the instance.

What to do next

To update the app after installation, use the **Configure JSON Payload** and **Settings on Slack** icons on the Notification Integration page. You can view the JSON payload, view app credentials, change app appearance, or delete the app.

Install the Now Actions messaging app for Microsoft Teams

Associate the Now Actions messaging app with your instance.

Before you begin

Role required: admin.

Procedure

1. Associate your Microsoft Teams app with your instance.

   a. Navigate to **System Notification** → **Messaging Integration Configuration**.

   b. Click the **Install** button next to Microsoft Teams.

   c. When a pop-up message appears to confirm redirection to Microsoft Teams to verify your identity, click **OK**.

   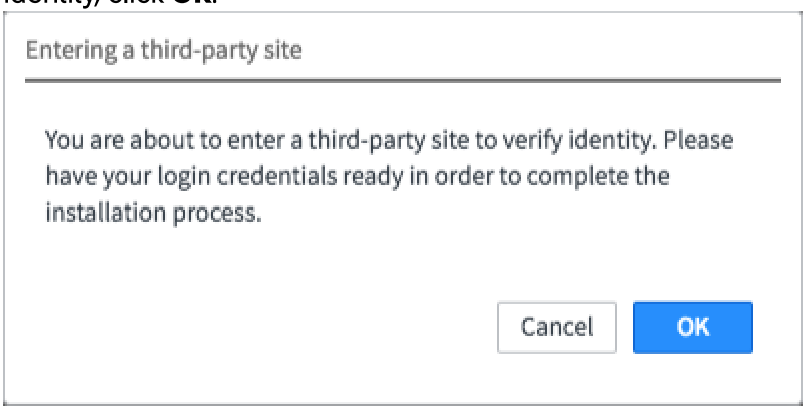

   d. When prompted, log in to Microsoft Teams with your Teams admin account.

   e. Click *Accept* to accept the permissions for the app.

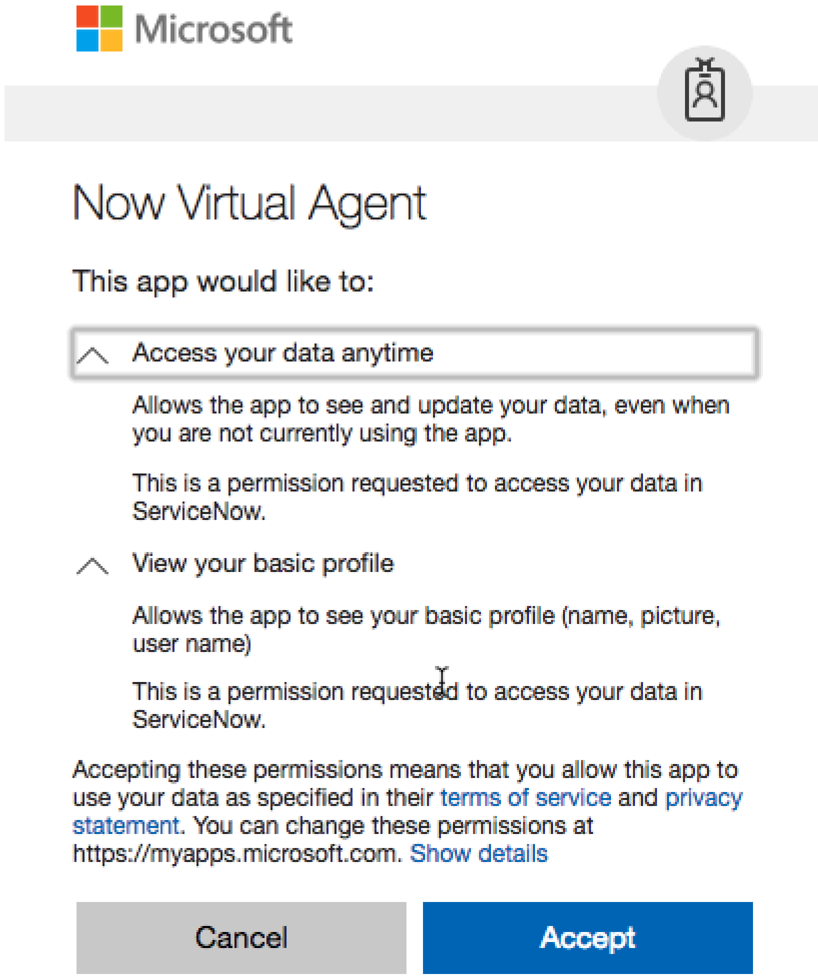

f. If the selected workspace has already been assigned, either click **Override** to change the workspace, or click **Cancel** to leave the current assignment in place.
After authentication, installation will start in the background. Once installation completes, a message appears confirming the installation.

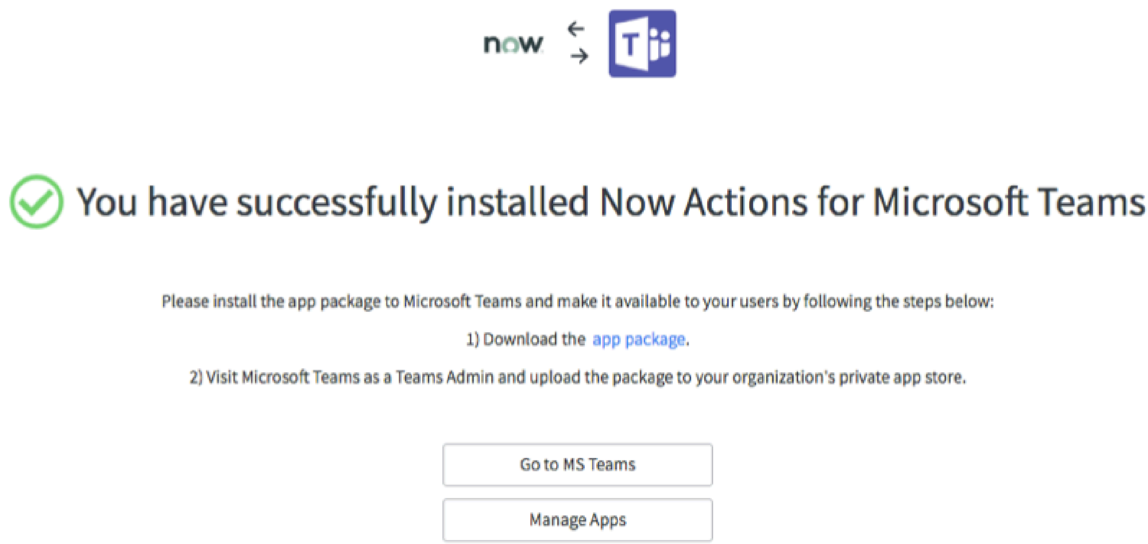

g. Click the **app package** link to access a knowledge base article KB0690098, which contains the application package for Microsoft Teams.

h. Follow the steps in the article to make this package available for your Microsoft Teams users.

2. Check for installation verification.

   During installation, the instance synchronizes with all public Teams channels. To add channels created after the install, open Microsoft Teams and @mention the app from the new channel. You can browse to **System Notification** > **Messaging Channels** on your instance to verify that the new channel appears on the instance.

   The new team should appear below the Teams heading on the Notification Integration page.

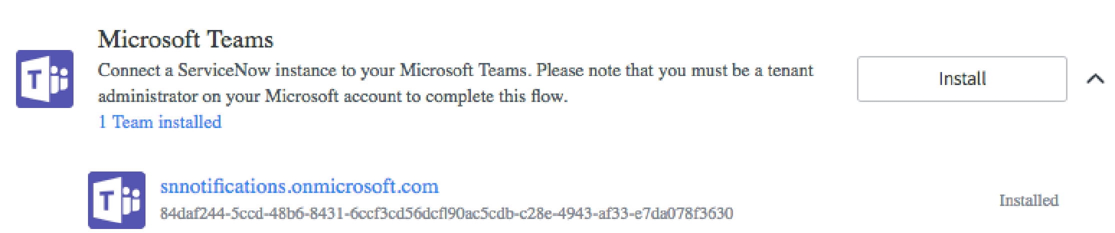

What to do next

To receive a ServiceNow notification in Microsoft Teams, create a message and a notification. See Configure message content and Create a messaging notification.

Configure message content

Create a message in the Message Contents [messaging_content] table that you can use in multiple notifications. Messages can be informative, or they can request action from an individual Slack or Teams user. For example, the Approval Request message includes Approve and Reject buttons to take the associated action in ServiceNow.

Before you begin

Role required: admin or messaging_admin

About this task

To enable users to take actions in Slack or Teams, use the Buttons message type and define scripted behavior. Only users that have linked their Slack or Teams account with their ServiceNow account can take action in ServiceNow from the third-party application.

Procedure

1. Navigate to **All** > **System Notification** > **Messaging Content**.
2. Open an existing notification record or click **New** to create a new record.
3. Fill in or modify the fields as appropriate.

Messaging content fields

Field	Description
Name	Enter a unique name for the message.
Type	Select a type of message. ◦ **Simple**: Send an informative message to an entire Team or Slack channel. ◦ **Buttons**: Send a message to an individual user in Slack or Teams and define the desired behavior in the **Script** field. Associate buttons in the message with actions in ServiceNow. Use the Approval Request as an example.
Target table	To use record content in your message, select the appropriate table. **Note:** Notifications running on the Task [task] table are not supported. Instead, use a table extended from task.
Fields	Select which fields from the **Target table** appear in a simple message. Changing the order of the fields changes the display order in the notification.
Message heading	Enter text to appear at the top of the message.
Message	Enter the content of the message. Use notification variables to include dynamic information in your message. For example, use `Incident ${number}` to include the incident number in the message. For available variables, see Notification variables.
Buttons	Select or create buttons to use on the form. After selecting or creating a button, you must define button behavior in the **Script** field. Only visible when Buttons is selected in the **Type** field
Script	This field is only visible when Buttons is selected in the **Type** field. Enter a script to define what happens in your instance when a user selects a button in Slack or Microsoft Teams. Use the `actions.get` method to return the name of the button selected by the user. `if (actions.get('button') == 'Approve') {` `  target.state = 'approved';`

Field	Description
	`target.update();` `}` In this example, the script conditionally executes when the "Approve" button is selected. Use the method `sn_notification.Messaging.send` to send information to the messaging window. `var content = new GlideRecord("messaging_content"); content.get("1bb10839572213007f004758ef94f9c4")); sn_notification.Messaging.send(application, userId, content, target);` In this example, the script stores a Message Content record in the `content` variable. It then posts a message from that Message Content record.
Application	Select the application scope for this notification.

4. Click **Submit**

Example:
This example shows the output of a message in Microsoft Teams. This message is of the buttons type, which shows buttons and performs scripted actions based on the user selection. In this case, the script uses the `sn_notification.Messaging.send`method to display the change in state for the record.

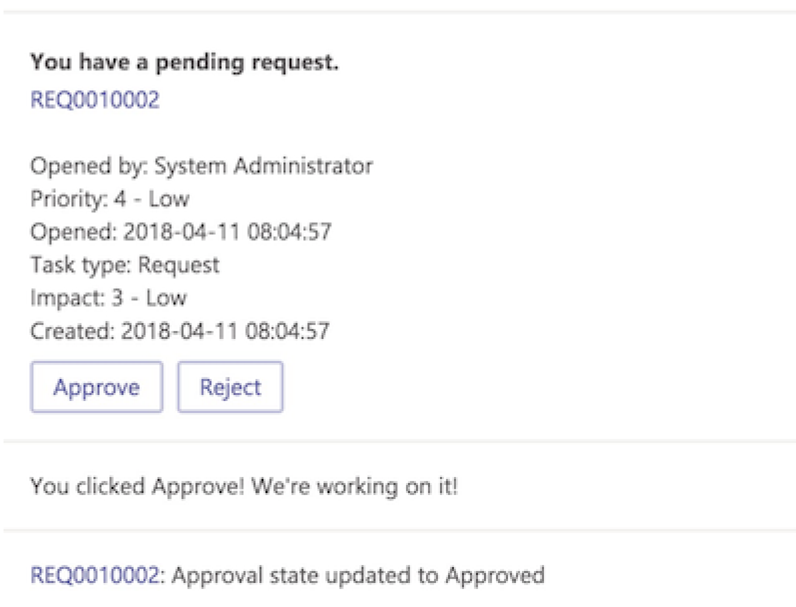

What to do next
Use the message in a notification. For more information about creating notifications, see Create a messaging notification.

Create a messaging notification

Notifications define when a message is sent, who it is sent to, and what it contains. To send a notification to Slack or Teams, add a messaging content record to the notification and configure the notification messaging channel. Notifications that require action must be sent to individual users instead of groups.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Messaging** > **Messaging Notifications**.
2. Click **New**.
3. Fill in or modify the fields as appropriate.

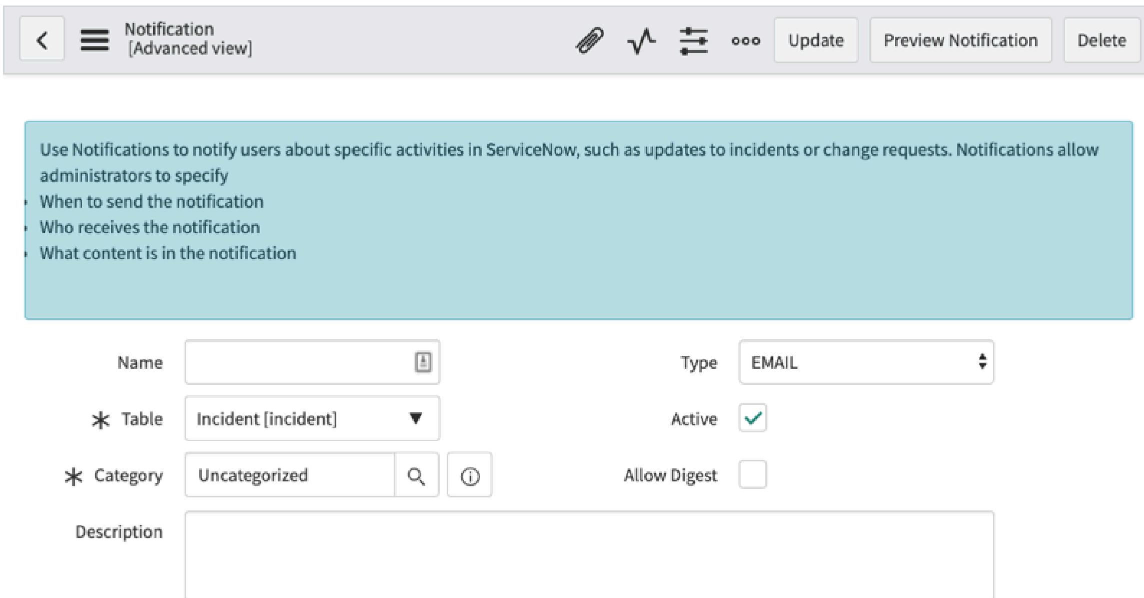

Messaging notification fields

Field	Description
Name	Enter a unique name for the notification. Use descriptive names to ensure that your users are able to easily distinguish between notifications.
Table	Select a table to trigger the message. Must be a table or database view within the same application scope. **Note:** Notifications running on the task [task] table are not supported. Instead, use a table extended from task.
Category	Select a category for the notification. A category groups related notifications in the system settings. **Note:** Do not leave the category as **Uncategorised**, as users may not be able to find the notification in their list of notifications.
Description	Enter a description for this notification.

Note:
Notification records are shared between messaging, email, and push notifications. Some fields on the form are not relevant to messaging notification records. Leave fields not described here at their default values.

4. Click**Advanced view** under **Related links** to switch the advanced view.

Note:
Switching to the Advanced view saves the current record.

5. Open the **When to send** tab and complete the form.

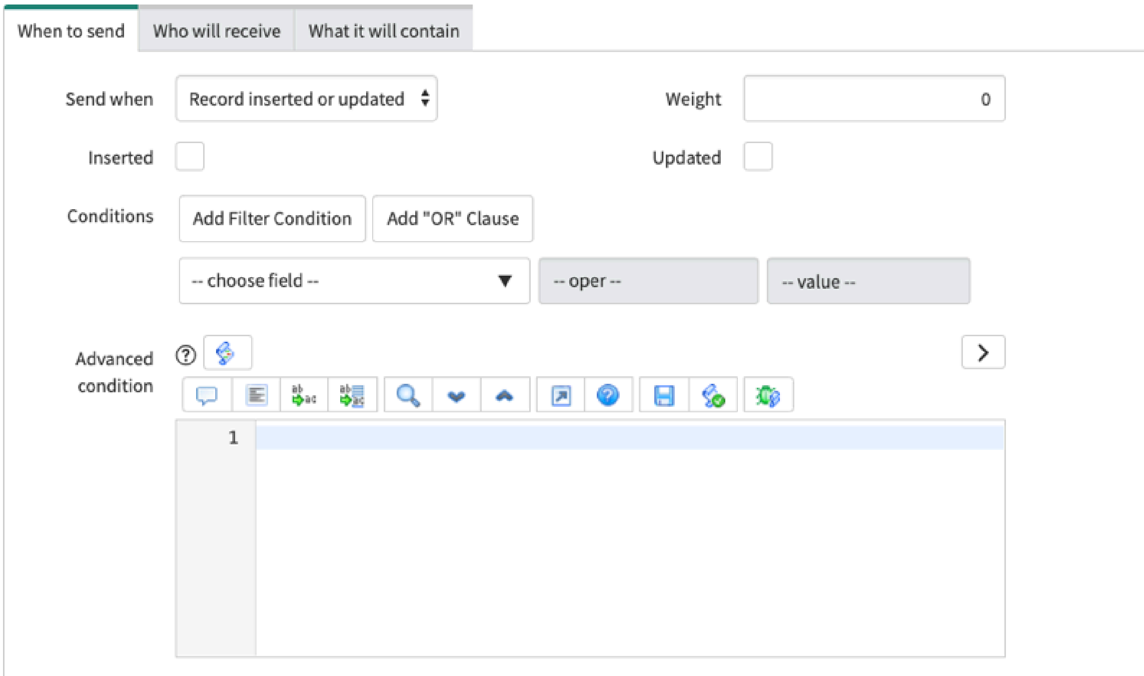

Note:
If the same trigger generates multiple notifications, the system only sends one notification. The system considers all other notifications, even if they have a different subject and body, as duplicates. The Ignore Duplicates business rule controls this functionality.

Field	Description
Send when	Select under what condition the notification is sent: ○ When a record is inserted or updated ○ When a particular event is fired ○ When Notification step in Flow Designer
Weight	Set a numerical value for the notification priority relative to other notifications with the same target table and recipients. The system only sends the notification with the highest weight. The default value 0 causes the system to always send the notification (assuming the conditions are met).
Conditions	Use the condition builder to select the conditions under which this notification is sent. For example, select **Priority** > **greater than** > **3 - Moderate** to send the notification only for High and Critical priority incidents.
Inserted	Select the check box to enable messaging notification when a record is inserted. This field appears when you set the **Send when** field to **Record inserted or updated**.

Field	Description
Updated	Select the check box to enable messaging notification when a record is updated. This field appears when you set the **Send when** field to **Record inserted or updated**.
Event name	Select the event that triggers this notification. This field appears when you set the **Send when** field to **Event is fired**.
Advanced condition	Create a script to define more filter conditions, like sending a notification based on the current messaging record, changing field values, or changing system properties. The advanced condition script must return **true** or set a global answer variable to **true** to send the notification. The advanced condition script uses the following global variables: ◦ *current*: Contains the current record from the table to which the notification is linked. ◦ *event*: Contains the event that triggered the notification. **Note:** The **Advanced condition** field is evaluated in addition to other conditions you set on the notification. Both the condition and advanced condition must evaluate to true to send the notification.

6. Open the **Who will receive** tab and fill in or modify the following fields.

Field	Description
Messaging Channels	Select the messaging channels to receive the messaging notification. Only complete this field if creating a simple notification to send to a channel. Simple notifications use a Messaging Content record of Type **Simple**.
Users	Select the users you want to receive the messaging notification. Only complete this field if creating a notification that requires user action. Notifications that require user action use a Messaging Content record of Type **Buttons**.
Users/ groups in fields	Select users or groups from reference fields. For example, if a notification uses the Incident [incident] table, you can select users or groups from incident fields like **Opened by** and **Assignment group**. **Note:** You can dot-walk to values in reference fields by clicking the plus sign in the field selector and then selecting the related field. If you address the notification to a user with an inactive record in the User [sys_user] table, the system does not send the notification to that user.
Groups	Select the groups you want to receive the messaging notification. You can search for groups with the reference lookup icon or by manually entering the group name. This list of groups is static.
Send to event creator	Select this check box to send the notification to the person who performed the action that started the notification process if the person is also a recipient. If the event creator is not specified in one of the recipient fields, the event creator does not receive a notification, regardless of the setting in this field.

Field	Description
	For new notifications, this option is selected by default.
Event parm 1 contains recipient	Select this check box if the event parameter 1 contains one or more notification recipients (in a comma separated-list). This field is visible only when the **Send when** field is set to **Event is fired**.
Event parm 2 contains recipient	Select this check box if the event parameter 2 contains one or more notification recipients (in a comma-separated list). This field is visible only when the **Send when** field is set to **Event is fired**.

7. Open the **What it will contain** tab and fill in or modify the following fields.

Field	Description
Messaging content	Select a messaging content record to be used in this notification.

8. Click **Update**.

Link your ServiceNow user account to your Slack or MS Teams account for Now Actions

Link your ServiceNow account to your Slack or Teams account to take action on ServiceNow notifications from within Slack or Teams. If you do not link accounts, you can still receive notifications within group channels. However, you cannot to approve ServiceNow requests or take action on notifications from within Slack or Teams.

Before you begin

Role required: none

Procedure

1. Open the Slack or Teams application
2. Navigate to the Now Actions bot, which can be found in the left panel of Slack under Apps, and in the left panel of Teams in the **Chat** tab.
3. Start a conversation with the Now Actions bot.
   The bot then presents a link to authenticate.

Authentication message

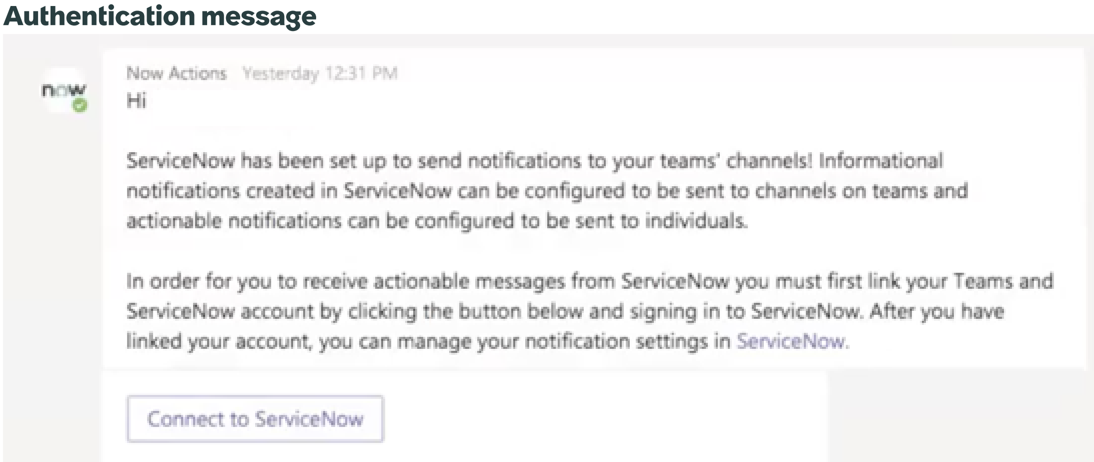

4. Click the **Connect to ServiceNow** link.
   A new browser tab opens and prompts the user to Deny or Confirm the connect between your Slack or Teams account and your ServiceNow account. This link will expire five minutes after it is displayed.
5. Click **Confirm** to link the accounts.

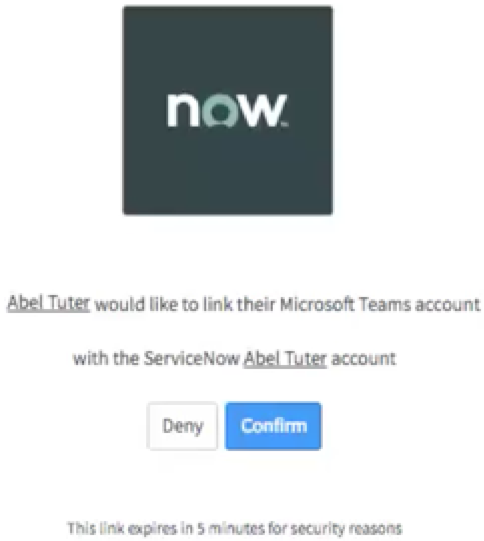

6. After confirmation, you are directed to the instance user record.
   A confirmation message is displayed at the top of the screen.

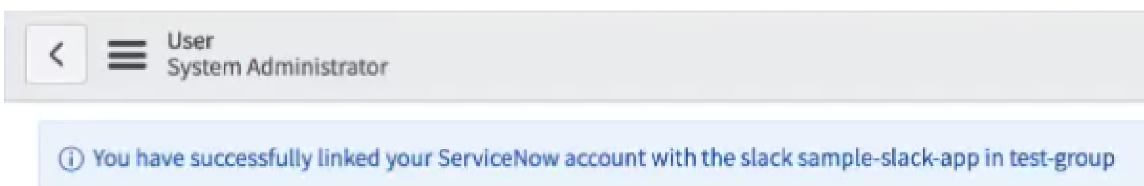

Configure system settings

Configure the system settings of your instance to set the UI experience that you want.

Before you begin

Role required: admin.

About this task

You can customize your general settings, theme, accessibility, lists, forms, notifications, and developer tools. You can also switch your instance UI to a previous version.

Procedure

1. Navigate to **User Menu** > **Preferences** > **Notifications**.
2. Select the **General** tab and change your system settings.

Tab	System Setting
General: General settings for your instance.	**Compact user interface**: Option to display a compact user interface.

servicenow

<table>
<tr><th>Tab</th><th colspan="2">System Setting</th></tr>
<tr><td rowspan="6"></td><td colspan="2">Compact list date/time: Option to display a compact list date and time.</td></tr>
<tr><td colspan="2">Keyboard shortcuts enabled: Option to enable keyboard shortcuts.</td></tr>
<tr><td colspan="2">Home: Option to select Homepages or Dashboards for your instance home.

Important:

The functionality found in homepages, arranging information from your instance to tell a story about your data, is found in dashboards on new instances. On upgraded instances with Next Experience enabled, users can view existing homepages if they have a direct URL, but they can't create or edit them. Responsive dashboards and Analytics Center dashboards take over homepage functionality.

Use the Homepage deprecation help tool to convert the homepages on your instance to responsive dashboards.

For more information, see:
◦ Dashboards in the Analytics Center.
◦ Working with responsive dashboards.</td></tr>
<tr><td colspan="2">Date/Time: Select Calendar, Time Ago, or Both.</td></tr>
<tr><td colspan="2">Time zone: Time zone displayed in your instance.</td></tr>
<tr><td colspan="2">Printer friendly version: Printer-friendly version of a notification.</td></tr>
<tr><td>Theme: Color scheme for your instance.</td><td colspan="2">Option to select a color scheme for your instance.</td></tr>
<tr><td rowspan="4">Accessibility: Accessibility features for your instance.</td><td colspan="2">Data visualization patterns enabled: Option to enable data visualization patterns.</td></tr>
<tr><td rowspan="3">Accessibility enabled: Option to enable accessibility.

Enable Accessibility to toggle these features.</td><td>Show tooltips on forms: Option to show tooltips on forms. Tooltips display when you point to a UI element.</td></tr>
<tr><td>Show Date & Time format in field label: Option to show date and time format in field labels.</td></tr>
<tr><td>Show Record Preview tooltips on Slushbucket items when available: Option to show record preview tooltips on slushbucket items when available. Tooltips</td></tr>
</table>

Tab	System Setting
	display when you point to a UI element.
Lists: List preferences for your instance.	**Wrap longer text in list columns**: Option to wrap text in list columns.
Forms: Form preferences for your instance.	**Tabbed forms**: Option to enable tabbed forms.
	Related list loading: Select **With the Form**, **After Form Loads**, or **On-demand**.
Notifications: Notification settings for your instance.	**Allow Notifications**: Option to enable notifications. To receive notifications for Simple type messages, enable the notification channel. You don't need to toggle individual notifications. For more information on message types, see Configure message content. **Note:** You can only enable notifications in notification channels that you have permission to access.
	Create Channel: Select to configure a channel.
	Primary email ○ Option to receive email notifications in your primary email. ○ Configure your primary email by clicking the configure icon (>).
	Create Personal Notifications: Select to configure a personal notification. Select a parent notification from the list and click the configure icon (>). Child notifications are displayed for each entry. Toggle a notification.
Developer: Developer tools for your instance.	**Application**: Application of your instance. Set if other than **Global**.
	Show application picker in header: Option to display the application picker in your instance header.
	Update Set: Update set for your instance. Select an update set if other than**Global**.
	Show update set picker in header: Option to display the update set picker in your instance header.
	JavaScript Log and Field Watcher: Option to display the JavaScript log and field watcher.
	Automated Test Framework Page Inspector: Option to display the automated test framework page inspector.

Unlink your ServiceNow user account from your Slack or MS Teams account for Now Actions

Unlink your ServiceNow account from your Slack or MS Teams account to stop receiving actionable ServiceNow notifications in Slack or Teams.

Before you begin

Role required: none

About this task

If you unlink your account to a messaging app for Now Actions, you must also disable the corresponding messaging channel in your notification preferences.

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.
2. Click the **View Linked Accounts** related link.
3. In the Linked Accounts page, check the selection box for the messaging integration to be unlinked.
4. Select **Actions on selected rows...**, then click **Unlink account**.
   Your ServiceNow account is unlinked from the Now Actions messaging integration (the link is inactive).

   Note:
   Even though you unlinked your account, you still receive notifications in any group channels to which you belong.

5. Disable the messaging channel in your Notification Preferences.
   a. Select the gear icon in the banner frame to access the System Settings window.
   b. Under Notification Channels, disable the appropriate messaging channel.

   Note:
   If you want to link your ServiceNow account again, follow the steps in Link your ServiceNow user account to your Slack or MS Teams account for Now Actions.

Set up notifications for a custom Slack app

Send ServiceNow notifications to users via your own Slack bot.

Before you begin

- Request Integration Hub
- Request the Messaging Notification plugin
- Create a custom application for Slack

Role required: admin or messaging_admin

Procedure

1. Create a JSON for your custom Slack application.

   Ensure that your JSON matches the following example format:

```
{
"bot_user_id" : "U9TLFR6TA",
"scope" : "bot,channels:read",
"client_id" : "294380102274.335446283495",
"client_secret" : "ed0da5fe1f844fdac15a9e1ac1b0b831",
"verification_token" : "YXSNsRhioCeU8h9Fm2k4ACu0",
"app_id" : "A9VD48BEK",
```

servicenow

```
"access_token" :
 "xoxp-294380102274-295238228534-335452648487-2fac31f923dadd7ff
30f2bc8c852809d",
"bot_access_token" :
 "xoxp-333695856928-6pNwn03Y6vNSGI7KCAIPwv6V",
"team_id" : "T8NB63082",
"team_name" : "sn-notif-ext-test",
"user_id" : "U8P706QFQ",
"bot_username": "empdmyers1snnotifext"
}
```

To get the Bot User ID, copy the member ID of the bot user in the Slack app.

The Scope is always `bot,channels:read`.

To get the Client ID, Client Secret, and Verification Token, see the Basic Information page in your Slack workspace.

To get the App ID, copy the 9-character identifier from the app URL in your Slack workspace.

To get the Access Token and Bot Access Token, see the OAuth & Permissions page in your Slack workspace.

To get the Team ID, view the HTML source code for your Slack workspace page.

The Team Name is the same as your subdomain in your Slack workspace.

To get the User ID, copy the member ID for your own account in the Slack app.

The Bot Username is the name that you assigned to the bot user in Slack.

2. In your ServiceNow instance, navigate to **System Notification** > **Messaging** > **Messaging Integration Configuration**.
3. Next to the Slack listing, click **Install**.
4. On the JSON Configuration window, paste the JSON that you created in Step 1.

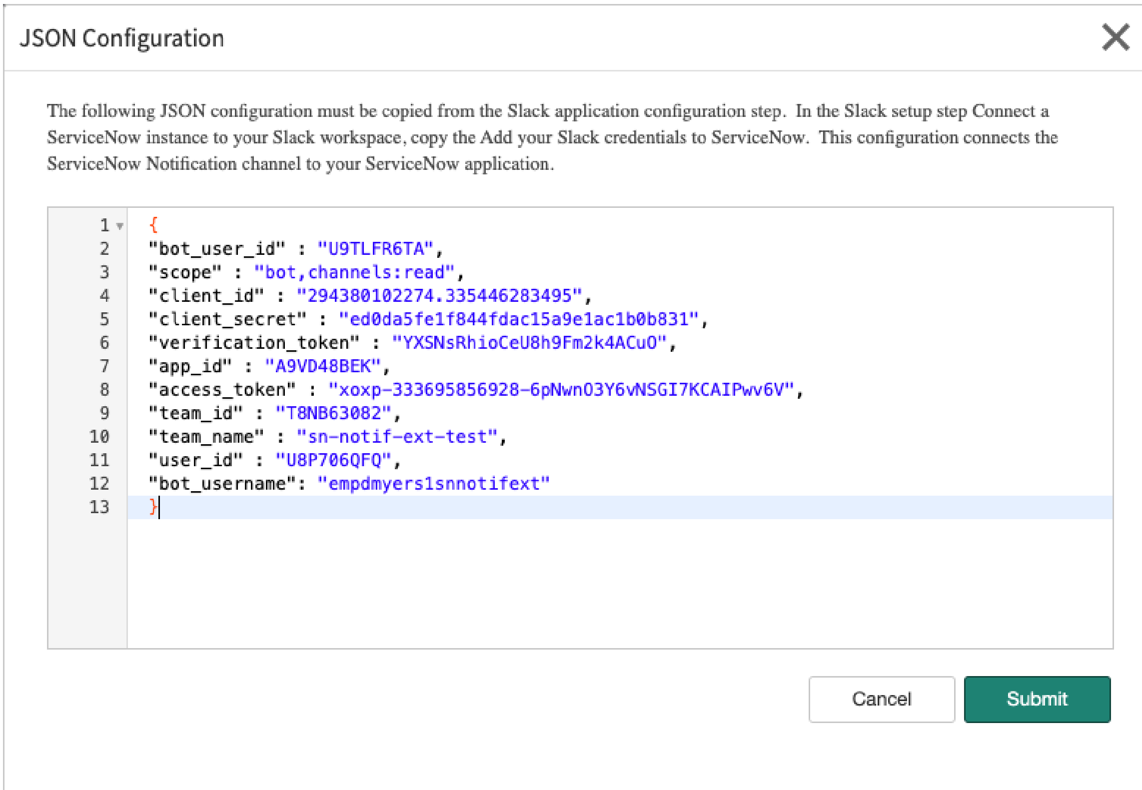

5. Click **Submit**.

6. On the confirmation window, click **Skip**.

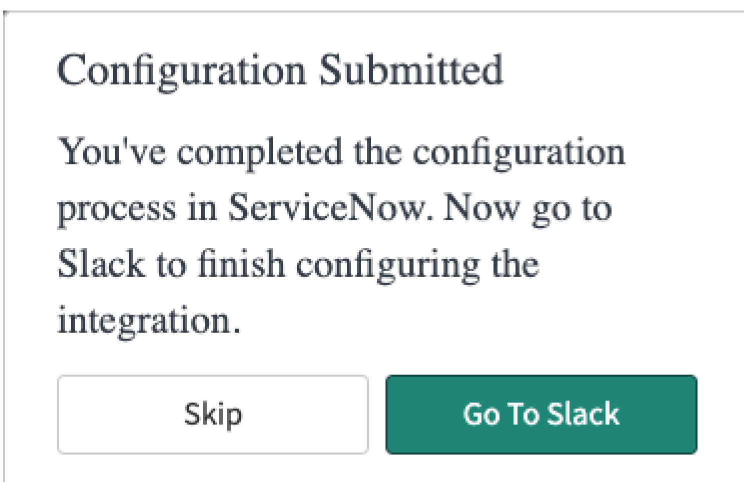

Result

The system creates a notification channel for your custom Slack app automatically.

What to do next

- Link ServiceNow user accounts to your custom Slack app so that users can take action on notifications that are delivered by your Slack bot.
- For each user account that you link, enable the notification channel using the system settings.
- Create notifications for your custom Slack app by following the steps in Create a messaging notification and Configure message content.

Related topics

Slack API documentation

Provider notifications

Send notifications directly to recipients via the provider on supported channels. Notifications can be simple messages for review, or actionable messages with buttons that users can select to perform certain actions.

Get started

Explore	Create
Explore and learn about provider notifications	Create provider notifications

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Provider notification

Send notifications directly to recipients via the provider on supported channels. Notifications can be simple messages for review, or actionable messages with buttons that users can select to perform certain actions.

Important terms

Provider

A provider is an implementation that delivers notifications to destinations of a specific channel.

Channel

A channel represents a particular communication mechanism with a recipient.

Destination Type

A destination type represents a logical way of categorizing a destination for the underlying channel and is required to send via a channel, as well as set a preference. As an admin, you can modify the notification **Auto Opt-in** behavior for a destination type.

Destination

A destination is a specific contact information of a recipient for a channel. It includes a destination type and an identifier suitable for use in the channel.

Destination Type sent by default is either set to true or false during the adapter installation. This determines if the destination type is opt-in or opt-out.

- If sent by default is set to true, notifications are delivered to the recipients if they have not explicitly opted-in/out to receive a notification on the destination type.
- If sent by default is set to false, notifications are not delivered to the recipients if they have not explicitly opted in to receive the notification on the destination type.

Destination preference precedence

Notifications are delivered to the destinations that honor the following preference precedence:

1. Notification preference of the recipient
2. Destination preference of the recipient
3. Destination Type send by default. You also have the option to modify **Auto Opt-in** behavior for a destination type.

Note:
Groups are not supported as recipients.

Common content for a notification

Provide common notification content so that different notification channels can use the same generic content.

Actionable notification

Enable provider admins to define an actionable notification and send it to users. You can also enable response actions.

Note:
A notification should either have Common Content or Provider Content in order to be delivered to the recipients.

All the above-mentioned capabilities are applicable only for Virtual Agent and Workspace providers.

Create provider notification

Create and update provider notifications, add and implement actions, and customize notifications

Create and update a provider notification for all users

Create a notification in the notification provider framework to send to both sys and non-sys users.

Before you begin
Role required: notifications provider admin and admin

Procedure

1. Navigate to **All** > **System Notification** > **Provider** > **Notifications**.
2. Select **New** to create a notification.
3. On the form, fill in the fields.

Notification form

Fields	Description
Name	Name of the notification.
Application	Scope of the notification record.
Category	Select the category to which this notification belongs. A category identifies and groups related notifications. This notification, if active and listed in the selected **Category** in the notification preferences for each user (Notifications tab in the System Settings window.
Active	Option to activate the notification.
Triggered By	Trigger used in the notification record. **Note:** The options in this field depend on the selected table. For example, if you select an incident table, the options are Events or Record Change.
Table	The table name on which the notification is triggered.
Inserted	Condition to trigger the notification.
Updated	Condition to trigger the notification.
Event	Event that triggers the notification. For example, to send a notification after an incident is closed or resolved, you would select the `incident.inactive` event. By default, this event is logged in the system each time an incident is resolved or closed. **Note:** You can select only an event that shares the same table as the notification. This field appears when you select **Event** in the **Triggered By** field.
Conditions	Additional conditions to trigger the notification.
Users	Names of the available sys users.
Recipients in fields	Field in the currently selected table that contains the users who receive an email notification.

Fields	Description
	Note: The **Recipients in fields** field under Recipients enables you to choose both sys and non-sys users to whom you can send notifications.
Is the originating user included?	Option to include the originating user to be included in the notification.
Event Parm1 contains recipient(s)?	Option to select if Parm 1 contains a list of comma-separated sys_ids.
Event Parm2 contains recipient(s)?	Option to select if Parm 2 contains a list of comma-separated sys_ids.
Resolve Parm1 recipients from table	Table that is used to resolve the sys_id of a Parm 1 recipient. **Note:** This field appears only when you select Event Parm 1.
Resolve Parm2 recipients from table	Table that is used to resolve the sys_id of a Parm 2 recipient. **Note:** This field appears only when you select Event Parm 2.
Advanced event condition	Option to select when the recipients are from more than 2 tables. These options are dynamic conditions.

4. Select **Submit**.
5. Open the recently created notification.
6. Update the form using any of the available related lists.
   See Create and add additional recipients for more information.

   Note:
   All the validation messages given by different providers are retained. Warning messages of one provider isn't overridden by warning messages of another provider.

Create and add additional recipients

Update the notifications form by adding additional recipients from the Additional Recipients related list to send notifications to.

Before you begin

Role required: admin and notifications provider admin

Procedure

1. Click Additional Recipients related list on the notification form to add new recipients.

   Note:
   An error message about adding content to the notification shows up. See create common content and create default content to create content for a notification.

2. Click **New** in the Additional Recipients related list to add additional recipients.
   The Additional Recipients form shows up.
3. Fill up the fields in the Additional Notification Recipients form.

Fields	Description
Notification	Name of the notification **Note:** The Notification name auto-populates and the **Recipient Table** field gets User[sys_user] as the default value.
Application	Scope of the recipient record
Recipient Table	Table name of the recipient **Note:** You can also select a non-sys users table.
Active	Option to activate the notification
Static Recipients	Names of static recipients of the notification
Dynamic Conditions	Dynamic filter conditions for the users **Note:** Click the refresh button next to the records matching condition link to update the exact number of records matching the latest filter condition. By default, it shows the total number of records in the current table. Only one active record can be added for a particular recipient table. If you want to add another record for the recipient table, you will have to deactivate the current record.

   Note:
   For one recipient table, we can use either static recipients or dynamic conditions.

4. Click **Submit**.

Create default content in a provider

Create and use the default content for a notification that doesn't have a content provider or some of the fields in the content provider are not present.

Before you begin

Role required: admin

Procedure

1. Enter sys_notification_content.list in the left navigator.
   A list of **Notification Contents** shows up.
2. Click **New Provider Content**.

   Note:
   Both shared content and provider content can be default content. Provider content has been considered for this example.

3. Select the type of Notification Content you would like to create.
4. Fill the **Notification Content** form for the type you have selected in the previous step.

   Note:
   The **is default** field is visible only if the **Notifications** field is not present.

5. Click **Submit**.
   A new default content has been created.
6. Go to **Notification Providers** and select one of the providers.
   The **Notification Providers** form for the selected provider opens.
7. Click the search icon in the **Default content** field to select the default content.

   Note:
   The **Default content** is associated with a provider through the **Default content** field.

   The **Select the document** modal opens up.
8. Select the default content in the **Document** field.
   You can select the recently created or any previous default content.
9. **Save** the changes.

Create custom notifications in Next Experience

Provide better context and information about users' work by creating custom content for in-product notifications in Next Experience.

Before you begin

Role required: admin

About this task

Create notifications with custom title and message content for your applications and use cases. Next Experience notifications appear as toast banners and in the Notifications menu.

Note:
Next Experience notifications aren’t supported in the legacy Agent Workspace.

Procedure

1. Navigate to **All** > **System Notification** > **Provider** > **Notifications**.
2. Select **New**.
3. On the form, fill in the fields.

Notifications form

Field	Description
Name	Name of the notification.
Application	Application scope of the notification.
Active	Option to activate the notification. The notification is active by default.
When to send	
Trigger	System action that triggers the notification. You can send the notification after a record is changed or after an event is triggered.
Inserted	Option to send the notification after a record is inserted. This field appears when you select **Record Change** in the **Trigger** field.
Updated	Option to send the notification after a record is updated. This field appears when you select **Record Change** in the **Trigger** field.
Event	Event that triggers the notification. For example, to send a notification after an incident is closed or resolved, you would select the `incident.inactive` event. By default, this event is logged in the system each time a user resolves or closes an incident. **Note:** You can select only an event that shares the same table as the notification. This field appears when you select **Event** in the **Trigger** field.
Table	Table to receive notifications about.
Conditions	Filter to specify the table records that users receive notifications about. For example, to send notifications about top-priority incidents, you would select **Incident [incident]** in the **Table** field and set the conditions to **[Priority] [is] [1 - Critical]**.
Who will receive	
Users	Users who receive the notification.

Field	Description
	Note: Notifications in messaging channels are sent only to users with ServiceNow accounts (sys_user profiles). Consumers and customer contacts are considered as guests and can't receive notifications in messaging channels.
Recipients listed in fields	Record fields that include users who receive the notification. For example, to send the notification to the record assignee, select **Assigned to**. **Note:** This field shows up only if you select a table.
Include the person whose action triggered the notification	Option to include the user who changed the record or triggered the event.
Recipient(s) listed in event parm1	Option to select whether Parm 1 contains a list of comma-separated sys_ids. For more information, see Classic Events.
Recipient(s) listed in event parm2	Option to select if Parm 2 contains a list of comma-separated sys_ids
Table containing recipients from event parm1	Table that is used to resolve the sys_id of the recipient **Note:** This field appears only when you select Event parm 1.
Table containing recipients from event parm1	Table that is used to resolve the sys_id of the recipient **Note:** This field appears only when you select Event parm 2.
Advanced event condition	Option to select when the recipients are from more than two tables. These are dynamic conditions.

4. From the form context menu, select **Save**.
5. Select the **Contents** tab and then select **New Provider Content**.
6. Select **Next Experience** as the content provider.
7. On the form, fill in the fields.

Notification Next Experience Content form

Field	Description
Name	An internal name for the notification content.
Notification	The notification trigger that sends the notification content.
Active	Option for whether the notification is sent when triggered.
Table	Table to receive notifications about.

Field	Description
Route	Next Experience Route
Message Heading	Header text at the top of the message. From the Select variables list, you can select variables to include your message heading.
Message	Secondary text of the message that follows the heading. From the Select variables list, you can select variables to include your message body.

Example

The following example shows the values for a notification with custom content that references the number and short description of an incident.

Field	Value
Name	Incident record update content
Notification	Incident records
Active	true
Message Heading	${number} changed
Message	Short description: ${short_description}

8. Select **Submit**.

Result

When triggered, users configured to receive the notifications see notifications with custom content as banners and in the Notifications menu in the Next Experience Unified Navigation.

The UI Notification Inbox [ui_notification_inbox] table lists all sent notifications and their recipients.

Trouble?

If users receive duplicate notifications, verify that the notification trigger has only one content provider. If the notification has both a Next Experience and Workspace content provider, remove the Workspace content provider from the notification trigger to send only the Next Experience notification with custom content.

Implementing notification action with a script

Use the **Script** field in the Notification Action form to update and implement a notification action.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Provider** > **Notifications**.
   A list of notifications shows up.
2. Select an existing notification.
3. Scroll down to the Notification Actions related list.

4. Open the Notification Action form and choose between the following options to reach the Notification Action form:
   - If there are no existing notification actions under Notification Actions related list, do the following.
     a. Click **New Provider Action** under Notification Actions related list.
     b. Select the type of notification action you would like to create from the Notification Action wizard.

        Note:
        Scriptable Action, Flow Action and Virtual Agent are the only types of notification actions available.
   - If there is an existing notification action, select one from the list under Notification Actions related list.

   The Notification Action form shows up.
5. On the form, update the fields as required.

Fields	Description
Name	Name of the actionable notification.
Application	Scope of the actionable notification record.
Notification	Name of the notification containing the actionable notification record.
Active	Option to activate the actionable prompt record.
Script	Script to implement the actionable notification.

6. Click **Update**.

Select provider notification action capabilities

Enable provider notification definitions to accept actions as scriptable action, flow actions, and virtual agent.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Provider** > **Notifications**.
   A list of notifications shows up.
2. Select an existing notification.
3. Scroll down to the related lists.
4. Select the Actions related list.
5. Click **New Provider Action**.
   The Notification Action wizard shows up.
6. Select one of the following types of notification actions you would like to create from the Notification Action wizard.

- Scriptable Action
- Flow Action
- Virtual Agent

Create a scriptable notification action

Select **Scriptable Action** from the Notification Action wizard to create a scriptable notification action.

Before you begin

Role required: admin

```
Before performing this task, you must complete Select provider notification action capabilities .
```

Procedure

1. Select **Scriptable Action**.
2. On the form, fill in the fields.

Fields	Description
Name	Name of the notification action.
Application	Scope of the notification action record.
Notification	Name of the notification containing the notification action record.
Active	Option to activate the notification action record.
Script	Script to implement the notification action.
Acknowledgement	Acknowledgement message after an action is selected from a notification.

3. Click **Submit**.

Create a flow notification action

Select **Flow Action** from the Notification Action wizard to create a flow notification action.

Before you begin

Role required: admin

Procedure

1. Select **Flow Action**.
2. On the form, fill in the fields.

Fields	Description
Name	Name of the notification action.
Notification	Name of the notification containing the notification action record.

Fields	Description
Async	Option to execute configured flow or action asynchronously.
Active	Option to activate the notification action record.
Choose Flow or Action	Option to select either Flow or Action. See Flow Designer for more information. **Note:** You can select either Flow or Action, but not both.
Inputs	Script used for the creation of flow notification action. **Note:** The script returns Json of key value pairs, which is passed as input in executing configured flow or action.
Acknowledgement	Acknowledgement message after an action is selected from a notification.

3. Click **Submit**.

Add actions to the notification content

Use the Link Actions to Content related list to add actions to the notification content, selection of the provider notification and ordering of actions of an actionable content.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Notification** > **Provider** > **Notifications**.
2. Select **New** to create a notification, or select an existing notification.

   Note:
   If the notification includes one or more actions, a message appears to remind you to ensure the actions are associated with content.

3. In the related lists section, select the Link Actions to Content related list.
4. Select **New** to create an actionable content.
5. On the form, fill in the fields.

Actionable Content form

Fields	Description
Name	Name of the action set.
Application	Scope of the notification.

Fields	Description
Notification	Name of the notification.
Actions	Option to select the actions. **Note:** You can also set the order of actions of an actionable content.
Content	Option to select the content to apply to the selected action.
Active	Option to activate the action set.
Action(s) listed in event parm1	Option to link actions from event parm1. Actions must be sys_ids from the **sys_notification_action** table.
Action(s) listed in event parm2	Option to link actions from event parm2. Actions must be sys_ids from the **sys_notification_action** table.

6. Select **Submit**.

Email client

Send email directly from any record, such as an incident, change request, or problem record.

Get started

Explore	Configure	Components Installed
Explore and learn about sending an email directly from any record	Configure and manage the behavior of your email client	View components installed with email client

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Email client

Send email directly from any record, such as an incident, change request, or problem record.

Using the email client

You can use the email client either from a form or the record view in a workspace.

- To open the email client from a form, open a record and click the more options menu (). Click the email icon ().
- To open the email client in a workspace, open a record and click the menu icon ( ). Click **Compose Email**.

The email client looks like a standard email interface, which contains a toolbar for text formatting and adding attachments.

The email client uses the multipart/mixed content type and supports HTML markup in the message body.

Note:
The **Subject** field on the email client allows a larger character count than the default setting for the **Subject** field on the Email [sys_email] table. If the subject text from the client is being truncated, increase the **Max length** value for the **Subject** field on the Email table.

For more information on using the email client in a workspace, see Communicating via email in workspace.

Activation information

The email client is activated with the Email Client plugin (com.glide.email_client), which is active by default on the ServiceNow AI Platform.

The email client is enabled by default on the incident table. To enable the email client for another table, add the `email_client` dictionary attribute on the table's collection record. For more information, see Enable the email client for a table.

The email client is enabled by default on the change table.

Give users access to the email client by assigning them the email_composer role.

More setup

You can configure the email client by creating an Email Client Configuration [sys_email_client_configuration] record. For more information, see Email client configurations.

You can also create email client templates or quick messages to use predefined content in email client messages. For more information, see the following:

- Create an email client template
- Composing emails with quick messages

Optionally, you can give users the option to send email client messages as SMS text messages. For more information, see Enable SMS delivery with the email client.

Configure email client

Enable and configure email client.

Configuration overview

Enable the email client for a table

Enable the email client for a table so that users can send emails directly from the table record.

Before you begin

Role required: admin

About this task

The email client is enabled by default on the Incident [incident] table. You can enable the email client for another table by adding the *email_client* dictionary attribute to the table.

Note:
This capability is not inherited by tables that extend the current table. For example, enabling the email client on the Task table does not enable it for the Incident or Problem tables.

Procedure

1. Open a record in the table that you want to enable the email client for.
   For example, to enable the email client for the Problem [problem] table, navigate to **Problem** > **Open**, and then open any problem record.
2. On the form, click the menu icon (⋮) and then click **Configure** > **Dictionary**.
3. On the Dictionary Entries list, open the first record.
   The first record has the record type **Collection** and does not have any entry for **Column name**.
4. On the form, in the Related Links section, click **Advanced view**.
5. In the **Attributes** field, enter `email_client=true`.
   If there are other values in the field, separate the attribute with a comma.
6. Click **Update**.

Email client configurations

Use email client configurations to manage the behavior of your email client. Each configuration consists of different email controls for setting allowable email recipients and email addresses.

In new and upgraded instances, the Email Client Configuration module is activated by default through the Email Client plugin (com.glide.email_client).

Note:
If you upgraded to the Madrid release, a default email client configuration is created automatically during upgrade. This default email configuration is used if you do not define any other email client configurations. Also, if you previously defined certain email client properties to control email address auto-complete, disambiguation, and the display of the From and Reply-To email addresses, those settings are retained in the default email client configuration.

How email client configurations work

Email client configurations determine how your email client behaves, such as whether the From address or Reply-To address should be displayed. A configuration applies to a specific ServiceNow table, which enables you to define email client behavior for a specific context, purpose, or application. For example, if you have email client users who send email about Customer Service case records, you can create a configuration that applies email client controls specific to customer service. You can also select the email client recipient qualifiers for controlling the auto-complete recipient results from a selected table, such as the Customer Contact table.

An email client configuration consists of the following components:

Email client recipient qualifiers

Recipient qualifiers control the email recipients returned in the auto-complete results for the To, Cc, and Bcc fields of the email client. You can provide additional qualifiers displayed in the auto-complete results to differentiate recipients that have the same names and narrow the results.

> **Note:**
> Use the **To is Optional** field in the Email Client Configuration form to make the **To:** field optional to send an email.

Email client From addresses

This control enables the display of the From and Reply-To addresses. Use the Email Client Template to create a set of allowable From addresses that can be used when a user sends an email with the email client. You can provide a specific From address to be used in the email client, or you can choose to hide it entirely.

Email client attachment handling

By default, attachments that are sent from the email client are attached to the target record. Instead, you can select to attach files to the Email [sys_email] record, or you can set conditions that define which email client attachments are added to the target record. To see the status of all email attachments, go to the Email Attachments [sys_email_attachment] table.

You can also use an Email Client Template to define the content of an email client message and set other message characteristics, such as an email sender (From address) configuration.

Email client configuration process

Users with the admin role define and edit email client configurations. Users with the email_client_admin role can edit configuration information in selected tables.

1. **Build the various email client controls that can be used to create an email client configuration.**
    - Define email client recipient qualifiers
    - Define email client from addresses
2. **Create the email client configuration.**

    Define an email client configuration that determines how your email client behaves when users create an email message.

Define email client recipient qualifiers

Create a configuration that controls the auto-complete list of recipients displayed in the email client.

Before you begin

Role required: email_client_admin

About this task

You can specify recipient qualifiers that display additional fields (from a selected ServiceNow table) in the auto-complete list. These fields differentiate email recipients who have the same first and last names.

Use the following tabs in the Email Client Recipient Qualifier form to define a recipient configuration.

- **Display Configuration** – Set up the email recipient auto-complete behavior and optionally select additional fields to differentiate recipients who have the same name. The additional fields ensure that users select the proper recipient for an email.
- **Query Configuration** – Specify a condition or script that queries the selected table and filters the recipient results returned.

You can define different recipient configurations, which can be used in an email client configuration.

Procedure

1. Navigate to **Email Client** > **Email Client Recipient Qualifier** and click **New**.
2. Fill in the fields at the top of the Recipient Qualifier form.

Field	Description
Name	A unique name for this email client recipient configuration.
Application	The type of scoped application.
Table	The ServiceNow table to be queried for recipients.
Description	A brief description of this recipient configuration.

3. Fill in the fields in the **Display Configuration** tab to control the auto-complete display in the email client.

Field	Description
Email Address	The field on the table that contains the recipient email address.
Display Name	The field on the table used for the recipient name displayed.
Order	The order of results returned relative to other recipient qualifiers defined.
Additional Display Fields	A content list for choosing additional fields from the table to be displayed in the auto-complete list. Differentiates recipients who share the same name.
Search Additional Fields	A check box for enabling additional fields to be searched. Expands the recipients returned in the auto-complete list by querying (searching) the **Additional Display Fields**.

4. Fill in the fields in the Query Configuration tab.

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td>Prefix Search</td><td>Option to find the desired recipient using prefix.

Note:
Selecting the Prefix Search option displays the Last Name Field drop-down menu with options to filter the recipient's name.</td></tr>
<tr><td>Type</td><td>Select the method for filtering the recipients returned in the auto-complete list:
◦ Condition
◦ Script</td></tr>
<tr><td>Conditions</td><td>If you selected Condition for the Type, use the condition builder to specify the conditions that must be met to return the appropriate recipients.</td></tr>
<tr><td>Script</td><td>If you selected Script for the Type, enter a script that uses these variables to return the appropriate recipients:
◦ <code>recipientQuery</code>: GlideRecord for the table being queried.
◦ <code>targetRecord</code>: GlideRecord for the target record from which the email client was opened.

For example, this recipientQuery is a GlideRecord that represents the Customer Contact table. The targetRecord variable is a Customer Service case from which the email client was opened. This query limits the email client auto-complete results to contacts who belong to the same account as the Customer Service case.
<pre>(function (recipientQuery, targetRecord) {

  // Limit results to contacts belonging to the account for the case.
  var account = targetRecord.getValue('account');
  recipientQuery.addQuery('account', account);

})(recipientQuery, targetRecord);</pre></td></tr>
</table>

5. Click **Submit**.
   The recipient qualifier configuration is added to the Recipient Qualifier [sys_recipient_qualifier] table and is available for use in an email client configuration.

Define email client from addresses

Set an allowable email address that is displayed in the From address of a message sent from the email client.

Before you begin

Role required: admin

Procedure

1. Navigate to **Email Client** > **Email Client From Address** and click **New**.
2. In the **Display name**, enter the name to be displayed in the From field of the email client message.
3. Enter a valid **Email address** for the From address in the email client.
4. Click **Submit**.
   The From address is added to the Email Client From Address [sys_email_client_from_address] table.

Create an email client configuration

Define a configuration that controls the display and behavior of the email client that is based on a selected ServiceNow table.

Before you begin

Role required: email_client_admin

About this task

The configuration determines the recipient auto-complete results that are displayed, email addresses that can be entered, and a set of allowable From addresses to be used when sending a message with the client.

You can create one email client configuration per ServiceNow table.

Procedure

1. Navigate to **Email Client** > **Email Client Configuration** and click **New**.
2. On the form, fill in the fields.

Field	Description
Name	Unique name for this email client configuration.
Application	Type of scoped application.
Active	Option that indicates that the configuration is active. The configuration is active by default.
Table	ServiceNow table that this client configuration applies to. Select a table from the list.

3. In the Recipient Configuration section, select **Recipient Qualifiers** to be used for this client configuration.

a. Click the lock icon ( ).
By default, two types of recipients are available: Active Users with email accounts and Active Groups with email accounts.

b. Click the search icon ( ) and select a recipient configuration from the list of available recipient configurations.

c. Click the lock icon ( ) to lock it.

4. In the Display Configuration section:
   a. Select the **Display From** option to display the From email address in the email client, or clear the check box to hide the From email address in the email client.
   b. Select the **Display Reply-To** option to display the Reply To address in the email client, or clear the check box to hide the Reply To email address in the email client.
5. In the Attachment Handling section, select from one of three choices for the **Attachment Send Action** field.
6. In the Filter Configuration section, select **Email Address Filters** to be applied to the email client.
   a. Unlock the lock icon ( ).
   b. Click the search icon ( ) and select one or more Email Address Filters.
   c. Click the lock icon ( ) to lock **Email Address Filters**.

   Each filter that you select evaluates a given email address and determines whether or not the given address is valid to apply to an outbound email.
7. Click **Submit**.
   The configuration is added to the Email Client Configuration [sys_email_client_configuration] table.

Enable SMS delivery with the email client

Give users the option to send an email client message as an SMS text message.

Before you begin

Role required: admin

Procedure

1. In the application navigator, enter `sys_properties.list`.
2. From the list, find and open the record for the *`glide.email_client.show_sms_option`* property.
3. On the form, change **Value** to **true**.
4. Click **Update**.

Result

In the email client, user can select to send the message as an SMS text message.

Email client SMS delivery

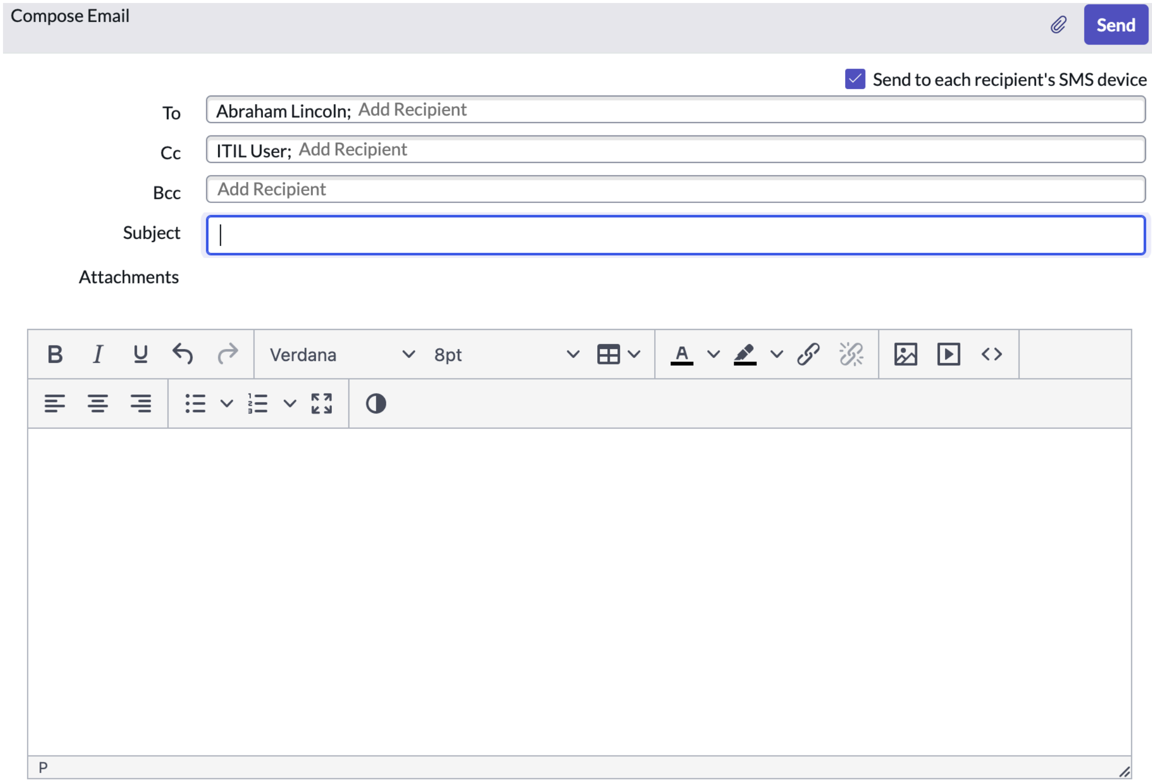

If the recipient doesn't have a registered SMS device, the email client sends the message to the primary email device.

Related topics

Add a system property

Create an email client template

Create a different template for each table that uses the email client to pre-populate data for recipients and email content.

Before you begin

Role required: admin

About this task

The email client uses its own email templates to define default values for fields. Use the following sections in the Email Client Template form to build a client template:

- Content: The message body.
- Recipients: The email addresses of users receiving the email message. The email addresses are displayed in the To, Cc, and Bcc fields of the message.
- Sender Configuration: The method used to generate the email sender (From email address) of the message.

Procedure

1. Navigate to **All** > **Email Client** > **Email Client Templates**.
2. Select **New**.
3. On the form, fill in the fields.

Email Client Template form

Field	Description
Name	Unique template name.
Table	The table that the template applies to. Enable the email client for the same table.
Conditions	The conditions that determine when this client template is used. Use the condition builder to identify the target record that must match before this template is applied to the email client.
Application	The type of scoped application.
Execution Order	A number that indicates the order in which the template conditions are evaluated. The template with the lowest execution order is considered.
Include attachment from email	Option for including attachments from email to the **replyall-received**, **reply-received**, **forward-sent**, **reply-sent**, **replyall-sent**, and **forward-received** response templates. **Note:** This option is applicable only for response email in a configurable workspace.

4. On the **Content** tab, fill in the fields.

Content tab

Field	Action
Subject	Enter a description of the email. **Note:** The **Subject** field on the email client allows a larger character count than the default setting for the **Subject** field in the Email [sys_email] table. If the subject text from the client is being truncated, increase the **Max length** value for the **Subject** field in the Email table.
Content Type	Select the format of the message body content: **HTML** or **plain text**.
Body HTML	If you've selected **HTML** for the **Content Type** field, enter the content of the message body using the HTML editor toolbar to format the HTML. To reference field values, select variables from the variables list or manually type variable references using the syntax `${table_name.variable_name}`. To reference the user who launches the email client, enter the variable `${current_user}`. **Note:** Line breaks don't appear for multi-line fields such as ${description} and ${comments} in the email client template.

Field	Action
	You can call a mail script by using `${mail_script:script name}` in the **Body HTML** field.
Body text	If you selected **plain text** for the **Content Type** field, enter any text or mail script that you want to appear in the message body. You can insert a mail script in the **Body text** field by using the following syntax: `<mail_script> [code] </mail_script>`

5. On the **Recipients** tab, fill in the fields.

Recipients tab

Field	Action
To	Enter a comma-separated list of either field names that contain user email addresses or just email addresses. To reference an email address using a script, create a script include and then call the script include in the **To** field. Your script must start with the `javascript:` prefix. The script must return email addresses in a comma-separated string.
Cc	Enter a comma-separated list of either field names that contain user email addresses or just email addresses. To reference an email address using a script, create a script include and then call the script include in the **Cc** field. Your script must start with the `javascript:` prefix. The script must return email addresses in a comma-separated string. This field can't have the same addresses as the **To** field.
Bcc	Enter a comma-separated list of either field names that contain user email addresses or specific email addresses. To reference an email address using a script, create a script include and then call the script include in the **Bcc** field. Your script must start with the `javascript:` prefix. The script must return email addresses in a comma-separated string. This field can't have the same addresses as the **To** or **Cc** fields.

6. On the **Sender Configuration** tab, select the **From Generation Type** method to determine how the sender (From address) in the email client message is generated.
   Use this method only if you want a different From address than the one defined in your SMTP email account.
   - None: The From address isn't generated in the email client message.
   - SMTP Email Account: Use the From address of the SMTP email account for the instance as the sender.
   - Select From List: Choose from a list of allowable From addresses defined in the Email Client From Address [sys_email_client_from_address] table.
   - Script: Run a GlideRecord query on the Email Client From Address [sys_email_client_from_address] table.

For example, the following script sets the From address based on the location of the incident caller:

```
(function (fromAddressQuery, targetRecord) {
    // targetRecord is incident for this template
    var location = targetRecord.caller_id.country;

    if (location == 'us')
        fromAddressQuery.addQuery('email_address',
 'servicedesk.us@example.com');
    else if (location == 'japan')
        fromAddressQuery.addQuery('email_address',
 'servicedesk.jp@example.com');
    else if (location == 'uk')
        fromAddressQuery.addQuery('email_address',
 'servicedesk.uk@example.com');

})(fromAddressQuery, targetRecord);
```

- Text: Enter the email From address to be used in the client.

Note:

The From email address doesn't appear in the email client unless you configure an email client to display the From address. For more information on managing the behavior of an email client, see Email client configurations.

If the domain for the From address isn't the domain of the SMTP email account server, the owner of the From domain address must configure the SPF record for the domain. The owner changes the domain settings to enable the instance to send an email as if from that domain. For more information on using SPF records, see ServiceNow KB0535456.

7. Select **Submit**.
   The template is added to the Email Client Templates [sys_email_client_template] table.

Related topics

Script includes

Apply an email client template

After you create an email client template, you can apply it to an email.

Before you begin

Role required: admin

About this task

The following steps are implemented only after creating an email client template

Procedure

1. Click the ellipsis (...) icon at the top of the template to apply the template to an email.
2. Select **Email**.
3. Fill in the fields to compose an email.

See Create an email client template for field details.

Note:
The body of the email is filled up with the details from the incident record. If you modify one or more fields in the incident record and are not saved, the unsaved changes still reflect in the email body.

Note:
2000 characters is the total limit for all the changes in the incident record. If some of your changes are not reflected and get truncated in the email body, this is because the total changes is more than 2000 character limit.

Create an email client response template

You can now create custom email client response templates that support conditions.

Before you begin

Role required: admin

About this task

Email client template is applied automatically in the email client based on the specified table, conditions, and execution order.

Procedure

1. Navigate to **All** > **Email Client** > **Email Client Templates**.
   The following options for **Response Template - Received** and **Response Template - Sent** show up.
2. Click one of the available options to create a template for that selected option.
3. Fill in the fields at the top of the Email Client Template form.

Field	Description
Name	Auto-filled template name based on the selected template option. **Note:** This is a read-only field.
Table	Auto-filled table name. **Note:** This is a read-only field because for response templates, the table name is always sys_email.
Conditions	The conditions that determine when this client response template is used. Use the condition builder to identify the target record that must match before this template is applied. **Note:** If the conditions are not true, the email client gets the default template. Select the target table name under conditions that the template applies to. Enable the response template for the same table.

Field	Description
	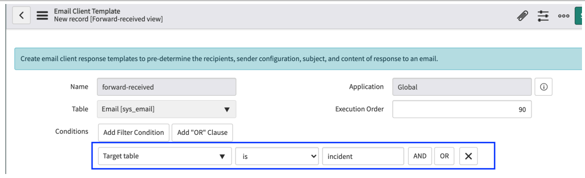
Application	The type of scoped application.
Execution Order	A number that indicates the order in which template conditions are evaluated.

4. Fill in the fields in the **Content** tab.

Field	Description
Subject	Enter a description of the email. **Note:** The **Subject** field on the email client allows a larger character count than the default setting for the **Subject** field on the Email [sys_email] table. If the subject text from the client is being truncated, increase the **Max length** value for the **Subject** field on the Email table.
Content Type	Select the format of the message body content: HTML or plain text.
Body HTML	If you selected HTML for the **Content type**, enter the content of the message body using the HTML editor toolbar to format the HTML. To reference field values, select variables from the variables list or manually type variable references using the syntax `${table_name.variable_name}`. To reference the user who launches the email client, enter the variable `${current_user}`. **Note:** Line breaks don't appear for multi-line fields such as ${description} and ${comments} in the email client template. You can call a mail script by using `${mail_script:script name}` in the **Body HTML** field.
Body text	If you selected plain text for the **Content type**, enter any text or mail script that you want to appear in the message body. You can insert a mail script in the **Body text** field by using the following syntax: `<mail_script> [code] </mail_script>`

5. Fill in the fields in the **Recipients** tab.

Field	Description
To	Enter a comma-separated list of either field names that contain user email addresses or specific email addresses. To reference an email address using a script, create a script include and then call the script include in the **To** field. Your script must start with the `javascript:` prefix. The script must return email addresses in a comma-separated string.
Cc	Enter a comma-separated list of either field names that contain user email addresses or specific email addresses. To reference an email address using a script, create a script include and then call the script include in the **Cc** field. Your script must start with the `javascript:` prefix. The script must return email addresses in a comma-separated string. This field cannot have the same addresses as the **To** field.
Bcc	Enter a comma-separated list of either field names that contain user email addresses or specific email addresses. To reference an email address using a script, create a script include and then call the script include in the **Bcc** field. Your script must start with the `javascript:` prefix. The script must return email addresses in a comma-separated string. This field cannot have the same addresses as the **To** or **Cc** fields.

6. In the **Sender Configuration** tab, select the **From Generation Type** to determine how the sender (From address) in the email client message is generated.
   Use this field only if you want a different From address than the one defined in your SMTP email account.
   - None: The From address is not generated in the email client message.
   - SMTP Email Account: Use the From address of the SMTP email account for the instance as the sender.
   - Select From List: Choose from a list of allowable From addresses defined in the Email Client From Address [sys_email_client_from_address] table.
   - Script: Run a GlideRecord query on the Email Client From Address [sys_email_client_from_address] table.

     For example, the following script sets the From address based on the location of the incident caller:

```
(function (fromAddressQuery, targetRecord) {
    // targetRecord is incident for this template
    var location = targetRecord.caller_id.country;

    if (location == 'us')
        fromAddressQuery.addQuery('email_address',
 'servicedesk.us@example.com');
    else if (location == 'japan')
        fromAddressQuery.addQuery('email_address',
 'servicedesk.jp@example.com');
    else if (location == 'uk')
        fromAddressQuery.addQuery('email_address',
 'servicedesk.uk@example.com');
```

```
})(fromAddressQuery, targetRecord);
```

- Text: Enter the email From address to be used in the client.

Note:
The From email address does not appear in the email client unless you configure email client to display the From address. For more information on managing the behavior of email client, see Email client configurations.

If the domain for the From address is not the domain of the SMTP email account server, the owner of the From domain must configure the SPF record for the domain. The owner changes the domain settings to allow the instance to send email as if from that domain. For details on using SPF records, see ServiceNow KB0535456.

7. Click **Submit**.

Composing emails with quick messages

Insert predefined content into the message body of emails that you send from the email client.

Using quick messages

In the email client, you can select a quick message to fill the email body with the content that is specified in the quick message. After you define one or more quick messages, the Quick Message selector appears in the email client so that you can construct an email with multiple quick messages.

By default, the quick message content is inserted at the place of the cursor. It doesn't replace the content that you created before selecting the quick message. To replace message content with a quick message, highlight the text to replace and then select a quick message.

Note:
To instead have quick message content replace all existing content in an email draft, set the *glide.email_client.quick_message.insert* property to **false**.

Adding quick messages to an email

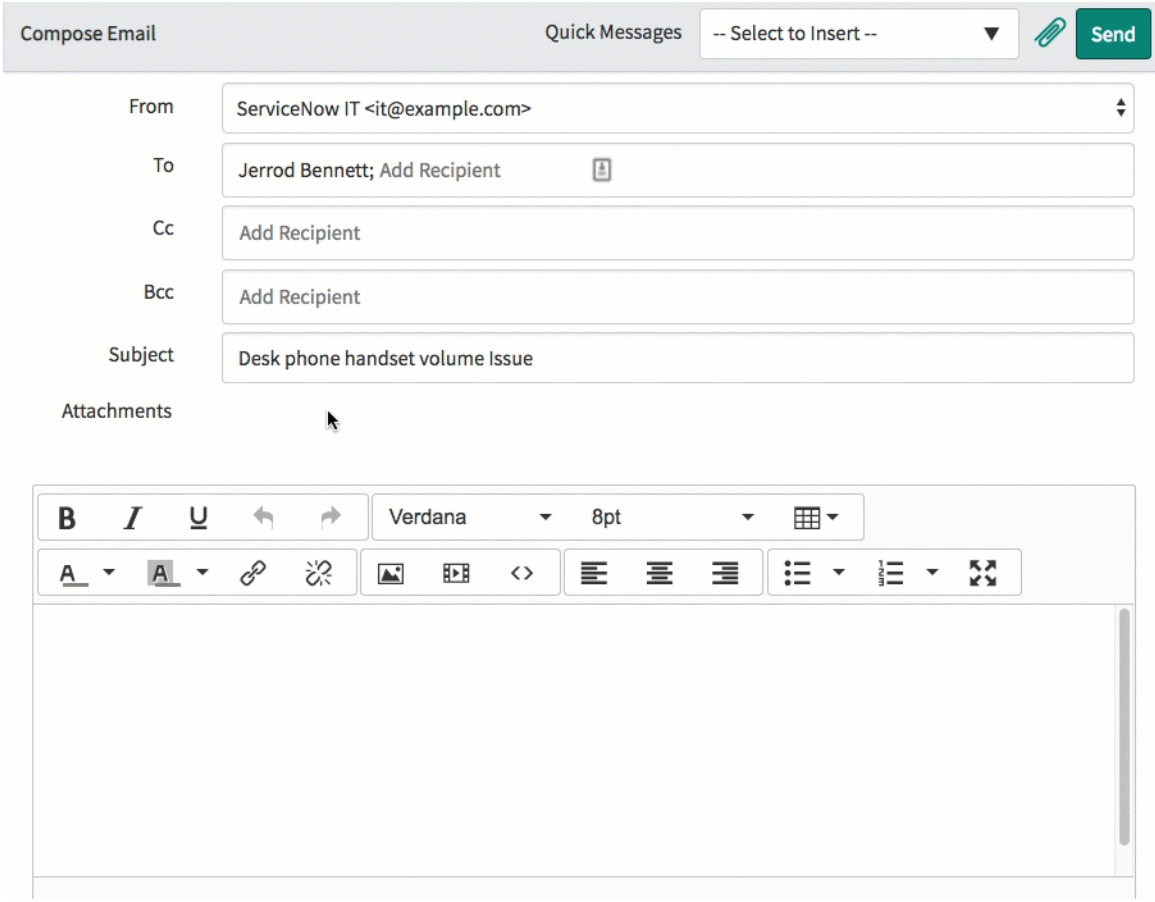

Creating quick message content

Define quick messages by creating records in the Email Client Canned Messages [sys_email_canned_message] table.

When you define a quick message, you can add any of the following types of content into the message body:

- Icons
- Logos
- Pictures
- Rich-text HTML
- Hyperlinks
- Variables
- Any other HTML constructs

Note:
If you add attachments to the quick message, the attachments aren't sent as part of the email distribution.

Assign the email_client_quick_message_author role to business managers so that they can create quick messages for users in their group and other groups.

Setting quick message conditions

When you define a quick message, you can limit the availability of the quick message according to:

- The user who launches the email client
- The group of the user who launches the email client
- The table of the record from which you launch the email client
- The target record from which you launch the email client

For example, if you designate Beth Anglin as the user for a certain quick message, the quick message is available in the email client only for Beth Anglin.

If you associate a quick message to the Incident [incident] table, the quick message is available in the email client only after you launch from an incident record. If you specify a target record from the Incident table for the quick message, the quick message becomes available in the email client only for matching incident records.

Define a quick message

Create predefined content to add in the email client so that users can write emails consistently and efficiently.

Before you begin

Role required: email_client_quick_message_author or admin

Procedure

1. Navigate to **Email Client** > **Quick Messages**.
2. Click **New**.
3. On the form, fill in the fields.

Field	Description
Title	Name that appears in the Quick Message selector.
Active	Option for activating the quick message. When selected, the quick message is available for selection in the Quick Message selector.
Application	Type of scoped application.
User	User who has access to this quick message. Selecting a user restricts access to that user only. Leave the field blank to have no user-based restrictions.
Group	Group whose members have access to this quick message. Selecting a group restricts access to members of that group only. Leave the field blank to have no group-based restrictions.
Table	Table to which the quick message applies. To make the quick message available for all tables, leave blank.
Conditions	Target record that must match the conditions before the quick message appears for selection in the email client.
Body	Content to insert in the **Message Text** field in the email client. By default, the field supports HTML format.

Field	Description
	To reference field values, select variables from the variables list or manually type variable references using the syntax `${table_name.variable_name}`. To reference the user who launches the email client, enter the variable `${current_user}`. You can use these values to include a predefined signature in email responses from agents to customers. **Note:** Line breaks don't appear for multi-line fields such as ${description} and ${comments} in the email client template.

4. Click **Submit**.

Components installed with the email client

Several types of components are installed with activation of the Email Client (com.glide.email_client) and Email Client Template (com.glide.email_client_template) plugins, including tables and user roles and tables.

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

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
Email Client Configuration [sys_email_client_configuration]	Controls for the email client. There's only one Email Client Configuration record for each table that uses the email client.
Email Client From Address [sys_email_client_from_address]	Email addresses that users can select as the From address of an email client message.
Recipient Qualifier [sys_recipient_qualifier]	Email addresses that users can send email client messages to.
Email Client Template [sys_email_client_template]	Templates that are applied automatically to qualified email client messages.

Email client administrator [email_client_admin]

Configures the email client to fit the needs of your organization and can read or write email client templates.

Contains Roles

List of roles contained within the role.

- email_client_quick_message_author
- email_client_template_write

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Email composer [email_composer]

Creates email client messages and can read email client templates.

Contains Roles

List of roles contained within the role: email_client_template_read.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Read-only role for email client templates [email_client_template_read]

Reads email client templates.

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.

Read/write role for email client templates [email_client_template_write]

Reads and writes email client templates.

Contains Roles

List of roles contained within the role: email_client_template_read.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Quick message author [email_client_quick_message_author]

Creates quick messages for users in their group.

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.

Inbound email

Save time from responding to emails manually when you configure your instance to send replies, create incidents, or update records automatically in response to inbound emails.

Get started

Explore	Configure	Actions
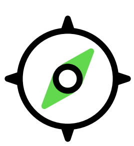	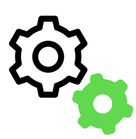	
Explore automated system responses to inbound email	Configure a processing order for inbound email actions	Run Inbound actions by comparing the inbound email type and inbound action conditions

Use	References and Examples	
Process inbound email actions	See examples for inbound email actions, inbound events integration and details about email objects, email matching, user impersonation and redirecting emails	

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Automating system responses to inbound email

Save time from responding to emails manually when you configure your instance to send replies, create incidents, or update records automatically in response to inbound emails.

You can define system responses to inbound emails in two ways:

- Create an inbound email flow in Flow Designer
- Script an inbound email action

Inbound email trigger in Flow Designer

With the inbound email trigger in Flow Designer, you can create flows that define the automated processes that your instance takes when it receives an email.

Inbound email flows take priority over inbound email actions. If you create flows with inbound email triggers, emails are first processed by the inbound email triggers before they are processed by inbound email actions.

The following diagram shows the processing order for inbound emails. When an email is sent to your instance, the system first classifies the email as a reply, forward, or new email. Then the system runs the inbound email through an inbound email flow. If the flow issues stop processing, the email is finished being processed. If the flow does not issue stop processing, the system tries to match the email to another inbound email flow. If at any point the email does not match with an inbound email flow, the system matches the email to an inbound email action instead.

Processing order for inbound emails

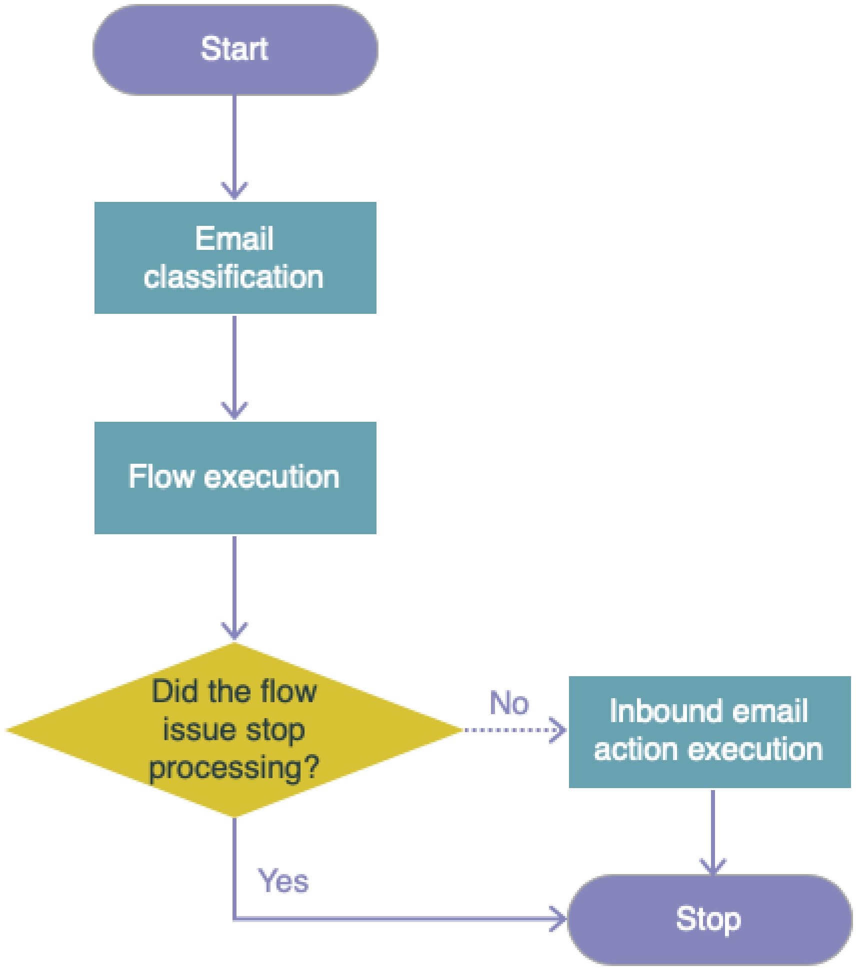

The benefits to automating system responses to inbound emails in Flow Designer are:

- Provides an easy and accessible interface that uses natural language.
- Consolidates configuration and runtime information into a single environment so process owners and developers can create, operate, and troubleshoot flows from a single interface.
- Reduces upgrade costs, with upgrade-safe ServiceNow AI Platform logic replacing complex custom script.
- Reduces development costs by providing a library of reusable actions.

For more information on creating inbound email flows in Flow Designer, see Flow Designer trigger types or follow the steps in Create a flow with an inbound email trigger. To view or manage your inbound email flows, navigate to **Flow Designer** > **Inbound Email Flows**.

Inbound email actions

Define an inbound email action and processing to script how the system responds to an inbound email.

Inbound email actions

Define an inbound email action to script how the system responds to an inbound email.

Note:
Inbound email flows take priority over inbound email actions. If you create flows with inbound email triggers, emails are first processed by the inbound email triggers before they are processed by inbound email actions.

Inbound email actions are similar to business rules: both use conditions and scripts that take action on a target table. An inbound email action checks the email for a watermark that associates it with a task and checks for other conditions. If the conditions are met, the system takes the inbound email action that you configure. The system can take two types of actions:

- Record action: setting a value for a field in the target table.
- Email reply: sending an email back to the source that triggered the action.

By default, if an email has no identifiable watermark, an inbound email action attempts to create an incident from the message. If the email has a watermark of an existing incident, an inbound email action updates the existing incident according to the action's script.

Watch this six-minute video to learn about the actions that the instance can take in response to messages from users and shows how to create or modify email notifications to users. Configuring inbound email actions and email notifications

Inbound email receive types

The system classifies all incoming email into one of three types: forward, reply, or new.

Inbound action classifications

Order	Type	Criteria
1	Forward	The system classifies an email as a forward only when it meets all these criteria: • The subject line contains a recognized forward prefix such as **FW:**. • The email body contains a recognized forward string such as **From:**. The system classifies any email that meets these criteria as a forward, even if the message contains a watermark or record number that otherwise classifies it as a reply.
2	Reply	The system classifies an email as a reply when it fails to match it to the forward receive type and it meets any one of these criteria: • The subject line or email body contains a recognized watermark such as **Ref:MSG0000008**. • There is no watermark and the Reply-To header contains a recognized message ID of an email with watermark. • There is no watermark and the subject line contains a recognized reply prefix such as **RE:** and a recognized record number such as **INC0005574** **Note:** The Reply-To header for the associated email should have a watermark associated with it, if there is no watermark then it will be classified as new.
3	New	The system classifies an email as new when it fails to match it to the forward and reply receive types.

Note:
From Paris release and beyond, if an email body has multiple watermarks, the last watermark in the email body is considered.

Determining the type of incoming email

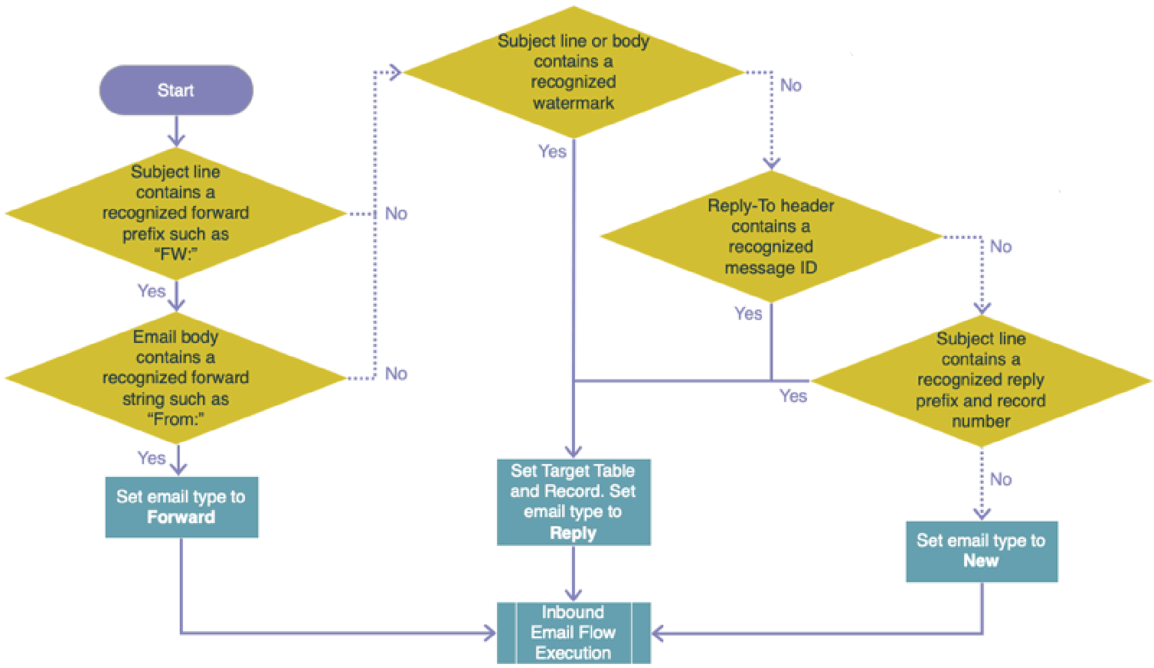

Attachments

If an inbound email contains one or more email attachments, the inbound email action adds the attachments to the first record the action produces.

Character encoding

- If the email encoding is ASCII-7 or UTF-8, inbound email actions preserve the character encoding in any associated task records they produce.
- If the email encoding is ISO-8859-1, the inbound email action attempts to convert the email to Windows 1252.
- Inbound email actions convert any other encodings (for example, Mac OS Roman) to plain text, which may or may not be readable.

See the System email log and mailboxes for examples of what you might see if a notification or inbound email action is not processed.

Note:
The state of all incoming emails that have been run against inbound email actions, even if there is no matching action, is changed to **Processed**.

Domain separation

The system ignores the domain that the inbound email action record is in when it creates a record based on the inbound email action. Keep inbound actions in the global domain. For example, if your inbound email action creates an incident, the system creates the incident in the same domain as the user in the **Caller** field. If that user is not in the User [sys_user] table, the incident is in the global domain.

Related topics

Notification variables

Watermarks on notification emails

Inbound email action processing

The system determines which inbound actions to run by comparing the inbound email type and inbound action conditions to the incoming email message. Certain properties are available to set the reply and forwarding prefixes in the email subject lines that your instance recognizes when processing inbound emails.

Note:
Inbound email flows take priority over inbound email actions. If you create flows with inbound email triggers, emails are first processed by the inbound email triggers before they are processed by inbound email actions.

The system follows this processing flow to determine whether to run an inbound action.

Inbound action processing work flow

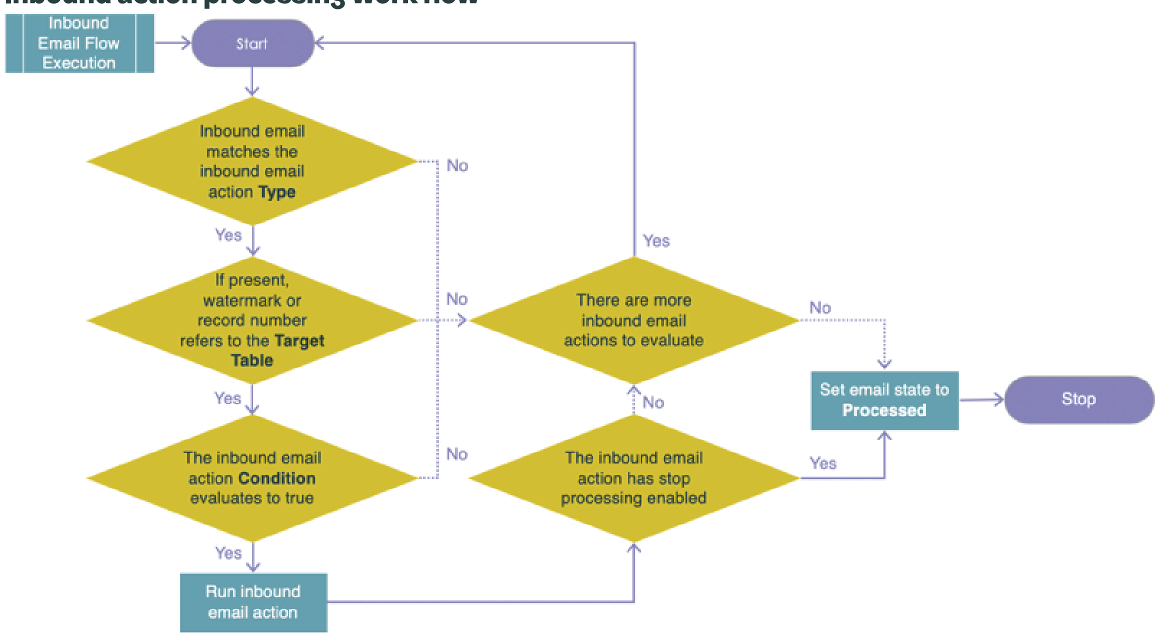

The system only runs an inbound action when:

- The incoming email type matches the inbound action **Type**.
- If present, the watermark or record number refers to a record in the **Target table**.
- The inbound action **Conditions** evaluates to true.

If any of these criteria are not met, the system skips the current inbound action and evaluates the next active inbound action. The system processes inbound actions from the lowest to highest **Order** value. If the inbound action has **Stop processing** enabled, the system updates the **State** of the email record to **Processed** after running the inbound action **Script**.

The following video shows how an inbound action condition prevents an incident from being created.

Troubleshooting when inbound email isn't processed

Prefixes recognized in email subject lines

Email reply prefixes

When no watermark is present or the In-Reply-To email header is present, the instance recognizes email containing a prefix from the *glide.email.reply_subject_prefix* property as reply email. You can use this property to set non-standard reply prefixes in your email system.

Property	Description
glide.email.reply_subject_prefix	Specifies the comma-separated list of prefixes in the subject line that identify an email reply. • Type: string • Default value: re:,aw:,r:,Accepted:,Tentative:,Declined: **Note:** The case of the reply prefix in the email, for example RE:, must exactly match the case of the prefixes defined in this property. If, for example, an email contains the Re: prefix and only RE: is defined in the property, the email will not be recognized as a reply. Therefore, it is a best practice to define multiple versions of the prefix, including mixed-case versions, such as RE:, Re:, and so on.

Email forward prefixes

Emails with certain prefixes trigger the forward type of inbound email action. The instance recognizes any email whose subject line contains a prefix from the *glide.email.forward_subject_prefix* property as forwarded email. Emails with these prefixes trigger inbound email actions of the type forward. Use this property to set non-standard forward prefixes in your email system or you want email forwards to behave like replies. If the value of the system property is empty, then the system reverts to using the values **fw:** and **fwd:**.

Property	Description
glide.email.forward_subject_prefix	Specifies the list of prefixes (comma-separated) in the subject line that identify a forwarded email. • Type: string • Default value: fw:,fwd: • Location: Add to the System Properties [sys_properties] table **Note:** Prefixes are case insensitive.

Email forwards as replies

Properties are available to force inbound actions to process forwarded mail as replied mail. These properties control the subject prefix that the inbound actions use.

Property	Value needed
glide.email.reply_subject_prefix	re:,Re:,RE:,aw:,r:,fw:,fwd:,Fwd:,FWD:
glide.email.forward_subject_prefix	[any text that is not a forward prefix]

These properties cause the Update Incident inbound action to process all forwarded and replied-to mail.

Note:
The *glide.email.forward_subject_prefix* property must contain some text so that the forwarded email can be processed as a Reply. It can be any text except a forward prefix (that is, fw:,fwd:,Fwd:,FWD:).

Matching a sender email address to a user

The instance matches a senders email address to an active user in the User [sys_user] table using inbound actions.

When processing an email, the instance sets the current user to the user whose email address matches **email.from**. Inbound actions can then reference that current user. For example, the base system inbound action Create Incident sets the **caller_id** of the incident to the value returned by `gs.getUserID()`.

If multiple users have the same email address, the instance first searches for an active user with the email address. The instance does not match inactive users.

Note:
Each user record must have a unique email address so that the instance can reliably match the email to the correct user.

If a unique email address for each user is not possible, assign a shared email address to only one active user so that the instance always matches incoming email from that address to the active user.

Matching watermarks in the Subject line or Body

The following examples illustrate how the instance matches randomized watermarks in an email subject line or body.

Note:
For instances upgraded from a release before Jakarta, the system can recognize both randomized and non-randomized watermarks during a watermark transition period.

Examples of matching watermarks in the Subject line or body

Subject Line or Body Contents	Matching Results
Ref:MSG0000008_ aLJc130zDhCVuh3spXmt	The instance recognizes this string as a watermark and searches the Email Watermarks [sys_watermark] table for a record with the number MSG0000008_

Examples of matching watermarks in the Subject line or body (continued)

Subject Line or Body Contents	Matching Results
	aLJc130zDhCVuh3spXmt. If this watermark exists, the instance matches the email to the associated record. If this watermark does not exist, the system processes inbound email messages as described in Criteria for matching email to inbound actions.
Ref:MSGWTR0000008_wfLLz42IxCgUvG2JIYnh	The instance recognizes this string as a watermark and searches the Email Watermarks [sys_watermark] table for a record with the number MSGWTR0000008_wfLLz42IxCgUvG2JIYnh. If this watermark exists, the instance matches the email to the associated record. If this watermark does not exist, the system processes inbound email messages as described in Criteria for matching email to inbound actions.

Matching record numbers in the Subject line or Body

The following examples illustrate how the instance matches record numbers in the subject line of an email to an existing record when no watermark is present.

Examples of matching record numbers in the Subject line

Subject line contents	Matching results
RE: Example INC0005574	The instance recognizes this subject line as a reply and recognizes the INC prefix as belonging to the Incident table. The instance searches the Incident table for an existing record `INC0005574`. If this incident exists, the email is associated with this incident. If this incident record does not exist, the instance uses the inbound action for new emails to create an incident and associates the new incident with the email.
RE: Example "INC0005574" RE: Example *INC0005574	The instance recognizes this subject line as a reply but does not recognize the `"INC` prefix as belonging to the Incident table because of the quotation mark. The same error occurs for any character other than a space before the record number. The instance instead uses the inbound action for new emails to create an incident and associates the new incident with the email.
RE: "Example INC0005574" RE: Example INC0005574*	The instance recognizes this subject line as a reply and recognizes the `INC` prefix as belonging to the Incident table. The instance searches the Incident table for an existing record `INC0005574"`, which it cannot find because of the quotation mark. The same error occurs for any character other than a space at the end of the record number. The instance instead uses the inbound action for new emails to create an incident and associates the new incident with the email.
RE: CHG0008593	The instance recognizes this subject line as a reply and recognizes one, but not both, of the number prefixes. There is no way to predict which prefix the instance matches first. Whichever prefix it matches, it searches the

Examples of matching record numbers in the Subject line (continued)

Subject line contents	Matching results
and INC000576	corresponding table for a matching record. If the record exists, the email is associated with the table. If the record does not exist, the instance uses the inbound action for new emails to create an incident and associates the new incident with the email. **Note:** The instance does not support processing email with multiple numbers in the subject line because there is no way to predict which record the instance matches first. For this reason, do not include more than one $number variable in your notifications.
FW: Example INC0005574	The instance recognizes this subject line as a forward because of the `FW:` prefix. It uses the inbound action for forwarded emails to create an incident and associates the new incident with the email.
Example INC0005574	The instance recognizes this subject as a new email because it does not contain a matching reply or forward prefix. It uses the inbound action for new emails to create an incident and associates the new incident with the email.

Criteria for matching email to inbound actions

The system matches incoming email to the conditions of the active inbound actions.

The default inbound actions create or update task record under these conditions.

Default matching criteria

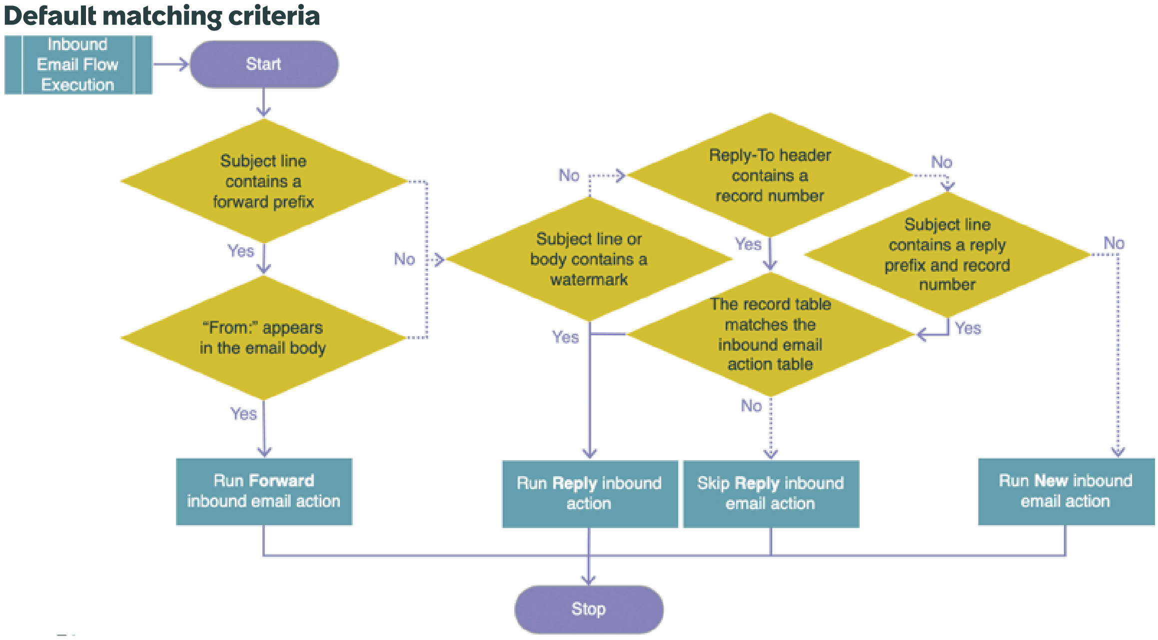

If you customize or deactivate the default inbound actions, the system checks the conditions of the active inbound actions. If the system cannot find an inbound action with matching conditions, it sets the state to **Processed**.

servicenow

Custom matching criteria

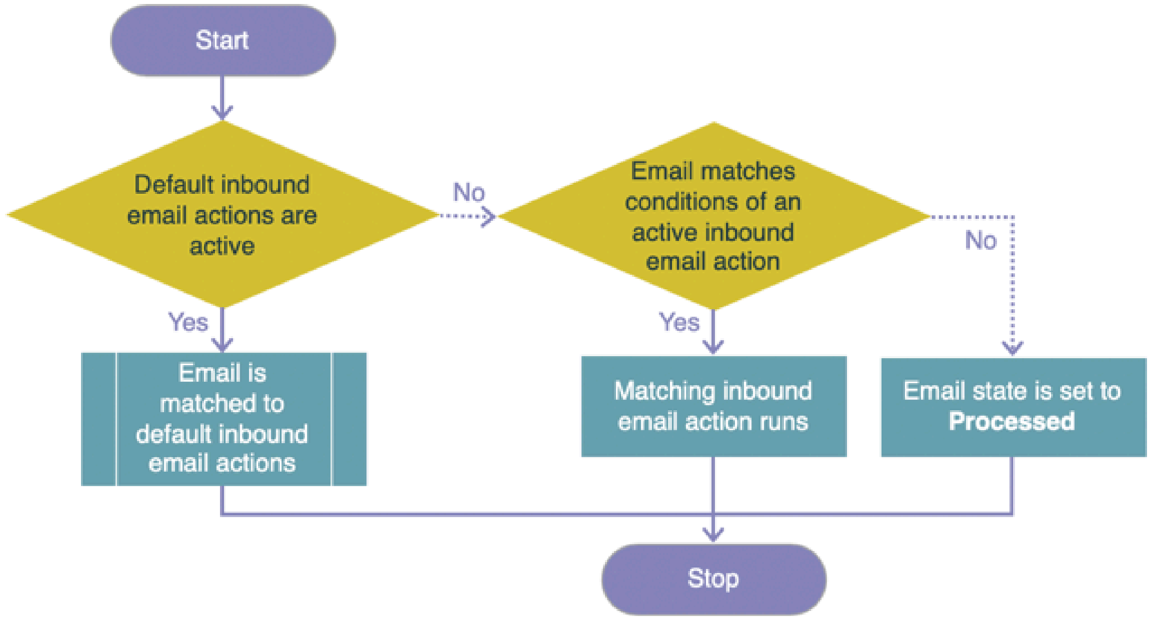

Inbound action type criteria

Inbound email action type	Required matching criteria	Name of default action (Incident table)	Result of default action
Forward	The email contains the following conditions: **1.** A subject starting with a recognized forward prefix (even if a watermark or an In-Reply-To header is present). **2.** From <user email> appears anywhere in the email body.	Create Incident (Forwarded)	Create new record
Reply	The email contains one of the following conditions and the table specified in the email matches the table of the inbound action: **1.** A valid watermark that matches an existing record. **2.** An In-Reply-To email header (when no watermark is present) that matches an existing record. **3.** A subject line starting with a recognized reply prefix (when neither a watermark nor an In-Reply-To header is present) and a valid record number that matches an existing record.	Update Incident (BP)	Update existing record
New	The email does not meet the conditions for either a reply or forward type inbound email action	Create Incident	Create new record

If more than one inbound action is available for a particular type, the instance uses the Table field to match the email to a particular table. If there is also more than one action for the inbound action's table, the instance uses the **Order** field to determine the order in which the actions run.

Configure inbound email actions

Configure a processing order for inbound email actions and use the order to manage multiple filters in inbound email actions.

Enable automatic user creation

An administrator can set an email property to automatically create users from incoming email. The administrator provides a list of trusted domains to prevent untrusted users from being automatically created.

Before you begin

Role required: admin

About this task

For example, you can prevent email from users outside your company domain from creating incidents. When an instance receives a message and there is no matching email address from the sender, the instance can create a user with the User ID [sys_user.user_name] set to the sender's full email address.
Users in your instance must still have write and update access to the records that they create or update through inbound email actions.

Note:
To learn more about the `glide.user.trusted_domain` property, which affects automatic email creation, see Restrict emails by domain for user creation (instance security hardening) in the Instance Security Hardening Settings.

Creating users from incoming email

Value of email.from Variable	User ID Created	Email Address	Name
new.user@company.com	new.user@company.com	new.user@company.com	New User
"New User" <new.user@company.com>	new.user@company.com	new.user@company.com	New User
"User, New" <new.user@company.com>	new.user@company.com	new.user@company.com	New User
"User" <nuser@company.com>	nuser@company.com	nuser@company.com	User

Procedure

1. Navigate to **All** > **System Properties** > **Email Properties**.
2. Select the check box for **Automatically create users for incoming email from trusted domains** (`glide.pop3readerjob.create_caller`).

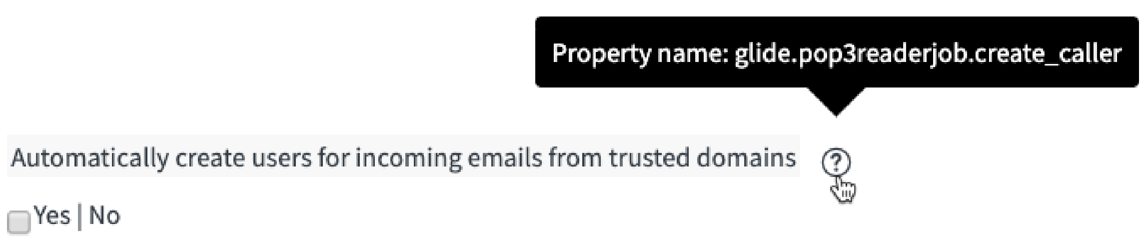

3. Enter the list of trusted domains in **Trusted domains for creating users from incoming emails** (`glide.user.trusted_domain`).

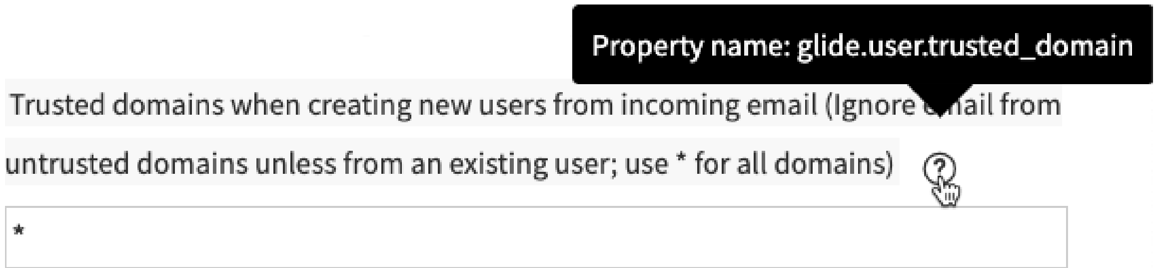

Note:
The *glide.user.trusted_domain* property prevents user creation if the sender is not from a trusted domain. However, the system may still process inbound actions for emails that are received from the domain. To have the system ignore these emails, set up a system address filter. For more information on setting up system address filters, see System address filters. You can also prevent untrusted users from triggering inbound actions by locking out the guest user.

4. Click **Save**.
5. **Optional:** Complete the following steps to lock out the guest user.
   a. Navigate to **User Administration** > **Users** and select the user **guest**.
   b. Select the **Locked out** field to disable the guest account.

What to do next

When the property `glide.pop3readerjob.create_caller` is set to `false`, the instance runs inbound actions from users who do not match an existing user by impersonating the guest user.

If the property `glide.pop3readerjob.create_caller` is set to `true`, but a user has a valid email address associated with a non-primary device, the instance creates a new user record for that email address if there is no matching email address in the Users [sys_user] table. The instance does not validate non-primary email addresses against the Notification Devices [cmn_notif_device] table.

The method the instance uses to create users can be upgraded to use the full email address by activating the Email Automatic User Creation plugin.

The plugin sets the property `glide.email.create_userid_from_email` to **true** by default. After activating the plugin, enable automatic user creation from email.

Note:
The `glide.email.create_userid_from_email` is not used by the system when the record is created using record producer or created manually.

Warning:
Review your existing user records to reconcile any that contain identical email addresses. If you activate the plugin prior to reconciling email addresses, your instance cannot distinguish between users with identical email addresses and randomly selects one of the users with the matching email address.

Setting field values from the email body

Values in an inbound email can set field values in a task record.

Any name:value pair in an inbound email body gets parsed into a variable/value pair in the inbound email script. The name:value pair must be on its own line. Note that most email clients

limit the number of characters allowed per line and may truncate excessively long name:value pairs.

Tip:
To prevent unexpected parsing, ensure that all the names in the name:value pairs are unique.

See Redirecting Emails for an example of using *setDisplayValue()* in an inbound email action.

Note:
The action always generates a lowercase variable name. Also, this functionality does not work on reference fields.

For example, if an email body contains this line:

```
Foo:bar
```

The inbound email script creates the variable *email.body.foo* with the value of **bar**. You can use these variables to create conditions such as:

```
if(email.body.foo!=undefined){
   current.[field]=email.body.foo;}
```

In this example, the script sets the value of *[field]* to the value **bar**.

Note:
Spaces are rendered as underscores when a name:value pair gets parsed into a variable/value pair. For example, if an email body contains a line with spaces like `my variable:data`, then the inbound email script creates the variable *email.body.my_variable*. The value of the variable is **data**.

Specifying the inbound email processing order

You can configure a processing order for inbound email actions and use the order to manage multiple filters in inbound email actions.

The Ordered Email Processing (com.glide.email_ordered_processing) plugin is enabled by default for new instances. Users with the admin role can activate the plugin for upgraded instances. The plugin adds the **Order** column to the Rules [sysrule] table, which the instance uses to determine when to process emails. Admins can also add a command to an action script that halts processing after the script runs.

Configure the processing order

Configure the processing order for inbound email actions to force them to run in a prescribed order.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Policy** > **Email** > **Inbound Actions**.
2. Open an existing inbound action or create one.
3. Complete the form and assign an order number to the **Order** field to establish when this inbound rule should run in relation to other rules.
   If you upgraded and activated the plugin, the **Order** field might be named **Execution Order**.

Note:
Ensure each inbound action has a unique **Order** value to ensure the system stops processing as expected. If multiple inbound actions have the same **Order** value, the system might evaluate all of the inbound actions, even if one of them contains the `event.state="stop_processing";` script or has the **Stop processing** option selected.

4. To stop rule processing when an inbound email action runs successfully: add the following line to the bottom of the script:
    - Select the **Stop processing** check box.
    - Add the following line to the bottom of the **Actions** script:

    ```
    event.state="stop_processing";
    ```

Manage multiple filters in an inbound email action

Use process ordering and the *`stop_processing`* command to manage multiple filters in inbound email actions.

Before you begin

Role required: admin

About this task

In this example, you can create new problem records when `prb:` appears in the subject line and new change requests when `chg:` appears in the subject line. All other emails are used to create an incident. The actions are set up as follows:

Procedure

1. Create an action with a condition of `Subject > starts with > chg:` and the `event.state="stop_processing";` command appended to the script.
2. Assign this action an Order value of `100`.
3. Create an action with a condition of `Subject > starts with > prb:` and the `event.state="stop_processing";` command appended to the script.
4. Assign this action an Order value of `200`.
5. Create an action for incident with no conditions and an Order value of `300`.

    It is not necessary to add the *`stop_processing`* command to the script for the incident action unless you want processing to stop at this rule to avoid continuing to another action.

Result

If either a change request or a problem is created, the *`stop_processing`* command stops processing, and no incident record is created. If neither a change request nor a problem is created, the inbound email action for incident creates a record.

Use Inbound email actions

Create inbound email actions and reprocess received emails.

Create an inbound email action

You can create inbound email actions to define the actions that the system takes when an email is received.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Policy** > **Email** > **Inbound Actions**.
2. Click **New**.
3. Fill in the fields as described in the table.

   Note:
   You might need to configure the form to see all fields.

Inbound Email Actions form fields

Field	Description
When to run	
Name	Enter a descriptive name for this email action.
Target table	Select the table where the action adds or updates records.
Action type	Select the type of action the instance takes. Select **Record Action** to modify a record in the instance, or select **Reply Email** to have the instance send an email back to the source of the inbound email.
Active	Select the check box to activate the inbound email action. Clear the check box to disable the action.
Stop processing	Select this check box to prevent the system from running additional inbound email actions after this action runs.
Type	Select the message type required to run the action. The action runs only if the inbound email is of the selected type. Available types are: ◦ *New*: An email that is not recognized as a reply or forward. ◦ *Reply*: An email with a watermark with an In-Reply-To email header, or whose subject line begins with a recognized reply prefix. ◦ *Forward*: An email whose subject line begins with a recognized forward prefix, even if the email also contains a watermark or In-Reply-To header. **Note:** By default, inbound emails of the *Forward* type always generate new incidents regardless of the presence of a watermark. If this behavior does not match your business logic, you can change the recognized reply and forward prefixes to treat forwards like replies.
Required roles	Specify required roles the sender must have to trigger the inbound action.
Order	Enter a number that specifies when this inbound action runs relative to other inbound actions that use the same target table. The instance processes the action with the lowest order number first.
From	Select the user required to run the action. If a user is selected, the action runs only when the email sender matches the user name. Leave this field blank to perform the action for all users.

Field	Description
	Note: If the selected user is later archived or deleted, the restriction is removed and anyone can trigger the inbound email action.
Condition	Specify the condition that must evaluate to true to trigger the inbound action. Build a condition with the choice lists or enter a statement that determines when the inbound email action runs. For example: `email.subject. startsWith ( "chg:" )`
Actions	
Field actions	This field appears if the **Action type** is **Record Action**. Specify how information in the email is inserted into the record. For example, select **[Created by] [From email] [Sender]**, so that when the request is inserted, you can see who it is for.
Reply email	This field appears if the **Action type** is **Reply Email**. compose the email message to send to the source that triggered the inbound email action.
Script	Enter the script the action runs. Typically, this script uses the validators script include and email variables. A template is provided: `(function runAction(/*GlideRecord*/ current, /*GlideRecord*/ event, /*EmailWrapper*/ email, /*ScopedEmailLogger*/ logger) {` `// Implement email action here` `})(current, event, email, logger);` The following objects are available: ◦ **current**: access the record referred to by the inbound email. For example, `current.assigned_to` accesses the person assigned to the task. ◦ **event**: access one of the parameters of the originating event. For example, `event.parm1` accesses the first parameter of the event or `event.parm2` for the second parameter. See Events for more information. ◦ **email**: access the inbound current email record. For example, `email.subject` accesses the content in the subject line of the email. See Accessing email object variables for more information. ◦ **logger**: add a message to the log file with the source set to `email.<Sys ID of incoming email>`. For example: `logger.log ("Some information")`
Description	
Description	Enter a detailed explanation of what this inbound email action does.

Field	Description
Other fields	
Order	Enter a number to define the order in which this email action should be processed. Actions with lower numbers are processed first. The Ordered Email Processing plugin does not install this field.

Reprocess received emails

When a email has been processed incorrectly, fix the underlying problem, then reprocess the email. You can reprocess emails individually or in a batch.

Before you begin

Role required: admin

About this task

Various scenarios can result in improperly processed emails. After taking corrective action, reprocess the emails to get correct email records. Two of the most common scenarios are:

- The inbound action failed to trigger correctly. After you update or fix the inbound action, reprocess the affected emails to create the correct records.
- Emails have accidentally ended up on a development instance, usually due to the username or password of a production or test instance being reused on the development instance. In this case, you first export the emails from the development instance as XML and import them to the correct instance. On the correct instance, reprocess the imported emails to obtain correct email records.

You can reprocess emails with the types `received` or `received-ignored`. Such emails are processed to the system mailboxes **Received** and **Junk**, respectively.

Take the following steps only after fixing the underlying problem. This procedure only reprocesses the emails.

Procedure

1. Navigate to the relevant **System Mailbox** or navigate to **All** > **System Logs** > **Email**.
2. Decide whether to reprocess emails individually or in a batch.
   - To reprocess an individual email, open that email and click **Reprocess email**.
   - To reprocess several emails at once, select them from the mailbox or the email log and apply the **Reprocess received emails** action.

   Remember that you can reprocess only `received` or `received-ignored` emails.

References for Inbound email

Email object variables, user impersonations and inbound actions and examples for processing inbound email actions.

Accessing email object variables

An inbound email action script contains the email object to access various pieces of an inbound email through variables. You can use the global variable *`sys_email`* with inbound email actions.

Accessing email objects with variables

Variable	Contents
email.to	Contains a comma-separated list of email addresses in the To: and Cc: boxes.
email.direct	Contains a comma-separated list of email addresses in the To: box.
email.copied	Contains a comma-separated list of email addresses in the Cc: box.
email.body_text	Contains the body of the email as a plain text string.
email.body_html	Contains the body of the email as an HTML string.
email.from	Contains an email address that depends on the following conditions: • If the address listed in the email **Headers** field matches an existing user's **Email address**, this variable contains the user's email address. • If the address listed in the email **Headers** field does not match an existing user's **Email address**, this variable contains the address listed in the email **Headers** field.
email.from_sys_id	Contains the Sys ID of the user who sent the email to the instance.
email.fromAddress	If system property *glide.email.inbound_action.extract_from_hea* property is set to true, *origemail* is computed from the headers. The default value is false if the property does not exist.
email.origemail	Contains the address of the email sender as listed in the email **Headers** field.
email.subject	Contains the subject of the email as a plain text string.
email.recipients	Contains a comma-separated list of recipient addresses as a plain text string, in the To: box.
email.recipients_array	Contains the recipient addresses as an array.
email.content_type	Contains the MIME content type of the email (for example, *text/plain; charset="us-ascii"* or *text/html; charset="us-ascii"*).
email.headers	Contains details about the sender, route, and receiver as a plain text string in the format of the sending email client.

Accessing email objects with variables (continued)

Variable	Contents
email.importance	Contains an indication from the sender about how important a message is. The value can be **High**, **Low**, or empty.

Note:
The instance follows RFC 2822 (Internet Message Format), which requires multiple email addresses in a group to be separated by commas, not semicolons. The instance can set the values of the *email.to*, *email.direct*, and *email.copied* variables only if emails addressed to groups follow the expected RFC format.

Inbound email.recipient variables

The recipients variables (*email.recipients*, *email.recipients-array*) allow processing of inbound email based on the email recipients. For example, you can create a script to process email based on the array values:

```
var rarray  = email.recipients_array ; for ( var i  = 0 ; i  <
 rarray.length ; i ++ ) { var recipient  = rarray [i ] ; // do
 something with it }
```

The sys_email variable

This variable lets you access the received sys_email record that triggered the inbound email action. It can be used to reference fields on the email record, such as **uid**, **sys_id**, **content_type**, and so on.

Email user matching

When the instance receives an email message, the system searches for an existing user record with the same email address as the sender.

Matching email to existing users

Value of email.from Variable	Matching User ID	Email Address	Name
michael.tossi@company.com	michael.tossi@company.com	michael.tossi@company.com	Michael Tossi
"Michael Tossi" <michael.tossi@company.com>	michael.tossi@company.com	michael.tossi@company.com	Michael Tossi
"Tossi, Michael" <michael.tossi@company.com>	michael.tossi@company.com	michael.tossi@company.com	Michael Tossi
"Tossi" <mtossi@company.com>	mtossi@company.com	mtossi@company.com	Tossi

Note:
This functionality requires that you activate the Email Automatic User Creation plugin.

Inbound Email Action scripts no longer support the `gs.createUser()` method. Use either the automatically-generated email variables or the `gs.GetUserID()` method instead.

User impersonations and inbound actions

When the instance receives an email, it can take a variety of actions by impersonating the sender.

If the sender of an incoming email matches an existing user, the instance impersonates the matching user to complete any inbound email actions. If the sender does not match an existing user, the instance impersonates the Guest user to complete any inbound email actions. If the impersonated user is locked out , the inbound email action fails.

Note:
If inbound email comes from an untrusted domain, the instance impersonates the Guest user unless you explicitly prevent users from untrusted domains from triggering inbound actions. For more information on filtering domains, see Designate untrusted and trusted email domains.

Allowing locked out users to process inbound email actions

A property is available to allow locked out users to trigger inbound actions. For example, enabling the property can allow locked out users to reset their password and send email to the instance asking for assistance.

Property allowing locked out users to trigger inbound email actions

Property	Description
glide.pop3.process_locked_out	Enables (true) or disables (false) the ability for locked out users to trigger inbound actions. • Type: true | false • Default value: false • System Properties [sys_properties] table

Warning:
Enabling this property (*glide.pop3.process_locked_out*) also enables users from untrusted domains to trigger inbound actions. The user must still be active.

Redirecting email to the instance POP3 account

Configure other mailboxes to forward email to the instance's POP3 account.

By default, the *POP Reader* scheduled job checks for new email every two minutes. It connects to the mail server and account specified in email properties. The *POP Reader* downloads any email waiting on the mail server and creates *email.read* events. After the instance processes the events, the inbound email actions run.

The *POP Reader* shows the number of emails processed during the reader's last run. The message shows the number of emails the reader processed or `0 processed` if no emails were available. The reader resets the status each time it runs.

It is not possible to specify more than one POP3 account for an instance, but you can forward other mailboxes to the designated POP3 account. This script can be added to the Create Incident inbound email action to differentiate the content based on the original recipient, and then set an *assignment_group* value.

There can be multiple active POP3 accounts with different email addresses but only one instance POP3 account.

```
if(email.direct.indexOf('facilities@anycorp.com')>-1)
current.assignment_group.setDisplayValue('Facilities
 Management');
```

Inbound email action examples

Various examples of inbound email actions are available to help you build your own inbound email actions. These examples show how to set up inbound email actions to handle email replies, create (log) a problem record, request a change, and update an incident.

Inbound email action example: handling email replies

This example shows you how to set up inbound email actions to handle replies that users send back to the instance.

Before you begin

Role required: admin

About this task

The inbound email action parses the email and responds using a script. By default, an email received by the instance creates a new incident, and the body of the email is added to the **Additional Comments** text box. More refined Inbound Email Actions can create incident tickets with more data, thus saving the incident management team valuable time.

Normally, when a user responds to an email sent by the instance, the inbound email action matches the watermark to an existing incident, and updates the incident rather than creating a new record. However, if the watermark is missing, this inbound email action attempts to match a reply to the original incident.

This inbound email action is replicated in Flow Designer as the sample flow **Inbound Email Flow Example: handling email replies**. To view the sample flow, navigate to **Flow Designer** > **Designer**.

Procedure

1. Navigate to **System Policy** > **Inbound Actions** and click **New**.
2. Populate the form as follows:

Inbound action field values

Field	Value
Name	Update Incident
Type	Reply
Target table	Incident [incident]

3. In **Script**, enter this code.

```
gs.include('validators');

//Note: current.caller_id and current.opened_by are already set
 to the first UserID that matches the From: email address

if (current.getTableName() == "incident") {
  current.comments = "reply from: " + email.origemail + "\n\n"
 + email.body_text;

  if (email.body.assign != undefined)
```

```
    current.assigned_to = email.body.assign;

  if (email.body.priority != undefined &&
 isNumeric(email.body.priority))
    current.priority = email.body.priority;

  if (email.body.category != undefined)
    current.category = email.body.category;

  if (email.body.short_description != undefined)
    current.short_description = email.body.short_description;

  current.update();
}
```

Inbound email action example: logging a problem

This example shows you how to set up inbound email actions to a create a problem record.

Before you begin
Role required: admin

About this task
Inbound email actions allow users to log or update incidents on an instance via email. The inbound email action parses the email and responds using a script.

This inbound email action is replicated in Flow Designer as the sample flow **Inbound Email Flow Example: logging a problem**. To view the sample flow, navigate to **Flow Designer** > **Designer**.

Procedure

1. Navigate to **System Policy** > **Inbound Actions** and click **New**.
2. Populate the form as follows:

Field	Entry
Name	Log Problem
Type	New
Active	True
Target Table	Problem [problem]
Condition	`email.subject.indexOf("Problem: ") == 0`
Script	`current.description = email.body_text; current.short_description = email.subject.toString().substring(9);` `current.assignment_group.setDisplayValue("Development");`

Field	Entry
	``if (email.body.assign != undefined) current.assigned_to = email.body.assign; current.insert();```

Inbound email action example: requesting a change

This example shows you how to set up inbound email actions to create a change request record.

Before you begin

Role required: admin

About this task

Inbound Email Actions allow users to log or update incidents on an instance via email. The inbound email action parses the email and responds using a script.

Procedure

1. Navigate to **All** > **System Policy** > **Inbound Actions** and click **New**.
2. Populate the form as follows:

Field	Entry
Name	Request Change
Type	New
Active	True
Target Table	Change Request [change_request]
Condition	`email.subject.indexOf("Change Request: ") == 0`
Script	```current.comments = email.body_text; current.short_description = email.subject; current.notify = 2; if (email.body_text.assign != undefined) current.assigned_to = email.body_text.assign; if (email.body_text.priority != undefined) current.priority = email.body_text.priority;```

Field	Entry
	`if (email.body_text.category != undefined) current.category = email.body_text.category; current.insert();`

Values automatically set from incoming email

The default inbound action for the Incident table automatically sets the following field values when it receives an incoming email.

Values automatically set from incoming email

Field value set	Value used from incoming email
`current.caller_id`	User ID of the first user whose email address matches the `email.from` variable.
`current.opened_by`	User ID of the first user whose email address matches the `email.from` variable.

If multiple users have the same email address, the instance first searches for an active user with the email address. Use unique email addresses for each user record whenever possible. If not, having only one active user with the shared email address guarantees that the instance always matches incoming email from this address to the active user.

Integrate inbound events

This example illustrates how to create a notification from an inbound JSON request.

Before you begin

Role required: admin

About this task

When complete, you will be able to:

- Send a JSON request to the [imp_notification] web service import set with the JSON processor
- Create a new import set in the [imp_notification] table in the instance using data from the JSON request

The following example steps assume you have your own demonstration instance.

Procedure

1. Activate the JSON Web Service plugin.
2. Install the RESTClient Firefox plugin.
3. Open the RESTClient.
4. Create the following JSON request.
   - **Method**: *POST*
   - **URL**: `http://<instance name>.service-now.com/imp_notification.do?JSON`

- **Headers**: Authorization: Basic
- **Body**:

```
{"sysparm_action":"insert","message":"this is an
event","uuid":"abc"}
```

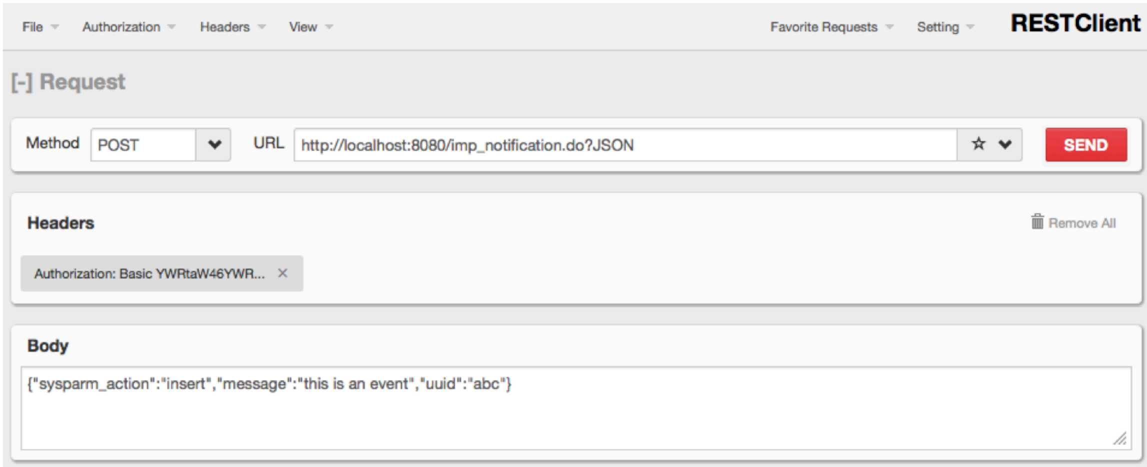

5. Click **Send**.
6. Navigate to **Response** > **Response Body (Raw)**.
7. Verify that the instance sends back a response with a `sys_id`.

8. Login to your development instance.
9. In Navigation filter, enter `imp_notification.list`.
10. Verify that the import set table has an event matching your JSON request.

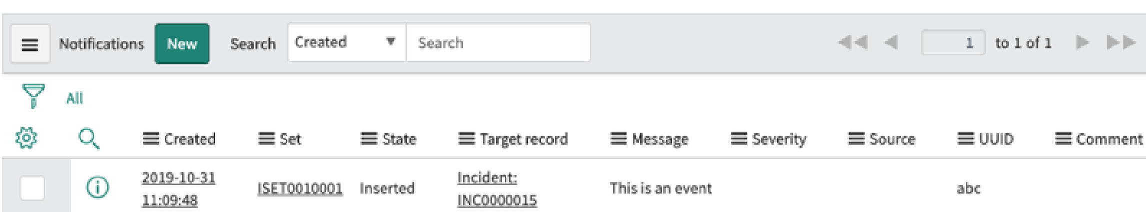

Notification Preferences

With notification preferences you can manage and choose the notifications you want to receive.

Get started

Explore	Notification Preferences in Next Experience	Notification Preferences in Core UI
Learn about preferences and choose the notifications you want to receive	Manage and set notification preferences, including customized notifications and channels	Set notification preferences, including personal subscriptions and channels

Troubleshoot and get help

- Troubleshooting Notification emails
- Ask or answer questions in the Platform community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Notification preferences

With notification preferences you can manage and choose the notifications you want to receive.

Notification preferences enable you to set global preferences, search for notifications, enable and disable notifications, and customize the way you receive system notifications.

Notification preferences for Next Experience

You can manage your own notification preferences for Next Experience, including personal notifications and channels for receiving them. With **Advanced Preferences** managing notification preferences in Next Experience is more efficient and also supports provider notifications.

For more information, see System and custom notification and delivery channel preferences in Next Experience.

Notification preferences for Core UI

In the Core UI, you set personal notification preferences, including personal subscriptions and channels for receiving them, through the **Notifications** tab of the System Settings window.

> **Note:**
> To switch to the Core UI from the Next Experience UI, set the property `glide.ui.polaris.core.notification_preference.enabled` to true.

For more information, see Setting notification preferences in Core UI.

System and custom notification and delivery channel preferences in Next Experience

You can manage and set your own notification preferences, including customized notifications and channels for receiving them.

With notification preferences, you can control the system notifications and custom notifications that you receive and apply conditions to restrict notification delivery and choose how the notifications are delivered to you.

System notifications

System notifications alert you of record changes that might concern or interest you. You can be notified via email, SMS text messages, push notifications, or messaging applications.

With system notifications, you can:

- Search for a specific system notification in the list of your notifications to be able to set your preferences for that notification.
- Set preferences at a global level.
- Easily enable or disable all notifications or individual notifications.
- Add or modify a notification schedule for individual delivery channels for each system notification.
- Set a schedule for all notifications, individual system notifications, or delivery channels.

  For more information, see Apply notification conditions

Custom notifications

A custom notification is a system notification that you can subscribe to and customize by adding conditions.

With custom notifications, you can:

- Create and modify custom notifications, which are subscriptions to notifications that are important to you.
- Set preferences at a global level.
- Easily enable or disable all notifications or individual notifications.
- Enable or disable notification channels.
- Add or modify a notification schedule for individual delivery channels for each system notification.
- Set a schedule for all notifications, individual system notifications, or delivery channels.

  For more information, see Apply notification conditions

Note:

Custom notifications for Next Experience do not support notifications that require a related affected record. For these notifications, you have to set the preference in the Core UI. For more information, see Setting notification preferences in Core UI.

Note:
Subscription-based notifications are not domain aware and cannot support domain-specific settings.

Delivery Channels

Delivery channels is a list of your chosen channels such as email, SMS, instant message and voice channels for receiving notifications.

With Delivery Channels, you can:

- Enable or disable a particular channel for receiving notifications.
- Create, edit, or delete notification channels for instant messages, email, SMS, and voice.

  Note:
  Only admins can delete channels from the cmn_notif_device table.

- Choose how notifications are delivered.
- Set a schedule for the record in the channel and apply to the notifications you want to receive.

Set notification preferences

View, enable, and disable general settings for notification preferences, and delivery channels.

Before you begin

Role required: none

Procedure

1. From your user profile, navigate to **Preferences** > **Notifications**.
2. Select the **General** tab.
3. Set global notification preferences.

Customize system notifications

Create custom notifications, which are subscriptions to notifications of importance to you, and apply conditions that control specific content included in your custom notification. You can also enable or disable the channels for delivery.

Before you begin

Role required: none

About this task

Note:
Custom notifications for Next Experience do not support the setting for the affected record the notification is based on. For notifications that require an associated affected record, see Create personal notifications.

Procedure

1. From your user profile, navigate to **Preferences** > **Notifications** and then select **Advanced Preferences**.
2. Select the **Custom Notifications** tab.
3. Select **Create notification**.
4. Select the notification you want to receive.

   a. Enter a name for the notification in the **Personal notification name** field.
   b. Select the system notification you want to receive from the **Notification** drop-down.
   c. Select **Next**.
5. Select the channels from the delivery channels through which you want to receive the notifications.
6. Select **Next**.
7. Set a schedule or conditions for receiving notifications.
   a. Choose to have notifications delivered only during weekday working hours by selecting **Monday thru Friday (9 to 5)**.
   b. **Optional:** Set a filter with as many conditions as desired to limit the notifications you receive.

      (Optional) For example, you can choose to be notified only when an incident communication plan is created for a specific configuration item.

      Note:
      All conditions are set to **[Active][is][true]** by default.
8. Select **Save**.

Modify a notification

Modify a notification by establishing a schedule or setting conditions to control the notifications you receive.

Before you begin

Role required: none

Procedure

1. From your user profile, navigate to **Preferences** > **Notifications** and then select **Advanced Preferences**.
2. Determine the type of notification you want to modify by selecting either the **System Notifications** or **Custom Notifications** tab.
3. Access the Add/Edit Schedule dialog box.
   - In the **System Notifications** tab, select and expand the notification and select **Schedule**.
   - In the **Custom Notifications** tab, select **Schedule** in the notification entry.
4. Choose to have notifications delivered 24x7 by selecting **None** or only during weekday working hours by selecting **Monday thru Friday (9 to 5)** in the **Schedule** field.
5. **Optional:** Set a filter with as many conditions as desired to limit the notifications you receive.
   - In the **System Notifications** tab, select the **Advanced Filter** checkbox and set the filter in the **conditions** field.
   - In the **Custom Notifications** tab, set the filter in the **conditions** field.

   (Optional) For example, you can choose to be notified only when an incident communication plan is created for a specific configuration item.

Note:
All conditions are set to **[Active][is][true]** by default.

6. Select **Save**.

Delete a custom notification

Delete custom notifications that you don't need any more.

Before you begin
Role required: none

Procedure

1. From your user profile, navigate to **Preferences** > **Notifications** and then select **Advanced Preferences**.
2. Select the **Custom Notifications** tab.
3. In the notification entry, select the more actions icon (⋮) and select **Delete**.

Add a new notification delivery channel

Add new channels for email, instant message, SMS, and voice to receive notifications and set a schedule for the channel for the notifications you want to receive.

Before you begin
Role required: none

Procedure

1. From your user profile, navigate to **Preferences** > **Notifications** and then select **Advanced Preferences**.
2. Select the **Delivery Channels** tab.
3. In the relevant delivery channel type, select **+Add**.
4. In the dialog box, select the **Channel Info** entry.
5. Enter a channel name and the relevant details.
6. Select **Set Schedule** and select a schedule from the **Schedule**drop-down list to schedule notifications for the channel.
7. Select **Save**.

Edit a delivery channel

Edit an email, instant message, SMS, or voice channel for receiving notifications, and schedule when the channel can receive notifications.

Before you begin
Role required: none

Procedure

1. From the user profile, navigate to **Preferences** > **Notifications** and then select **Advanced Preferences**.
2. Select **Delivery Channels**.
3. Expand the delivery channel that you want to edit.

4. Select the more actions icon for the channel and select **Edit**.
5. Make the changes to the channel name, email address, or schedule.
6. Select **Save**.

Setting notification preferences in Core UI

You can set your own notification preferences, including personal subscriptions and channels for receiving them. All users can set these preferences through the **Notifications** tab of the System Settings window.

With the **Notifications** tab, you can:

- Search for a specific notification in the list of your notifications.
- Use a global switch to enable or disable all your notifications.
- Enable or disable a particular channel for receiving notifications, as well as create, edit, or delete channels.
- Control the notifications that you receive and apply conditions to restrict notification delivery.
- Create personal notifications, which are subscriptions to notifications that are important to you.

Comparison between notification preferences in Core UI and earlier interfaces

In new and upgraded instances that use the Core UI interface, the System Settings window is the central location for managing your notification preferences for Next Experience. The Notification Preference User Interface plugin (com.glide.notification.preference.ui) is activated by default.

Notification preferences in Core UI and UI15

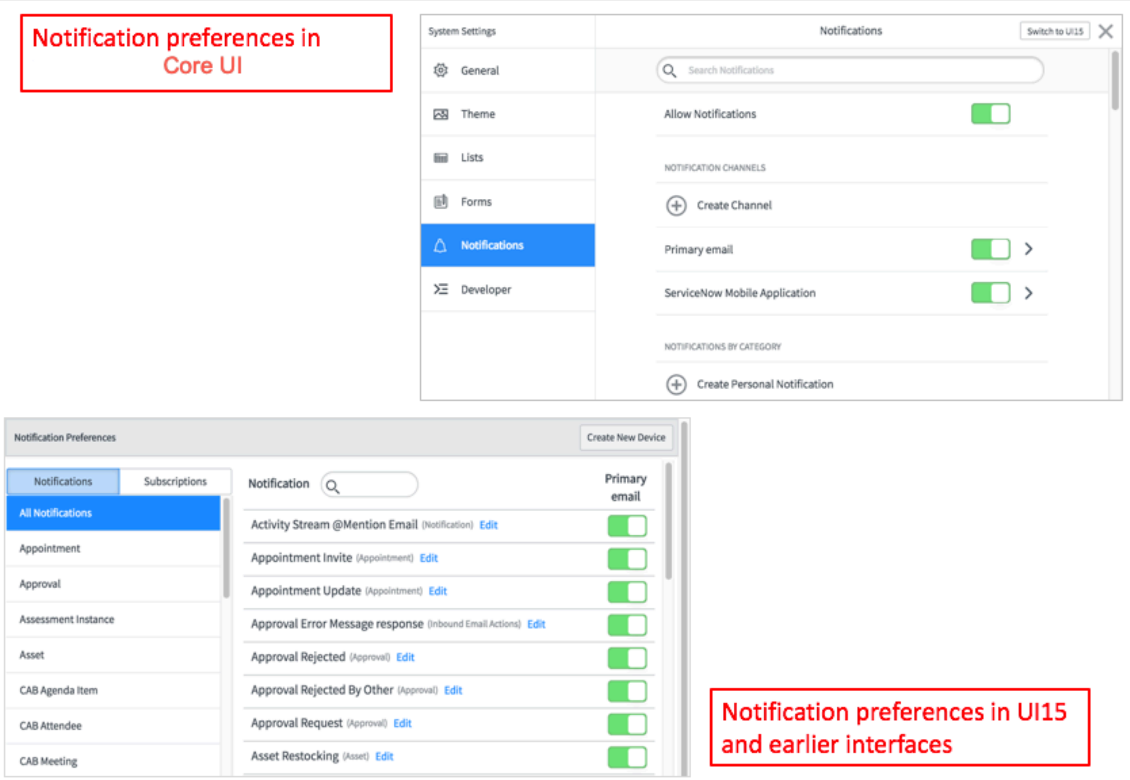

In UI15 and earlier interfaces:

- Users set notification preferences through the **Notification Preferences** link in the User form (**Self service** > **My Notification Preferences**).
- Admins created user devices and set notification preferences for users through the **Notification Preferences** link in the User form, accessed through the User Administration module (**User Administration** > **Users**).

> **Note:**
>
> Admins can revert to the UI15 notification preferences interface by setting the *`glide.notification.preference.UI.enabled`* property to false.

Manage your notification preferences in Core UI

> **Note:**
>
> To switch to the Core UI from Next Experience, set the property *`glide.ui.polaris.core.notification_preference.enabled`* to true.

To access your notification preferences, click the gear icon () in the banner frame, and in the System Settings window, click the **Notifications** tab. You can also access the System Settings by navigating to **Self Service** > **My Notification Preferences**.

Notifications tab in System Settings window

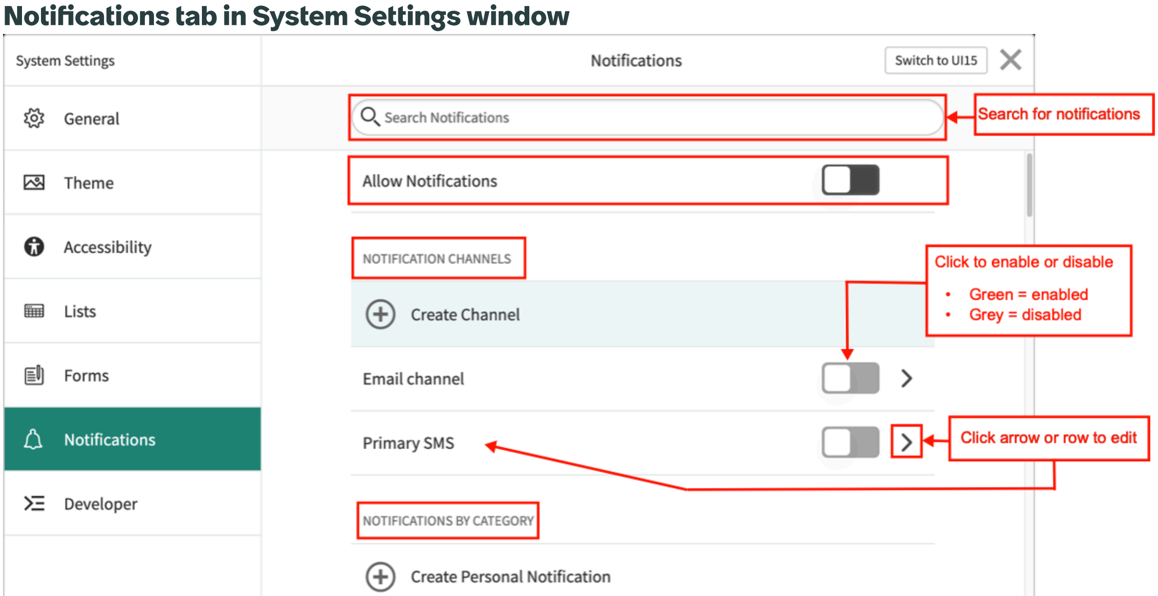

In the **Notifications** tab, use the settings in the following table to manage your notifications.

Setting	Description
Search	Search bar for finding notifications. To find a notification: • Enter the name of the notification you are searching for. When you type the first three characters, the system automatically returns a list of notifications matching the characters entered. • Enter two asterisks to return a list of all notifications in alphabetic order (the notifications are not organized by category).
	Global switch for enabling or disabling all notifications.

Setting	Description
Allow notifications	
Notification channels	List of your channels for receiving notifications. Use this section to: • Enable or disable a channel by using its toggle switch. If you disable a channel, notifications are not delivered through that channel. • Add channels using the **Create Channel** option. • Modify channel information. Click the channel row or the right arrow (**>**) next to the appropriate channel to edit or delete it.
Notifications by category	List of notification categories that identify and group related notifications. Each category contains the notifications that you can subscribe to. To view the notifications in a given category, click the category row or the right arrow (**>**) next to the appropriate category. You can: • Edit settings for a notification: ◦ Enable or disable channels for the selected notification. ◦ Apply notification conditions, such as schedules and filters, that affect the delivery of the notification. • Create personal notifications, which are subscriptions to specific notifications that matter to you.

Note:
Subscription-based notifications are not domain aware and cannot support domain-specific settings.

Subscription-based notifications

Subscription-based notifications enable users to proactively subscribe to items that interest them and unsubscribe from messages that are not mandatory.

Users can also specify additional notification channels that each of their notifications can be configured to use.

Before users can manage the notifications that are sent to them, administrators must create email notifications to which users can subscribe. Administrators can also make subscription-based notifications mandatory so users cannot unsubscribe to them. Then users can subscribe or unsubscribe to the notifications, and add schedules and filters to the subscription to limit the notifications that can be received.

Notifications that administrators mark as subscribable are automatically available in user notification settings. Users are limited to one subscription per notification.

Administrators should create subscription-based notifications when they do not want to specify users for a notification and want to let users proactively subscribe to the notification.

Note:
Subscription-based notifications are not domain aware and cannot support domain-specific settings.

Subscriptions 2.0 plugin

The Subscription Based Notifications 2.0 plugin must be active to use subscribable notifications. This plugin is active by default on all new and upgraded instances.

The plugin installs the Notification Subscription [sys_notif_subscription] table, which holds user subscriptions to all notifications.

Subscriptions and notification preferences

Users can subscribe to notifications available to them through their notification preferences. In the Core UI interface, all users can set and modify their notification preferences through the **Notifications** tab of the System Settings window.

Create personal notifications

You can create personal notifications, which are subscriptions to notifications of importance to you. You can apply conditions that control specific content included in your personal notification, and also enable or disable the channels for delivery.

Before you begin
Role required: none

About this task
A personal notification is a subscribable notification that you can customize by adding conditions. You can assign the notification a name that is meaningful to you and filter the information received in that notification.

Note:
Users are limited to one subscription per notification.

Procedure

1. Click the gear icon ( ) in the banner frame to open the System Settings window, and click the **Notifications** tab.
2. Click **Create Personal Notification**.
3. Complete the New Personal Notification form (see table).

Field	Description
Name	A descriptive name for your personal notification. Use a meaningful name to distinguish it from system notifications.
Notification	The notification to subscribe to. You can subscribe only to notifications that are configured to allow subscriptions. Depending on the notification you select, the **Table**, **Affected record**, and **Send when** fields are displayed.
Table	The database table that the notification is linked to.
Affected record	The specific record that the notification is based on. This field appears when you select a notification that has an affected record (for example, a notification for a problem record that has multiple affected CIs). Click the lookup icon, select the table and specific document (record) in that table, and click **OK**.
Send when	Another condition that must be met to send the notification. For example, you might create a filter whose conditions send notifications when an incident with a priority of **1 - Critical** is opened for a network issue.

4. **Click Save.**
   In the following example, notice the category in which the system placed the personal notification.
   The system creates the personal notification and adds it to the appropriate notification category.
5. Enable or disable the channels for your personal notification.
6. Navigate back to the list of notifications for the category by clicking the left back arrow (**<**) at the top of the form.
   Notice that your Personal Notifications are listed above the System Notifications in the given category.
7. Navigate back by clicking the left back arrow (**<**) at the top of the form or by clicking the **Notifications** tab.

Create a service provider

Administrators can configure service providers for devices that use SMS.

Before you begin

Role required: admin

About this task

Administrators also have the option of configuring how a device's service provider affects the construction of the device's email address.

Procedure

1. Navigate to the Notification Service Provider [cmn_notif_service_provider] table by typing `cmn_notif_service_provider.list` in the application navigator filter.
2. Fill out or change the fields on the form (see table).
3. Click **Update**.
4. Click **Save** on the Notification Preferences page.

Notification Service Provider form fields

Field	Description
Name	Configured name of the service provider.
Type	Type of device, in this case **SMS**.
Active	Enables or disables this notification device.
Advanced notification	Removes all the previous options and displays the **Advanced script** field.
Advanced script	Used for custom notifications that run a script rather than construct a traditional SMS/email. This is generally used when all outbound SMS messages must run through a central SMS hub, as opposed to being sent directly from the instance to the SMS provider. Advanced users can construct a script in this field that will send a notification to an old style numeric pager that cannot receive SMS communications. This field is rarely used and employs advanced scripts. Contact your representative to assist you with any advanced SMS scripting.
Construct address manually	Removes the prefix and suffix options and displays the **Construction script** field.

Field	Description
Construction script	Allows you to script the email address construction as you would in a business rule. For example, `abc + current.phone_number + def@text.att.net` would construct an email address of `abc2223334444def@text.att.net`. **Note:** `Current` is a reference to the device, not the service provider, hence the *`current.phone_number`* variable that uses the device's `phone_number` value.
SMS Provider Email Prefix	Places the provided text before the device's specified phone number; for example: `PREFIX2223334444@text.att.net`
SMS Provider Email Suffix	Places the provided text after '@' sign; for example: `2223334444@SUFFIX`
Notification Device Variables	Optional attributes of an SMS device used inside an SMS service provider's **Advanced script**. This is generally used to deliver SMS notifications to an internal SMS distribution technology. This is an advanced scripting procedure and is not necessary for configuring external SMS providers.

Select a service provider

You can configure how a device's service provider affects the construction of the device's email address.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Users** and open any user's record.
2. Under *Related Links*, click **Notification Preferences**, and then click an SMS notification device.
3. If no SMS device is present, click **New Device** and configure one.
4. Select the appropriate service provider, and click the reference icon for the *Service provider* field.
   The service providers are saved in the Notification Service Provider [cmn_notif_service_provider] table. Only active providers are visible.

SMS notification advanced scripting

You can use these objects and their attributes in your advanced notification scripts on the SMS Notification Service Provider form.

Objects and attributes

Object	Example attributes
current	• current.number • current.assigned_to

Objects and attributes (continued)

Object	Example attributes
	• current.company • current.state
email	• email.notify • email.sourceTable • email.contentType • email.attachments • email.HTML • email.dataVersionHeader • email.UID • email.allowSavingNoRecipientEmail • email.recipients • email.ignore • email.save • email.headers • email.sysID • email.attachmentLimits • email.class • email.sysId • email.textBody • email.hashCode • email.weight • email.equals • email.logEmail • email.reset • email.wait • email.body • email.SMSText • email.watermark • email.textBodyLegacy • email.sourceHeader • email.subject • email.instance • email.importance

Objects and attributes (continued)

Object	Example attributes
device	• device.service_provider • device.group • device.order • device.sys_id • device.sys_updated_by • device.sys_created_by • device.primary_email • device.schedule • device.name • device.sys_created_on • device.email_address • device.active • device.phone_number • device.sys_mod_count • device.sys_updated_on • device.user • device.sys_meta • device.type

Make a notification mandatory

To prevent users from turning off or deleting a subscription to a notification, make the notification mandatory.

Before you begin

Role required: admin

About this task

When you make a notification mandatory, the notification is locked in user notification preferences, preventing the user from removing or unsubscribing to the notification, filtering it, or changing the schedule.

Users receive mandatory notifications even if they disable notifications.

Note:
Mandatory notifications apply to only the primary device of the user. You cannot make a notification mandatory for secondary devices.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Open the appropriate notification.
3. Configure the form and add the **Mandatory** field.

Note:
This field does not display by default.

a. Click the context menu icon (≡) and select **Configure** > **Form Layout.**
b. Select the **Mandatory** field and the order in which you want the field to appear.
c. Click **Save**.

4. Select the **Mandatory** check box.

   When a user checks their notification preferences, the control button for the mandatory notification is on and is read only.

Force a notification to be sent

To force a notification to be sent to the specified users, enable forced delivery.

Before you begin
Role required: admin

About this task
Forcing a notification means that the relevant users receive the notification, even if they have not subscribed to the notification or have turned off the subscription. In addition, the users receive the notification even if the **Notification** field is set to **Disable**.

Note:
Forced delivery applies to only the primary device of the user. You cannot force a notification to be sent to secondary devices.

Procedure
1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. Open the appropriate notification.
3. Configure the form and add the **Force delivery** field.
4. Select the **Force delivery** check box.
5. Click **Update**.

   Unlike using the **Mandatory** option, forcing the delivery does not lock the user's preference or prevent the user from unsubscribing from the notification.

Create a notification filter

Notification filters enable a user to control the delivery of messages by creating special conditions on multiple tables in a single, reusable filter.

Before you begin
Role required: admin

About this task
For example, you can create a filter that controls message delivery when active incidents, problems, and change requests for network issues reach a critical state. For Core UI, filters for notifications or channels are set through the Notifications tab of the System Settings window. For details, see Apply notification conditions.

Note:
The system applies the user's filter conditions after the administrator's conditions have been evaluated. If the administrator's conditions fail, the system ignores notification filters.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notification Filters** and create a record.
2. In the *Notification conditions* related list of the new record, create and submit filter conditions on one or more tables.
3. Repeat the procedure to create additional conditions on other tables for this filter.

Filter device notifications using a schedule

You can associate devices, such as Email, SMS, and Voice, to schedules that define when the devices can and cannot receive notifications.

Before you begin

Role required: admin

About this task

Notifications that are triggered outside of the scheduled days and times for the device are not queued up for delivery at a later time. For example, if an administrator selects the **Weekdays** schedule for an email device, the device receives email notifications triggered between Monday and Friday. If notifications are triggered on Saturday, they are not delivered to the device.

Procedure

1. Define schedules as needed using **System Scheduler** > **Schedules** > **Schedules**.
2. Add or edit a device.
3. Configure the New Device for System Administrator form and add the **Schedule** field.
4. In the **Schedule** field, select the schedule for the device.
5. Click **Submit**.

Edit the schedule or filter of an existing notification message

You can update a schedule or filter that was previously created for an email notification.

Before you begin

Role required: admin

About this task

To edit the schedule or filter of an existing notification message:

Procedure

1. In the Notification Preferences screen, click the message to edit.
   You cannot edit any attributes of a mandatory message.
2. Make the appropriate changes in the form.
3. Click **Update**.

Create notification channels

You can add channels to receive your notifications. A notification channel is an email account or voice message system that you have access to.

Before you begin

Role required: none

About this task

Notification channels include email addresses, service providers for SMS messages, and mobile applications. You can create voice notification channels to support applications like Notify.

Note:
If you are using the ServiceNow mobile application or a custom push application, you do not need to create a push channel for your mobile device. The system automatically creates a channel for the mobile app after you initially log in to your instance from your mobile device.

Procedure

1. Click the gear icon ( ) in the banner frame to open the System Settings window, and click the **Notifications** tab.
2. Click **Create Channel**.
3. Complete the fields on the New Channel form.

Field	Description
Name	A descriptive name for the channel, such as the device or email account.
Type	The type of channel: ◦ **Email**: for email messages. **Note:** All users with an email address have a primary email channel, which is created automatically after a notification is sent to them. ◦ **SMS**: for SMS messages. ◦ **Voice**: for phone messages. If you are using the ServiceNow mobile application or a custom push application, the system automatically creates a channel for the mobile app after you initially log in to your instance from your mobile device.
Email address	The email address of the channel.
Phone number	The phone number for SMS messages or for voice messages.
Service provider	The service provider for SMS messages.

4. Click **Save**.
   The system creates and enables the channel, and adds it to the list of notification channels.

What to do next

To receive notifications on your new notification channel, you must enable the channel for individual notifications. After you enable the channel for a notification, you can set conditions to further control the notifications that you receive on the channel. For more information, see Apply notification conditions.

Modify notification channels

You can update channel information for your notifications.

Before you begin

Role required: none

Procedure

1. Click the gear icon (⚙) in the banner frame to open the System Settings window, and click the **Notifications** tab.
2. In the Notification Channels section, click the channel row or the right arrow (**>**) next to the channel name.
3. In the Edit Channel form, update the fields (see Create notification channels for descriptions of the form fields).
4. Click **Save**.

Result

The system updates and saves the channel information that you modified.

Delete a notification channel

Delete a notification channel so that you no longer receive notifications through the channel.

Before you begin

Role required: admin

Procedure

1. Click the gear icon (⚙) in the banner frame to open the System Settings window, and click the **Notifications** tab.
2. In the Notification Channels section, click the channel row or the right arrow (**>**) next to the channel name.
3. On the Edit Channel form, click **Delete**.

Result

The system deletes the channel from the list of notification channels and no longer delivers notifications to that channel.

Apply notification conditions

Set various conditions and filters to control the notifications you receive. You can also enable or disable an email digest and the delivery of individual notifications by channel.

Before you begin

Role required: none

About this task

- Enable or disable an email digest if a digest is available for the notification. An email digest is a single email that summarizes the activity for the notification during a time interval that you specify. Admins configure the email digest content for a notification.
- Enable or disable a channel on which the notification is received.
- Set preconfigured schedules and filters that determine when you receive your notifications.
- Set more advanced filter conditions to limit the notifications delivered to you.

Procedure

1. Select the system settings gear icon (⚙) to open the System Settings window.
2. Select the **Notifications** tab.
3. Select the notification:
   a. In the Notifications By Category section, select the row or right arrow (**>**) of the category that you want to view.
   b. In the list of notifications for the category, select the row or right arrow (**>**) next to the notification that you want to edit.
4. **Optional:** Enable or disable an email digest for the notification, if an email digest is available:
   a. Select the **Email Digest** switch.
   The email digest is enabled when the switch is green and turned off when the switch is grey.
   b. To enable the digest, select the **Interval** (length of time) during which the notifications are accumulated.

   The interval begins with the first occurrence of notification activity during the specified interval. For example, if you selected the Daily interval and the first notification occurs at 07:00, the system begins accumulating notifications at 07:00 and stops at 07:00 the next day. Soon after the interval ends, the system sends the email digest instead of sending the individual notifications that were triggered during the digest interval. If you turned off the digest, the system immediately stops accumulating the notifications and doesn't send the email digest.
5. Toggle the channel switch enable or disable a channel on which the notification is received.
6. To apply conditions for notification delivery, select the row or the right arrow (**>**) next to the channel for the notification.
   a. Fill the Apply Conditions form.

Field	Description
Schedule	Determines when the notification can and can't be received.
Filter	Preconfigured filter with the criteria that determines when the notification can be sent. For example, you might select a filter whose conditions send notifications when an incident with a priority of **1 - Critical** is opened for a network issue.
Advanced filter	Check box to use the condition builder to create additional criteria. When you select the check box, the **Table** and **Conditions** fields are displayed.
Table	Table for the notification. For example, if you selected an incident-related notification, you might select the Incident [incident] table.
Conditions	Define as many conditions as needed to limit the notifications you receive. For example, you can choose to be notified only when an incident communication plan is created for a specific configuration item.

   b. Select **Save**.

System Events

System Events in ServiceNow are crucial for automating processes and maintaining platform efficiency. They act as triggers for business rules, notifications, work-flows, and other actions, ensuring timely and consistent execution of tasks. By monitoring and responding to these events, organizations can streamline operations, improve response times, and enhance user experiences.

Get started

Exploring system events Learn about how system events function as triggers for actions and enable process automation.	Configuring System events Set up and manage system events to trigger actions and automate processes.	
Managing system events Connect system events with external systems and applications to enable seamless data exchange and automated work-flows.	System events reference More information about the predefined system events available and their usage.	

Troubleshoot and get help

- Ask questions and explore other resources for System events in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring system events

Events are special records that the system uses to log when certain conditions occur and to take some kind of action in response to the conditions.

System events overview

Learn about how system events function as triggers for actions and enable process automation.

Realize the benefits of leveraging system events by doing the following tasks in your instance:

- Create events
- Log events
- Implement automated job scheduling
- Move an event
- Register events
- Reprocess events
- Pass event parameters from a work-flow to a notification

System events users

Users

User	Description
Administrator	Administrators leverage system events to trigger automated actions and work-flows based on specific occurrences in the platform, enabling asynchronous processing and responsive automation. They can generate events via scripts or business rules, register them in the Event Registry, and handle them using notifications, script actions, or flows.
Business analysts	Business analysts analyze trends and processes by leveraging events to capture specific system activities, generate logs, or send data to analytic tools for better decision-making and process improvement.
Compliance officers	Compliance officers help to ensure adherence to regulatory requirements by using events to log critical activities, generate compliance reports, or initiate corrective actions when deviations from standards are detected.
Customer support	Customer support employees use system events to enhance customer service by triggering automated responses to customer requests, creating follow-up tasks, or escalating unresolved cases to appropriate teams.
Developers	Developers create and manage system events to integrate modules, trigger work-flows, and implement custom logic in response to user actions or record changes, helping to ensure extensibility and dynamic system behavior.
HR administrators	HR professionals use system events to automate HR processes, such as notifying employees of updates, generating on-boarding work-flows when a new hire is recorded, or triggering document reviews during off-boarding.
ITOM/ITSM	IT service managers use system events to monitor and automate ITSM processes, such as sending notifications for SLA breaches, escalating incidents, or triggering task work-flows, helping to ensure efficient IT service delivery and compliance with SLAs.
Marketing teams	Marketing teams automate customer engagement by using events to trigger campaign actions, such as sending promotional emails or updating CRM systems when specific conditions (e.g., customer interaction) are met.

Users (continued)

User	Description
Operations teams	Operations teams monitor and respond to operational alerts or system changes by triggering events that initiate corrective actions, such as restarting services, updating CMDB records, or notifying relevant teams of maintenance schedules.
Project managers	Project managers use system events to trigger notifications for project updates, create task dependencies automatically, or update stakeholders when milestones are achieved, enabling better project tracking and communication.
Security analysts	Security analysts leverage system events to trigger immediate responses to security incidents, such as creating alerts, generating logs, or initiating containment work-flows in response to detected threats or vulnerabilities.

System events benefits

System events benefits

Benefit	Feature	Users
Allows users to define new events that can be triggered by specific conditions or actions, enabling customized work-flows and automated responses.	Create an event	Developers, Administrators, IT Service Managers
Captures event activity in logs for auditing, debugging, and monitoring system behavior, providing insights into event performance and bottlenecks.	Event logs	Administrators, Compliance Officers, Security Analysts
Schedules jobs to be executed automatically at specific times or intervals, ensuring routine tasks are performed without manual intervention.	Implement automated job scheduling	Operations Teams, IT Service Managers, Administrators
Transfers events between queues or categories, helping optimize processing priority and load balancing in event-driven systems.	Move an event	Administrators, Developers
Registers custom or system-defined events in the platform to make them available for triggering specific actions or work-flows.	Register an event	Developers, Administrators
Provides the ability to reprocess failed or incomplete events, ensuring system reliability and minimizing the impact of errors.	Reprocess an event	IT Operations Teams, Developers, Administrators
Enables seamless transfer of contextual data from work-flows to notifications, ensuring messages contain relevant and actionable information for recipients.	Pass event parameters from a work-flow to a notification	Developers, IT Service Managers, HR Professionals

What to explore next

To learn more about configuring and using System events, see:

- Configuring system events
- Managing system events
- System events reference

System Events

Events are special records that the system uses to log when certain conditions occur and to take some kind of action in response to the conditions.

The system uses business rules to monitor for system conditions and to generate event records in the Event [sysevent] table, which is also known as the event log or event queue.

Event-generating business rules typically use this script logic:

If [some condition is true for the current record], then [add a specific event to the queue].

For example, here are some of the conditions in the **incident event** business rule:

- If a user adds a comment to an incident record, add an `incident.commented` event.
- If a user adds an incident record, add an `incident.inserted` event.
- If a user updates an incident record, add an `incident.updated` event.

Event-generating business rules use the GlideSystem `eventQueue` method to insert event records, which typically contain this information:

Event fields

Field	Description
Name	Unique name of event. Baseline event names include the record affected and the triggering action, such as `incident.commented`.
Parm1	Event-specific parameter the system uses to pass record information to other parts of the system, such as a record Sys ID or a field value.
Parm2	Event-specific parameter the system uses to pass record information to other parts of the system, such as a record Sys ID or a field value.
Table	Table to which the event applies. This is the same table on which the business rule ran.
Instance	Sys ID of the record to which this event applies.

Scheduled jobs periodically read the event queue and forward them to the appropriate handler for processing. The handler uses information from event records to take some kind of action such as:

- Run a script action
- Send a notification
- Trigger a workflow activity
- Trigger an inactivity monitor

By default, the system provides events covering a broad view of application activity. If existing events do not meet your needs, you can create your own events to watch for specific changes to records.

For developer training, see Scheduled Script Executions and Events Objectives on the ServiceNow® Developer Site.

Use the System Events and Jobs Dashboard to monitor the system event processing system and the scheduled jobs processing system.

Event processing during platform upgrade

Determine which events are to be processed during a platform upgrade by configuring the following properties:

- **glide.event_processor.all_events_upgrade_safe**: Set to **true** to process all events during platform upgrade. Default = **false**.
- **glide.event_processor.upgrade_safe_events**: Configure a comma-separated list of event names to indicate that only specified events are processed. Default = <empty>.

  This property is relevant only when *`glide.event_processor.all_events_upgrade_safe = false`*.

Event registry

The events registry lists the events the system recognizes. Use registered events to automate other activities, such as script actions or notifications.

After you create a new event and a business rule that uses the event, you must register it.

Registration lets other parts of the system, such as Email and SMS notifications and Script Actions, see the event in their list of available events and react to the event when it occurs.

Related topics

Register an event

Configuring System events

Plan, configure, and implement system events without having to complete any extra configurations first. Follow the task listed in the configuration overview to implement automated job scheduling.

Configuration overview

Implement automated job scheduling

Implement the message processing framework (automated jobs scheduling) by using the Queue Registration link on the Event Registration form.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Performance Analytics** > **System** > **Event Registry**.
2. Select **New** to register a new event.

Note:
You can also select an existing event to update the configurations.

3. Click the Queue Registration link in the banner message on the Event Registration form.

   Note:
   You can also disable automated job scheduling by selecting **Disable Automatic Job Scheduling** button on the Queue Registry list page.

   If you have started a new event, the Queue Registry list shows up. If you have selected an existing event, the Queue Registration new record shows up.
4. Select **New** to create a new queue.
   The Queue Registration form shows up.
5. On the form, fill in the fields.

Field	Description
Queue	Name of the queue
Application	Scope of the application
Automatic Job Scheduling	Option to implement automated job scheduling
Event Processing Order	Option to select the way of processing events ◦ Parallel: Process multiple events simultaneously ◦ Sequential: Process only one event at a time. The events are inter-dependent and can process only after the previous event is completed **Note:** Once you have configured the processing order and submitted the queue, it can't be edited. You can delete the queue and create a new queue.
Job configuration type	Type of job configuration **Note:** This field is disabled when you select Sequential as the Event Processing Order. If you select Parallel as the Event Processing Order, you have 2 options. ◦ Constant: Creates the required number of jobs ◦ Scale with node: Multiplies jobs by the number of nodes. You can use this option if you are unaware of the number of available nodes

Field	Description
Scale factor	Total number of jobs based on the Job configuration type selected
Poll interval	Determines how frequently jobs should poll and process events of this queue
Description	Optional field to describe about this queue related information

6. Click **Submit**.
   The Queue Registry list shows up. It also shows the update made to the recent queue.
7. Select the recently added queue.
   The Queue Registration form shows up.
8. Scroll down to the Related lists section.
9. Select the **Status** entry in the Queue Details related list.

   Note:
   Since the new queue has been submitted, the Status entry is Active.

   The Queue Details form shows up.
10. **Optional:** Update the number of jobs for this queue as per the requirement.
    You can now manually scale up or down as per the requirements without having to configure everything and can eventually avoid misconfiguration.
11. **Optional:** Click Rollback related link on the Queue Registration form if you want to roll back a queue to its previous configurations without automated job scheduling.

    Note:
    The Rollback related link shows up only if you are in an automated jobs scheduling configuration.

12. **Optional:** Select **Retrieve Configurations** if you want to retrieve any existing job configurations for this queue.

    Note:
    This process is possible only if you are in an automated jobs scheduling configuration.

    If any queue with your suggested queue name already exists in the manual jobs scheduling configuration, the configurations of the existing queue are retrieved. You can then validate and update them as required.

Managing system events

Managing system events involves creating, monitoring, and handling events to drive automation and streamline processes. Administrators and developers can define custom events, register them in the system, and set up actions or work-flows that respond to these events. This functionality helps ensure efficient task execution, helps improve system monitoring, and enables more robust error handling for enhanced operational reliability.

Overview of managing system events

- Create an event

  Create custom system events.

- Move an event

  Accelerate the execution process by moving high priority events from the default queue to an isolated queue.

- Pass event parameters from a workflow to a notification

  Pass two event parameters that send information about a record or related records from a work-flow to a notification.

- Register an event

  Register an event for a specific table and a business rule that fires the event.

- Reprocess an event

  Re-fire an event for testing or diagnostic purposes.

System events reference

For more information, see System events reference.

Create an event

If you do not find a suitable existing event, you can create your own.

Before you begin
Role required: admin

About this task
The gs.EventQueue function works directly with the backend and therefore business rules that are called by gs.EventQueue() are not invoked.

Procedure

1. Navigate to **All** > **System Policy** > **Events** > **Registry**, and then select **New**.
2. On the form fill in the fields.

Event Registration Form Completion

Field	Description
Event name	Name of your new event.
Table	Database table for this event. **Note:** Only tables and database views that are in the same scope as the event appear on the listing.
Queue	Name of the queue that the event is placed into when triggered. Use only lowercase letters, no spaces, and no special characters except underscore (_). For example, `my_queue`. See Using custom queues to process events.
Priority	Order in which messages will be processed. Lower values have higher priority.

Field	Description
	Note: The com.glide.sysevent.priority.enabled system property is enabled by default. If you disable it, events processing is not done by priority.
Caller Access	Restricted caller access settings. ◦ Caller Restriction: calls to the resource must be manually approved. ◦ Caller Tracking: Calls to the resource are automatically approved.
Fired by	Name of the business rule that runs the event. This field is for reference only and is not used by any process. Make sure that there is enough information to locate your event again.
Description	Short description of the purpose of the event.

3. Click the **Business Rules** related link.
4. If you are creating an event for a base system table, select the existing event business rule for the table.

   Example
   For example, select the *sc request events* business rule to create a custom Request event.
5. If you are updating an existing event business rule, add a new condition to the **Script**.

   Example
   The following sample script adds a *request.commented* event with the user's Sys ID as parm1 and the user's user name for parm2.

```
if (current.operation() != 'insert' &&
 current.comments.changes()) {
gs.eventQueue('request.commented', current, gs.getUserID(),
 gs.getUserName());
}
```

6. If you are creating an event for a custom table, create a new business rule that runs after database operations.

   Example
   For example, this business rule defines several events for a custom application called Marketing Events.

Sample event business rule

Field	Value
Name	Attendee Events
Table	Attendee [x_snc_marketing_ev_attendee]
Application	Marketing Events
Advanced	Selected
When	after
Insert	Selected

Field	Value
Update	Selected
Delete	Selected
Script	Add custom script that: ◦ Checks for one or more conditions on the current record. ◦ Calls the gs.eventQueue() method and specifies an event name. See code sample.

Note:
If you add **Filter Conditions**, **Role conditions**, or a **Condition** value, verify it runs the business rule when expected.

```
(function executeRule(current, previous /*null when async*/) {
   //This function will be automatically called when this rule
 is processed.
        //Add event when attendee inserted
        if(current.operation() == 'insert' &&
 current.marketing_event.changes()) {

 gs.eventQueue('x_snc_marketing_ev.attendee.added', current,
current.marketing_event, current.email);
        }
        //Add event when marketing event changes
        if(current.operation() == 'update' &&
 current.marketing_event.changes()) {

 gs.eventQueue('x_snc_marketing_ev.attendee.deleted', previous,
previous.marketing_event, previous.email);

 gs.eventQueue('x_snc_marketing_ev.attendee.added', current,
current.marketing_event, current.email);
        }
        //Add event when attendee deleted
        if(current.operation() == 'delete') {

 gs.eventQueue('x_snc_marketing_ev.attendee.deleted', current,
              current.marketing_event, current.email);
}
})(current, previous);
```

7. Register the event.

What to do next

Create a script action or notification to process the event.

Related topics

Register an event

Script actions

Move an event

Accelerate the execution process by moving high priority events from the default queue to an isolated queue with just one click.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Event registry**.
   A list of registered events shows up.
2. Select the event you want to move.
   The Event Registration form of the selected event shows up.
3. Select **Move to Default Queue** to move the event from the current queue and accelerate its execution.

   Note:
   If you are currently in the default queue, you will have the **Move to Adaptive Event Queue** option to select.

   The processor can execute only one event at a time. So, if you move an event to a queue, it gets executed only when the processor is available.

4. **Optional:** Navigate to **System Diagnostics** > **Stats** > **Adaptive Events** to track the average execution time of the events.
   You can select any slow moving event and change the queue as required. You can also register an event to show up on the event registry list.

Pass event parameters from a workflow to a notification

Pass two event parameters that send information about a record or related records from a workflow to a notification.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Policy** > **Events** > **Registry** and define a new event to call.
2. Create the activity step that calls the event from your workflow and assign the two possible parameters.
   These parameters can be references or fields on the record that triggered the workflow, such as `current.number` for the request item number. You can also dot-walk to records in other tables. The event then sends the parameters information to the notification that it triggers.
3. Retrieve the parameters in the notification with email scripts by using `event.parm1` and `event.parm2`.
   For example:

```
var string = event.parm1.toString();
template.print(string);
```

Example:
Another example using Sys IDs gives you access to anything in the record that triggered the workflow. Use a script like this to send information about a change or request to support personnel, for example.

```
var id = event.parm1.toString();
var gr = new GlideRecord('sc_req_item');
gr.get('sys_id', id);
if (gr.next()){
   // Do something.
}
```

Related topics

Scripting for email notifications

Register an event

You can register an event for a specific table and a business rule that fires the event.

Before you begin

Role required: admin

Procedure

1. To register an event, browse to **System Policy** > **Events** > **Registry**, and then click **New**.
2. On the form fill in the fields.

Event Registration Form Completion

Field	Description
Event name	Name of your new event.
Table	Database table for this event. **Note:** Only tables and database views that are in the same scope as the event appear on the listing.
Queue	Name of the queue that the event is placed into when triggered. Use only lowercase letters, no spaces, and no special characters except underscore (_). For example, `my_queue`. See Using custom queues to process events.
Priority	Order in which messages will be processed. Lower values have higher priority. **Note:** The com.glide.sysevent.priority.enabled system property is enabled by default. If you disable it, events processing is not done by priority.
Caller Access	Restricted caller access settings. ◦ Caller Restriction: calls to the resource must be manually approved. ◦ Caller Tracking: Calls to the resource are automatically approved.

Field	Description
Fired by	Name of the business rule that runs the event. This field is for reference only and is not used by any process. Make sure that there is enough information to locate your event again.
Description	Short description of the purpose of the event.

Reprocess an event

You can fire the event again for testing or diagnostic purposes.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Logs** > **Events**.
2. Open an event.
3. Under **Related Links**, click **Reprocess Event**.
   The event returns to the event queue.

System events reference

Reference topics provide information about system events.

Event logs

The event log records all system events that occur within the ServiceNow AI Platform.

This log provides the following information for all events that occur:

Event log

Field	Description
Created	Date and time of the event for the locale of the machine running the instance.
Name	Name of the event as listed in the Event Registry.
URI	HTTP query that generated the event.
Parm1	Event-specific value that depends on the event and the recipient.
Parm2	Event-specific value that depends on the event and the recipient.
Table	Database table acted on for this event.
Processed	Date and time the event started processing. This time reflects the locale of the machine running the instance.
Processing time	Time taken to process this event, in milliseconds.
Queue	Processor queue name.

Event states

The event state describes where in the life cycle the event is.

Event states

State	Description
Ready	The system created the event and it is in the queue waiting to be processed.
Processed	The event successfully ran. An event does not necessarily trigger any further action when processed. Additional functionality must make use of the event.
Error	The event encountered an error during processing. This state is often caused by invalid event parameters. Reprocessing the event may resolve the error.
Transferred	The event was rotated to a different shard of the Event [`sysevent`] table. When an event is rotated, a duplicate record is created in an active shard to be processed. A scheduled job processes the event when it is next in the queue, but it is not possible to predict when this will happen as because several events may need to be processed before it. Therefore, you can reprocess the event. See Reprocess an event.

The incident events business rule

The incident events business rule comes with the system and defines a number of events that can be triggered by different actions in the Incident table.

Incident events business rule

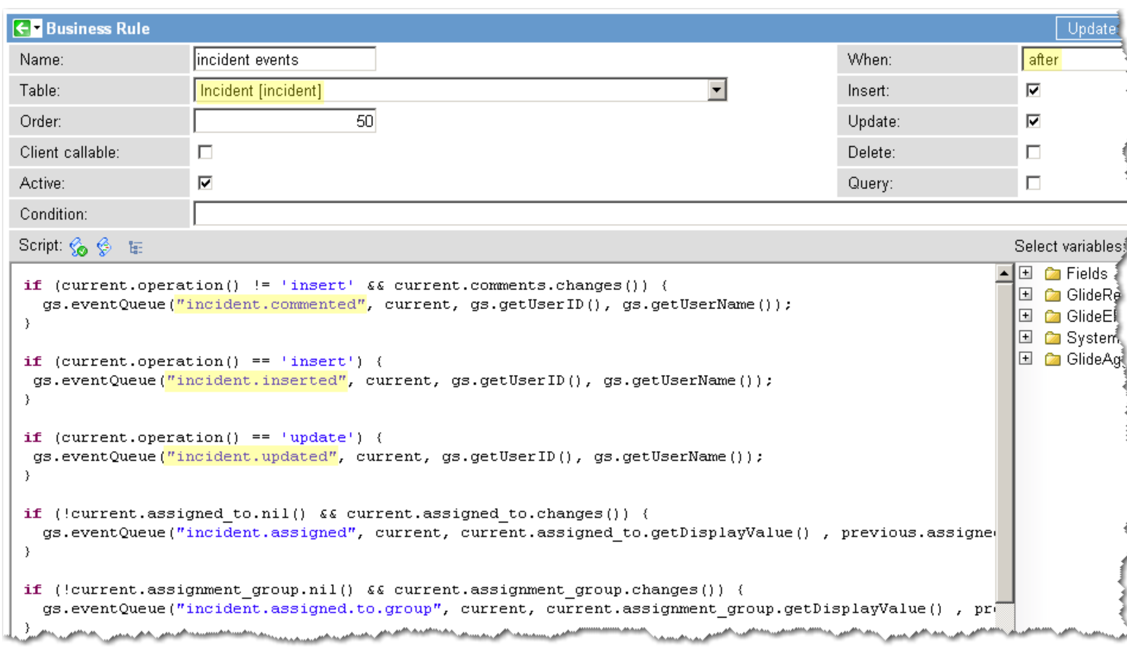

This business rule defines several events, three of which are triggered after a record in the Incident table is inserted or updated. The first script is:

```
if (current.operation() != 'insert' &&
 current.comments.changes()) {
gs.eventQueue("incident.commented", current, gs.getUserID(),
 gs.getUserName());
}
```

The condition in this script requires that a change be made to the **Comments** field in an existing (not inserted) incident record. If this condition is true, then the platform adds the `incident.commented` event to the event queue.

The second condition requires that a record be inserted before the event is added to the queue.

```
if (current.operation() =='insert') {
```

The third condition is true whenever the incident record is updated (including updates to the **Comments** field, as specified by the first script).

```
if (current.operation() == 'update')
```

The then part of each script, the *`gs.eventQueue`* function, adds the event to the event queue. This statement uses the following syntax, set off with braces:

```
gs.eventQueue("incident.updated", current, gs.getUserID(),
 gs.getUserName());
```

The *`gs.eventQueue`* function takes the following parameters:

Field	Input Value
Name	The name of the event triggered, set in quotation marks
Record	The record referenced when the condition in the script evaluates to *true*. Usually this is expressed as *current*, meaning the current record the business rule is working on. If the business rule is being triggered as part of a scheduled job, use a GlideRecord argument in its place.
Parm 1	An optional parameter you can use to pass system or record information with the event. For example, the GlideSystem API call *`gs.getUserID()`* passes the Sys ID of the user who acted on the current record as a string value. Other scripts can reference this string value as parm1 using the format `${event.parm1}`.
Parm 2	An optional parameter you can use to pass system or record information with the event. For example, the GlideSystem API call *`gs.getUserName()`* passes the user name of the user who acted on the current record. Other scripts can reference this string values as parm2 using the format `${event.parm2}`.

Note:
The gs.EventQueue function works directly with the backend and therefore business rules that are called by gs.EventQueue() are not invoked.

Global events

Your instance has a global function called *`global_events()`* that triggers from a business rule when certain conditions occur.

This function triggers when your instance is:

- Inserting new records
- Updating existing records
- Adding comments to an existing record
- Assigning a record to a user
- Exceeding a record's inactive timer

For example, if you add the script global.events(current) to a business rule on the `change_request` table, the instance automatically configures the following events:

- change_request.inserted
- change_request.updated
- change_request.commented
- change_ request.assigned
- change_ request.inactive

Script actions

You can use script actions to create server-side scripts that perform a variety of tasks, such as modifying a configuration item (CI), or managing failed login attempts. Script actions are triggered by events only.

Configuration

To create a new script action, navigate to **System Policy** > **Events** > **Script Actions** and click **New**.

Script actions

Field	Description
Name	Unique name for your script action.
Application	Application that contains the script.
Event name	Event to use for this script. If you do not find an event for your script action that suits your purpose, you can create a new one.
Active	Check box that enables or disables the script action. Select **true** to enable the script action.
Execution order	Order in which the script will be executed.
Condition script	Statement for a condition under which this script should execute. The system only parses the **script** field if the condition evaluates to true. If you decide to include the condition statement in the script, leave this field blank.
Script	Script that runs when the condition you define evaluates to true. Two additional objects are available in this script:

Script actions (continued)

Field	Description
	• *event*: a GlideRecord - the sysevent that caused this to be invoked. If you want this first parameter on the event, use *event.parm1* or *event.parm2* for the second parameter. For the date/time of the event, use *event.sys_created_on*. To get the user ID that created the event (if there was a user associated), use *event.user_id*. • *current*: a *GlideRecord*- the event scheduled on behalf of (incident for example).

Sample of a script action that creates an email notification for Workflow activity:

Script action workflow

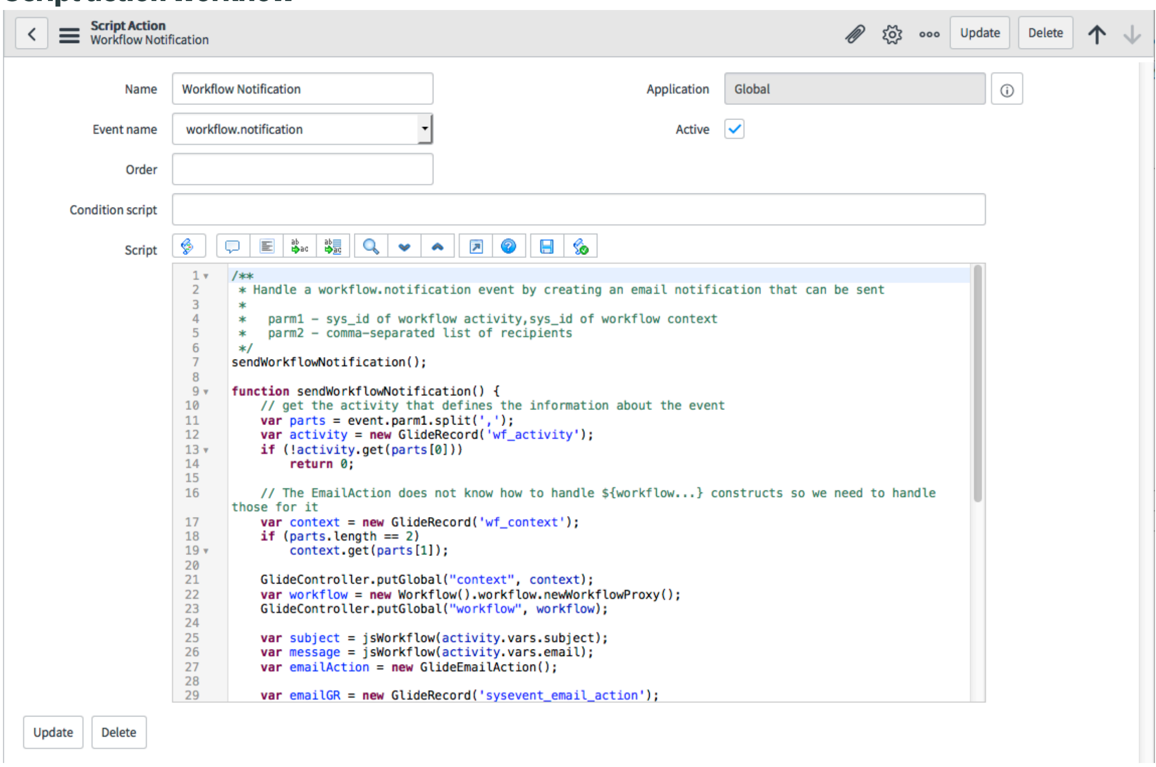

Sample scripts from the change events business rule

Several scripts are found in the baseline change events business rule.

This business rule defines events that fire after a change request is inserted or updated.

```
if (current.operation() == 'insert') {
 gs.eventQueue("change.inserted", current, gs.getUserID(),
 gs.getUserName());
}

if (current.operation() == 'update') {
 gs.eventQueue("change.updated", current, gs.getUserID(),
 gs.getUserName());
}
```

```
if (!current.assigned_to.nil() && current.assigned_to.changes())
 {
  gs.eventQueue("change.assigned", current,
 current.assigned_to.getDisplayValue() ,
 previous.assigned_to.getDisplayValue());
}

if (current.priority.changes() && current.priority == 1) {
  gs.eventQueue("change.priority.1", current, current.priority,
 previous.priority);
}

if (current.risk.changes() && current.risk== 1) {
  gs.eventQueue("change.risk.1", current, current.risk,
 previous.risk);
}

if (current.start_date.changes() || current.end_date.changes()
 || current.assigned_to.changes()) {
  if (!current.start_date.nil() && !current.end_date.nil()
 && !current.assigned_to.nil()) {
    gs.eventQueue("change.calendar.notify", current,
 current.assigned_to, previous.assigned_to);
  }

  // Remove from previous assigned to, due to assigned_to
 changing
  if (!previous.assigned_to.nil()) {
      if (!current.assigned_to.nil() &&
 current.assigned_to.changes() &&
         (!previous.start_date.nil()
 && !previous.end_date.nil())) {
        gs.eventQueue("change.calendar.notify.remove", current,
 current.assigned_to, previous.assigned_to);
      }
   }
  // Remove old calendar from current assigned to, due to date
 changing
  else if (!current.assigned_to.nil()) {
     if ((current.start_date.changes()
 && !previous.start_date.nil()) ||
         (current.end_date.changes()
 && !previous.end_date.nil())) {
       gs.eventQueue("change.calendar.notify.remove", current,
 current.assigned_to, current.assigned_to);
     }
  }
}
```

State Management

State Management enables an administrator to define State Models and State Transitions that control how a record is allowed to transition through a predefined list of states.

An example of a state transition is when the **State** field in a facilities request is moved from the **Assigned** state to the **Work In Progress** state.

State Management is active for all instances.

What is a state model?

A state model is a list of states that describe an expected record workflow through the lifecycle of the record. State models can be defined for any table that extends the task table. State models simplify defining the state transitions allowed for a specific task type.

In the State Model [sys_state_model] table, define the name of the state model and which task table the state model is applied to. Use the condition builder to specify any conditions for applying the state model to records and any required condition for moving between states.

For example, you could define a state model for a new custom application for airline reservations. The custom application has a Reservation Request [reservation_request] table with 4 states: **Held**, **Confirmed**, **Completed**, and **Canceled**. You could define the state model to target the Reservation Request table, and then define the state transitions for each of the 4 states. When you enable the state model, the choice list for the **State** field includes only the choices allowed by the conditions in the state transitions.

Note:
State Management includes example state models that are copies of the normal, emergency, and standard change request state models. By default, these examples are not enabled. Use them only as examples to develop a state model and transitions for a task table that does not have a state model. Do not enable these example state models for change requests and then make changes to them. Doing so breaks existing transitions for change requests.

What is a state transition?

State transitions are a list of conditions for entering or exiting each state defined for a table. In the State Transitions [sys_state_transition] table, use the condition builder to build a list of conditions required for entering or exiting each state.

To prevent users from choosing an invalid state, any attempt to update a record's state is denied if it violates the state transitions, whether the attempt is through user input, a script, a Web API such as REST or SOAP, or any other source.

State transitions control the choice list for the **State** field on the target task table and prevent you from choosing any state value that does not adhere to the underlying process or does not meet the defined conditions for the transition.

For example, if the enter condition for the **Completed** state is **State is Confirmed**, only records in the Confirmed state can transition to the Complete state. When a record is in the Confirmed state, the only choice in the **State** field choice list is **Completed**.

Add a state model and transitions

Add a state model and transitions to specify conditions for moving between states.

Before you begin

Role required: state_model_admin or admin

About this task

Develop and test your state model in a non-production instance before deploying it in the production instance.

Procedure

1. Navigate to **All** > **State Management** > **State Models**.
2. Click **New**.
3. Fill in the fields on the form.

State model fields

Field	Description
Name	Descriptive name for this state model.
Table	Table this state model applies to.
Application	Defaults to **Global**.
Order	Order in which this state model is evaluated for the target table. State models are evaluated in numerical order, from the lowest number to the highest number. Leave space between the numbers so that you can insert a new state model later, if necessary. For example, enter 10 for the first state model to be evaluated, 20 for the next state model to be evaluated, and so on.
Active	Indicates that this state model is enabled. Clear this check box until you define the state transitions for the state model.
Description	Description of this state model.
Condition	Condition builder that sets the conditions for applying this state model to records. For example, you could add the condition **Category is Networking** versus **Category is Security** if there are different state models for networking requests and security requests.
Common Exit Condition	Condition builder that sets any common condition required for moving from one state to another.

4. Open the form context menu and click **Save**.
   The **State Transitions** related list appears.
5. Click **New** on the **State transitions** related list.
6. Fill in the fields on the form.

State transition fields

Field	Description
State Model	State model that uses this state transition.
State	State this transition controls, for example **Ready** or **In Progress**.

Field	Description
Application	Defaults to **Global**.
Order	Order in which this state transition is evaluated. State transitions are evaluated in numerical order, from the lowest number to the highest number. For example, if the first predefined state is **New**, give that state the lowest order number so that the conditions for entering and leaving the state are evaluated first. Leave space between the numbers so that you can insert a new state transition later, if necessary. For example, enter 10 for the first state transition to be evaluated, 20 for the next state transition to be evaluated, and so on.
Terminal State	Selected if there are no transitions from this state.
Description	Description of this state transition. For example, the **Description** for the Assess state could be `New state can move only to Assess state`.
Enter Condition	Condition builder that sets the conditions for entering this state.
Exit Condition	Condition builder that sets the conditions for exiting this state.

7. Click **Submit**.
   The state model form reopens, and the new transition is added to the related list.
8. Continue adding state transitions for the remaining states.
9. When finished adding state transitions in the state model form, select the **Active** check box and click **Update**.

Implement process flow and UI actions with a state model

You can implement a process flow and UI actions with a state model.

Before you begin

Role required: state_model_admin or admin

About this task

State models provide a way to limit the choices for moving between states in a form. If you set up only a state model, users manually change states in the form. Additional steps are required to add the process flow or UI actions, as illustrated in the following example.

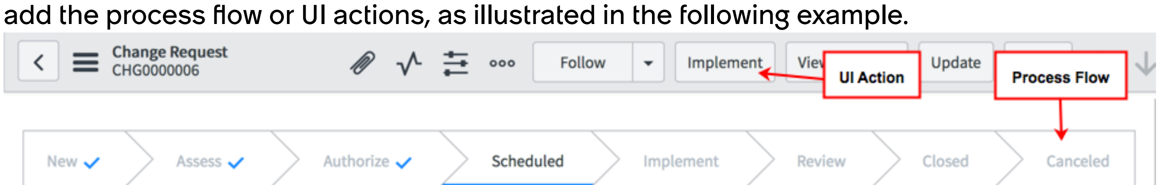

Procedure

1. Verify choices for the **State** field for the table and ensure that you created transitions for them.
   For more information about how to add **State** field choices, see Best practices for state field choice values.
2. After creating the state model, define the process flow.
   For more information about process flows, see Process flow formatter.
3. Define UI actions as desired to move between states.
   For more information about UI actions, see UI actions.

Result

After the state model is enabled, only the defined state transitions are included in the choice list for the **State** field, and the process flow and UI actions are implemented.

Installed with State Model

Several types of components are installed with the State Model.

Tables installed with State Model

State Model adds the following tables.

Table	Description
State Model [sys_state_model]	Defines the name of the state model and which table the state model targets. Also defines the condition for applying the state model to records and any common condition for moving from any state to any other state.
State Transition [sys_state_transition]	Defines how each record in the target table transitions to (the enter condition) and/or from (the exit condition) each state predefined for that table.

Roles installed with State Model

State Model adds the following role.

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

State Model Administrator	Description	Contains roles
[state_model_admin]	Allows users to add, modify, and delete state models for task tables.	None

State model Administrator [state_model_admin]

Add, modify, and delete state models for task tables.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Script includes installed with State Model

State Model adds the following script include.

Script include	Description
TableChoiceUtils	Utility class to get the choice list of a table field.

Client scripts installed with State Model

State Model adds the following client script.

Client script	Table	Description
Populate State Choices	State Transition [sys_state_transition]	Populates the **State** field choices in the targeted record.

Business rules installed with State Model

State Model adds the following business rules.

Business rule	Table	Description
Set the display value for state	State Transition [sys_state_transition]	Sets the **State** field value based on the state model for the table.
Set the table to pick choices	State Transition [sys_state_transition]	Caches the **State** field value from the first table in the hierarchy for further rule processing.

Metrics

A metric measures and evaluates the effectiveness of IT service management processes.

For example, a metric could measure the effectiveness of the incident resolution process by calculating how long it takes to resolve an incident.

A metric can often be easily obtained from the data. To find the number of incidents created today, a report counts the number of incidents in the incident table with a Created date of today. However, metrics must be gathered as data is updated. For example, determining how long an incident was assigned to a certain group requires collecting information about assignment changes and calculating the duration of each assignment.

The Metric plugin provides an easy, declarative way of defining metrics. Once defined, the data for the metric is gathered, and instances of the metric are calculated and stored. For example, the "Assigned to Duration" metric measures the amount of time an incident is assigned to an individual. To define this metric, you create a metric definition of the type "Field value duration" and select the "Assigned to" field from the Incident table. A metric instance is created for each incident assignment showing its duration. Reporting on the duration of incident assignments becomes easy.

You can report on a metric by using a Creating database views for reporting that links the metric to the table on which it is defined.

To install the Metric Definition plugin (com.glide.metrics), navigate to **System Definition** > **Plugins**, search for Metric Definition, and click **Install**.

Define a metric

Create a metric definition for a task table.

Before you begin

Role required: metric_admin

Procedure

1. Navigate to **All** > **Metrics** > **Definitions**.

   > **Important:**
   > Don't confuse these metric definitions with the metric definitions for the Assessments feature. The Assessments feature is concerned with questionnaires. It isn't related to the Metrics feature.

2. Select **New**.
3. Complete the Metric definition form then select **Submit**.

   > **Note:**
   > In the base system, metrics are configured to work on the task table only. To apply metrics to cmdb_ci tables, duplicate the metric events business rule that currently runs on the task table for the cmdb_ci table. Without the events created, no metric processing can occur.

New metric definitions

Field	Description
Number	The system generates a unique record number for the metric definition.
Name	Enter a unique name to describe what metrics are collected.
Table	Select the table that you want to collect metrics for. A metric can only apply to one table. **Note:** The list shows only tables and database views that are in the same scope as the metric definition
Field	Select the table column that you want to monitor for changes. Metrics only work on audited fields. Creating a metric for a non-audited field produces unreliable metrics.

Field	Description
	Note: Fields starting with 'sys_' don't work for a metric definition, except for 'sys_domain'.
Type	Select what values this metric generates. ◦ **Field value duration**: This type of metric measures the duration of time from when the value of the specified field is set until it is changed. A Field value duration metric can optionally specify a script. The script can either return a duration value or set the answer variable to false to stop processing the metric. For example, the baseline incident metrics stop calculating duration when the Active field of an incident is set to false. The script can also carry out any other action such as closing the duration of other metrics defined on the same record. For more information, see Sample field value duration script. ◦ **Script calculation**: This type of metric creates a metric instance using a script. The script has access to the current row in the table (for example an incident) and the metric definition. The script must perform metric calculations and insert data into the metric_instance table. The calculation doesn't have to result in a duration. It can calculate any type of value and store it in the metric instance value.
Timeline	Check this box to display the metric on a timeline page. For more information, see Timeline pages.
Active	Check this box to monitor changes for this metric.
Description	Specify what data the metric monitors and its conditions. A good description helps other metrics administrators understand the purpose of your metric.
Script	Enter the script you want the metric to run to either calculate a duration or perform some calculation on the metric_instance table.

Sample field value duration script

Review the existing Incident Open metric definition to see how you can create your own custom metric.

This script either provides a duration value or stops processing durations (sets the answer variable to *false*) when an incident is closed.

```
// script can set answer to false to terminate processing of the
 metric
    // mi - MetricInstance
    // answer
    if (!current.active) {
    answer = false;
    mi.endDuration();
    gs.log("Closing field durations");
    closeDurations(mi.current);
    }

    function closeDurations(current) {
    var now_GR = new GlideRecord('metric_instance');
    gr.addQuery('id', current.sys_id);
    gr.addQuery('calculation_complete', false);
    gr.addQuery('definition.type', 'field_value_duration');
```

```
    gr.query();
    while (gr.next()) {
    gs.log("closing: " + gr.definition.name + " for: " +
current.number);
    var definition = new GlideRecord('metric_definition');
    definition.get(gr.definition);
    var mi = new MetricInstance(definition, current);
    mi.endDuration();
    }
    }
```

Metric instance

A metric instance is a record in the metric_instance table. A record holds one instance of a metric.

Metric instance records get created and updated in one of two ways: 1) If the metric is a duration, the system automatically populates the metric instance table with duration values. 2) If the metric is calculated from a script, the script itself must populate the metric_instance table.

Metric plugin

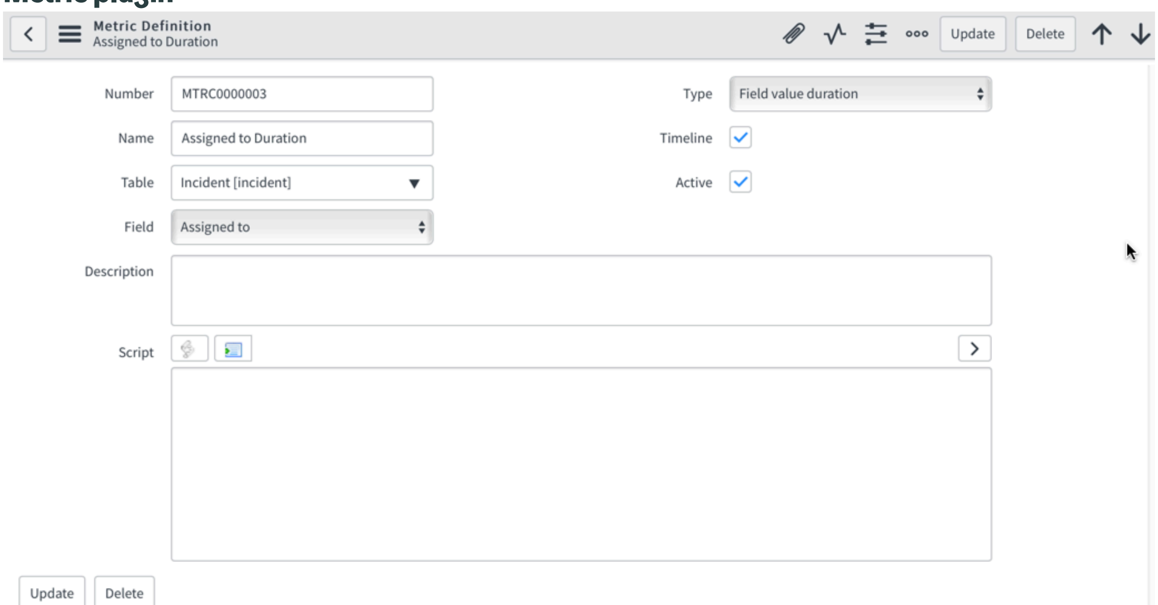

Some of the notable fields in the metric_instance table are:

- **Metric definition**: the metric definition for which this metric instance was gathered.
- **Value**: For a "Field value duration" metric this is the value of the table field for which duration is calculated. For example, for the "Assigned to Duration" metric, the Value is the name of the person assigned to the incident. For other metrics, the value can be any value calculated by the metric.
- **ID**: Identifies the specific record for which the metric is gathered. For example, the specific incident.
- **Duration**: Time duration for a Field value duration metric.

Maintaining and monitoring the ServiceNow AI Platform

Keep your instance running smoothly, improve performance, and find issues before they arise, all with this suite of tools.

Instance Scan	Events	Platform performance
Find potential health issues and opportunities for best practice improvement in your instance.	Get in-depth analysis for all displays and user-based events in your application.	Optimize your application with diagnostics and performance monitoring.

Instance Scan

Get detailed analysis of your instance configurations, and stop potential system health issues before they arise. Use ServiceNow's Instance Scan any time and ensure system stability, whether you are developing, releasing, or upgrading your application. Schedule, execute, and track scans from single-point to full-suite tests, and review the data they provide with dashboards.

See Instance Scan.

Events

Track every occurrence in your app, from screen loads to user clicks. Usage Insights provides detailed charts and tables for every event, to show you where the most activity goes on in your app. Analyze core processes in your business and improve user efficiency and workflow.

See System Events.

Platform performance

Keep your system optimized and running smoothly. Use diagnostics to track processing and response times, network speeds and behaviors, and browser settings. Monitor platform performance, set quotas and control system bandwidth to ensure stable operations for users, set up test modules, and more.

See Platform performance overview.

Instance Scan

Use ServiceNow® Instance Scan to interrogate your instance for configurations that indicate health issues and identify opportunities to address best practices. Instance Scan checks your existing configurations and helps you avoid creating future configuration issues. Instance Scan is a tool that can be used as a part of your development operations, release management as well as pre- and post-upgrades.

servicenow

Get started

Explore	Configure	Use
Learn about Instance Scan records and components.	Execute scans to keep a track on the health of your instance.	Use Instance Scan dashboard for a system wide visual representation of the health of your instance.

Exploring Instance Scan

If you are new to Instance Scan, read this overview to learn what the tool can do. Follow the tutorial to create checks and execute scans that uses most basics of Instance Scan features.

Note:
Instance Scan doesn't fully support domain separation. Findings are visibly domain separated based on the domain of the source record. For more information see Domain separation.

Instance Scan record and components

Checks

Checks are singular focused rules that detect anomalies or opportunities in an instance. These checks can run against tables, records, or metadata. Checks are defined to identify security, upgrade best practices, manageability, user experience and performance vulnerabilities. See Getting started with checks for more information.

Results

An Instance Scan result reports the status and type of the scan. See Results for more information.

Findings

A finding is a reference to a record that has violated a rule from a check on the instance. See Findings for more information.

Dashboard

The Instance Scan dashboard is a system-wide visual representation of the health of your instance. The dashboard helps you manage and analyze the full scan results against your instance. See Instance Scan dashboard for more information.

Quota rule

A quota rule determines the execution threshold of a scan. The quota rule prevents the instance from running long scans. For example, any scan running longer than the threshold set by the quota rule will result in a failure. See Quota rules for more information.

Scan types

Instance Scan deals with the following types of scans.

Full scan

Execute a scan for the entire instance by selecting **Execute Full Scan**. Implementing a full scan runs all the active checks present in your instance.

Point scan

Execute all applicable checks against a single record, update set, or an application by selecting **Run Point Scan**. For example, if you execute a point scan against a business rule, only the checks that are applicable to the business rule table run, and only that single target record is scanned. If you execute an update set scan or an application scan, all records related to that update set or application are scanned. See Execute an app scan and Execute an update set scan for more information.

Test scan

Execute a test scan to verify if the check works as expected. The test scan enables you to test a single check instead of a full scan by selecting a single check and selecting **Test Check** on the Check form.

Roles

Instance Scan has the scan_user role that can run different types of scans and view the findings and results.

Getting started with checks

Checks are singular focused rules that detect anomalies or opportunities in an instance. These checks can run against tables, records, or metadata. Checks are defined to identify security, upgrade best practices, manageability, user experience and performance vulnerabilities.

Check types

You can create your check by selecting one of these types.

Table Check

Create a check by selecting **Create a new Table Check** if you know which specific table and conditions you want to test. This check type is applied on only one table at a time. You can also include your own script for more complex capabilities by selecting the **Advanced** option on the form.

Column Type Check

Retrieve all records containing a specific column field type from all tables in an instance by selecting **Create a new Column Type Check**. The **Column Type Check** type implements the rule you created to iterate all records matching the target column field type.

Script Only Check

Create a check without specifying a table or a column type by selecting **Create a new Script Only Check**. You can verify meta data, configurations, and execute complex checks by writing your own script.

Linter Check

Create a linter check to identify any issues in a script. When a linter check is run on a record, an abstract syntax tree for its code is generated. You can use the abstract syntax tree to analyze issues with the code.

Create a check

Create your own checks by implementing the following procedure.

Before you begin

Role required: scan_admin

Procedure

1. Navigate to **All** > **Instance Scan** > **Checks**.
2. Select **New** from the Checks context menu.
   The Check Interceptor list appears.
3. Select one of the available types of checks.
   See Check types for more information.

   Note:
   If you want to update an existing check, the **Version** field increments itself. The new version provides you a change history. You can also know which particular version of the check ran with a given scan.

Create a table type check

Create a check by selecting **Create a new Table Check** if you know which specific table and conditions you want to test. This check type is applied on only one table at a time. You can also include your own script for more complex capabilities by selecting the **Advanced** option on the form.

Before you begin

Role required: scan_admin

```
Before performing this task you must complete Create a check.
```

Procedure

1. Select **Create a new Table Check**.
2. On the form, fill in the fields.

Fields	Description
Name	Name of the check record.
Application	Application that has the check record.
Category	Category of the check record.
Priority	Priority at which the findings should be resolved.
Version	Current version of the check record. **Note:** If you want to update an already existing check, the **Version** field increments itself.
Active	Option to activate this check record.
Short Description	Mandatory description about the check records.
Description	Additional information about the check records.

Fields	Description
Resolution Details	Guideline to avoid check to go off against one or more tables.
Table	Table queried for this check.
Conditions	Conditions that filter records on a table. Only the records that meet the criteria are scanned.
Advanced	Option to execute the defined script against each record that matches the condition and control when a finding is generated if **Advanced** is true. If **Advanced** is false, a finding is generated for every record that matches the condition in the defined table.
Documentation URL	Documentation link for the check details. **Note:** A scan_user can't edit the **Documentation URL** field.
Run Condition	Boolean field that determines if the check should be run.
Table	Table from which records will be scanned
Conditions	Conditions that determine which records will run
Script	The script that executes against each record that matches the condition in the defined table. **Note:** This field shows up only if the **Advanced** check box is selected.

Create a column type check

Retrieve all records containing a specific column field type from all tables in an instance by selecting **Create a new Column Type Check**. The **Column Type Check** type implements the rule you created to iterate all records matching the target column field type.

Before you begin

Role required: scan_admin

```
Before performing this task you must complete Create a check.
```

Procedure

1. Select **Create a new Column Type Check**.
   A New record form appears.
2. On the form, fill in the fields.

Fields	Description
Name	Name of the check record.
Application	Application that has the check record.
Category	Category of the check record.
Priority	Priority at which the findings should be resolved.
Version	Current version of the check record. **Note:** If you want to update an already existing check, the **Version** field increments itself.
Active	Option to activate this check record.
Short Description	Mandatory description about the check records.
Description	Additional information about the check records.
Resolution Details	Guideline to avoid check to go off against one or more tables.
Column Type	Column type queried for this check.
Documentation URL	Documentation link for the check details. **Note:** A scan_user can't edit the **Documentation URL** field.
Run Condition	Boolean field that determines if the check should be run.
Column Type	Type of columns that will run.
Script	Script that executes against each record in a table containing a field of the defined column type.

Create a script only check

Create a check without specifying a table or a column type by selecting **Create a new Script Only Check**. You can verify meta data, configurations, and execute complex checks by writing your own script.

Before you begin

Role required: scan_admin

```
Before performing this task you must complete Create a check.
```

About this task

The script only check runs only during a full scan and executes only once during a scan.

Procedure

1. Select **Create a new Script Only Check** from the **Check Interceptor** list.
2. On this form, fill in the fields.

Fields	Description
Name	Name of the check record.
Application	Application that has the check record.
Category	Category of the check record.
Priority	Priority at which the findings should be resolved.
Version	Current version of the check record. **Note:** If you want to update an already existing check, the **Version** field increments itself.
Active	Option to activate this check record.
Short Description	Mandatory description about the check records.
Description	Additional information about the check records.
Resolution Details	Guideline to avoid check to go off against one or more tables.
Documentation URL	Documentation link for the check details. **Note:** A scan_user can't edit the **Documentation URL** field.
Run Condition	Boolean field that determines if the check should be run.
Script	Option to write a custom script to generate findings.

Create a linter check

Create a linter check to identify any issues in a script. When a linter check is run on a record, an abstract syntax tree for its code is generated. You can use the abstract syntax tree to analyze issues with the code.

Before you begin

Role required: scan_admin.

`Before performing this task you must complete` Create a check.

Procedure

1. Select **Create a new Linter Check**.
2. On the form, fill in the fields.

Linter Check form

Fields	Description
Name	Name of the check record.
Application	Application that has the check record.
Category	Category of the check record.
Priority	Priority at which the findings should be resolved.
Version	Current version of the check record. **Note:** If you want to update an already existing check, the **Version** field increments itself.
Active	Option to activate this check record.
Short Description	Mandatory description about the check records.
Description	Additional information about the check records.
Resolution Details	Guideline to avoid check to go off against one or more tables.
Documentation URL	Documentation link for the check details. **Note:** A scan_user can't edit the **Documentation URL** field.
Run Condition	Boolean field that determines if the check should be run.
Script	The script that executes against each record that matches the condition. By default, this field has an engine object to use. This engine object contains other objects like `engine.finding` which you can call `engine.finding.increment()` to manually increment the **Count** field on that specific finding. Another object is `engine.current` which is the current

Fields	Description
	GlideRecord that the check would be viewing when running. **Note:** This is new to Quebec release, where previously you would call `finding` and `current` directly.

Advanced linter check scripts

Linter check scripts helps you in writing checks that look for issues in scripts. When a linter check is run on a record, it provides an abstract syntax tree for its code. You can use this abstract syntax tree to analyze issues with the code such as too many nested if statements or usages of slow API in a while loop.

Linter check usage

Linter checks have a unique object to use called `rootNode`. This is the root node of the parsed abstract syntax tree (AST) for the script of the current record. This object has many functions, but the `visit` function is the most important. The `visit` function takes a callback function as a parameter which gets called on every node of the tree. The callback function then takes a node as a parameter which represents the current node during its iteration. You can return false from the callback function to stop iterating the tree early, otherwise it keeps visiting every subnode in the tree of the node you called `visit` on.

For example, if you want to scan for a deprecated API called `badFunction()`. You can write a Linter Check with a script in the following example:

```
(function(engine) {
 engine.rootNode.visit(function(node) {
  if (node.getTypeName() === "NAME" &&
    node.getNameIdentifier() === "badFunction" &&
    node.getParent().getTypeName() === "CALL") {
   engine.finding.incrementWithNode(node);
  }
 });
})(engine);
```

When you run a scan with this Linter Check, it checks every record in the scan with a script field. For example, if a record in that scan has a script that looks like the following, the scan picks up a finding for this record.

```
/*
badFunction()
*/
// badFunction()
function badFunction() {
 return;
}
var GoodClass = Class.create();
GoodClass.prototype = Object.extendsObject(Object, {
 badFunction: function() {
  // actually good
 }
});
var badFunction = ["badFunction", "badFunction()"];
```

```
badFunctionButSometimesGood();

badFunction();
```

Node functions

The following is a list of available functions to call on the node object.

- `getRootNode()`: Get the root node of the tree for this node
- `getParent()`: Get the parent of this node
- `getLineNo()`: Get the line number location for this node
- `getColumnNo()`: Get the column number location for this node
- `toSource()`: Get the source for this node. The source is based on this node and its children
- `getTypeName()`: Get the node type name of this node. For example, a function call in source is tokenized as a node with the type name of "CALL"

  Note:
  The values of `node.getTypeName()` are from different types of nodes in an Abstract Syntax Tree (AST). The possible values come from Rhino's AST parser itself.

- `getNameIdentifier()`: If the node type of this node is "NAME", then return the identifier, which is the string value of the name itself. A "NAME" node represents a simple name that is not a keyword, like a function name or a variable name.

  Note:
  If this node is not a NAME node, then the result is null.

- `getAbsolutePosition()`: Get the absolute position of this node. The absolute position is the number of characters from the start of the script to this node
- `compareTo(other)`: Compare this node with another node. The node with a greater absolute position and length will be the larger in comparison. `other` is a node to compare with this node
- `visit(callbackFunction)`: Visit each node in the subtree starting from this node and execute the given callback function on each node. `callbackFunction` is a function that will be executed on each node in the subtree of this node. This callback function takes a LinterCheckAstNode as a parameter which is will be the node being visited
- `debugPrint()`: Returns a string representation of the abstract syntax tree, starting from this node. Each line contains information about a node, and the indentation represents the hierarchy relationship between the nodes. The information in each line is arranged as follows

  1. Absolute position
  2. Node type name
  3. Position relative to parent
  4. Length
  5. Name identifier (if a NAME node)

For more information, see LinterCheckAstNode API - Scoped, Global.

Suites management

A suite is a collection of individual checks and suites that can be used for a scan. Suites can be created for specific business needs. All child suites are executed when a parent suite is used in a scan.

Navigate to **Instance Scan** > **Suites**. Select one of the suites from the list. Select **Execute Suite Scan** to run a suite scan. See Execute a suite scan for more information.

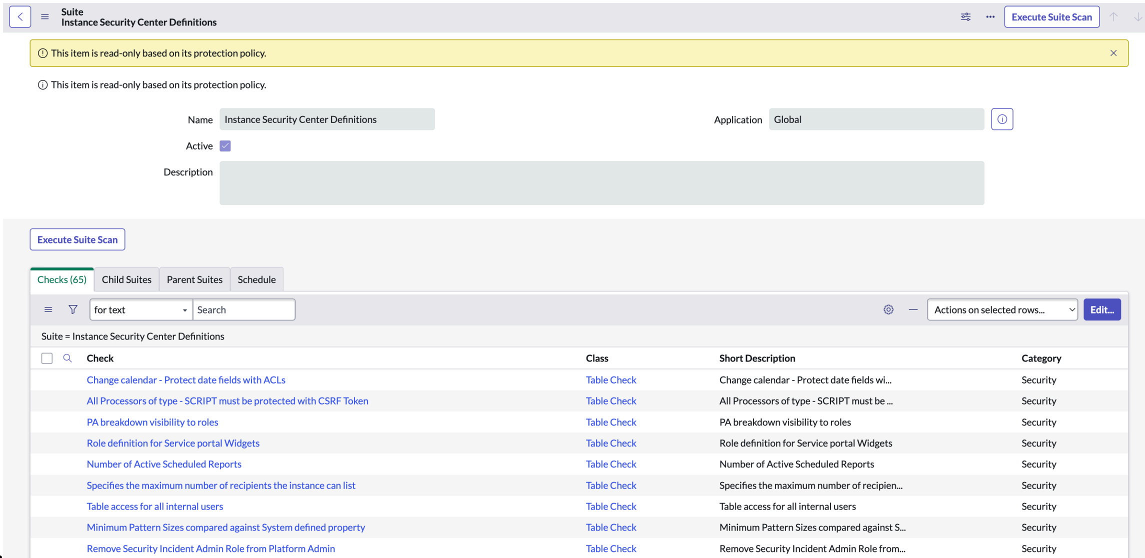

Note:
Some ServiceNow provided suites are protected suites and the checks in the suite are not modifiable. For example, if you have a suite from the app store, the suite is unmodifiable. You can't add a child suite to the unmodifiable suite. But you can add a parent suite because it doesn't affect the suite.

Create a check suite

Create a check suite to bundle a group of checks into a suite and execute it.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **Instance Scan** > **Suites**.
   A list of suites shows up.
2. Select **New** to create a new suite.
3. Fill in the details on the new suite record.

Fields	Description
Name	Name of the suite
Application	Scope of the suite
Active	Option to activate this suite
Description	Information about the suite

4. Select **Submit** to create a new check suite.

   Note:
   The suites that you create are completely modifiable. You can modify the checks in the suite and can also add child and parent suites.

Configuring Instance Scan

Execute the different kinds of scans that Instance Scan offers to keep a check on the health of your instance. You can also schedule a full scan so you can ensure the health of your instance even without an active session.

Note:
Instance Scan scans only customized records, but point scan scans all records including the default records.

1. Execute scans: Execute the different types of scans available as per your requirements.
2. Schedule scans: Create a schedule to regularly trigger either a full or a suite scan even when you don't have an active session.
3. Monitor scans: Track the progress status of the executed and scheduled scans.
4. Review scans: Review your executed scans with the findings and results components.

Instance Scan roles

Instance Scan is installed with these roles.

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

Instance Scan user [scan_user]

The scan_user role can run different types of scans and view the findings and results.

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.

Execute a point scan

Execute all applicable checks against a single record, update set or an application by selecting **Run Point Scan**.

Before you begin
Role required: admin

About this task

For example, if you execute a point scan against a business rule, only the checks that are applicable to the business rule table run, and only that single target record is scanned. If an update set scan or an application scan is executed, all records related to that update set or application are scanned.

Procedure

1. Navigate to **All** > **Instance Scan** > **Checks**.
2. Navigate to the record in a table applicable for point scan.

   Note:
   The **Run Point Scan** related link appears in the UI only for applicable records.

3. Scroll down to the **Related Links** section.
4. Select the **Run Point Scan** related links.

   Note:
   The **Run Point Scan** related link is available only if all the following conditions are true.

   - Available checks that are applicable to the record
   - The user has read access to the record
   - The record is on a table that extends sys_metadata
   - The role of the user must be scan_user
   - The system property *glide.scan.enable_point_scan_ui_action* must not be false

   The progress tracker appears showing the status of the scan.
5. Select **Go to Result**.
   The Scan Result record appears.
6. Select **Checks** related list to look at the findings.
   If you want to review the failure reasons, select the **Failures** related list.

Result

A scan of all applicable checks against only a single record is executed.

Execute an update set scan

Use update set scan to execute applicable checks against records that are affected by the update set.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Update Sets** > **Local Update Sets**.
2. Select the desired update set from the list.
3. Click **Scan Update Set** from the **Related Links** list.

Execute an app scan

Scan the installed files of an application as well as the application record itself with applicable checks by executing an application scan.

Before you begin
Role required: admin

About this task
For example, if the application installs a business rule, checks that are relevant to sys_script table are executed against that business rule. To trigger an application scan, go to the record of the application you desire to scan and select **Scan Application** form link UI action.

Procedure

1. Navigate to **System Application** > **All Available Applications** > **All**.
   A list of all available applications and plugins appears.
2. Select the menu icon of the desired application.
3. Select **Go to application** option from the menu icon.

   Note:
   The **Go to application** option appears only if the selected application is not a plugin.

4. Select **Scan Application** to run app scan.

Execute a test scan

Execute a test scan to verify if the check works as expected. The test scan enables you to test a single check instead of a full scan by selecting a single check and selecting **Test Check** on the Check form.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Instance Scan** > **Checks**.
   A list of checks appears.
2. Open the check record you want to test from the list.
3. To start the test, select **Test Check**.
   The status of the scan appears in the progress tracker. The **Findings**module in the application navigator shows the number of records found from running the check test.

Execute a full scan

Execute a scan for the entire instance by selecting **Execute Full Scan**. Implementing a full scan runs all the active checks present in your instance.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Instance Scan** > **Checks**.
2. To run all the active checks in your instance, select **Execute Full Scan**.
   The progress tracker shows the status of the scan.

Note:
Running of multiple full scans simultaneously is not allowed. If two or more checks overlap during full scan execution, only the first check succeeds. The first check keeps running while the other overlapping checks fail immediately.

3. To find the **Result Number** and the number of checks that ran as part of the full scan, select **Go to Result**.
4. **Optional:** From the related list, select **Checks** to view the list of all the checks that ran as a part of the full scan.
5. **Optional:** From the related list, select **Errors** to review the errors encountered during the full scan.

Result
A scan of the whole instance is executed.

Schedule a full scan

Create a schedule to regularly trigger a full scan even when you don’t have an active session. A schedule creates an established baseline for the health of the instance and provides a comparison to configuration and instance health over time.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Instance Scan** > **Checks**.
2. Select **Schedule Full Scan**.
3. On the Trigger full scan form, fill in the fields.

Field	Description
Name	Name of the check record.
Active	Option to execute the scan at a scheduled time.
Application	Application where the job belongs.
Conditional	Option to select conditions under which the scan runs.
Run	List used to select the scheduling of full scan.
Run this script	Selection to run the default script.
Update	The script that is updated for full scan.

Field	Description
Execute Now	The default script that is executed for full scan.
Delete	The script is deleted for full scan.

4. To run the full scan, select **Execute Now**.
   A list of findings is displayed.
5. From the application navigator, select **Results**.
   The status of the execution is displayed along with the **Result Number**.
6. To display the list of findings, select the **Result Number**.
7. **Optional:** From the application navigator, select **Dashboard** to display data about the findings.

Result

A schedule to trigger a full scan regularly even without an active session is created.

Execute a suite scan

Execute suites of available checks against a specific target to have targeted scan results by selecting **Execute Suite Scan**.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **Instance Scan** > **Suites**.
   A list of all suites displays.
2. Select one of the suites from the list.
   The form for the selected suite displays.
3. Select **Execute Suite Scan**.
   The **Scan Suites Now** modal displays.
4. Select a target for the scan.

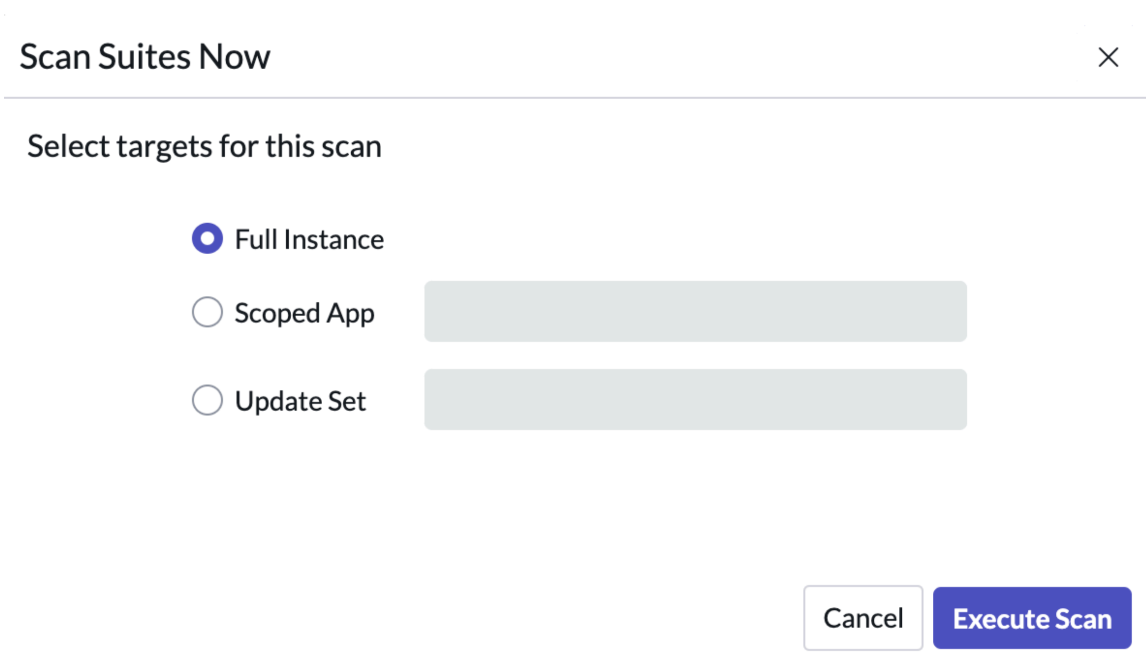

- **Full Instance**: Scans all the available records in the instance.
- **Scoped App**: Scans selected scoped apps. You can select multiple scoped apps.
- **Update Set**: Scans multiple update sets.

5. Select **Execute Scan**.

> **Note:**
> If a record is inactive, the checks don't generate any findings for that record.

Schedule a suite scan

Create a schedule to regularly trigger a suite scan even when you don’t have an active session.

Before you begin
Role required: admin.

Procedure

1. Navigate to **All** > **Instance Scan** > **Suites**.
   A list of suites displays.
2. Select a suite from the list to schedule a suite scan.
3. Select Schedule related list.
   A list of scheduled scans displays.
4. Select **New** to create a schedule for the suite scan.
   The **Schedule Suite Scan** modal displays.
5. Select a target for the scan.
   - **Full Instance**: Scans all the available records in the instance.
   - **Scoped App**: Scans selected scoped apps. You can select multiple scoped apps.
   - **Update Set**: Scans multiple update sets.
6. Select **Schedule Scan**.

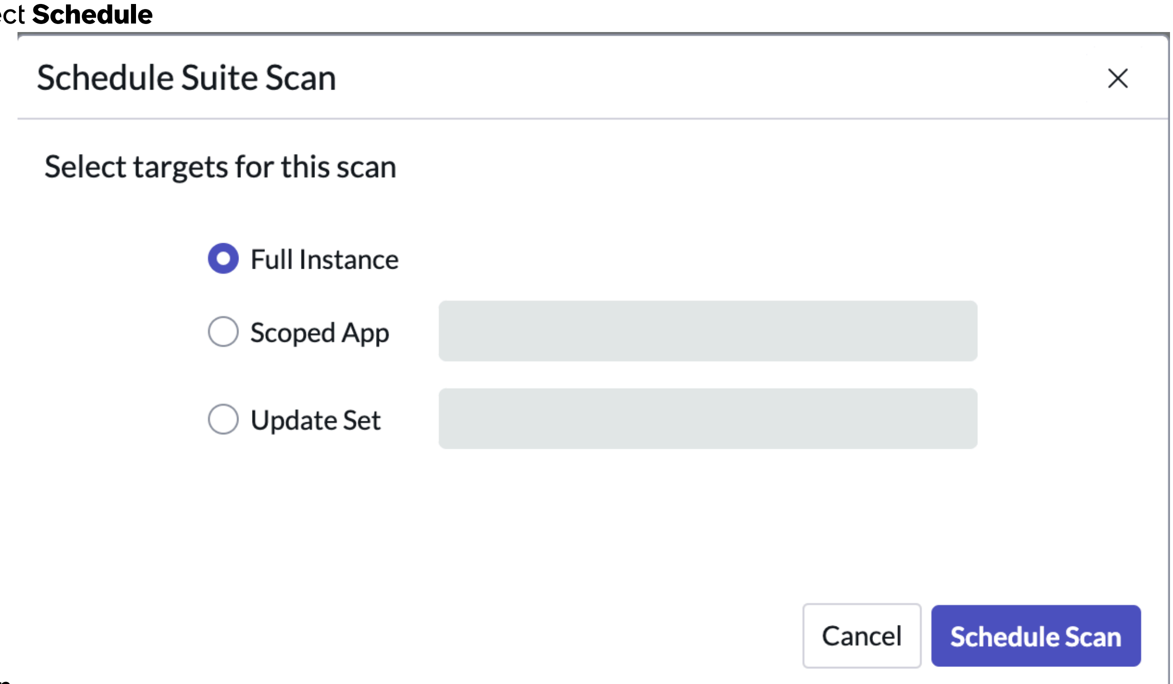

Note:
If a record is inactive, the checks don't generate any findings for that record.

The **Scheduled Scan** form displays.

7. On the form, fill in the fields.

Linter Check form

Fields	Description
Name	Name of the scheduled scan.
Application	Application that has the record.
Combo	References a record which has a list of sources and a list of targets to run the scan against.
Run	Frequency at which the scan should be scheduled to run. ◦ Daily: Scan is scheduled to run daily. Time: The time at which the scan is scheduled to start daily. ◦ Weekly: Scan is scheduled to run weekly. ▪ Day: The day in the week on which the scan is scheduled to start weekly. ▪ Time: The time at which the scan is scheduled to start. ◦ Monthly: Scan is scheduled to run monthly. ▪ Day: The day in the month on which the scan is scheduled to start monthly. ▪ Time: The time at which the scan is scheduled to start. ◦ Periodically: Scan is scheduled to run periodically. ▪ Starting: Date and time at which the scan is scheduled to start. ▪ Repeat interval: Date and time at which the scan has been scheduled to repeat. ◦ Once: Date and time at which the scan is scheduled to start. ◦ On Demand: Scan is not scheduled to run and runs only manually when **On Demand** option is chosen for the **Run** field. ◦ Business Calendar: Business calendar time of the user to schedule the scan run either at the beginning of that period or at the end, depending on their selection.
Active	Option to activate the record.

Fields	Description
Conditional	Option to write your own script condition, which is always verified before scheduled scan runs.
Run as	User role used to run the scan.
Run as tz	Timezone used to run the scan.

8. Select **Submit**.
   The scan that has been scheduled displays under Schedule related list.

Progress status of a scan

You can check the progress status of a scan by checking the progress tracker.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Instance Scan** > **Checks**.
   A list of checks appears.
2. Run the required scan.
   The progress tracker displays the status of the scan run. You can close the progress tracker and view it later without having to stay in the modal for the execution to complete. To see the status later, go to **Results** in the application navigator. Select the required scan results record to review its status.
3. To see the status of the scan, select **Go to Result**.
   If the scan completes, a list of findings is displayed. You can also find all the checks that ran as a part of the scan by selecting the Checks related list. If any of the checks fails, click Failures related list to review them. If the full scan takes some time to complete, you can check the status by selecting **Results** from the application navigator.

Execute a reactive scan

Execute a reactive scan to set up a scan that runs automatically on an execution tracker failure.

Before you begin

Role required: admin.

About this task

For many failed processes on your instance, you can find a corresponding record on the execution tracker. The progress window shows the status of the test. If a process fails on an execution tracker, Instance Scan can watch for that failure on the progress status and run a scan automatically.
The **Trigger Type** in the Scan Trigger form should match the **Source Table** field on the sys_execution_tracker record of the failed process to enable the connection between the two. For example, if the failed test is in the Update Sets list, you can set up the **Trigger Type** as sys_update_set. This enables Instance Scan to be connected to the execution tracker test failure.

Procedure

1. Enter `scan_trigger_list.do` in the navigation filter.
   A list of scan triggers shows up.
2. Select the required trigger from the list to update it.

Note:
The **Trigger Type** is based off of table names. You can also create a custom trigger type in your codes.

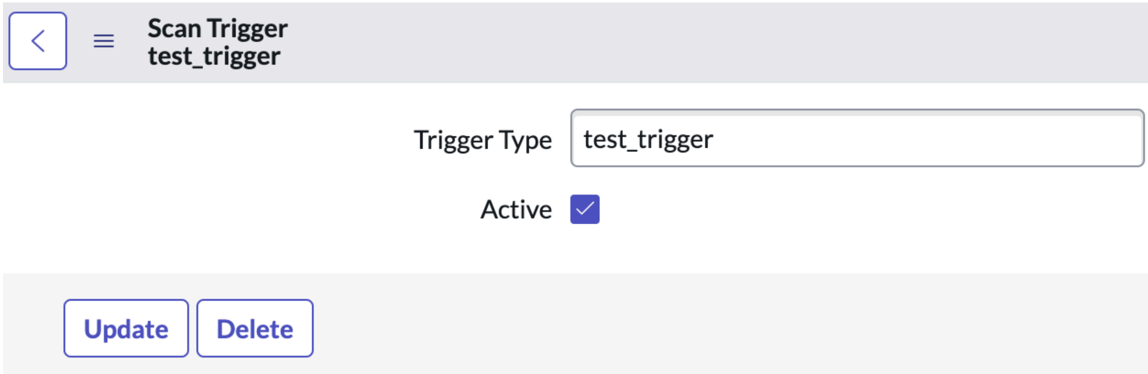

The reactive scan is now set up to automatically scan the execution tracker of the failed process which appears on the scan results table.

Reviewing of scans

After the scans have been executed, you can review them with the findings and results components.

- Results: An Instance Scan result reports the status and type of scan. You can also see all the checks that ran as part of the scan and all other information related to the scan such as errors and scan logs. Use Scan Results dashboard to get detailed information about an executed scan.
- Findings: A finding is a reference to a record that has violated a check which was used in a scan. You can also mute a finding if you don't want a finding to be reported in the next scan.

Results

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

On selecting **Go to Result** from the progress tracker of a scan, the Scan Result is identified by a **Result Number**. Select **Results Dashboard** to view the scan results.

The **Scan Type** field states the type of scan that was executed. The **Status** field shows the status of the scan which can either be **In progress**, **Complete** or **Failed**. The **Execution Time** field shows the duration of the scan in milliseconds.

There are the following related lists in the **Scan Result** form.

Scan Findings

The Scan Findings related list shows up all the findings encountered during the execution of the checks.

Suites

The Suites related list shows up all the suites that ran as part of the scan.

Checks

The Check related list shows all the checks that ran as part of the scan.

Failures

The Failures related list shows up all the checks that failed during the execution of the scan. It also shows up the reason of its failure in the form of an error message.

Scan Log

The Scan Log related list shows up messages encountered during the entire scan.

Scan Statistics

The Scan Statistics related list shows up

Target

The Target related list shows all the targets against which the checks have executed.

Scan results

Scan Results dashboard helps you with an overview of all details of an executed scan.

Navigate to **All** > **Instance Scan** > **Results**. Select a scan result from the Scan Results list. Click Results Dashboard link on the Scan Result form.

Select **Rescan** to execute the same scan again. If you want to see all the findings grouped by checks, select **Findings** at the top right corner.

Fields	Description
Previous Scan	Details about the previous scan. • Result number: Link to the previous scan result record • Outstanding Findings: Total findings whose status has not been marked as Resolved.

Fields	Description
	• Resolved Findings: Total findings whose status has been marked as Resolved. **Note:** Not all findings that have been marked Resolved are fixed. If a finding has not been fixed and has been marked Resolved, it still adds to the count of total findings. • Date: Date and time on which the previous scan started • Duration: Total time needed to complete the previous scan
Latest Scan	Details about the recent completed scan. • Result number: Link to the latest scan result record • Findings: Total number of findings found in the latest scan • Target: Target against which the scan is executed ◦ Full Instance: Scans all the available records in the instance ◦ Scoped App: Scans selected scoped apps. You can select multiple scoped apps. ◦ Update Set: Scans multiple update sets. • Date: Date and time on which the scan started • Duration: Total time needed to complete the scan
Scan Tasks By Priority	Scan tasks sorted by priority. The priority comes from the checks to which these tasks belong. **Note:** The states come from the states of the tasks associated with the findings.
Average Findings Over Time	Reporting of the average number of findings in a day.
Category Breakdown	Category breakdown of the checks that have run as part of the scan
Product Breakdown	Product family of each of the scan findings records.

Fields	Description
	Note: Incident records don't have product family.
Findings By Developer	Name of the user whose changes generated the findings. **Note:** Only the top 5 user names appear on the list. Click **View All** to view all the findings grouped by users.

If a scan is in progress and at least one of the checks fails, the following warning message shows up. 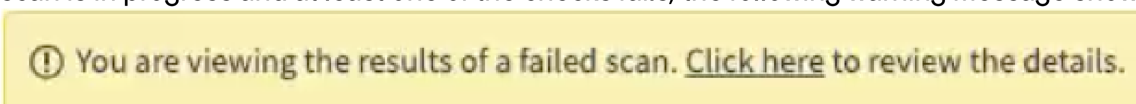

Findings

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

Each finding is specific to a result and each result is specific to a scan that has been executed.

Note:
When you run a scan, you definitely get a result of the scan but might not get any findings. There might be no findings encountered during the execution of the checks.

Note:
If a record is inactive, the checks don't generate any findings for that record.

The Scan Finding gives information about the check that has been violated by a record. The **Source Table** and **Source** fields point to the record that has violated a rule from a check during the execution. You can find more information about the record by clicking the information icon. The **Count** field shows the number of times the record has violated the check rules. The finding also shows the version of the check that generated the finding. The **Product Family** field shows you the product family for each check. The **Task** field helps you to assign a task to an individual or a group. See Creating scan tasks for more information. If you want to avoid a finding from being generated from any future scans, use the **Mute Rule**.

Note:
If you use the **Mute Rule**, you must also select a reason for muting the finding.

Mute a finding

Mute a finding on the Scan Finding form to avoid the finding from being reported in the later scan results.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **Instance Scan** > **Results**.
   A list of Scan Results shows up.

2. Select the scan result you ran from the list.
   The Scan Result form for the scan shows up.
3. Select the Scan Findings related list.
   A list of findings for the scan shows up.
4. Use the information icon to open the finding you want to mute.
5. Select **Mute** to avoid the finding from being reported in the next scan.

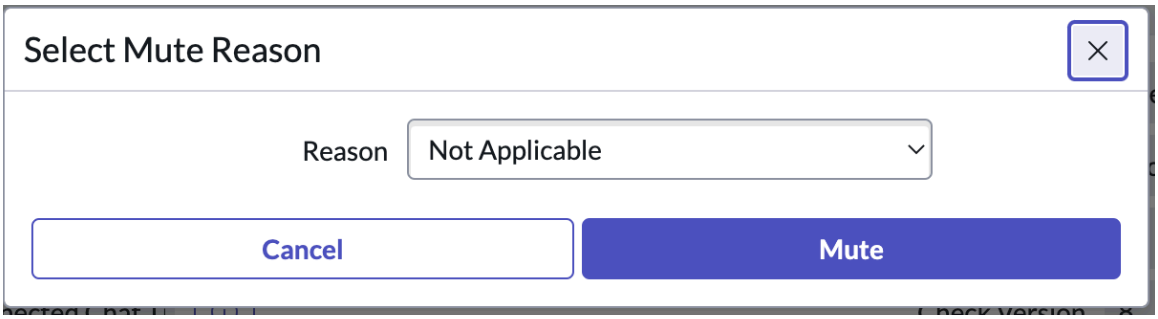

   Note:
   The Select Mute Reason modal shows up only if you select **Mute** on the Scan Finding form.

6. Select the reason for muting the finding and then select **Mute** on the modal.

   Note:
   Selectc **Unmute** if you want to the finding too be reported in the later scans. The **Unmute** button is visible only when you have muted a finding.

Creating scan tasks

Create a scan task to facilitate task assignments from the finding of a record.

Before you begin

Role required: scan_user

Procedure

1. In the navigation filter, enter `scan_task_list.do`.
2. Select **New** to create a new scan task.
   The **Scan Task** form appears.
3. Fill in the following fields in the form.

Scan Task form details

Field	Description
Number	Auto-generated task number
Assigned to	Person responsible to work on the task
Assignment group	Group responsible to work on the task
Active	Option to activate task
State	State of the task.

Field	Description
	○ **New** ○ **In Progress** ○ **On Hold** ○ **Resolved**
Short description	Brief description about the task
Description	Additional information about the task
Work notes	Additional notes

Note:
The scan_user can assign the task only to a non-scan_user. The user to which the task has been assigned also gets the same privilege as the scan_user.

Using Instance Scan dashboard

The Instance Scan dashboard is a system wide visual representation of the health of your instance. The dashboard helps you to manage and analyze the results of full scan against your instance.

The dashboard offers several options to filter your search results to find the exact check finding that might be causing some issues. The **Instance Scan Category** is the filter picker that shows all the categories of all the checks that were scanned. You can select the type of scan category that you want the dashboard to display as a part of the scan results. **Instance Scan Priority** helps you to select the priority level of the results to decide which scan results must be addressed first. This priority level is often referred to as the rating level of the scans.

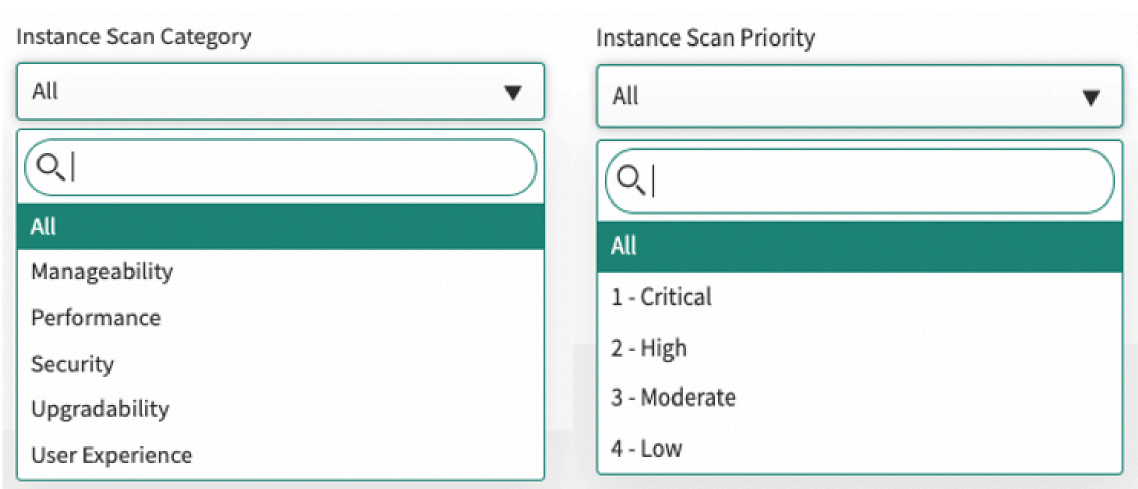

Note:
The dashboard is only applicable for full scans. The **Instance Scan Category** is only available if you have data or a finding in the result set that are applicable to the categories.

For example, you can use the dashboard to compare the results of the check findings. If you see a sudden change in the number of check findings, select the particular scan result from the Scan Results List.

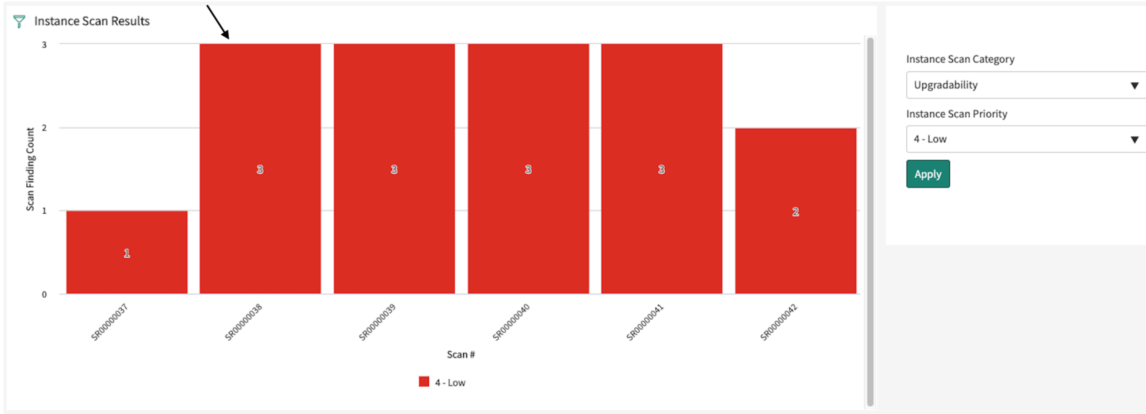

You can review all the findings for the scan that generated an unexpected number of findings. This list displays all the findings that are related to one or more of all the scans.

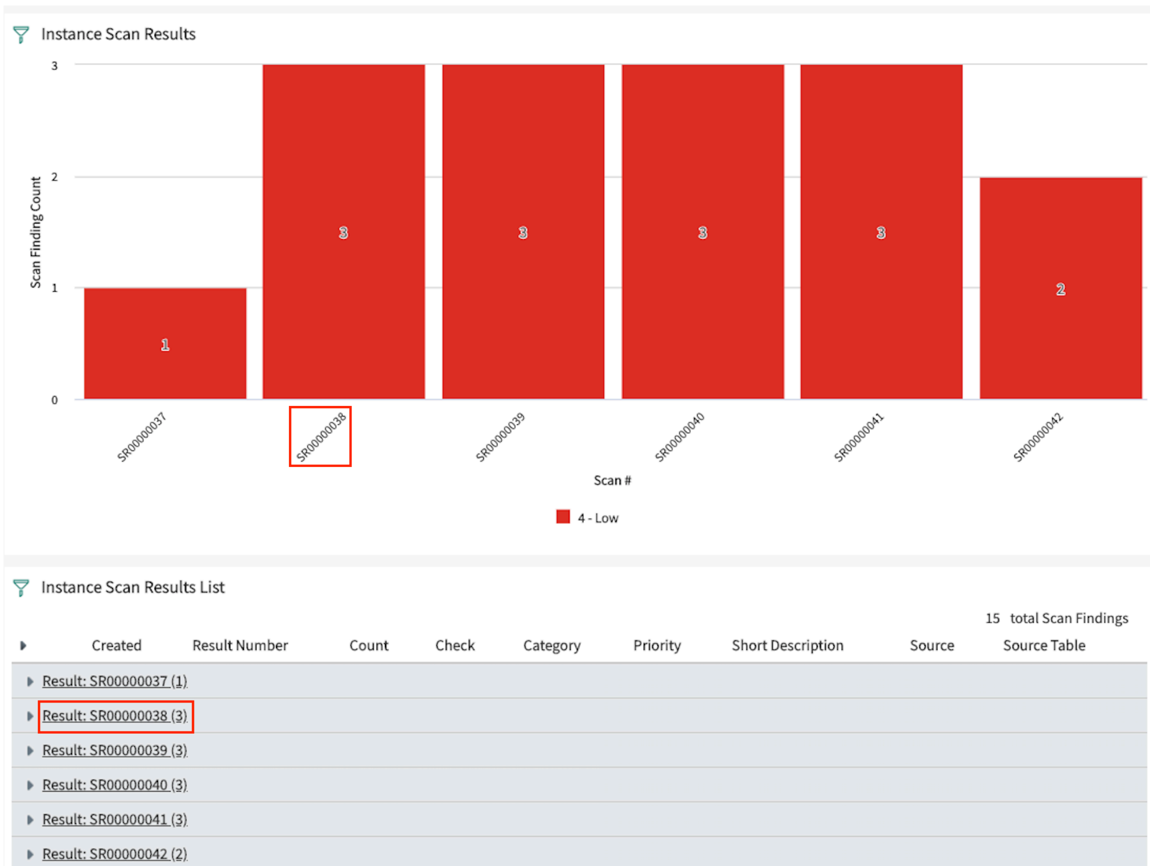

You can also directly select the column with the required findings to display the related list for a particular scan.

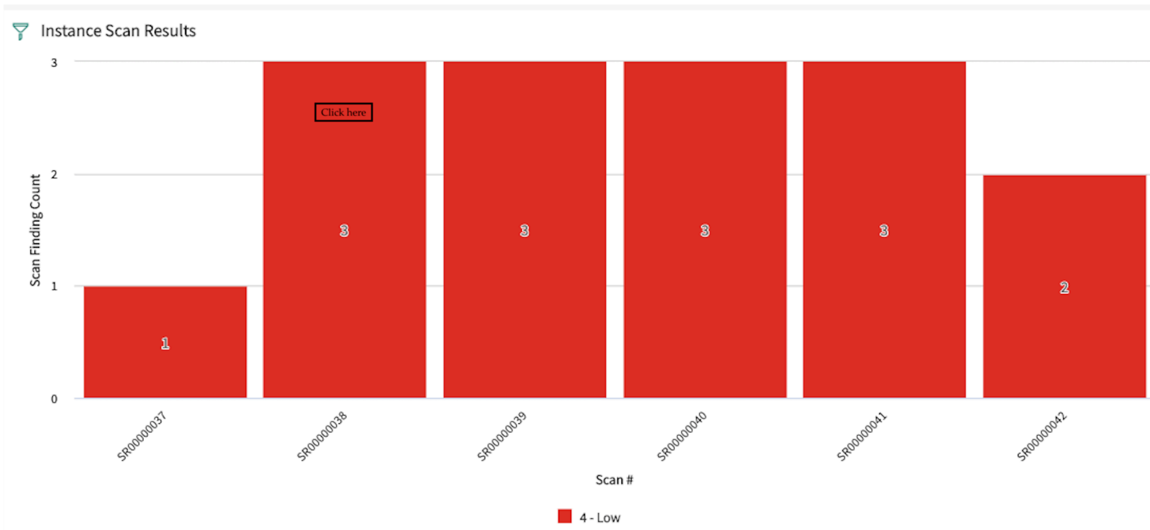

The list also indicates the source from which the finding has been retrieved. You can export the list view data to an external source such as a .csv file or an .xlsv file for further performance comparison and analysis.

Implement the Instance Scan dashboard

Compare and analyze the results of the scans in your instance by representing it on the premium Performance Analytics (PA) dashboard.

Before you begin

You've created some checks to execute the full scan. See Create a check and Execute a full scan for more information.

Role required: admin.

Procedure

1. Navigate to **All** > **Instance Scan** > **Dashboard**.
2. Select an option from the **Instance Scan Category** or **Instance Scan Priority** list.

   Note:
   The dashboard is applicable only for full scans. The **Instance Scan Category** appears only if you have data or a finding in the result set that are applicable to the categories.

3. Select the scan result from the Scan Results list that generates an unexpected number of check findings.
   A related list of findings related to only that scan is expanded.

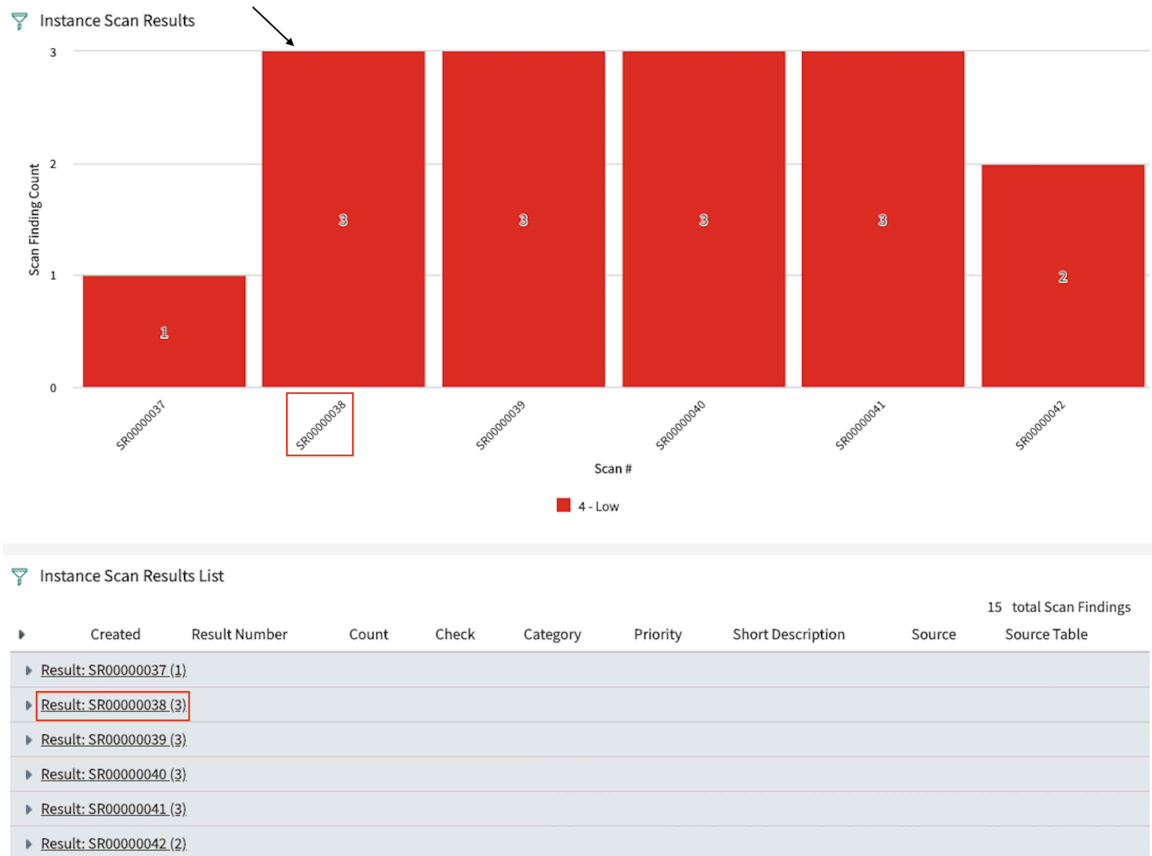

4. **Optional:** Select the scan column directly from the dashboard to display the related list of checks.

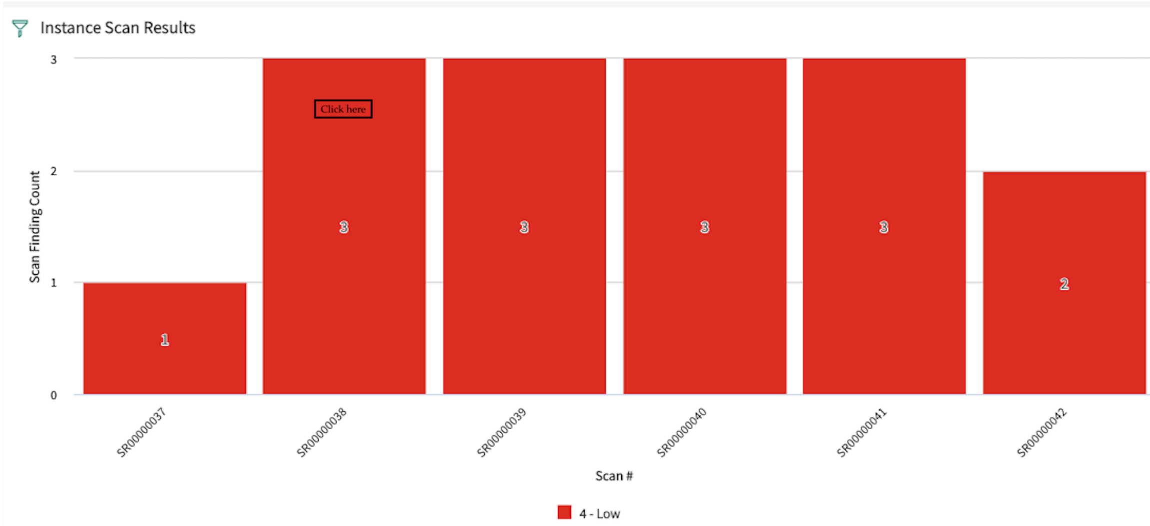

The related list of checks appears. The list also indicates the source that the finding has been retrieved from.

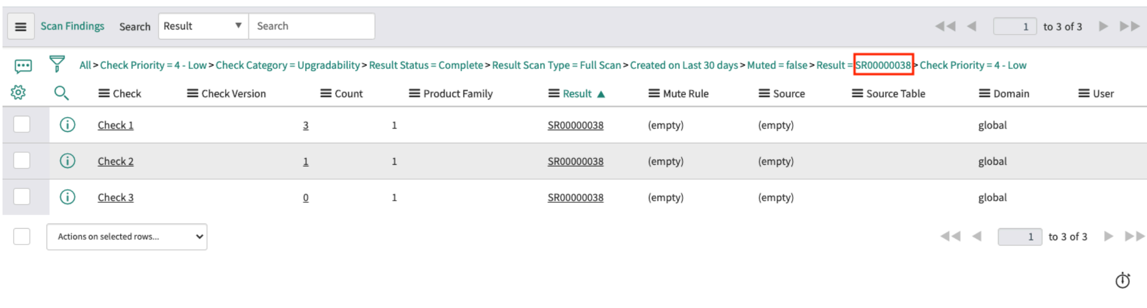

Result

A dashboard showing the results of scans is created.

Instance Scan table cleanup policies

Instance Scan offers table cleanup policies to erase previous scan results.

The cleanup policies are based on the time that the scans are created. The timer starts when you run a scan. **Table Cleanup** is run against the Auto Flushes [sys_auto_flush] table and cleans up the previous scan results.

- If the result records are older than 90 days (**7,776,000** seconds), the first policy clears all scan result records and related records.
- If the following conditions are true, the second policy deletes all scan result records and related records.
  - The **Scan type** is **Test Scan**.
  - The records are older than 14 days, and **Age in seconds** is equal to or greater than **1,209,600** seconds.

You can also modify these policies by entering any customized time period. See Customize table cleanup policies for more information.

Customize table cleanup policies

Clean up previous scan results at a regular interval by customizing the table cleanup time period.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Instance Scan** > **Table Cleanup**.
2. Select one of the cleanup policies from the Auto Flushes list.
3. Customize the **Age in seconds** field as required.
4. To save the customized table cleanup policy, select **Update**.

Timeout threshold

A timeout determines the execution threshold of a scan from running long scans. For example, any scan running longer than the set time period results in a failure.

If the scan or the check takes longer than the set time period set by the quota rule, an error message appears.

> **Note:**
> By default the scans are limited to three hours and individual checks are limited to 10 minutes.

To learn how to configure these thresholds, see Implement a scan timeout threshold and Implement a check timeout threshold.

Implement a scan timeout threshold

Set the execution time of a single scan by implementing **Transaction Quota Rules**. Setting of a quota rule prevents your instance from running a long scan.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Transaction Quota Rules**.
2. To set a definite execution time, select **Scan timeout**.
3. Enter the definite execution time in the **Maximum Duration (seconds)** field and select **Update**.

   > **Note:**
   > The minimum allowed timeout threshold is 5 seconds. If you set the timeout to anything less than 5 seconds, the system still considers it to be 5 seconds. By default, it has been set to 10,800 seconds.

   The scan fails if the execution time exceeds the set timeout. For more information on the type of scan you want to run, see Configuring Instance Scan .

Implement a check timeout threshold

Set the execution time of an individual check by implementing timeout system property. Setting of a timeout threshold prevents your instance from running a long check.

Before you begin
Role required: admin

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. From the System Properties list, select `glide.scan.process_check.time_out`.

   Note:
   If `glide.scan.process_check.time_out` is not present in the list, see Create a new check timeout system property for more information.

3. In the **Value** field, set the execution time for the check in seconds.

   Note:
   The minimum allowed timeout threshold is 5 seconds. If you set the timeout to anything less than 5 seconds, the system still considers it to be 5 seconds. By default, it has been set to 10,800 seconds.

Result
A timeout threshold for a check is set by the user. Individual checks will be cancelled after the set time threshold exceeds.

Create a new check timeout system property

Create a new timeout threshold property for a check if the `glide.scan.process_check.time_out` system property is not present. Setting of a timeout threshold prevents your instance from running a long check.

Before you begin
Role required: admin.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Click **New** to create a new timeout system property.
3. In the **Name** field, enter `glide.scan.process_check.time_out`.
4. In the **Value** field, enter the check execution time in seconds.

   Note:
   The minimum allowed timeout threshold is 5 seconds. If you set the timeout to anything less than 5 seconds, the system still considers it to be 5 seconds. By default, it has been set to 600 seconds.

Platform Performance

Ensure system optimization through performance monitoring and diagnostics.

Get started with Platform Performance

Performance Dashboards

View a wide range of platform performance metrics for your instance and for the machine on which your instance is running.

Performance Metrics

You can view a wide range of platform performance metrics for your instance and for the machine on which your instance is running.

Transaction Quotas

Transaction quotas allow you to define a quota policy for different types of transactions.

System Diagnostics Stats Tool

To aid in performance evaluation, the Stats Tools records statistics for system activities that affect performance such as the execution of queries, scripts, and transactions.

Platform performance overview

Ensure system optimization through performance monitoring and diagnostics.

Term	Description
Application server response	Time for the application server to process a request and render the resultant page.
Network latency and throughput	Time for the network to pass your request to the server and the response back.
Browser rendering and parsing	Time for your browser to render the HTML and parse/execute Javascript.
Instance Cache	Amount of system resources available for processing.

Instance cache effects on performance

There is a performance degradation whenever you purge and rebuild the instance cache. During core business hours, avoid or minimize any of the following actions that will cause a purge and rebuild of the instance cache.

- Adding or updating system properties
- Adding or updating dictionary entries
- Committing update sets
- Adding or updating translations

Transaction log response times

The instance automatically logs the vital statistics of every transaction it processes, and that information is available to you as an administrator.

To look at the log, navigate to **System Logs** > **Transactions**.

To see the average response time of all listed transactions, right-click the column **Response time**. Select **Configure** > **List Calculations**, and then select the **Average value** check box.

It is a good practice to limit the list to those transactions that took place during the time period of interest. The default filter returns transactions from today.

Transaction log

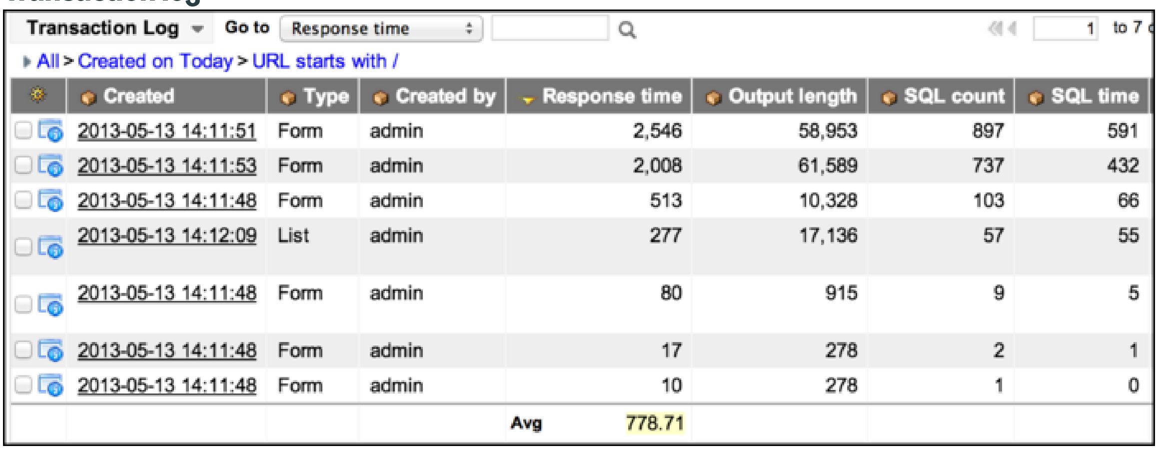
Transaction Log Go to Response time 1 to 7

All > Created on Today > URL starts with /

Created	Type	Created by	Response time	Output length	SQL count	SQL time
2013-05-13 14:11:51	Form	admin	2,546	58,953	897	591
2013-05-13 14:11:53	Form	admin	2,008	61,589	737	432
2013-05-13 14:11:48	Form	admin	513	10,328	103	66
2013-05-13 14:12:09	List	admin	277	17,136	57	55
2013-05-13 14:11:48	Form	admin	80	915	9	5
2013-05-13 14:11:48	Form	admin	17	278	2	1
2013-05-13 14:11:48	Form	admin	10	278	1	0
			Avg 778.71			

For each completed transaction, the following information is available (times are in milliseconds):

- Date/time, User ID, IP address, and URL of the transaction.
- Total response time, which does not include the browser time because the server does not have that information.
- **Network time**: network transmission time, both from and to the user.
- **SQL time**: time spent executing SQL commands.
- **SQL count**: number of SQL commands executed.
- **Business rule time**: time spent processing business rules.
- **Business rule count**: number of business rules executed.
- **Output length**: how many bytes the transaction returned, after any compression.

Response time on forms

A response time indicator (clock) appears at the bottom right of many forms and lists.

This indicator provides the processing time for a completed transaction, including the total time and the time for each step. Click the icon to show and hide the response time details. Point to the icon to see a tooltip with the details. The following example shows the response time for retrieving a filtered list in a demo instance.

Response time

In this example, the transaction took the following amount of processing time:

- 1019 milliseconds total time
- 3 milliseconds moving data across the network
- 526 milliseconds on the server

> **Note:**
> Server time is calculated using the following formula:
>
> ```
> server time = (client response time - browser time - client
>  network time)
> ```

- 490 milliseconds in the browser, rendering the HTML and parsing and executing JavaScript

Response does not appear for simple operations, such as paging, changing a list sort order, or for the first transaction in a session.

To view a detailed breakdown of the browser processing time on forms, click **browser**.

Browser response time details

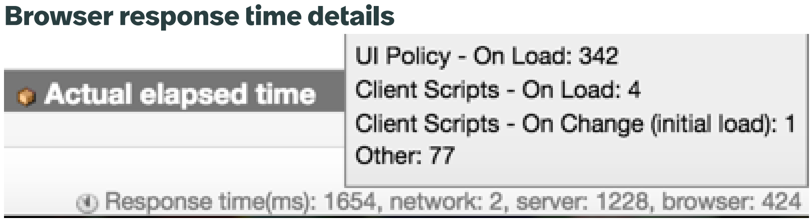

Administrators can disable the response time by setting the *`glide.ui.response_time`* property to **false** in the System Properties [sys_properties] table. In the Navigation filter, enter `sys_properties.list` and search for *`glide.ui.response_time`*.

When you review response times, look for the following issues.

1. A period where all transactions took an unusually long time. For example, transactions that normally took 1 second took 15 seconds between 11:00 AM and 11:20 AM. This issue can indicate that an unusual load, such as a large report, ran on that app server during that time.
2. A specific transaction which repeatedly took an unusually long time. For example, the list of closed incidents sorted by short description took 30 seconds each time it was displayed. This issue can indicate that a particular transaction put an unusual database load on the system, such as sorting 500,000 records on an unindexed field.

Try the following actions to improve performance.

1. Look for one or more transactions that span the entire window. For example, you observe that the response was slow for six minutes and one transaction ran the entire time. You can try adding additional indexing to the database to make the transaction faster. Certain types of queries are always going to run more slowly than others, regardless of indexing.
2. Ensure that a cache flush is not being run during business hours. Cache flushes are intended to prevent older data from interfering with changes and updates, and are performed automatically when using update sets. Scheduled cache flushes, using cache.do, can affect

overall performance, and degrade system response times. Do not run cache flushes during business hours, and do not trigger cache flushes automatically.

3. If you cannot find any specific issues when experiencing slow response time, contact support to see if there are global issues with the application server hardware.

Client transaction timings

The Client Transaction Timings plugin provides extra information about the amount of time spent on both the client and server side, and by the browser and network. This feature not only helps find long-running processes, but also provides information about where in the process the performance issue is caused.

Network response times

Troubleshooting a poor network response time can be difficult, but there are certain quick tests you can perform.

One clear indicator of a network issue is when you find that users in one location have good performance, and users in another location have poor performance. That tells you that the server and application are fine. Assuming that browser settings are identical, the only meaningful difference is the network.

Ping times

The coarsest measure of network response time is a ping. A ping measures the total time for a packet to make it from the source machine to the target and back again.

To perform a ping in Windows, bring up a command window (DOS prompt) and type:

```
ping -t <yourinstancename>.service-now.com
```

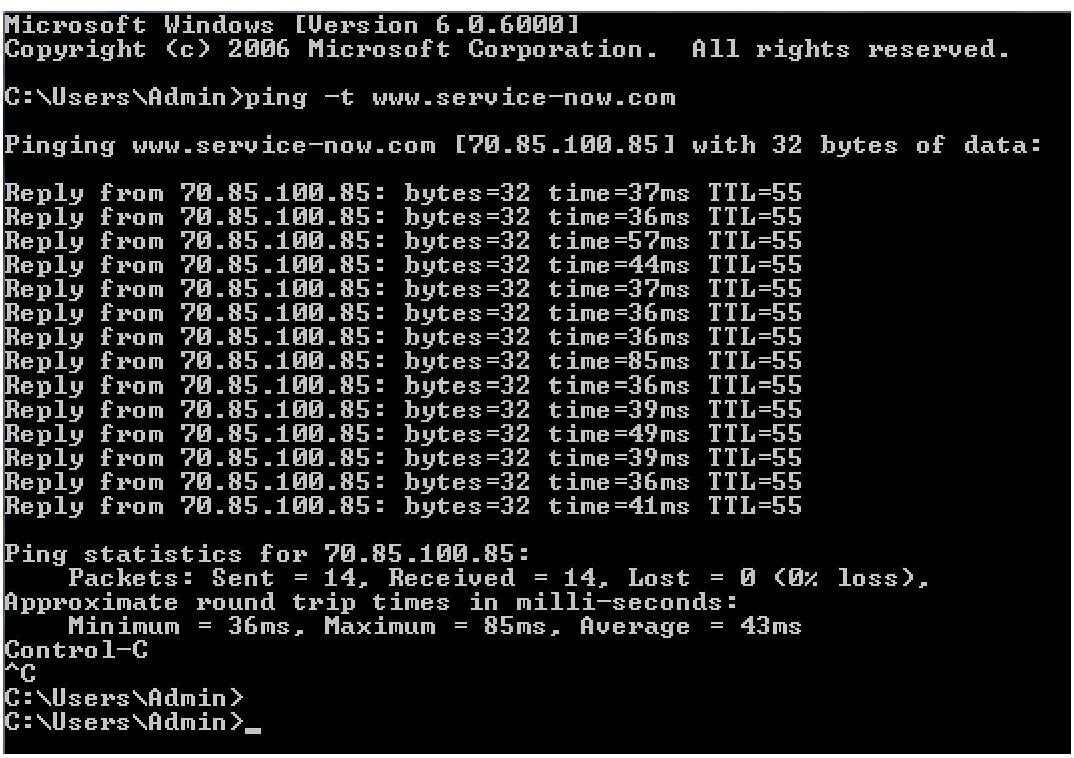

```
Microsoft Windows [Version 6.0.6000]
Copyright (c) 2006 Microsoft Corporation.  All rights reserved.

C:\Users\Admin>ping -t www.service-now.com

Pinging www.service-now.com [70.85.100.85] with 32 bytes of data:

Reply from 70.85.100.85: bytes=32 time=37ms TTL=55
Reply from 70.85.100.85: bytes=32 time=36ms TTL=55
Reply from 70.85.100.85: bytes=32 time=57ms TTL=55
Reply from 70.85.100.85: bytes=32 time=44ms TTL=55
Reply from 70.85.100.85: bytes=32 time=37ms TTL=55
Reply from 70.85.100.85: bytes=32 time=36ms TTL=55
Reply from 70.85.100.85: bytes=32 time=36ms TTL=55
Reply from 70.85.100.85: bytes=32 time=85ms TTL=55
Reply from 70.85.100.85: bytes=32 time=36ms TTL=55
Reply from 70.85.100.85: bytes=32 time=39ms TTL=55
Reply from 70.85.100.85: bytes=32 time=49ms TTL=55
Reply from 70.85.100.85: bytes=32 time=39ms TTL=55
Reply from 70.85.100.85: bytes=32 time=36ms TTL=55
Reply from 70.85.100.85: bytes=32 time=41ms TTL=55

Ping statistics for 70.85.100.85:
    Packets: Sent = 14, Received = 14, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 36ms, Maximum = 85ms, Average = 43ms
Control-C
^C
C:\Users\Admin>
C:\Users\Admin>
```

Look for a time under 100 ms if you are in the U.S., or 150 ms if you are in Europe or Asia. In practice, anything less than 250 ms is not of concern as it is not generally a major component in your perceived response time.

Traceroute

If you are seeing slow ping times, you can run a traceroute. Some networks refuse to forward ICMP, and your traceroute request may not work. If it does work, it is a great tool for identifying network bottlenecks. To run a traceroute on Windows, bring up a command window and run the following command.

```
tracert <yourinstancename>.service-now.com
```

Sample output:

```
 C:\dev\mysql5\bin>tracert mycompany.service-now.com
 Tracing route to mycompany.service-now.com [70.87.98.130]
 over a maximum of 30 hops:
 1     1 ms     1 ms     1 ms  12.192.116.193
 2     4 ms     4 ms     4 ms  12.116.227.37
 3    32 ms    32 ms    32 ms  gbr1-p90.sd2ca.ip.att.net
 [12.123.145.178]
 4    33 ms    33 ms    33 ms  tbr1-p013503.phmaz.ip.att.net
 [12.122.2.142]
 5    34 ms    33 ms    33 ms  tbr2-cl1521.phmaz.ip.att.net
 [12.122.10.194]
 6    32 ms    33 ms    33 ms  tbr2-cl1592.dlstx.ip.att.net
 [12.122.10.81]
 7    31 ms    50 ms    31 ms  gar1-p370.dlrtx.ip.att.net
 [12.123.16.173]
 8    31 ms    31 ms    31 ms  12.119.136.14
 9    31 ms    31 ms    31 ms  te9-1.dsr02.dllstx3.theplanet.com
 [70.87.253.22]
10    37 ms    37 ms    37 ms  vl41.dsr01.dllstx4.theplanet.com
 [70.85.127.83]
11    31 ms    37 ms    31 ms
 gi1-0-1.car16.dllstx4.theplanet.com [67.18.116.67]
12    32 ms    32 ms    32 ms  70.87.98.130
Trace complete.
```

Each line in the traceroute represents a network step between the source machine and the destination machine. In the sample traceroute, there were a total of 12 steps required to get the network traffic from the laptop to <yourinstancename>.service-now.com.

- The left column is the step number.
- The next three columns are latency estimates, performed three times to give an average.
- The last column is the machine you are hopping to.

For example, from rows #1 and #2 above, you can tell:

```
 1     1 ms     1 ms     1 ms  12.192.116.193
 2     4 ms     4 ms     4 ms  12.116.227.37
```

At the end of row 1, it was at 12.192.116.193. It then took 4 ms (on average) to get to 12.116.227.37.

Generally, with a traceroute, you are looking for individual steps that take a long time, like 500 ms for a particular hop. You are also looking for steps that show an asterisk (*) instead of a step time, for example:

```
 1     100 ms   *        500ms  12.192.116.193
```

The asterisk indicates that a particular packet failed to arrive, which can indicate network problems on that particular hop. You also see an asterisk if that particular router is set to not forward ICMP. This outcome is potentially a false alarm if all three latency times for a step are asterisks.

Browser settings and performance

Depending on the browser you use, browser settings can affect the performance of your instance.

Web pages in your instance can be large, for example, over 500 kb for a long list of incidents with many columns. To speed performance, most browsers can accept compressed data from an application server so that the full 500 kb does not have to be sent over the wire. Instead, the browser indicates that it can accept compressed data if the server can send it. The app server then compresses the response, which transforms the 500 kb data to about 20 kb.

Compression is enabled by default on all ServiceNow application servers. The application server always sends compressed data if your browser accepts it. There are browser settings that control whether your browser accepts compressed responses.

System Diagnostics Stats Tools

To aid in performance evaluation, the Stats Tools records statistics for system activities that affect performance such as the execution of queries, scripts, and transactions.

> **Note:**
> The Stats Tools plugin is activated by default. It requires the admin role to activate or upgrade, and it requires the com.snc.jrobin.

Stats Tools adds modules under **System Diagnostics** > **Stats**, including **Slow Queries**, **Slow Scripts**, and **Slow Transactions**. Each module accesses a table of activity patterns [sys_query_pattern], [sys_script_pattern], [sys_transaction_pattern]. Each pattern table represents a collection of unique activities. Each collection is an aggregation of executions of that unique activity over all time. Each record provides basic timing analysis with example identifiable details of the activity.

> **Note:**
> To aid in debugging, you can filter most of these logs by application scope, limiting the transactions (for example, slow scripts or events) to only those transactions originating in specific scopes.

Activity patterns are immediately recorded to a cache and are later persisted to their pattern table. If you flush server caches, then recorded activities that have not been persisted are cleared. The following are examples of pattern records.

- Each time a query is executed that meets the recording and persistence threshold it is aggregated and stored as a query pattern record.
- Each time a particular business rule is executed it aggregates to a script pattern record.
- Each time a particular background job runs it aggregates into a unique transaction pattern record.
- Each click of the **New** button on the Incidents list counts as a list type transaction pattern with specific form action.

Metrics

Metrics include total and average times of interest per unique execution pattern over the total execution count. Metrics are aggregated with each new instance of the unique activity and persisted to the pattern record.

Metadata

Example data from specific executions are included on each pattern to identify execution details.

Characterizations of each activity type

Transactions	Any transaction type includes server-side and related client-side transactions. Metrics include **Total server load time**, which aggregates the total server-side time excluding semaphore and session wait times. It also aggregates relevant server transaction times that are found on the syslog_transaction table. Transaction types: • An HTTP Request transaction is identified by a URL, transaction type, processor, form/list action, URL query (filters), and related table name. • Any other transaction is identified by its transaction URL/page/name, transaction type, and processor or thread name.
Scripts	Any script activity type includes scripts evaluated by GlideScopedEvaluator. Script Types: • A Jelly Script is identified by the sys_jelly_file table, jelly file path, line number, and script that was executed. • Any other script is identified by the table and sys_id.
Queries	Any query activity includes prepared statements executed by GlideDBI. Query Types: All queries are identified by MongoDB query or insert, update, or select statements, as well as other components of the statement like selected columns, where clause, unions, column sets, and limits.

Slow mutex locks record detail

Administrators can use slow mutex logs to gain insight into how mutex locks are affecting platform performance.

Field	Description
Mutex name	The name of the mutex lock for this record.
Is Fast Lock	If **false**, the lock was placed on a record in the ServiceNow platform. If **true**, the lock was placed on a record in the underlying database.
Average execution time	The average time for which the mutex lock was held.
2-hour moving average (ms)	The average of the moving average values of execution time, calculated over periods of two hours.
Day moving average (ms)	The average of the moving average values of execution time, calculated over periods of one day.
Month moving average (ms)	The average of the moving average values of execution time, calculated over periods of one month.
Execution count	The number of similar occurrences that are aggregated.
Last sighting	The time and date the last occurrence was noted.
First sighting	The time and date the first occurrence was noted.
Example stack trace	A stack trace for an individual mutex lock.
Total execution time	The sum of execution time for the aggregated occurrences.
Has	The hash value for this record.
Label	The mutex lock label.
Example URL	The URL for an individual mutex lock.
Event Execution Time / Event Count Trend graphs	The Mutex Execution Time Trend graphs show the total execution time of these mutex locks over the most recent period of two hours, one day, or one month. The Mutex Count Trend graphs show the mutex lock counts over the most recent period of two hours, one day, or one month.

Slow events log record detail

As an administrator, you can use Slow Events logs to gain insight into how events are affecting platform performance. To aid in debugging, you can filter slow event log detail by application scope, limiting events to only those events originating in specific scopes.

Field	Description
Event name	Name of the event for this record.
Queue	Event queue containing this event.
Event application	Application scope the slow event originated in. Global appears if the slow event originated in the global scope.
Execution count	Number of similar occurrences that are aggregated.
Average execution time (ms)	Average duration to execute the event in a single occurrence, in milliseconds.
Day moving average (ms)	Average of the moving average values of execution time, calculated over periods of one day.
Last sighting	Time and date the last occurrence was noted.
First sighting	Time and date the first occurrence was noted.
Hash	Hash value for this record.
Month moving average (ms)	Average of the moving average values of execution time, calculated over periods of one month.
Label	Label for the event.
Example URL	URL for an individual event.
Example stack trace	Example stack trace for an individual event.
Total execution time	Total sum of the event execution times for the aggregated occurrences.
Event Execution Time / Event Count Trend graphs	Event Execution Time Trend graphs show the total execution time of these events over the most recent period of two hours, one day, or one month. Event Count Trend graphs show the event counts over the most recent period of two hours, one day, or one month.

Slow scripts log record detail

As an administrator, you can use Slow Scripts logs to gain insight into how events are affecting platform performance. To aid in debugging, you can filter a slow script log by application scope, limiting scripts to only those scripts originating in specific scopes.

Field	Description
Label	Label assigned to the script.
Script source	Source of the script.
Script application	Application scope the slow script originated in. Global appears if the script originated in the global scope.
Average execution time (ms)	Average duration to execute the script in a single occurrence, in milliseconds.
Execution count	Number of similar occurrences that are aggregated.
Day moving average (ms)	Average of the moving average values of execution time, calculated over periods of one day.
Original application	Total sum of the script execution times for the aggregated occurrences.
Last sighting	Time and date at which the script was last executed.

Slow interactions log record detail

Administrators can use Slow Interactions logs to gain insight into how interactions are affecting platform performance.

Field	Description
Label	Label assigned to the interaction.
Average execution time (ms)	Average duration to execute the interaction, in a single occurrence, in milliseconds.
Execution count	Number of similar occurrences that are aggregated.
Total execution time	Total sum of the execution times for the aggregated executions of the interaction.
First sighting	Time and date the first occurrence was noted.
Hash	Hash value for this record.
Last sighting	Time and date the last occurrence was noted.
Example Java stack trace	Example stack trace for an individual event.
Window start	Start time
Window end	End time

Slow transactions log record detail

As an administrator, you can use Slow Transactions logs to gain insight into how transactions are affecting platform performance. To aid in debugging, you can filter transaction log detail by application scope, limiting transactions to only those transactions originating in specific scopes.

Field	Description
Average execution time (ms)	Average duration to execute the transaction, in a single occurrence, in milliseconds.
Execution count	Number of similar occurrences that are aggregated.
Total execution time	Total sum of the execution times for the aggregated executions of the transaction.
Average client time (ms)	Average duration to execute the transaction on the client, in a single occurrence, in milliseconds.
Total client time (ms)	Total sum of the transaction execution times for the aggregated occurrences on the client, in milliseconds.
Origin application	Application scope the transaction originated in. Global appears if the transaction originated in the global scope.
Example type	Example transaction type executed.
Example processor	Example processor for an individual transaction.

Use a slow query log

Administrators can use slow query logs to gain insight into how queries are affecting platform performance. The Slow Queries log aggregates data for similar queries. Use slow query data to evaluate the need for new indexes, changes to existing indexes, or changes to frequent queries.

Before you begin

Role required: admin

About this task

Queries are similar when they select from the same table and query the same field in the where clause, but search for different values in the field. For example, these queries are aggregated as similar queries.

```
SELECT * FROM sys_user WHERE name="ITIL User"
SELECT * FROM sys_user WHERE name="System Administrator"
```

Aggregating similar query data allows administrators to monitor the performance impact of slow queries and of queries that occur frequently. The log reports data for similar queries where the total execution time exceeds 5 seconds.

Procedure

1. Navigate to **All** > **System Diagnostics** > **Stats** > **Slow Queries**
2. Open a query record for more details.

Query record fields

Field	Description
Example	An SQL statement for an individual query, before being canonicalized to aggregate it with similar queries.
Hash	The hash value for this record.
Example stack trace	A stack trace for an individual query. Lines referencing script code (includes both custom and base system code) appear in the following format: `table name.sys_id:line number` For example, this stack trace indicates a function call from line 119 of a script include. `sys_script_include.105f70abc0a8010300d4d79ed1b93eb0:119` For more information, see stack traces.
Example URL	The URL for an individual query, depending on how the query was called. ◦ User transactions: lists the transaction parameters. For example, if a user navigates to the incident list, the URL is `/incident_list.do sysparm_query=active=true` ◦ Scheduled jobs: lists the name of the scheduled job. ◦ Any other method: lists an empty field.
First sighting	The first occurrence of a similar query.
Last sighting	The last occurrence of a similar query.
Total execution time	Total execution time.
Execution Count	The number of similar query occurrences that are aggregated.
Average execution time (ms)	The average duration to execute one of these similar queries.

3. To see what the database is doing to retrieve the data, click **Explain Plan**.
   The query plan is reported in the **MySQL Explain Plan** related list. Use the query plan to evaluate the need for new indexes or changes to existing indexes. For more information, see MySQL explain plan documentation.

   This feature works on other supported databases, including MongoDB and Oracle.

4. To generate an index suggestion for a slow query, click **Index Suggestion**.
   For details, see Generate an index suggestion for a slow query.

Index suggestions for slow queries

The Index Suggestion Engine (ISE) can generate an index suggestion for a selected slow query. When you request an index suggestion for a slow query, the ISE analyzes the query and recommends an index that can improve the query execution time.

If you choose to use the index suggestion and create the index, the ISE continues to review the effectiveness of that index during a 14-day evaluation period. The ISE provides details on the index during the evaluation, including recommendations for managing the index.

Administrators use the ISE to:

- Generate an index suggestion for a slow query.
- Review index suggestions for slow queries in your instance.
- Export an index suggestion to a non-production instance for evaluation and testing.
- Schedule an index for creation.
- Monitor the effectiveness of an index during the index evaluation period.
- Test index performance (this test is an immediate performance assessment of the index).
- Drop an index that does not optimize query performance, as recommended by the ISE.

In new and upgraded instances, the Index Suggestion Engine plugin (com.glide.index_suggestion) is activated by default.

Note:
The ISE supports MySQL databases only.

How index suggestions work

You start the index suggestion process by requesting an index suggestion for a selected slow query. The ISE runs a daily job that collects column statistics from tables in the slow query, gathering data such as cardinality (unique columns in a table) and null/not null count.

Next, the ISE aggregates and analyzes the information collected, applies a weighted column ranking algorithm to the slow query, and generates an index suggestion for the query.

After an index suggestion is generated, you review the suggestion and determine whether to create the index for the slow query. When you create the index, the ISE provides information on the index as it moves through its life cycle. You can track the index suggestion through three main processing stages:

Index suggestions to review

During this initial stage, you can review index suggestions that the ISE generated for your slow queries. You can choose to ignore a suggestion, export the index suggestion to a non-production instance for further testing, or schedule the index for creation. If the ISE successfully generates an index suggestion and you choose to schedule the index for creation, the index suggestion moves to the next processing stage. However, if the database cannot use the suggestion or the suggestion degrades query performance, the ISE recommends that you drop the index suggestion.

Index in progress

In this stage, the ISE creates the index and the 14-day evaluation period begins. The ISE does an hourly evaluation to determine whether the index improves or degrades the query execution time. The ISE updates the index state, including recommended actions that you can take. For example, if the index does not improve the performance of the slow query, the ISE advises that you drop the index. You

can then schedule the index to be dropped from the database. During this stage, you can also choose to test index performance or accept an index, even if the ISE recommends dropping it.

Index done

In the last processing stage, the ISE describes the final state of the index and related processing activity. If the index improved the slow query time, the ISE changes the index state to Created and the database continues to use the index. If the index did not improve the query time and you chose to drop the index, the ISE drops the index from the database and changes the index state to Dropped.

Index suggestion life cycle

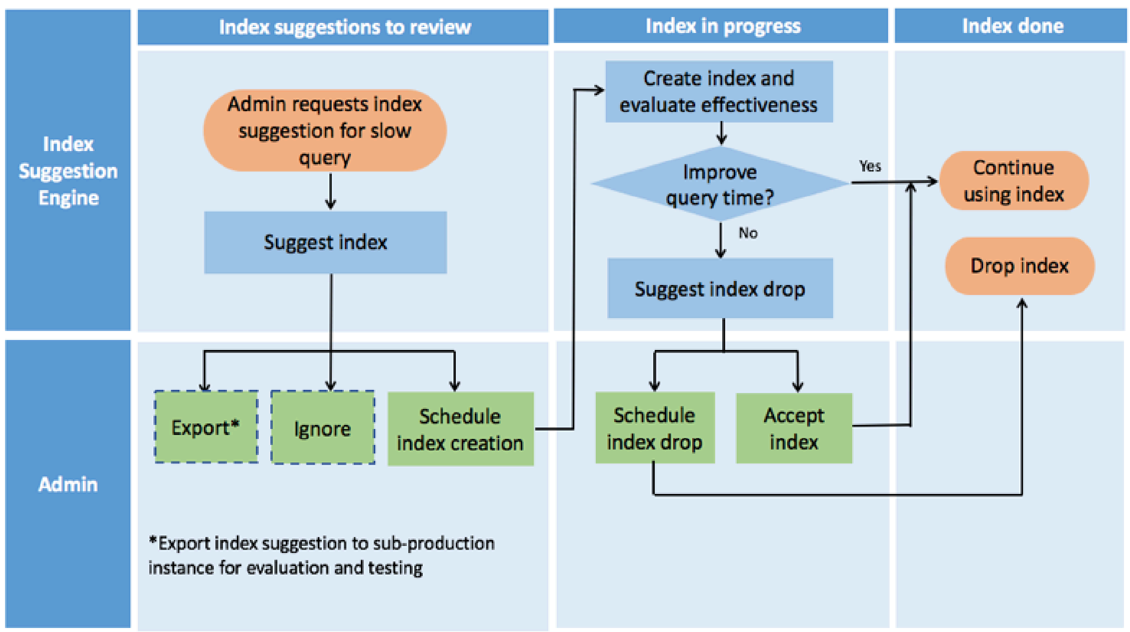

Processing states for index suggestions

The Index Suggestions [sys_index_suggestion] table provides state information on your indexes as they move through the three main processing stages:

- **Index Suggestions > To review**
- **Index Suggestions > In Progress**
- **Index Suggestions > Done**

The ISE uses the following states to describe the processing activity for an index.

Index suggestion states

State	Description
Index suggestions to review	
Suggested	ISE generated an index suggestion for the slow query.
Drop Suggested—Unused	ISE recommends that you drop the index, since the database is not using the index for the slow query.

Index suggestion states (continued)

State	Description
Drop Suggested —Performance Degradation	ISE recommends that you drop the index because the index did not improve the query time or made the performance worse.
Index in progress	
Creation Scheduled	You scheduled the index for creation.
Creation in Progress	ISE is creating the index.
Creation Failed	ISE could not create the index.
Evaluating Effectiveness	ISE created the index and is assessing index performance during the 14-day index evaluation period.
Drop Suggested—Unused	ISE recommends that you drop the index from the table for which the index was created, since the database is not using the index for the slow query.
Drop Suggested —Performance Degradation	ISE recommends that you drop the index because the index did not improve the query time.
Drop Scheduled	You scheduled the index to be dropped from the database.
Drop in Progress	ISE is dropping the index from the database.
Drop Failed	ISE could not drop the index. Contact Customer Service and Support for assistance.
Index done	
Created	After the 14-day evaluation period, the ISE determined that the index improved query performance. Indicates that the database continues to use the index.
Ignored	You chose to ignore the index suggestion.
Dropped	ISE successfully dropped the index.
Accepted	You chose to keep the index even though the ISE recommended dropping it.
Superseded	A recent index suggestion replaced the index for the same table and slow query.

Generate an index suggestion for a slow query

Use the Index Suggestion Engine (ISE) to generate an index suggestion for a slow query.

Before you begin

Role required: admin

About this task

When the ISE generates an index suggestion for a slow query, it reviews query metrics and information from tables in the query. If the ISE does not have the information required to generate an index suggestion, the system displays messages informing you about the data needed. These messages typically provide links for obtaining more information, such as updated query metrics

or gathering column statistics from tables in the query. Sometimes the ISE may be unable to suggest an index for the query and lets you know why the suggestion cannot be generated.

Procedure

1. Navigate to **All** > **System Diagnostics** > **Stats** > **Slow Queries**.
2. In the Slow Queries table, click the query record for which the index suggestion is to be generated.
3. In the Slow Queries record, click **Suggest Index**.
   If the ISE needs more information, the system may display one or more messages with details for obtaining that information. After you respond to each message, click **Suggest Index**. Repeat this process until the ISE has the information required to generate the index suggestion.

   For example, the ISE may need:

   - Recent metrics—If the system displays a message indicating that query metrics are outdated, the message provides a link for obtaining the most recent metrics.
     - Click the link to get the updated query metrics.
     - Click **Suggest Index**.
   - Column statistics—If the column statistics do not exist for the query, the system displays a message asking you to schedule the Collect Column Stats job. You can choose to run the job at the default time when the system is less busy, or you can have the system collect the statistics immediately.
     - If you schedule the Column Stats job to run later, the ISE generates the index suggestion after it collects the column statistics. After the column stats job runs, return to the Slow Queries record and click **Suggest Index**.
     - If you choose to collect column stats immediately, the system displays the Collect Column Statistics progress indicator while it gathers column statistics. In the Slow Queries record, click **Suggest Index**.

   When the ISE successfully generates an index suggestion, the Index Suggestion record for the query shows the index **State** as **Suggested**.
4. If you want to continue working with the index suggestion, select the next processing step

Review index suggestions for slow queries

You can review index suggestions and indexes created by the Index Suggestion Engine (ISE). The ISE provides three main views of index suggestions as they move through their life cycle: index suggestions in review, index in progress, and index done.

Before you begin

Role required: admin

About this task

Access index suggestions from the Application Navigator or the Performance homepage.

Procedure

1. Access the index suggestions for your instance

2. In the Index Suggestions table, click the **Table Name** associated with the index suggestion you want to view.

3. Review the index suggestion for the selected table.

Index Suggestion fields

Field	Description
Slow query	The slow query for which the index suggestion was generated.
Example	An actual SQL statement for an individual query, before being canonicalized to aggregate it with similar queries.
Table name	Name of the table used in the query.
Column names	Names of the columns used in the query.
Index name	Name of the index, assigned by the ISE after the index is created.
State	Processing state of the index. For details, see the Index suggestion states table.
Imported	This box, if checked, indicates that the index suggestion was imported from a different instance for evaluation and testing.

4. Choose the next processing step for the index suggestion.

Export an index suggestion to a non-production instance

You can export an index suggestion from a production instance to a non-production instance for evaluation and testing, before implementing the index on your production instance.

Before you begin

Role required: admin

About this task

This task involves exporting an index suggestion as an .xml file from your production instance and importing it to a non-production instance. You can then create the index and evaluate its performance on the non-production instance.

Procedure

1. In your production instance, navigate to the Index Suggestion to be exported:

   a. Access the Index Suggestions table to review the index suggestions for slow queries.

Location	Description
From the Application Navigator	Navigate to **System Diagnostics** > **Index Suggestions** > **Index Suggestions > To review**.
From the Performance homepage	If the suggestion icon indicates that one or more index suggestions are available ( ), click it to open the Suggested Indexes window, and click **View all suggested indexes**.

Location	Description
	Note: In the Suggested Indexes window, you can click **View** next to a specific table name to open its index suggestion record.

b. In the Index Suggestions table, click the **Table Name** to open the index suggestion.

2. In the Index Suggestions form, click **Export**.
   The ISE downloads the index suggestion as an XML file with the file name `sys_index_suggestion_<alphanumeric identifier>.xml`. The file name has a unique alphanumeric string to differentiate it from other index suggestions.
3. In your non-production instance, navigate to **System Diagnostics** > **Index Suggestions** > **Index Suggestions > To review**.
4. In the Index Suggestions table under Related Links, click **Import Suggestions**.
5. In the Import XML form, choose the index suggestion file that you exported and click **Upload**.

   After the import, the imported index suggestion appears in the Index Suggestions table on your non-production instance.

6. To continue the index evaluation process on your non-production instance, you can schedule the index suggestion for creation andtest index performance.

Schedule an index suggestion for creation

Use the Index Suggestion Engine (ISE) to schedule an index suggestion for creation.

Before you begin

Role required: admin

Procedure

1. Navigate to the index suggestion for the slow query.

2. In the Index Suggestion table, click the **Table Name** for which the index is to be created.
3. In the Index Suggestion form, click **Schedule Creation**.
4. In the Schedule Index Creation window, select the time that the index should be created and click **Create**.

   Note:
   If the index to be created involves large tables, such as the Incident or CMDB table, index creation may take up to one hour or longer.

   - **Create now** —The Schedule Index Creation progress indicator window shows the progress of the index creation. When index creation is complete, click **Close**.
   - **Create later** —If you select this option, the state changes to Create Scheduled, and the index is created at the scheduled time. After the index is created, the index state changes to Evaluating Effectiveness.

   When the ISE creates the index at the selected time, it assigns an index name based on the tables used in the slow query.

5. Verify that the index was created by reviewing the index **State** in the Index Suggestion record or in the Index Suggestion table (**System Diagnostics** > **Index Suggestions** > **Index Suggestions > In Progress**).
   The Index Suggestion record shows the state, Evaluating Effectiveness. The ISE begins the 14-day index evaluation period.
6. During the evaluation period, you can monitor the index **State** in the Index Suggestion record or in the Index Suggestion table (**System Diagnostics** > **Index Suggestions** > **Index Suggestions > In Progress**).

Test index performance

After the Index Suggestion Engine (ISE) creates an index for a slow query, you can test the index to see if it improves the query time.

Before you begin

Role required: admin

About this task

This test is a quick assessment of the index and is not part of the 14-day index evaluation process. If you are evaluating an index suggestion on a non-production instance, you can run this test after creating the index to get immediate results on index performance.

Note:
Index test results from your non-production instance may not necessarily predict similar performance results on your production instance.

Procedure

1. Open the Index Suggestion record for the index to be tested:
   a. Navigate to **System Diagnostics** > **Index Suggestion** > **Index Suggestions > In progress**.
   b. In the Index Suggestions table, click the **Table name** for the index to be tested.
2. In the Index Suggestion record, click **Test Performance** to start the performance test.
3. When the Test Index Suggestion progress indicator shows that testing is complete, click **Go to Result**.

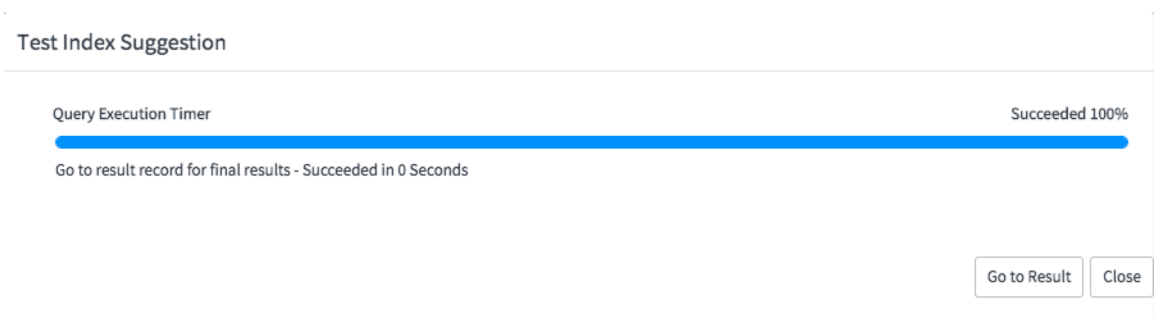

4. In the Query Performance Test Results form, compare the **Average with index** and **Average without index** fields to see if the index improved the query time.
   In the following example, the index improved the query time. The **Forced** check box is unchecked, indicating that the database is using the index. If the check box is checked, the database was forced to use the index during the test but is not using the index outside of the test.

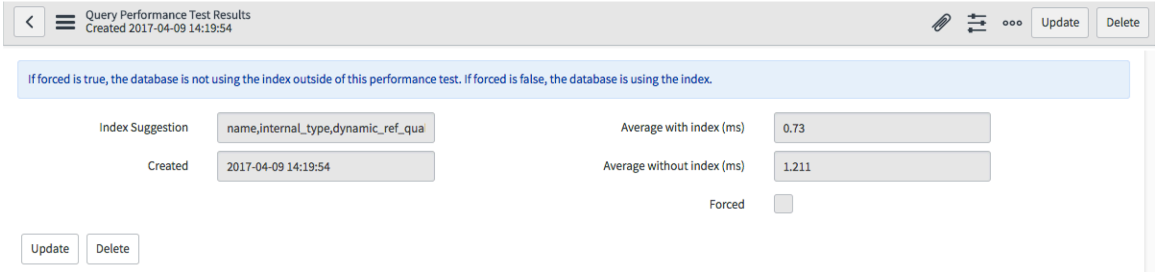

Query Performance Test Results

Field	Description
Index Suggestion	Name of the index suggestion.
Created	Date and time this performance test was run.
Average with Index (ms)	Query time using the index suggestion.
Average without Index (ms)	Query time before using the index suggestion.
Forced	Check box indicating whether the database is using the index. If this box is checked, the database was forced to use the index suggestion, but only for this test. If the box is unchecked, the database is using the index.

5. Click the back arrow to navigate back to the Index Suggestion record for the table.
   The result of this index performance test does not change the index state recorded in the Index Suggestion record.

Schedule an index to be dropped

You can drop an index for a slow query when the index is not used by the database or does not improve query performance.

Before you begin

Role required: admin

About this task

During the index evaluation period, the ISE may recommend that you drop the index when the database does not use the index or the ISE finds that the index does not improve the query time. You can drop the index immediately or schedule the drop for a later time.

> **Note:**
> If you want to retain an index even though the ISE recommends dropping it, you can use the **Accept** option in the Index Suggestion record to continue using the index. Before using this option, contact Customer Service and Support for guidance.

Procedure

1. Navigate to **All** > **System Diagnostics** > **Index Suggestions** > **Index Suggestions > In Progress**.
2. In the Index Suggestions table, click the Table Name for the index to be dropped.
3. In the Index Suggestion record, click **Schedule Drop**.
4. In the Schedule Index Drop window, select the time when you want the index dropped and click **Drop**.

Note:
If the index you want dropped involves large tables, such as the Incident or CMDB table, dropping the index may take up to one hour or longer.

- **Drop now** —The Schedule Index Drop progress indicator window shows the progress of the index drop. When the index drop completes, click **Close**. If the drop is successful, the Index Suggestion record shows the index state as Dropped.
- **Drop later** —When you schedule the drop, the index state changes to Drop Scheduled. When the system drops the index at the scheduled time, the index state changes to Dropped.

5. Verify that the index was dropped by reviewing the index state in the Index Suggestion record or in the Index Suggestion table (**System Diagnostics** > **Index Suggestions** > **Index Suggestions > Done**).

Transaction call chains

Identify and debug the root causes for a problematic transaction by recording, and then reviewing, its call chain. A call chain consists of the ordered series of scripts executed during the life cycle of a transaction.

Transaction Call Chain tracking enables an administrator to register the URL of a problematic transaction for recording of the call chain during the next URL occurrence. It is helpful for diagnosing transaction performance issues.

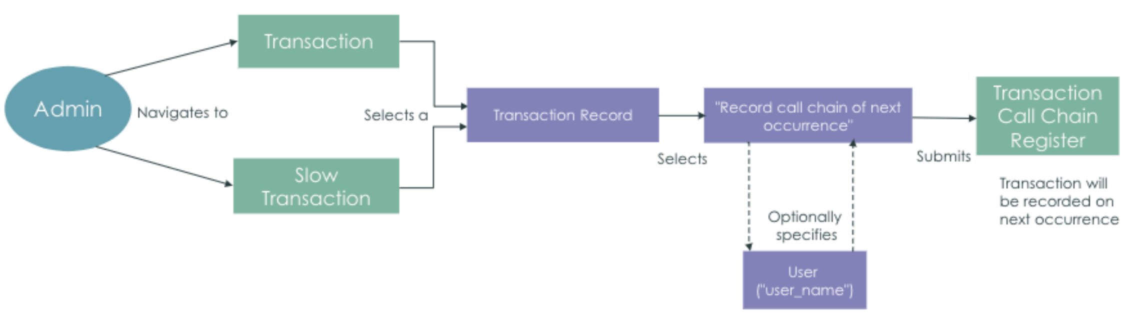

To enable transaction call chain tracking, perform the following tasks:

1. When you find a problematic transaction on the Transaction Log or Slow Transaction list that requires debugging, select it to view its detail record. To aid in debugging, filter and group these lists by application scope, limiting transactions that appear to only those transactions originating in specific scopes.

   Note:
   To learn more about these lists, see Transaction logs and Slow transactions log record detail.

2. In the record detail form, click **Record call chain of next occurrence**, located under Related Links.

- It designates that the next occurrence of the transaction URL should be persisted with all of its call chain information.
- You can optionally designate a user ID so call chain recording only takes place during the next transaction occurrence for a specific user.

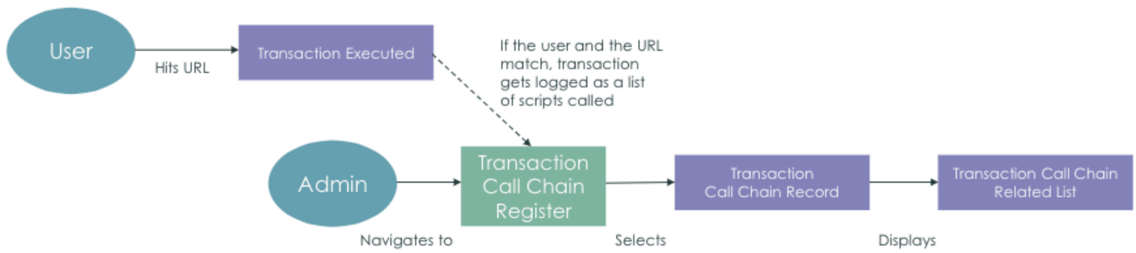

3. The next time a user accesses the URL, the Transaction Call Chain Register record is updated to contain a listing of the script resources executed to service the request.
4. In Transaction Call Chains, review the recorded call chain for the transaction, and then take needed the actions required to debug it.

Create a transaction call chain register

Create a call chain register to record call chain information for a problematic (or slow) transaction during the next occurrence of that URL. You specify that the next occurrence of the transaction URL should be persisted with all of its call chain information.

Before you begin

Role required: admin

Procedure

1. Navigate to the Transaction Log or Slow Transactions list, select the problematic transaction, and enable recording of the call chain.

2. When the Transaction Call Chain Register appears, fill in the **User** field.

Transaction Call Chain Register form

Control	Description
URL	URL for the selected transaction.
Transaction Log	Transaction log ID created during the next occurrence of the URL. This field is empty when you first create the call chain register for a transaction.
Recorded	Check box that indicates if call chain recording has already taken place for this transaction. ◦ This field is empty (**false**) when you first create the call chain register for a transaction. ◦ **true** appears when the call chain is recorded during the next occurrence of the URL.
User	(Optional) If call chain recording should only take place when a specific user executes the next transaction occurrence, select the user ID.

Control	Description
	Leave it empty if you want call chain recording to occur regardless of user.

3. Click **Submit**.

Result

A call chain register is created that designates that the next occurrence of the transaction URL should be persisted with all of its call chain information.

> **Note:**
> Once the call chain is recorded (**true** appears in the **Recorded** field), you can view it in the Transaction Call Chain Register form itself. To learn more about transaction call chains, see Review a transaction call chain.

Review a transaction call chain

Use Transaction Call Chains to review the resulting call chain for a selected transaction.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Diagnostics** > **Transaction Call Chain** > **Transaction Call Chains**.
2. To view call chain detail, click the **Created** field for the selected call chain component or script.

   > **Note:**
   > You can view transaction log detail by clicking the **Transaction Log** field for the selected call chain component or script. To learn more, see Transaction logs.

Transaction Call Chain form

Control	Description
Source Table	Source table for the script resource executed in the call chain.
Source	Script resource executed in the call chain.
Persisted Sequence	Order in which the script executed in the call chain.
Call Depth	Depth of the executed call chain. For example, a script calling a script include, which in turn calls another script include would have at least three rows. One row has a depth of 1, another has a depth or 2, and the third one has a depth of 3.
Trigger Table	Name of trigger table, if any, for the executed call chain component.
Duration(ms)	Observed duration for the script execution, in milliseconds.
Transaction Log	Transaction log ID created for the recorded call chain.

Transaction cancellation

A user can cancel a transaction that takes longer than expected to load.

Causes of slow-loading transactions include the following scenarios.

- Sorting a large number of records by an unindexed string field.
- Grouping a large table on a field that has predominantly distinct values.
- Exporting all rows from a large table.
- Testing a poorly scripted business rule that has an infinite JavaScript loop.

Note:
You cannot cancel an import with these controls.

Transaction cancel options

After the transaction has been cancelled, you can use the following options for different

purposes.

- Go Back: Returns to the previous page in the navigation history
- Reload: Reloads the targeted page
- Load List With Filter Only: Shows the list (without the records) only with the applied filters
- Export: Export of list view records

Note:
The Export option appears only if the following conditions are met.

- The glide.ui.export_on_tx_cancel.enable property is set to True.
- The list view is canceled.

You can use one of the following options to export the list view records.

- Excel
- CSV
- JSON
- PDF

The exported list view records is delivered to the provided email address. Only the PDF **File Type** option can have the **Orientation** option.

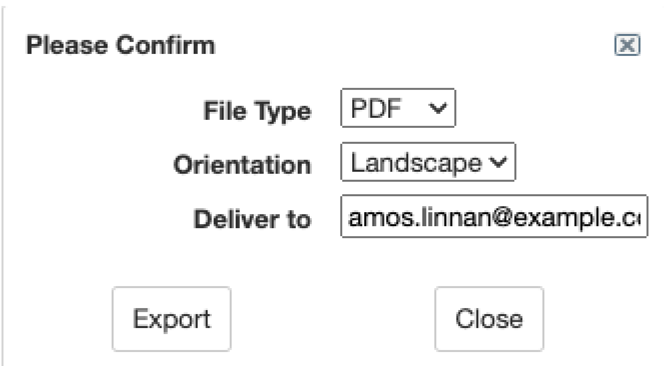

Transaction cancel timer

During a long-running transaction, a timer and a red cancellation button appear in the banner frame.

Cancel request

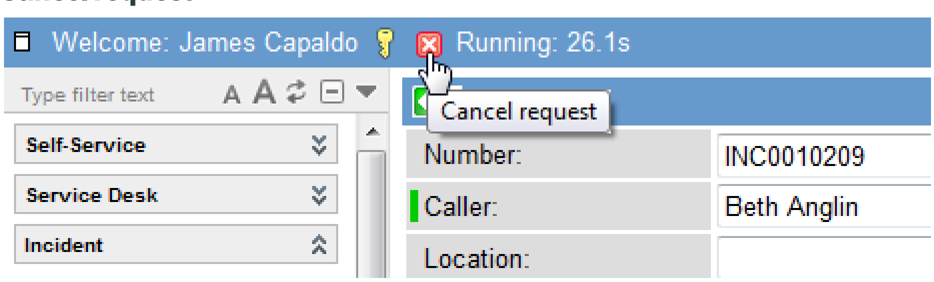

To cancel the transaction, click the cancellation button. The timer indicates that the transaction is being canceled. When it is finished, you see a message that the transaction was canceled.

If the transaction completes successfully, the timer indicates how long it took for the transaction to complete.

Transaction completed

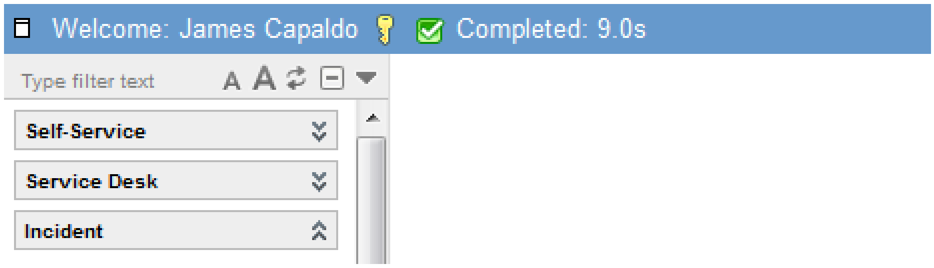

Canceled transaction logs

Canceled transactions appear in the transaction log with **CANCELLED** appended to the URL. Transactions canceled by a user are logged differently from transactions canceled by Transaction quotas.

Cancel transaction properties

Administrators can configure the behavior of the transaction cancel capability using the following properties.

Property	Definition	Location
com.glide.request_manager.active	Let users cancel long running transactions (enabled by default).	Open the sys_properties table.
glide.ui.transaction.long_response.time	Delay in seconds before the cancel transaction button appears for a long-running transaction.	Navigate to **System Properties** > **UI Properties**.

View and kill active transaction

Administrators can view active transactions and kill long-running transactions.

Before you begin

Role required: admin

About this task

Note:
Functionality described here requires the **Admin** role.

This action is intended to stop background processes, such as scheduled jobs, and not user-initiated transactions.

Procedure

1. If High Security Settings is enabled, elevate privileges to security_admin.
2. Navigate to the appropriate **Active Transactions** module.
   - To view and kill transactions on the current node for your instance, navigate to **User Administration** > **Active Transactions**.
   - To view and kill transactions on all nodes for your instance, navigate to **System Diagnostics** > **Active Transactions (All Nodes)**.
3. Select the transaction you want to stop.

The process to display the transactions is included in the list. Do not kill that process.

4. Either right-click the record and select **Kill**, or select the check box next to the record and select **Kill** from the **Actions on selected rows** drop-down at the bottom of the list.

Canceled-transaction logging to a table

The system logs canceled transactions to a table in addition to the transaction log.

Transactions canceled manually or by a quota rule are logged to the Transaction Cancellation Log [syslog_cancellation]. The Transaction Cancellation Log is an extra debugging feature, and will not contain information on all cancellations.

Note:
There are safeguards in place to prevent infinite loops and issues related to synchronicity, which mean it is not always possible to log a cancellation to a database table.

You can find the Transaction Cancellation Log by clicking **System Logs** > **System Log** > **Transaction Cancellations**.

You can turn off canceled-transaction logging to a table by changing a property.

Canceled-transaction log property

Property	Description
glide.quota.manager.log.cancellation	Indicates whether canceled transactions are to be logged to the Transaction Cancellation Log table. Type: Boolean Default value: true Location: Add to the System Properties [sys_properties] table

Import set performance

The algorithm transforms import sets from their staging table into their final destination.

Importing data via an import set requires two phases:

1. The data is loaded from a data source into a staging table
2. The data is transformed from the staging table into a target table

The transform algorithm operates in "blocks" of 100 records at a time as opposed to the previous algorithm which transformed a single record at a time. The newer approach allows the application server to pre-fetch a variety of information relevant to each block of records, reducing the number of unique interactions with the database and improving throughput.

Who should expect to see a performance improvement?

Any customer using import sets should expect to see a performance improvement in large transformations.

What kinds of transformations benefit the most from these changes?

Transformations with a large number of reference or choice type columns see the largest improvement.

What kind of transformations benefit the least?

Transformations that with complex or unkeyed coalesce conditions see a proportionately smaller benefit.

Thread performance monitoring

The instance tracks the performance of individual threads, which administrators can view in the transaction log or System Overview graph in the Performance homepage.

Thread performance can be monitored in two places as described in the following sections.

Transaction log

The amount of time a thread waits for an available semaphore or session synchronization is included in the **Response Time** field in the transaction log. To view the transaction log, navigate to **System Logs** > **Transactions**.

Performance graph

Every second, the system looks at all active threads (both UI and background) and places them into one of the following categories.

- **CPU**: The thread is active, but is not executing any of the steps. This condition typically means non-business rule compute time, although in this case a few other internal wait states are categorized as CPU. Therefore a 1:1 correlation between threads in a CPU count and hardware CPU utilization is not expected.
- **Database**: Waiting for information from the database.
- **Business Rule**: The system is running a business rule (synchronous or asynchronous) and is not currently executing a query (which would be database).
- **Network**: Writing data out to the network or waiting for an outbound network buffer to flush.
- **Concurrency**: Cannot run because they are waiting on a semaphore or session synchronization.

Every minute, the system averages these transactions and records them in the database. These averages are displayed in the **System Overview** graph available in the ServiceNow Servlet metrics. For details, see ServiceNow Servlet.

The **System Overview** chart can be added to any homepage. For more information, see Add existing reports to a homepage.

Important:

The functionality found in homepages, arranging information from your instance to tell a story about your data, is found in dashboards on new instances. On upgraded instances with Next Experience enabled, users can view existing homepages if they have a direct URL, but they can't create or edit them. Responsive dashboards and Analytics Center dashboards take over homepage functionality.

Use the Homepage deprecation help tool to convert the homepages on your instance to responsive dashboards.

For more information, see:

- Dashboards in the Analytics Center.
- Working with responsive dashboards.

ServiceNow Performance dashboards

View a wide range of platform performance metrics for your instance and for the machine on which your instance is running. The ServiceNow Performance dashboards provide central access to different graph sets for monitoring the performance of your instance.

Navigate to **Performance Analytics** > **Dashboards**. From the Dashboards list, select **ServiceNow Performance**.

ServiceNow Performance dashboards

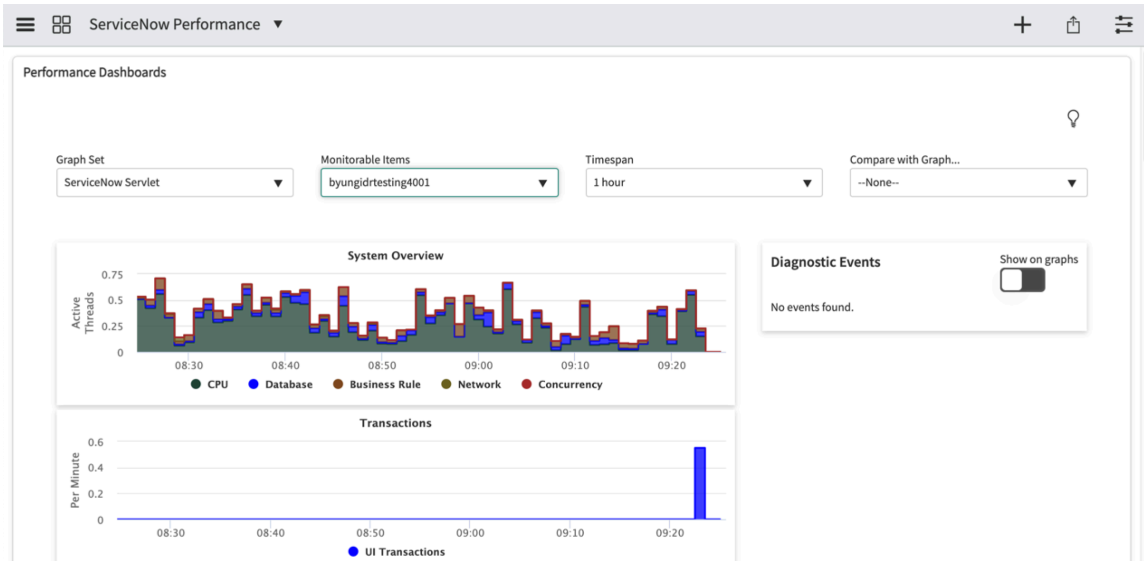

The ServiceNow Performance dashboards feature time series graphs for analyzing system behavior and patterns over selected time periods. The dashboard interface provides:

- Aggregate summaries of time series graph data
- Graph overlay capability to compare two different data sources within a single graph
- Diagnostic event overlay on each graph to highlight performance anomalies
- Zoom controls to expand each graph and view time series details

The graph sets provide performance data for the following functional areas:

- Asynchronous Message Bus (AMB)
- Database
- Instance View
- MySQL Global Status
- ServiceNow Servlet
- Slow Pattern

Filter the data in each graph using different measurements, such as:

- Maximum and minimum values
- Means
- Medians

Customer Service and Support uses some of these graphs to troubleshoot performance issues or help you tune your system for maximum efficiency.

ServiceNow Performance dashboards controls

Control	Description
Graph set	Graph controls: • Graph collections for viewing performance in functional areas: ◦ AMB ◦ Database ◦ Instance View ◦ MySQL Global Status ◦ ServiceNow Servlet ◦ Slow Pattern • **Monitorable items**: Name of the instance under evaluation. • **Timespan**: Period over which the data is captured: 1 hour, 2 hours, 1 day, 7 days, or 30 days • **Compare with graph**: Graph overlay feature that enables you to select a graph from a different data source and compare it to the current graph set.
Index suggestions icon	Index suggestions icon ( ). Opens the Suggested Indexes window, which provides links to index suggestions for slow queries. For details, see Review index suggestions for slow queries.
Display controls	• Information tooltip: Provides details about the area on the graph that you point to. • Refresh icon ( ): Refreshes the data in the graphs. Appears when you point to the upper-right corner of the page. • Help icon ( ): Opens the documentation page for the graph set. Appears when you point to the right corner of a graph. • Fullscreen icon ( ): Opens an enlarged view of the graph, which includes a legend that identifies the graph data. Appears when you point to the right corner of a graph. To zoom in on an area and view details over a smaller time, click and drag across a portion of the graph.
Show datatable	Show datatable icon ( ): Displays a table of data for each thread in the graph.
Diagnostic Events	Lists system events that occurred during the selected timespan. Provides controls for displaying diagnostic events on graphs.

To look for trends or patterns between different data sets, use the graph overlay feature. For example, you can overlay the Database Response graph over the ServiceNow System Overview graph to determine possible correlations.

To enable or disable the event overlay on the graph set, click **Show on graphs**. The event overlay is enabled when the switch is green. The following severity flags identify diagnostic events on the graphs:

- I: Information
- W: Warning
- E: Error

To limit the events that appear on the graphs, click the filter icon ( ).

- To see event information, point to an event in the graph.
- To access details on a particular event, click the event listed in the Diagnostic Events widget.

Platform performance metrics

You can view a wide range of platform performance metrics for your instance and for the machine on which your instance is running. The ServiceNow Performance dashboard provides central access to different graph sets for monitoring the performance of your instance.

To open the ServiceNow Performance dashboard:

1. Navigate to **Self-Service** > **Dashboards**.
2. Search for the **ServiceNow Performance** dashboard.

Note:

Beginning in the San Diego release, the ServiceNow Performance homepage is deprecated. After upgrading to the San Diego release, you can't edit the ServiceNow Performance homepage. You can Create a dashboard version of the Performance homepage. By default, new instances have ServiceNow Performance dashboards to replace the ServiceNow Performance homepage.

The ServiceNow Performance dashboard features time series graphs for analyzing system behavior and patterns over selected time periods. The interface provides:

- Aggregate summaries of time series graph data
- Graph overlay capability to compare two different data sources within a single graph
- Diagnostic event overlay on each graph to highlight performance anomalies
- Zoom controls to expand each graph and view time series details

Users must have an assigned pdb_user role to view performance metric graphs. For more information, see Add existing reports to a homepage and Customize your homepage content.

The graph sets provide performance data for the following functional areas:

- Asynchronous Message Bus (AMB)
- Database
- Instance View
- MySQL Global Status
- ServiceNow Servlet
- Slow Pattern

You can also filter the data in each graph by using different measurements, such as:

- Maximum and minimum values
- Means
- Medians

Customer Service and Support uses some of these graphs to troubleshoot performance issues or help you tune your system for maximum efficiency.

ServiceNow Performance dashboard features

The ServiceNow Performance dashboard provides the following graph and display controls.

Performance dashboard UI elements

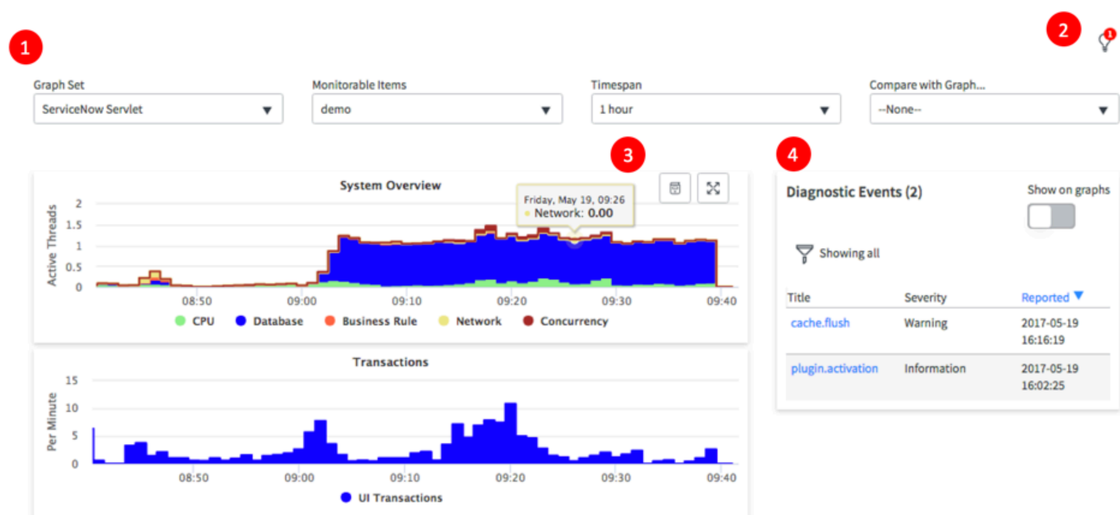

ServiceNow Performance dashboard elements

Number	Description
1	Graph controls: • **Graph set**: Graph collections for viewing performance in functional areas: ○ AMB, Database ○ ServiceNow® Discovery ○ MySQL Global Status ○ ServiceNow Servlet ○ Slow Pattern • **Monitorable items**: Name of the instance under evaluation. • **Timespan**: Period over which the data is captured: 1 hour, 2 hours, 1 day, 7 days, or 30 days. • **Compare with graph list**: Graph overlay feature that lets you select a graph from a different data source and compare it to the current graph set.
2	Suggestion icon ( ): Opens the Suggested Indexes window, which provides links to index suggestions that have been created for slow queries. For details, see Review index suggestions for slow queries.

ServiceNow Performance dashboard elements (continued)

Number	Description
3	Display controls: • Information tooltip: Provides details about the area on the graph that you point to. • Documentation icon ( ): Opens the documentation page for the graph set. Appears when you point to the right corner of a graph. Opens the documentation page for the graph set. • Fullscreen icon ( ): Opens an enlarged view of the graph, which includes a legend that identifies the graph data. Appears when you point to the right corner of a graph.
4	Diagnostic Events widget: Lists system events that occurred during the selected timespan. Provides controls for displaying diagnostic events on graphs.

Display and zoom controls

Fullscreen graph view

Click the fullscreen icon ( ) to open an enlarged view of a graph. The full screen view includes a legend with check boxes for controlling the data displayed in the graph. To disable the display of a data item in a graph, clear its check box. On some graphs, you can also change the sort order of certain metrics in the legend.

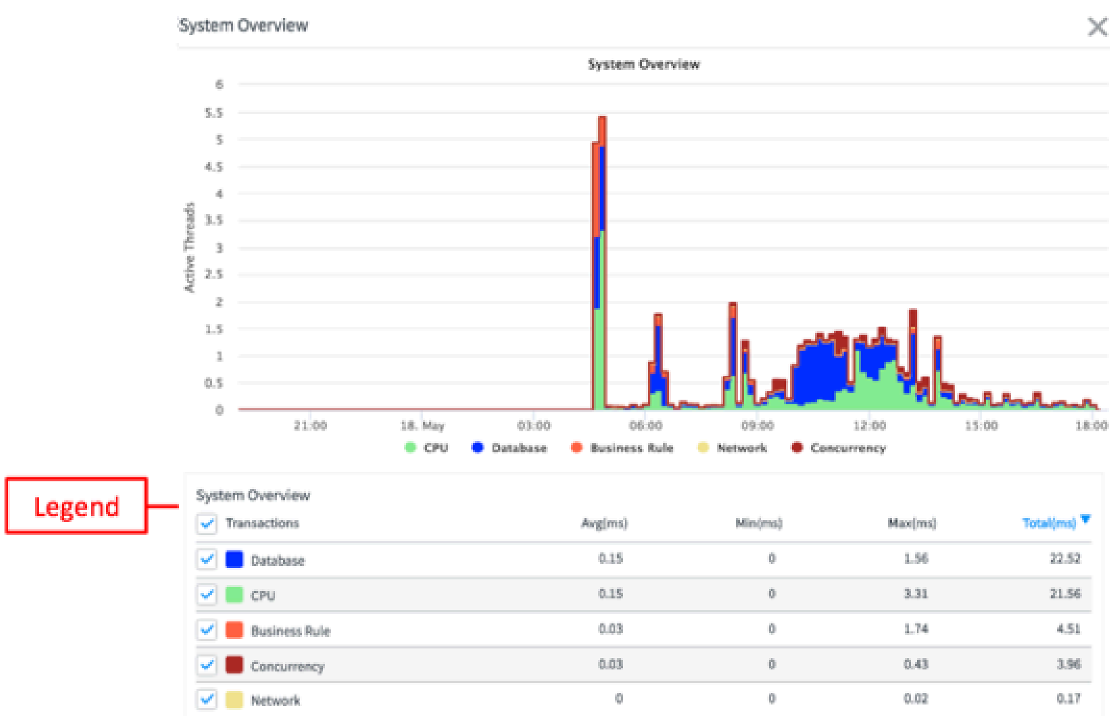

Zoom in on a selected area of the graph

Click and drag across a portion of the graph area to zoom in on that area and view details over a smaller time interval. To return to the original view, click **Reset Zoom**.

Zooming in a graph

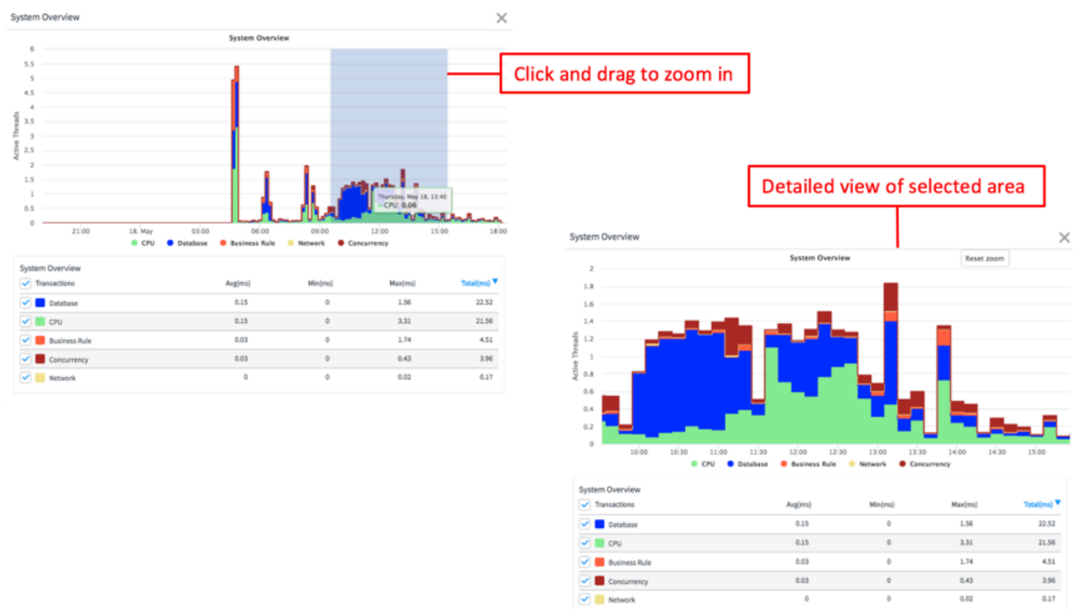

Graph overlays

Use the graph overlay feature to look for trends or patterns between different data sets. For example, you can overlay the Database Response graph over the ServiceNow System Overview graph to determine possible correlations.

- From the Compare with graph list, select a performance graph to overlay the current graph set.
- To remove the graph overlay, select **None** from the Compare with graph list.

Graphs with legends

Diagnostic event overlays

Display controls for diagnostic events

Use the following features in the Diagnostic Events widget to control the diagnostic events overlaid on graphs:

- **Show on graphs** switch: Click to enable or disable the event overlay on the graph set. The event overlay is enabled when the switch is green. The following severity flags identify diagnostic events on the graphs:
  - I: Information
  - W: Warning
  - E: Error
- Filter: Click the filter icon ( ) to limit the events that appear on the graphs. Clear the check boxes for events that you don't want displayed.
- **Clear filter**: Click to reset the filter to its default state (show all events).

Diagnostic events overlay

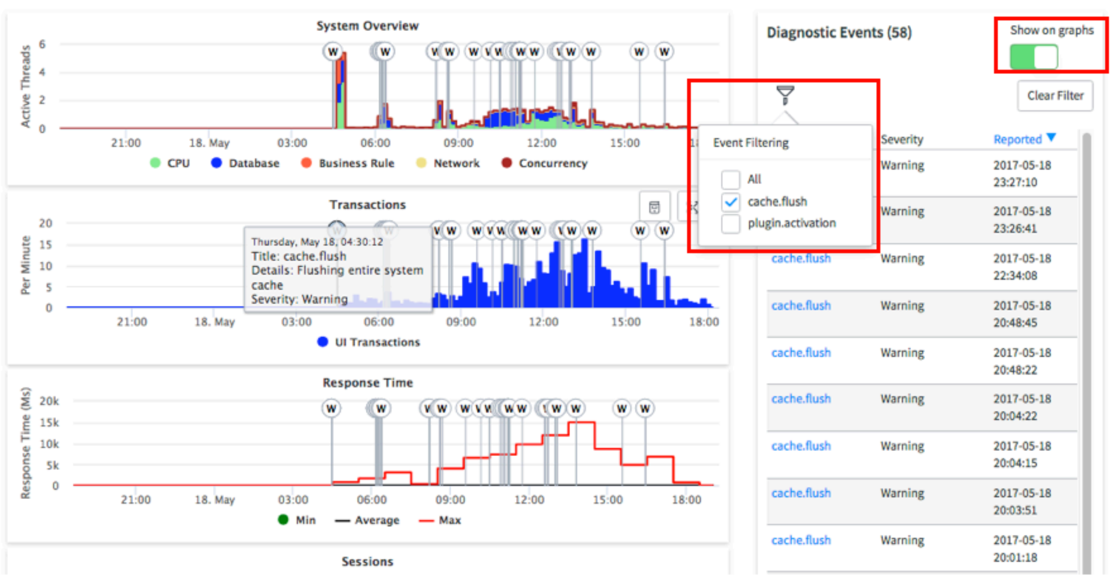

Diagnostic event information

- If you enabled the diagnostic event display, point to an event in the graph to see its information.
- To open the Event Details page, click the event listed in the Diagnostic Events widget.

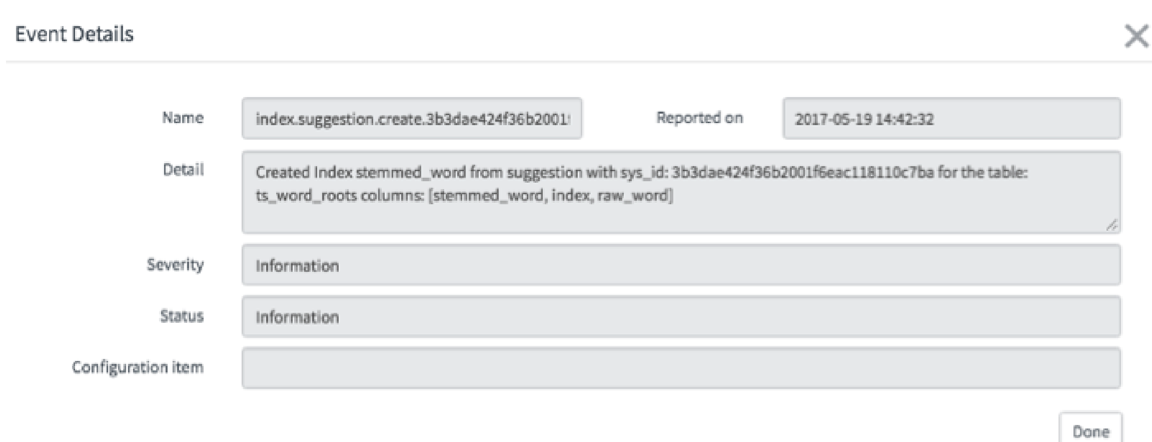

Event Details page

Field	Description
Name	Type of event.
Reported on	Date and time of event.
Detail	Additional information on the event, such as the type of plugin activated.
Severity	Type of diagnostic event: ◦ I: Information ◦ W: Warning ◦ E: Error
Status	Event state: ◦ Information ◦ Open ◦ Closed
Configuration item	Infrastructure component, if applicable.

Asynchronous Message Bus (AMB)

Asynchronous Message Bus (AMB) graphs, available to view on the ServiceNow Performance homepage, use asynchronous communication to monitor the transaction count and response times, which provide a performance analysis.

If you have an admin role, you can review performance metrics. The AMB graphs are intended for use by Customer Service and Support to troubleshoot performance issues.

To view AMB graphs on the ServiceNow Performance homepage:

1. In the **Graph Set** list, select **AMB**.
2. In the **Monitorable Items** list, select an instance.
3. In the **Timespan** list, select a time period.

For details on using the graph and display controls on the Customer Service and Support Performance homepage, see Platform performance metrics.

The following graph displays the maximum, median, and minimum response times in milliseconds over a user-selected timespan. The second graph shows the number of AMB transactions over a user-selected timespan.

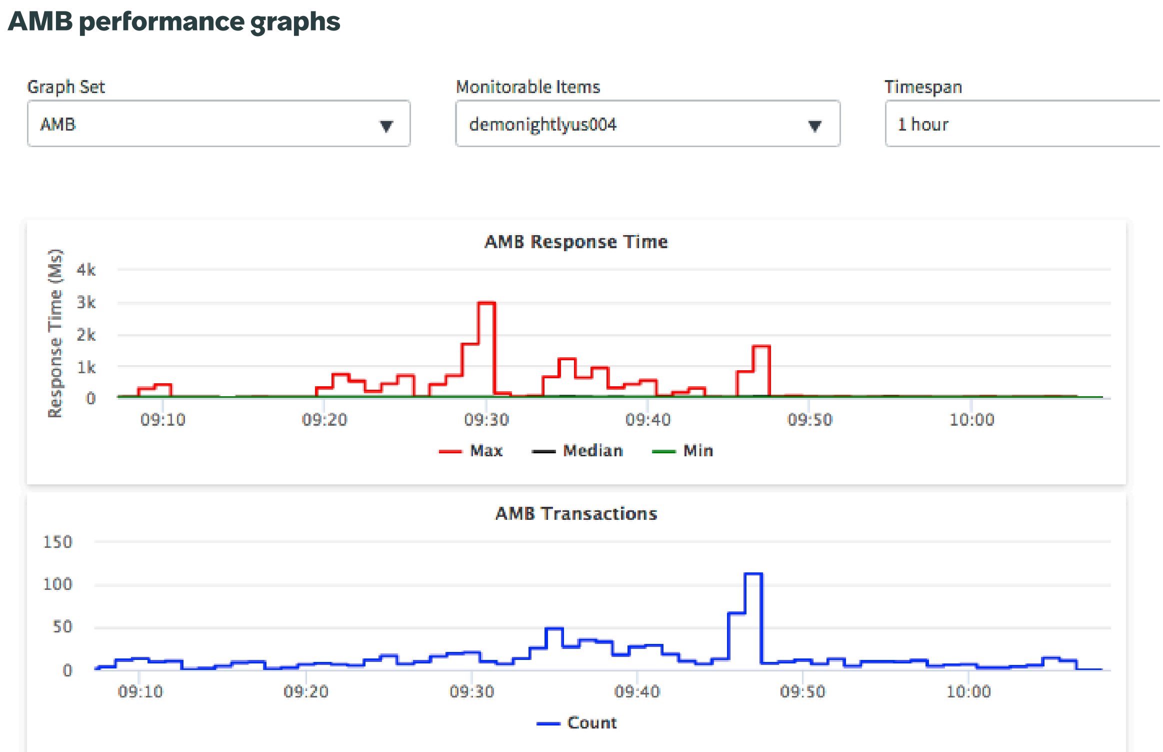

Database

Database graphs, available to view on the ServiceNow Performance homepage, display metrics for various database operations, for example, insertions and deletions. They also display a current count of database connections for the selected ServiceNow instance.

If you are running on a single node, which is typical, the graphs show everything. All the traffic shown in the graphs is specific to your cluster node or instance. If you are looking at the database graph for your Discovery node, it does not include traffic for your UI node or vice versa.

If you have the admin role, you can review performance metrics. To view database metrics on the ServiceNow Performance homepage:

1. In the **Graph Set** list, select **Database**.
2. In the **Monitorable Items** list, select an instance.
3. In the **Timespan** list, select a time period.

For details on using the graph and display controls on the ServiceNow Performance homepage, see Platform performance metrics.

- **Database Throughput**: Displays the count (per minute) for each type of database operation (inserts, deletes, and so on) over time.
- **Database Response**: Displays the response time (in milliseconds) for each type of database operation.
- **DB Connection Use**: Shows how many connections this instance has open to the database. View the maximum, median, and minimum number of database connections in use by the selected node.
- **Replication Throughput**: Measures the difference in the data (delta) between the production instance and the replicated instance as user activity changes the database. Replication is the process whereby an entire instance is replicated on a second machine for failover protection.

  Customer Service and Support uses this information to monitor the progress of replicating a customer instance.

Instance View

Review data points across all near nodes (instances) using the Instance View graph set. In a single view, you can review how all active nodes are responding and identify which nodes are experiencing issues.

If you have the admin role, you can review performance metrics. Follow this procedure to view instance performance graphs on the ServiceNow Performance homepage:

1. In the **Graph Set** list, select **Instance View**.
   - The **Monitorable Items** list and **Compare With** selections are both disabled.
   - The **Monitorable Items** list has a value **-All-** for all near (active primary) instances, and excludes far (primary offline or active secondary) instances.
   - Each metric displays data across all active nodes. For example, if you have three instances, all three appear in each metric.

     Note:
     You can view performance metrics for up to 49 near nodes. If you have more than 49 nodes, no data appears for the 50th instance and above.
2. In the **Timespan** list, select a time period.
3. To overlay diagnostic events that have taken place during the time span you requested, toggle **Show on graphs, which appear** next to Diagnostic Events.

For details on using the graph and display controls on the ServiceNow Performance homepage, see Platform performance metrics.

Instance View

- **Transactions**: Displays all UI transactions initiated by users. This graph can show large spikes in end-user traffic and identify when peak end-user activity occurs.
- **Response Time**: Displays the interval (in milliseconds) between the time the instance receives a transaction and the time the instance responds. Displays the time the server takes to complete a transaction, on average, during the given time span.

  An increase in average response time might indicate that there is a systemic issue or an influx of generally slower transactions. To identify possible performance problems, you can correlate response time with other areas, such as memory, database, or a CPU.
- **Sessions**: Shows active sessions, including those sessions initiated by the MID Server and external integrations. Large numbers of stale but active sessions can lead to memory and performance issues. Session counts larger than 10,000 can typically result in performance degradation. Consider reviewing integration session guidelines and limiting session timeouts.
- **Session Wait Queue**: Displays the number of transactions that are waiting on another transaction for the same user. Waiting sessions occur when a user submits a duplicate request before the prior request completes. Can indicate a slow page or a transaction that requires further investigation.

  To identify the transactions that are waiting, check the Transaction Log Entry [syslog_transaction] and view the Session wait time to find the transactions that are waiting. Next, find the Transaction Number that the user is waiting on.
- **Semaphore Use**: Shows the number of semaphores in use by the selected instance. Semaphores control the number of user transactions that can run in parallel.

  Long-running transactions on a semaphore can back up all semaphores, causing transactions to wait. The platform manages semaphores, requiring no customer administration. The semaphore graph is used only by Customer Service and Support for troubleshooting.
- **Semaphore Wait Queue**: Shows the wait queue for a semaphore. Use this graph with the **Semaphore Use** graph. A high wait queue indicates long-running transactions on the semaphore. A high and persistent semaphore queue can indicate that the instance node is overloaded with work.

  Check the Transaction Log [syslog_transaction] to find the longest-running transactions during that time period and identify the problem. This graph is used only by Customer Service and Support.
- **Scheduler Queue Length/ Scheduler Completed Jobs**: Displays all scheduler activity for the selected instance, including Discovery probes. You can determine the backlog of scheduled jobs in the queue for a particular time period. You can then compare that against the rate at which the jobs are being processed during the same period.
- **Errors**: Shows any severe errors printed to the localhost logs or syslog_transactions. Multiple severe errors indicate a problem that requires further investigation.
- **Events Processed**: Shows the average number of events processed during the selected time period.
- **HTTP Transactions**: Displays all completed HTTP transactions, including UI, integration, and AMB traffic. This graph can show large spikes in HTTP traffic and can help identify when peak user activity occurs.
- **Semaphore Rejections**: Display the number of transactions rejected by the following semaphore queues:
  - Semaphore Default Rejected Execution
  - Semaphore Debug Rejected Execution

- Semaphore AMB Receive Rejected Execution
- Semaphore AMB Send Rejected Execution
- Semaphore API INT Rejected Execution
- Semaphore Presence Rejected Execution

Identify processing bottlenecks by monitoring rejection counts in the queues. Look for spikes in the semaphore rejection graphs. A high rejection count indicates an unexpected proliferation in calls. Reduce the number of calls or stop long-running transactions that might be backing up the queue.

Isolating data for specific instances

Using the Fullscreen View, you can isolate information for specific instances in a selected metric.

1. Select a metric, and then click the Fullscreen icon to open an enlarged view of the graph.

2. Clear the check box next to the instances you want to exclude.
   - For example, clear the check box next to Instance_8080 to exclude its data.
   - Only data for Instance_8081 and Instance_8082 appears.

MySQL Global Status

The MySQL Global Status graphs, available to view on the ServiceNow Performance homepage, provide an aggregate view of all traffic going to your database server. These graphs are distinct from the Database graph set, which monitor only queries originating at your particular application instance.

All the traffic shown in the MySQL Global Status graphs is specific to your cluster node and instance. If you have an admin role, you can review performance metrics.

The MySQL Global Status graphs are intended for use by Customer Service and Support to troubleshoot performance issues.

To view metrics for MySQL server status global variables on the ServiceNow Performance homepage:

1. In the **Graph Set** list, select **MySQL Global Status**.
2. In the **Monitorable Items** list, select an instance.
3. In the **Timespan** list, select a time period.

For details on using the graph and display controls on the ServiceNow Performance homepage, see Platform performance metrics.

MySQL Global Status performance graphs

Graph Set	Monitorable Items	Timespan
MySQL Global Status	demo	1 day

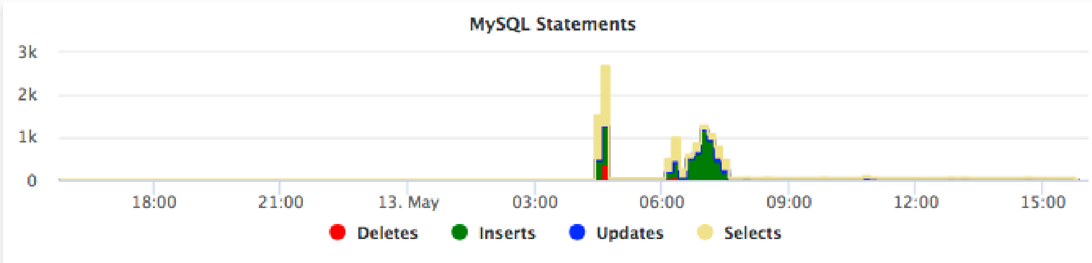

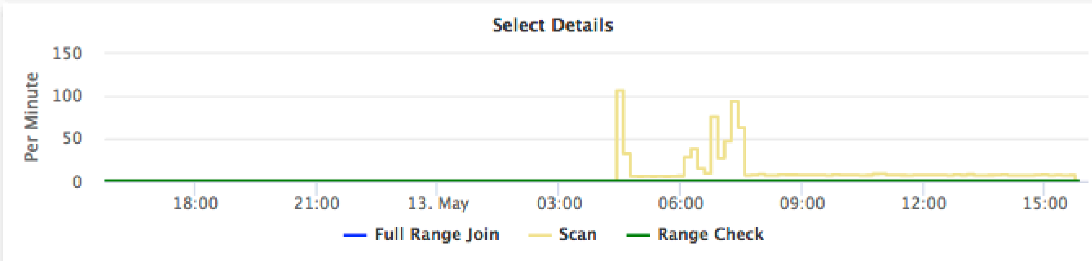

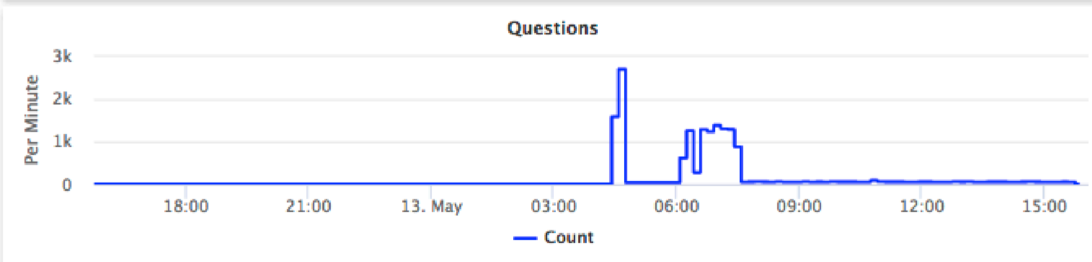

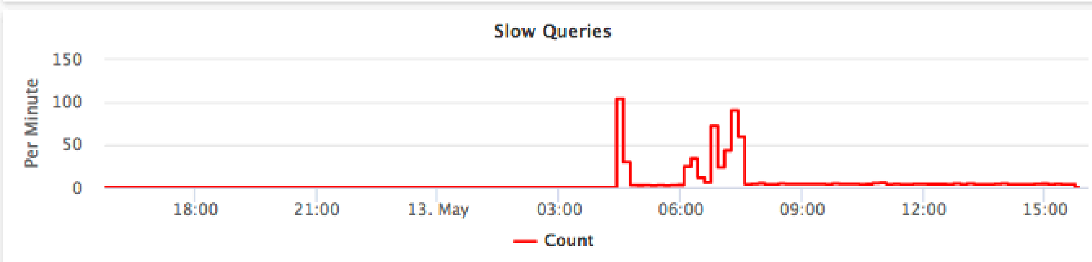

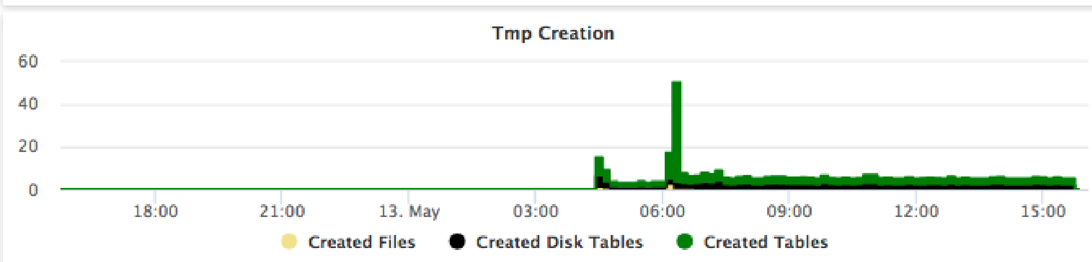

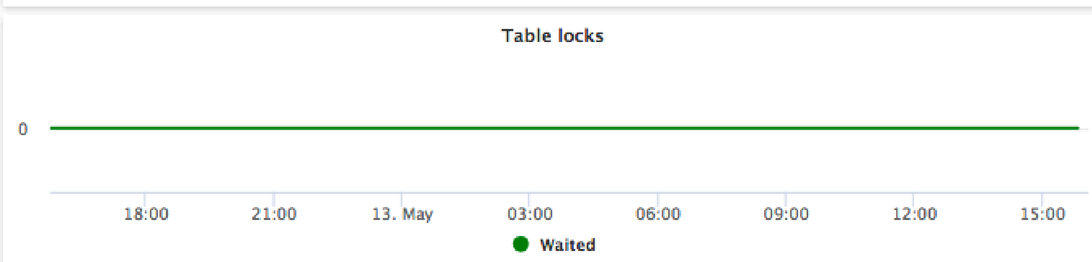

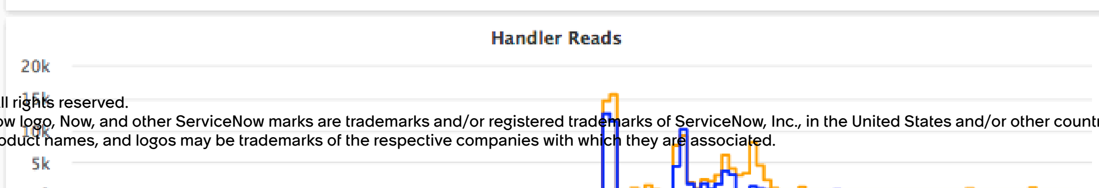

ServiceNow Servlet

Each instance has a servlet, and you can monitor its performance using the Servlet graph set on the ServiceNow Performance homepage.

If you have the admin role, you can review performance metrics. To view servlet graphs on the ServiceNow Performance homepage:

1. In the **Graph Set** list, select ServiceNow **Servlet**.
2. In the **Monitorable Items** list, select an instance.
3. In the **Timespan** list, select a time period.

For details on using the graph and display controls on the ServiceNow Performance homepage, see Platform performance metrics.

ServiceNow Servlet graphs

- **System Overview**: Provides thread performance information. Every second, the system looks at all active threads (both UI and background) and places them into one of the following categories:
  - **CPU**: The thread is active, but is not executing any of the steps. This condition typically means non-business rule compute time, although in this case a few other internal wait states are categorized as CPU. Therefore a 1:1 correlation between threads in a CPU count and hardware CPU utilization is not expected.
  - **Database**: The thread is waiting for information from the database.
  - **Business Rule**: The system is running a business rule (synchronous or asynchronous) and is not currently executing a query (which would be database).
  - **Network**: The thread is writing data out to the network or waiting for an outbound network buffer to flush.
  - **Concurrency**: The threads cannot run because they are waiting on a semaphore or session synchronization.

  The system averages these transactions every minute and records them in the database. The graph shows the averages for each category.

- **Transactions**: Displays all UI transactions initiated by users. This graph can show large spikes in end-user traffic and identify when peak end-user activity occurs.
- **Response Time**: Displays the interval (in milliseconds) between the time the instance receives a transaction and the time the instance responds. Displays the time the server takes to complete a transaction, on average, during the given time span.

  An increase in average response time might indicate that there is a systemic issue or an influx of generally slower transactions. To identify possible performance problems, you can correlate response time with other areas, such as memory, database, or a CPU.

- **Sessions**: Shows active sessions, including those sessions initiated by the MID Server and external integrations. Large numbers of stale but active sessions can lead to memory and performance issues. Session counts larger than 10,000 can typically result in performance degradation. Consider reviewing integration session guidelines and limiting session timeouts.
- **Session Wait Queue**: Displays the number of transactions that are waiting on another transaction for the same user. Waiting sessions occur when a user submits a duplicate request before the prior request completes. Can indicate a slow page or a transaction that requires further investigation.

  To identify the transactions that are waiting, check the Transaction Log Entry [syslog_transaction] and view the Session wait time to find the transactions that are waiting. Next, find the Transaction Number that the user is waiting on.

- **Semaphore Use**: Shows the number of semaphores in use by the selected instance. Semaphores control the number of user transactions that can run in parallel.

  Long-running transactions on a semaphore can back up all semaphores, causing transactions to wait. The platform manages semaphores, requiring no customer administration. The semaphore graph is used only by Customer Service and Support for troubleshooting.

- **Semaphore Wait Queue**: Shows the wait queue for a semaphore. Use this graph with the **Semaphore Use** graph. A high wait queue indicates long-running transactions on the semaphore. A high and persistent semaphore queue can indicate that the instance node is overloaded with work.

Check the Transaction Log [syslog_transaction] to find the longest-running transactions during that time period and identify the problem. This graph is used only by Customer Service and Support.

- **Scheduler Queue Length/ Scheduler Completed Jobs**: Displays all scheduler activity for the selected instance, including Discovery probes. You can determine the backlog of scheduled jobs in the queue for a particular time period. You can then compare that against the rate at which the jobs are being processed during the same period.
- **Errors**: Shows any severe errors printed to the localhost logs or syslog_transactions. Multiple severe errors indicate a problem that requires further investigation.
- **Events Processed**: Shows the average number of events processed during the selected time period.
- **Events Logged**: Shows the average number of events queued and added to the event log in the selected time period.
- **HTTP Transactions**: Displays all completed HTTP transactions, including UI, integration, and AMB traffic. This graph can show large spikes in HTTP traffic and can help identify when peak user activity occurs.
- **Semaphore Rejections**: Display the number of transactions rejected by the following semaphore queues:
    - Semaphore Default Rejected Execution
    - Semaphore Debug Rejected Execution
    - Semaphore AMB Receive Rejected Execution
    - Semaphore AMB Send Rejected Execution
    - Semaphore API INT Rejected Execution
    - Semaphore Presence Rejected Execution

  Identify processing bottlenecks by monitoring rejection counts in the queues. Look for spikes in the semaphore rejection graphs. A high rejection count indicates an unexpected proliferation in calls. Reduce the number of calls or stop long-running transactions that might be backing up the queue.

Slow Pattern

You can identify Slow Patterns (transactions, scripts, slow queries, events, mutexes, and cache) in your instance by using the Slow Pattern graph set on the ServiceNow Performance homepage.

If you have an admin role, you can review performance metrics. To view Slow Pattern graphs on the ServiceNow Performance homepage:

1. In the **Graph Set** list, select **Slow Pattern**.
2. In the **Monitorable Items** list, select an instance.
3. In the **Timespan** list, select a time period.

For details on using the graph and display controls on the ServiceNow Performance homepage, see Platform performance metrics.

In the fullscreen view of each Slow Pattern graph, the legend provides links to the Slow Pattern records represented in the graph. For example, the legend in the Slow Query Time graph lists the slowest queries for the instance during the selected timespan.

- Clicking a slow query item in the legend opens the Slow Queries record for that item. You can review query details and determine what action to take on the slow query.
- For convenience, associated Slow Pattern graphs appears below the record so you can review and compare the graphs with details from the Slow Pattern record.

Slow Pattern performance graphs

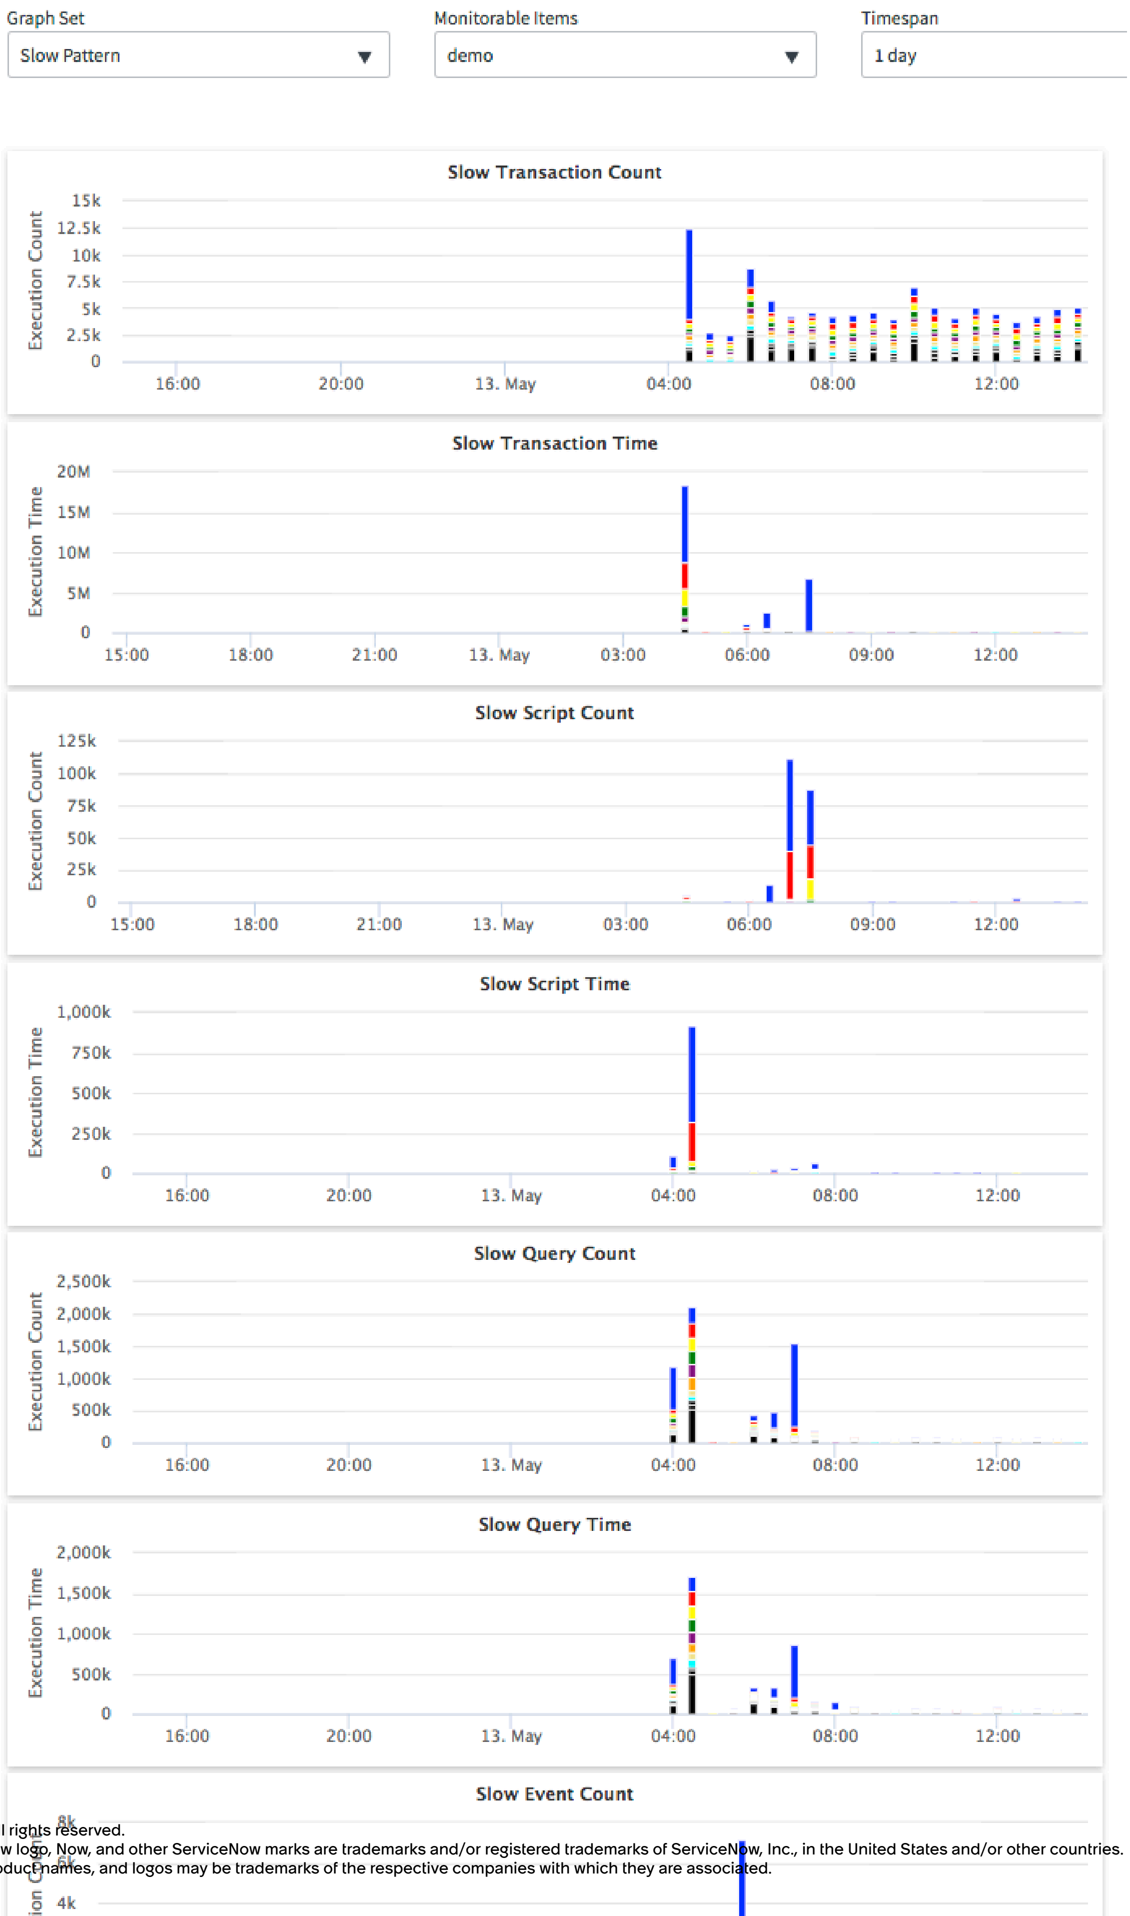

Application Insights

The ServiceNow® Application Insights application available in the ServiceNow Store, provides a centralized location where you can visualize and monitor system health.

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

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.

Overview

Application Insights enables you to proactively manage the health and performance of your instances. Application Insights provides time-series graphs and data that enable you to visualize performance trends, see correlated events, and access resources to take corrective action.

You can access Application Insights by logging in as a user with the admin or sn_app_insights.admin role and navigating to **All** > **Application Insights** > **Application Insights**.

Key benefits

- View the health and performance of all your instance nodes in one place.
- Detect anomalies or deviations in instance performance.
- Take corrective actions and resolve performance issues on your own without having to involve an administrator.
- Analyze the root cause of a performance issue by determining the context and correlation with system events such as an upgrade or application install.
- Proactively mitigate potential risks to the system. Identify and pre-empt risks that would otherwise lead to issues.

Application Insights

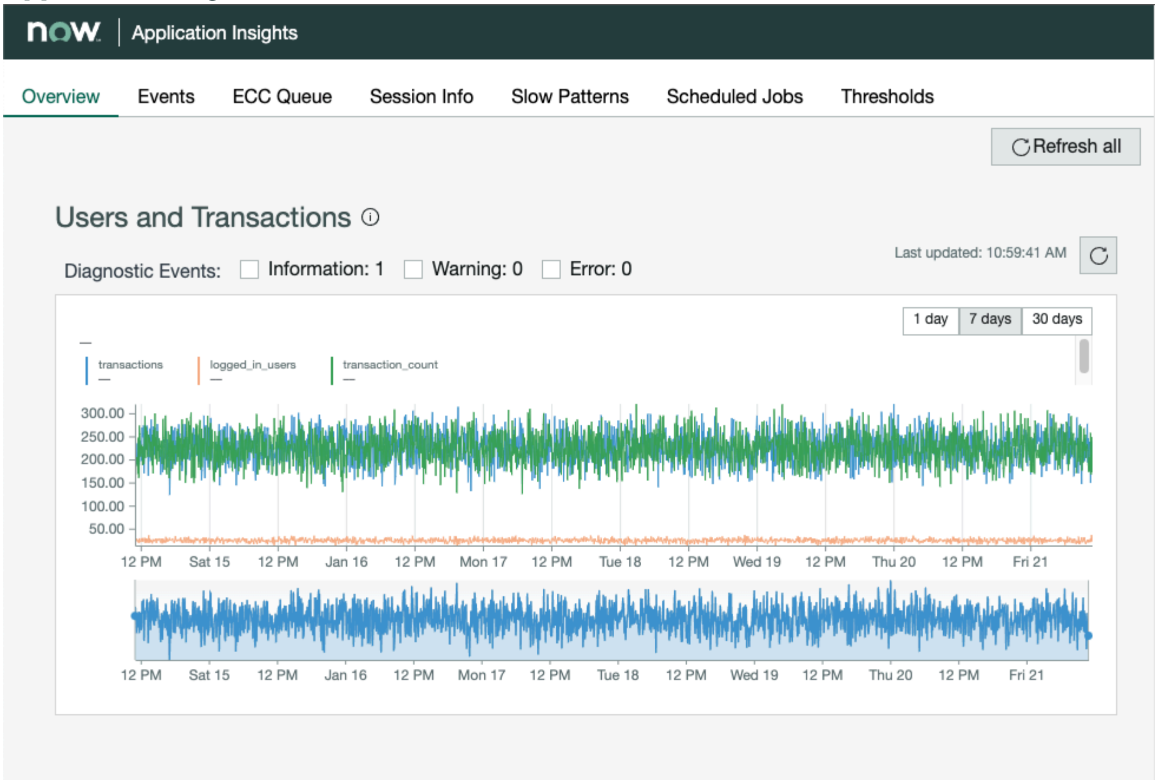

Working with Application Insights graphs

View key performance metrics for your instance using Application Insights graphs.

Application Insights provides time-series data in a collection of graphs that help you visualize the performance and health of your instance. The graphs are automatically refreshed every five minutes.

Overview graphs

You can view a summary of the overall health and performance of your instance on the **Overview** tab. It provides in an aggregated instance-level view by combining metrics from the **Events**, **ECC Queue**, and **Session Info** tabs.

Use the overview graphs to do the following:

- Analyze how metrics have been performing over time individually and in relation to each other.
- View trends or patterns between different data sets by overlaying metrics and diagnostic events.
- Closely and proactively monitor the instance when performing any impactful activity such as a system update or during periods of heavy user traffic.

Consolidated view of Users and Transactions graph

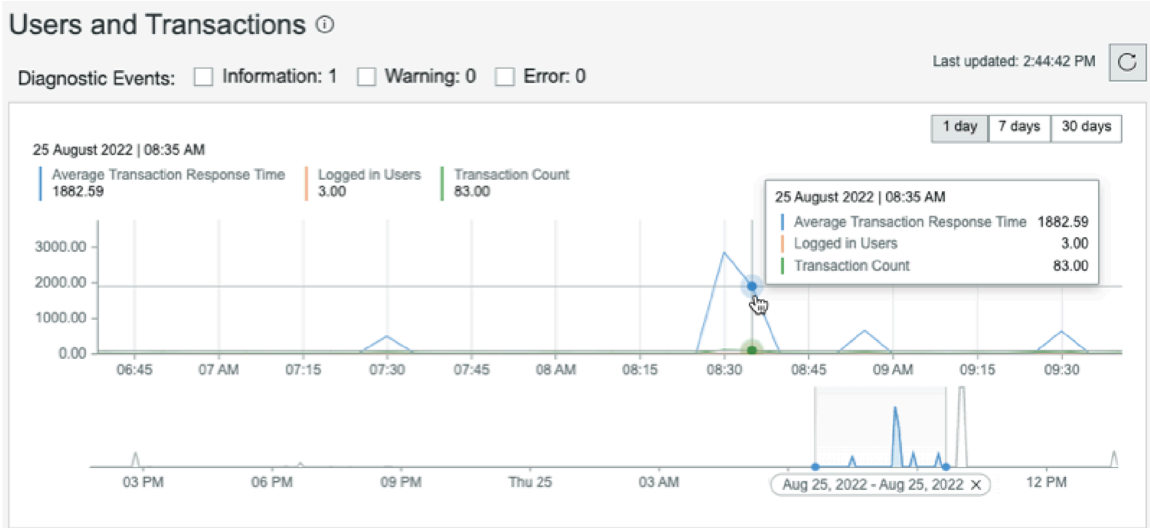

- Overlay diagnostic events to help you determine causation and correlation.
- Compare metrics or focus on a specific metric by selecting which graph lines to display from the legend above the graph.
- Change the time period by selecting a different day range from the date picker. The default value is 1 day, but you can also change the time period to 7 days or 30 days.
- View the details for all metrics at a point in time by pointing your cursor anywhere on the graph.

  View node-level details for a metric by pointing your cursor and selecting a data point. The corresponding detail graph on the **Events**, **ECC Queue**, or **Session Info** tab is displayed.
- Narrow the time range by selecting the left or right side of the time line below the graph and dragging across. Shift the time range by dragging this view along the axis.

Detail graphs

Each graph on the **Events**, **ECC Queue**, and **Session Info** tabs provides a detailed view of metrics for one or more nodes.

Use the detail graphs to do the following:

- Compare a specific metric to the median, 1-day moving average, or the 95th percentile value
- View trends or patterns by overlaying metrics and diagnostic events
- Determine whether performance degradation is happening on one node or across all nodes
- View detailed statistics for each node to monitor system health over time

Detail graphs include the same controls as overview graphs but provide additional functionality.

Detailed view of Logged in Users

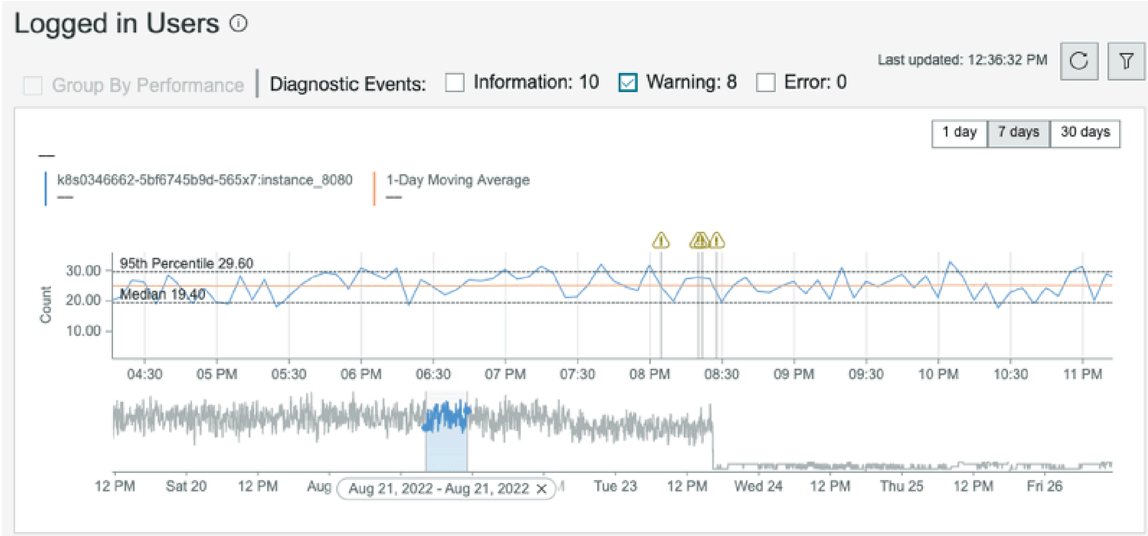

- See trends, issues, and system health for nodes and agents with similar performance values by selecting **Group By Performance**. If there are seven or more lines on the graph, this option is selected by default.
- Overlay diagnostic events to help you determine causation and correlation.
- Add a custom filter to display data for specific nodes. For example, display the average transaction time for nodes where the status is online.
- Compare the performance of different nodes to the 1-day moving average by selecting which graph lines to display from the legend above the graph.

  The **1-Day Moving Average** line represents the average value across all nodes over one day.
- Change the time period by selecting a different day range from the date picker. The default value is 1 day, but you can also change the time period to 7 days or 30 days.
- View the details for an agent or node by pointing your cursor anywhere on the graph.

  Compare the metric to the following horizontal lines on the graph:

  - **Median**: The middle value.
  - **95th percentile**: The 95th percentile value. A metric with a data point above this line is likely a cause for concern.

  View a list of associated events or transactions within a 10-minute period by pointing your cursor and selecting a data point.
- Narrow the time range by selecting the left or right side of the time line below the graph and dragging across. Shift the time range by dragging this view along the axis.

The median, average, and 95th percentile values for each node are displayed in the table below the detail graph.

- To sort records in the list, select the column header.
- To filter the list using a value in a field, right-click in the field and select **Show Matching** or **Filter Out**.

- To group records in the list by label, select the column options menu and then select **Group by**.
- To filter records in the list, select the column options menu and then select **Filter**. Enter a condition and value.

Application Insights overview graphs and metrics

The Application Insights overview graphs provide consolidated views of key metrics.

Application Insights Overview graphs and metrics

Graph	Key metrics	Description
Users and Transactions	• Average Transaction Response Time • Transaction Count • Logged in Users	Total number of logged-in users, transaction count, and average response time of transactions across all nodes
Events	• Events Processed • Events Logged	Total number of events logged and events processed across all nodes
Current ECC Queue Output	• Ready on Output • Processing on Output • Processed on Output	Total number of entries in the ECC output queue in the Ready, Processing, and Processed states
Current ECC Queue Input	• Ready on Input • Processing on Input • Processed on Input	Total number of entries in the ECC input queue in the Ready, Processing, and Processed states
Semaphore API_INT	• Semaphore API_INT Queue Depth • Semaphore API_INT Rejections	Average number of transactions in the API_INT semaphore queue and total number of transactions rejected across all nodes
Semaphore Default	• Semaphore Default Queue Depth • Semaphore Default Rejections	Average number of transactions in the Default semaphore queue and total number of transactions rejected across all nodes
Semaphore AMB Send	• Semaphore AMB Send Queue Depth • Semaphore AMB Send Rejections	Average number of transactions in the AMB Send semaphore queue and total number of transactions rejected across all nodes

Application Insights Overview graphs and metrics (continued)

Graph	Key metrics	Description
Semaphore AMB Receive	• Semaphore AMB Send Queue Depth • Semaphore AMB Send Rejections	Average number of transactions in the AMB Receive semaphore queue and total number of transactions rejected across all nodes

Application Insights detail graphs and metrics

The Application Insights detail graphs provide views of individual metrics at the node level.

Events detail graphs

Metric	Source	Description
Events Processed	`events_processed` key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of events processed
Events Logged	`event_logs` key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of events logged

ECC Queue detail graphs

Metric	Source	Description
Ready on Output	ECC Queue Statistics [ecc_queue_stats_by_ecc_agent] table	Average number of entries in the Ready state in the ECC output queue
Ready on Input	ECC Queue Statistics [ecc_queue_stats_by_ecc_agent] table	Average number of entries in the Ready state in the ECC input queue
Processing on Output	ECC Queue Statistics [ecc_queue_stats_by_ecc_agent] table	Average number of entries in the Processing state in the ECC output queue
Processing on Input	ECC Queue Statistics [ecc_queue_stats_by_ecc_agent] table	Average number of entries in the Processing state in the ECC input queue
Processed on Output	ECC Queue Statistics [ecc_queue_stats_by_ecc_agent] table	Average number of entries in the Processed state in the ECC output queue
Processed on Input	ECC Queue Statistics [ecc_queue_stats_by_ecc_agent] table	Average number of entries in the Processed state in the ECC input queue

Session information detail graphs

Metric	Source	Description
Average Transaction Response Time	*transactions* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Average transaction processing time in seconds
Transaction Count	*transactions* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Load of transactions on the instance
Logged in Users	*sessionsummary* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of users logged in to the instance over time
Semaphore API_INT Queue Depth	*semaphores* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions in the API Integrations semaphore queue
Semaphore Default Queue Depth	*semaphores* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions in the Default semaphore queue
Semaphore AMB Send Queue Depth	*semaphores* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of messages in the AMB Send queue being sent to the client
Semaphore AMB Receive Queue Depth	*semaphores* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of messages in the AMB Send queue being received from the client
Semaphore API_INT Rejections	*semaphore_api_int_rejected* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions rejected by the API Integrations semaphore queue
Semaphore Default Rejections	*semaphore_default_rejected* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions rejected by the Default semaphore queue
Semaphore AMB Send Rejections	*semaphore_amb_send_rejected* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions rejected by the AMB Send queue
Semaphore AMB Receive Rejections	*semaphore_amb_receive_rejected* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions rejected by the AMB Receive queue
Average Database Response Time	*sql_response* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Average response time (in milliseconds) for database operations
Database Throughput	*database_throughput* key in the Stats column in the Node stats [sys_cluster_node_stats] table	Number of transactions processed per second.

Monitoring users and transaction performance through Application Insights

Maintain the health and performance of your instance by monitoring key metrics related to users and transactions through the Application Insights Users and Transactions graph.

You can detect anomalies and performance deviations in your instance by monitoring the relationship between users, transactions, and average transaction response time in the Users and Transactions graph on the **Overview** tab.

Use the Users and Transactions graph to do the following:

- Monitor the performance of transactions on your instance.
- View the transaction load on your instance.
- Track the number of logged-in users.

You access the graph by navigating to **All** > **Application Insights** > **Application Insights** > **Overview**.

- Analyze system slowness by comparing the **Transaction Count** and **Logged in Users** figures. Note that a spike in the **Transaction Count** that has a corresponding spike in the **Logged in Users** count might be a false alarm.
- Determine whether a heavy transaction load is affecting system performance by comparing the **Average Response Time** figure to the **Transaction Count** amount. Look for spikes in the **Transaction Count** amount relative to the **Average Response Time** figure.

Dig deeper into a potential performance issue by drilling down to analyze issues at the node-level in the detail graphs on the **Session Info** tab.

- Look for spikes that indicate a bottleneck or loop condition in the **Transaction Count** detail graph. An upward trend line indicates the system is becoming overloaded. Select a data point on the graph to view a list of transactions in a 10-minute range.
- Look for spikes that indicate times of heaviest traffic on the instance in the **Logged in Users** detail graph. Monitor the trend lines to determine when the system is busiest.
- Determine whether performance degradation is happening on one node, across all nodes, or on a set of nodes that are consolidated using the **Group by Performance** option. Compare the node metrics to the **1-Day Moving Average** amount. If you don't see any obvious issues, look for anomalies in execution counts or high execution times on the **Slow Patterns** tab. If you see both, start your investigation there.

Monitoring semaphore queue efficiency through Application Insights

Monitor semaphore queue efficiency by tracking the queue depth and number of rejected transactions through the Application Insights semaphore graphs.

You can monitor performance of the semaphore queues in your instance by comparing the semaphore queue depth and rejection count in the semaphore graphs on the **Overview** tab.

Through the semaphore graphs, you can monitor the following metrics:

- The semaphore queue depth, which enables you to help prevent backlogs
- Rejection counts, which enables you to identify processing bottlenecks

Semaphores control the number of simultaneous transactions that can execute on a node. Low available semaphores indicate that the instance is running close to full transaction capacity. Incoming transactions will wait until the semaphores are available.

You access the semaphore graphs by navigating to **All** > **Application Insights** > **Application Insights** > **Overview**.

- Monitor the load level by comparing the queue depth to the queue depth limit that appears as a red line on each semaphore graph.
- Look for correlations or trends by comparing the queue depth metric to the rejection count metric over time.

Dig deeper into semaphore processing efficiency by drilling down to analyze issues at the node-level in the detail graphs on the **Session Info** tab.

- Look for spikes in the queue depth graphs. A high transaction count indicates the node might be overloaded. Select a data point in the detail graph to view a list of transactions. In the table, look for long-running transactions on the semaphore and stop them. Reduce the transaction count if the queue depth is consistently high.
- Look for spikes in the rejection count graphs. A high rejection count indicates an unexpected proliferation in calls. If the queue depth is reached, all subsequent requests are rejected, which results in a rejected requests (HTTP 429) error message. Reduce the number of calls or stop long-running transactions that might be backing up the queue.

Monitoring database performance through Application Insights

Maintain the health and performance of your database by monitoring the volume of database transactions and average response time through Application Insights.

You can detect anomalies and database performance issues by monitoring the database graphs on the **Session Info** tab.

Use the database graphs to monitor the following metrics:

- Average response time
- Volume of transactions by analyzing database throughput

You access the graphs by navigating to **All** > **Application Insights** > **Application Insights** > **Session Info**.

- Look for spikes that indicate delays in response time for each type of database operation (deletes, inserts, updates, and selects) in the **Average Database Response Time** graph. View a list of database transactions by selecting a data point in the spike, and then look for outliers in the response time. Review the source record to see whether the transaction can be optimized.
- Look for spikes or periods of heavy usage for each type of database operation (deletes, inserts, updates, and selects) in the **Database Throughput** graph. View a list of database transactions by selecting a data point in the spike, and then look for outliers in SQL count and response time. Review the source record to see whether you can optimize the transaction.

Monitoring event queue efficiency through Application Insights

You can monitor event queue performance in Application Insights by comparing and analyzing the rate at which events are logged and processed.

You can monitor the relationship between events logged and events processed by viewing the **Events** graph on the **Overview** tab.

Use the **Events** graphs to do the following:

- Monitor the rate of incoming events
- Monitor the rate of processed events
- Detect anomalies in processing events

You access the Events graphs by navigating to **All** > **Application Insights** > **Application Insights** > **Overview**.

- Look for bottlenecks by comparing the **Events Logged** and **Events Processed** totals. A spike in the **Events Logged** number without a corresponding spike in the **Events Processed** number indicates a problem with processing events.
- A high and consistent spike in both the **Events Logged** and **Events Processed** figures indicates the system is receiving and processing a large load of events. Look for loops or conditions that might be causing a continuous stream of incoming events.

Dig deeper into potential performance issues by drilling down to analyze issues at the node-level in the detail graphs.

- Analyze the processing of events over time by comparing processed events to the **1-Day Moving Average** shown on the **Events Processed** graph.
- Look for patterns in event logging over time. Determine whether the same spike is happening at the same time every week.
- Look for correlations between the count of events logged and system activities by overlaying diagnostic events on the graph. For example, if you see a spike in the Events Logged number and you notice that it coincides with the installation of an update set, you can investigate the update set to determine why it caused the spike in events logged.
- Investigate the cause of an issue by selecting a data point at the start of spike to view a list of events created 5 minutes before and 5 minutes after the issue.
- Analyze the rate of incoming events by comparing logged events to the **1-Day Moving Average** total on the **Events Logged** graph.
- Identify potential performance issues by comparing metrics for individual event queues in the Event Queue table.
  - Focus on a 1-day, 7-day, or 30-day period by selecting a day range.
  - Identify which queue had the most logged events by sorting on the **Logged events in range** column. An event is logged when it is inserted into the Events [sysevent] table.
  - Identify which queue had the most queued events by sorting on the **Queued events in range** column. An event is queued when it is assigned to a specific event queue in the Events [sysevent] table.
  - Identify which queue had the most unprocessed events by sorting on the **Unprocessed events in range** column.
  - Identify which queue took the most amount of time processing events in the selected day range by sorting on the **Processing duration in range** column.
  - Identify which queue processed the most events in the selected day range by sorting on the **Processed events in range** column.
  - Identify which queue took the most amount of time processing events on average in the selected day range by sorting on the **Average processing duration in range** column.

Monitoring MID server performance through Application Insights

You can monitor the performance of a MID Server by tracking the status of entries in the ECC queue.

Application Insights enables you to monitor the performance of the ECC queue, which is a communication log between your instance and an MID Server. For details, see MID Server ECC Queue.

You can monitor the relationship between the number of ready, processing, and processed records by viewing the ECC queue graphs. Use the ECC queue graphs to monitor the following:

- Load handling, viewing the number of entries in your ECC input and output queues
- Processing efficiency, by tracking entries in the ready, processing, and processed states in both of the ECC queues

You access the graphs by navigating to **All** > **Application Insights** > **Application Insights** > **Overview**.

- Look for a spike in the ready state without a corresponding spike in the processing and processed states. This pattern indicates a potential problem with reading the queue.
- Look patterns or trends that indicate worsening performance over time. A steady increase in ready entries with a flat or declining processing rate indicates a bottleneck or performance issue that merits further investigation.

Dig deeper into potential performance issues by drilling down to analyze issues for an agent or group of agents in the detail graphs.

- Make the detail graphs easier to read by grouping agents with similar performance metrics by selecting the **Group by Performance** check box. Look for groups with higher counts or spikes in the selected time range. View and compare trend lines for a group of agents in a separate graph by selecting a data point. Look for anomalies by comparing the individual agents to the **1-Day Moving Average** number.
- Analyze ECC queue efficiency by tracking the number of entries in the ready, processing, and processed states over time. Output metrics measure the number of jobs leaving the instance. Input metrics measure the number of jobs sent to the instance from the MID Server or another system.

Solving slow patterns

You can identify, prioritize, and troubleshoot performance issues related to slow patterns in your instance.

Identify performance issues related to slow patterns, for example, poorly performing scripts, queries, transactions, or events, using the information in the **Slow Patterns** tab.

- Identify slow patterns that are executed often within a given time range.
- Analyze the performance of slow events, transactions, queries, and scripts over time.
- Identify the source of a slow pattern.
- Prioritize potential performance improvements.

For example, the Slow Queries table shows which queries have high average execution times within the selected time range compared to the overall average execution time. A high average execution time indicates that an entry is performing poorly. Once you identify a slow query that needs investigating, select the query to view a graph of the performance over time, and then access the slow query record for details.

Troubleshoot a slow pattern

Identify the source of a slow pattern and prioritize potential performance improvements.

Before you begin

Role required: sn_app_insights.admin or admin

Procedure

1. Navigate to **All** > **Application Insights** > **Application Insights** > **Slow Patterns**.
2. Look for potential performance issues by viewing the Slow Events, Slow Transactions, Slow Queries, and Slow Scripts tables.
   - Focus on a 1-day, 7-day, or 30-day period by selecting a day range.
   - Identify patterns with consistently high execution times by sorting on the **Average execution time in range** column.
   - Find out which patterns are executed the most often by sorting on the **Execution count in range** column.
3. Select a slow pattern with a high execution time and high execution count.
   The **Average Execution Time** detail graph shows the slow pattern's execution time and its 1-day moving average.
4. View the performance for the slow pattern over time by analyzing the **Average Execution Time** detail graph.

   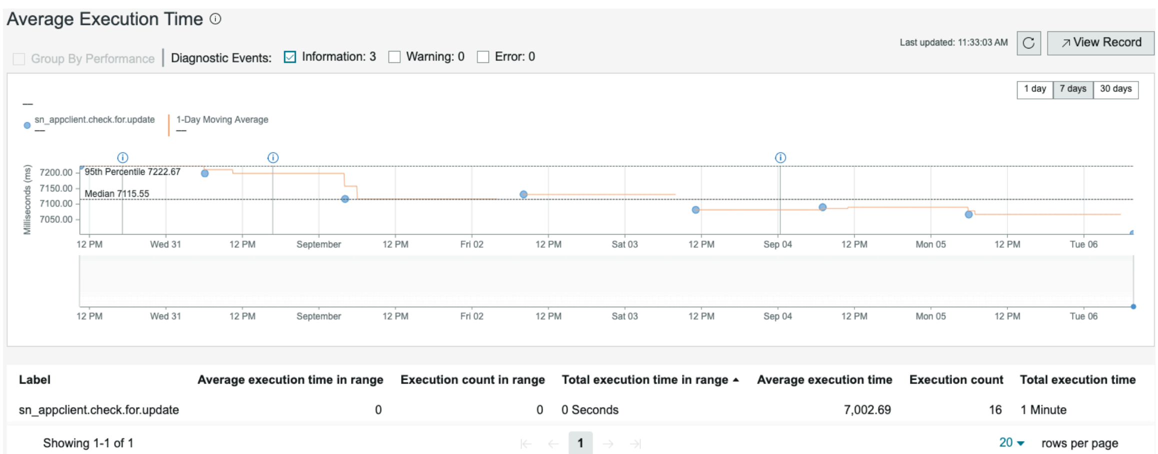

   - Look for spikes within the selected time range. Look for correlations to impactful system events that might indicate a false alarm by overlaying diagnostic events.
   - View the number of times the pattern was executed in the selected time range.
   - Determine whether performance is worsening over time by viewing the direction of the 1-Day Moving Average trend line.
5. Find out what is triggering the slow pattern by selecting **View Record**.
   The slow pattern record appears and provides additional details including the query or script contents, the first time it was executed, and the last time it was executed.
6. To access additional helpful information for troubleshooting, add the Referenced Scripts and Related Slow Patterns related lists.

Option	Description
Slow scripts	**a.** Select the form context menu ( ). **b.** Change the form view by selecting **View** > **Slow Script Insights**.

Option	Description
Slow queries	**a.** Select the form context menu ( ). **b.** Change the form view by selecting **View** > **Slow Query Insights**.

7. Determine the cause of the slowness.
   - When investigating a slow query, determine which script or business rule triggered the slowness by finding the entry with the highest calling order in the Referenced scripts related list. For example, suppose a slow query is triggered by a script whose calling order is 2 that is called by a business rule whose calling order is 1. That script directly triggered the slow query because it has the highest calling order.
   - When investigating a slow script, identify the slow patterns triggered by the script by viewing the patterns in the Related slow patterns related list. Determine which slow patterns to investigate first by sorting on the **Average Execution Time in Range** and **Average Execution Count in Range** columns. Investigate the slow patterns with the highest values in each column first.
   - Confirm a potential issue by viewing the Related slow patterns list, which provides a list of other slow patterns that have referenced scripts in common. If multiple slow patterns reference the same script include or business rule, you can be confident that is where the issue lies.
8. **Optional:** For troubleshooting multiple slow patterns, open each record, check the number of entries in the Related Slow Patterns related list, and prioritize debugging or resolving the slow pattern with the higher count.
   Debugging the slow pattern with the higher count is more likely to make a greater performance improvement.
9. Take actions to solve the performance issue.
   - Optimize or remove the offending script include or business rule.
   - Determine whether you can avoid using the slow query. If the query is required, try to optimize it with additional query conditions or a sys_id query so that it returns only the information that is needed.
   - Determine whether an index can optimize the performance of the slow query.

Example: Troubleshooting a slow query

After sorting the Slow Queries table by the **Execution count in range** field, you see a SELECT statement with a high execution count over the last 7 days.

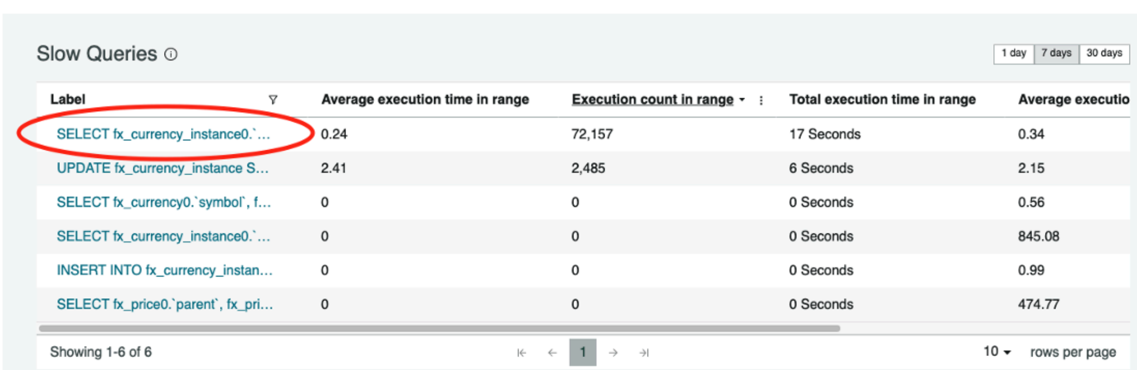

Slow Queries

1 day | 7 days | 30 days

Label	Average execution time in range	Execution count in range	Total execution time in range	Average executio
SELECT fx_currency_instance0.`...	0.24	72,157	17 Seconds	0.34
UPDATE fx_currency_instance S...	2.41	2,485	6 Seconds	2.15
SELECT fx_currency0.`symbol`, f...	0	0	0 Seconds	0.56
SELECT fx_currency_instance0.`...	0	0	0 Seconds	845.08
INSERT INTO fx_currency_instan...	0	0	0 Seconds	0.99
SELECT fx_price0.`parent`, fx_pri...	0	0	0 Seconds	474.77

Showing 1-6 of 6 | 1 | 10 rows per page

To find out what is triggering this slow query you would select the query name to open the detail graph, and then select **View record**.

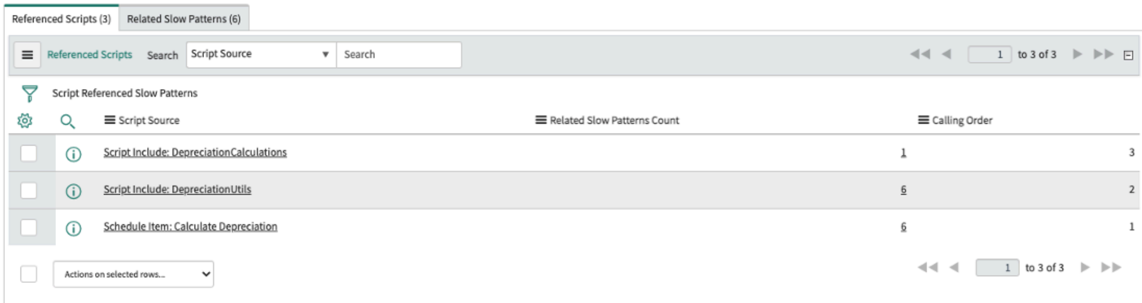

The Referenced Scripts related list displays three scripts. The DeprecationCalculations script has the highest calling order of all the referenced scripts, indicating that it directly triggered the slow pattern.

To remedy the issue:

- Examine the DeprecationCalculations script include and determine whether you can avoid using this query. In this example, the query is against the fx_currency table, which implies a currency calculation is happening. Try to circumvent this calculation or remove it if it isn't necessary.
- If the query cannot be avoided in DeprecationCalculations, determine whether the call to the other scripts with lower call orders could be avoided or decreased in execution count. In this example, DepreciationUtils needs to call into DeprecationCalculations, but perhaps you could set the scheduled job Calculate Depreciation to be run less frequently.
- If adjustments to the scripts lower in the call order are not viable, consider adding an index to potentially help with performance.

Troubleshoot a scheduled job through Application Insights

Identify a scheduled job that is causing slow performance or running more often than necessary through the Scheduled Jobs table.

Before you begin

Role required: sn_app_insights.admin or admin

Procedure

1. Navigate to **All** > **Application Insights** > **Application Insights** > **Scheduled Jobs**.
2. Identify potential performance issues by viewing the executed script includes in the Scheduled Jobs table.
   - Focus on a 1-day, 7-day, or 30-day period by selecting a day range.
   - Identify long-running jobs by sorting on the **Average processing duration in range** column.
   - Find out which jobs are running the most often by sorting on the **Run count in range** column.
   - Identify jobs with the most errors by sorting on the **Error count in range** column.
3. Investigate a potential issue by selecting a scheduled job.
   The **Processing Duration** detail graph shows the scheduled job's average execution time and its 1-day moving average.
4. View the performance for the scheduled job over time by analyzing the **Processing Duration** detail graph.

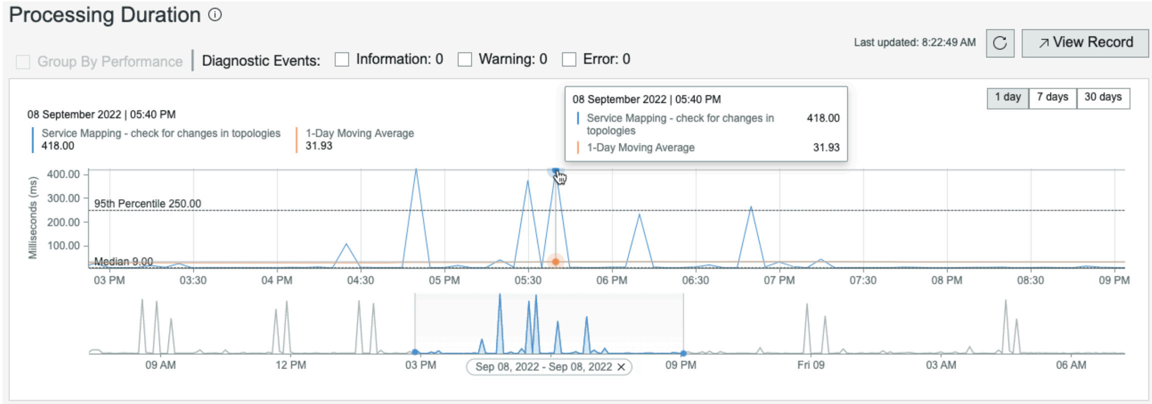

Determine whether the job is running unexpectedly, running during business hours, or running during peak usage time periods.

5. View the job's details and schedule to determine what improvements you can make by selecting **View Record**.
   - Change the job's schedule if the job is running too often or during times of peak usage.
   - Determine the best time to run a scheduled job by looking for periods of low user activity on the **Logged in Users** graph on the **Session Info** tab.
   - Review the scheduled script or report and determine whether it can be optimized or removed.
6. Select **Update**.
7. Check the performance of the scheduled job over time by viewing the **Processing Duration** detail graph.
8. Continue making adjustments to the schedule or job details until the average processing duration improves.

Monitoring API performance

Monitor the performance of your integrations with the Application Insights API metrics graphs. You can monitor how these integrations are performing over time and how those APIs perform over time.

> **Note:**
> admin and sn_app_insights.admin are the roles required to experience this improved visibility.

When you select the **API metrics** tab, the list of endpoints is displayed in tabular format. For each API endpoint and request method (GET, POST, PUT, PATCH, and DELETE), the following metrics are captured:

- Response time (ms)
- Requests per minute
- Payload size (bytes)
- Failure rate (%)

Configure API metrics

Configure the API metrics feature in Application Insights to enable instance administrators to proactively monitor instance performance.

Before you begin

The **API metrics** tab shows up under Application Insights if the following conditions are met.

- The API metric com.glide.rest.analytics.metricbase plugin is installed.

  Note:
  By default, the API metric plugin is installed along with the Application Insights version upgrade.

- Application Insights is installed on your instance.

Role required: admin

Procedure

1. Navigate to Application Manager and search for com.glide.rest.analytics.metricbase.
2. Select your target instance from the menu where Application Insights is installed and where the API metrics plugin should be activated.
3. Enter the following details.
   - Plugin ID: com.glide.rest.analytics.metricbase
   - Name: API Metrics

Application Insights p1 prediction model

Application Insights enables you to receive a warning when your instance is about to experience a priority 1 (p1) event.

When the p1 prediction feature is enabled, the system monitors system performance and models it to predict when a p1 condition is likely. If the system determines that a p1 condition is likely, a warning icon () is displayed on the graphs in the Application Insights application.

You must install the Predictive Intelligence (com.glide.platform_ml) plugin and the Application Insights application to enable the p1 prediction feature.

To see the warning message, point your cursor to the warning icon in any graph. A pane is displayed with a message similar to the following:

```
P1 Predicted: Entering a p1 alert state. We predicted a p1
might occur within 20 minutes with a confidence of 99.0%. It
might be due to exhaustion in the default semaphore set and/or
an unusually high number of transactions. <p>Performed on the
instance at 2021-07-28 11:12:26</P>
```

The warning enables you to take corrective action and avoid a p1 condition from occurring.

After you take corrective action to successfully avoid the p1 condition from occurring, the system displays a second warning indicator to signify the end of the event. Point your cursor to the warning icon. A pane is displayed with a message similar to the following:

```
P1 Predicted: Ending current p1 alert state. We no longer predict
a p1 might occur. <p>Performed on the instance at 2021-07-16
12:11:26</p>
```

The default system properties settings for the p1 prediction model are appropriate for most environments. You can change the properties for your environment. For more information about changing the default settings, see Application Insights properties.

Configure Application Insights thresholds

Configure conditional thresholds to trigger an alert that notifies you when one or more metrics, such as response time, is outside of the desired range.

Before you begin

Role required: sn_app_insights.admin or admin

Procedure

1. Navigate to **All** > **Application Insights** > **Application Insights**.
2. Select **Thresholds**.
3. Select **New**.
4. Enter a name.
5. Select a metric.
6. Select an operator.
7. Enter a value in the **Threshold** field.
8. Enter the number of minutes in the **Sustained for (min)** field.
9. To add an additional metric, select **and**.
10. To add an alternate metric, select **or**.
11. Select **Save**.
12. To edit a threshold, select it from the Thresholds list.
13. Edit the threshold and select **Save**.
14. To create a trigger action, select the threshold from the Thresholds list and from the ellipsis, select **Create flow**.
    The ServiceNow® Flow Designer opens with a default template in place. For more information about configuring triggers, see Configure Application Insights threshold triggers.
15. To delete a threshold, select it from the Thresholds list and from the ellipsis, select **Delete**.
16. To create a new threshold based on an existing threshold:
    a. Select the existing threshold from the Thresholds list.
    b. Edit the name and other values.
    c. From the ellipses, select **Insert and Stay**.

Result

The threshold is listed on the **Thresholds** tab and is shown as a horizontal dotted red line on the associated graph.

Configure Application Insights threshold triggers

Detect that a threshold has been exceeded and create a trigger to perform a sequence of actions.

Before you begin

Role required: sn_app_insights.admin or admin

About this task

Application Insights uses the ServiceNow® Flow Designer to create the trigger and action. You can configure one flow per threshold. You can use any of the actions available in the Flow Designer and the spokes from Integration Hub. This example configures a trigger to send an email.

Procedure

1. Navigate to **All** > **Application Insights** > **Application Insights**.
2. Select **Thresholds**.
3. Select an existing threshold from the Thresholds list.
4. From the ellipses, select **Create flow**.
   The Flow Designer opens. The system automatically populates a base template with a trigger and three actions.

   Modify the base template to fit your needs.

   For more information about using the Flow Designer, see Flow Designer.
5. Select **Send Email**.
6. In the To, CC, and BCC lines, enter the email addresses that you want to send the notification to.
7. In the **Subject** and **Body** fields, enter the text that you want.
8. Select **Done**.
9. Select **Save**.
10. To activate the trigger, select **Activate**.
11. To edit an existing trigger:
    a. Select **Thresholds**.
    b. From the Thresholds list, select an existing threshold.
    c. From the ellipses, select **Edit flow**.
    d. Change the settings.
    e. Select **Save**.

Install Application Insights

You can install the Application Insights application (sn_app_insights) 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.
- Application Insights requires the following plugin.

  Required ServiceNow plugins

  MetricBase (com.snc.clotho)

  You must request the MetricBase (com.snc.clotho) plugin from Now Support before you install Application Insights .

  > **Important:**
  > The MetricBase plugin must be activated by ServiceNow personnel via a Catalog Item request, in order to properly install and configure the ClothoDB. If you attempt to activate MetricBase on the instance yourself, the ClothoDB is not provisioned and Application Insights displays a 500: Internal Server Error.

  If you need this plugin activated in more than one instance, you must submit separate requests for each instance.

Role required: admin

About this task

The following items are installed with Application Insights :

- Roles
- Tables

For more information, see Components installed with Application Insights.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Application Insights application (sn_app_insights) 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.
4. Select **Install**.

Components installed with Application Insights

Several types of components are installed with the installation of the Application Insights application, including tables and user roles.

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.

Roles installed

Role title [name]	Description	Contains roles
Application Insights administrator [sn_app_insights.admin]	Users of the Application Insights application cannot: • Apply filters on graphs for events or transactions. • View syslog_transaction and sysauto records. • Create alerts.	None

Tables installed

Table	Description
Application Insights Aggregates [sn_app_insights_aggregates]	Temporarily stores aggregates for average, median, and 95th percentile of various statistics such as events processed and transactions.
Application Insights ECC Queue Aggregates [sn_app_insights_ecc_queue]	Temporarily stores ECC queue aggregates for average, median, and 95th percentile of various ECC-related statistics such as inbound and outbound processing.
Metric Trigger [sn_app_insights_metric_trigger]	Log of triggered metric threshold events.
Metric trigger log [sn_app_insights_metric_trigger_log]	Log of metric trigger events.
P1 Predict [sn_app_insights_p1_predict]	Stores the prediction model values to determine if a system is about to experience a P1 event.
Slow Patterns To Scripts Mapping [sn_app_insights_slow_patterns_to_scripts_mapping]	Stores mapping of slow patterns to scripts.
Script Referenced Slow Pattern [sn_app_insights_st_slow_pattern_referenced_script]	Temporarily holds the count of scripts referenced from the source of a slow pattern.
Slow patterns metrics [sn_app_insights_st_slow_patterns]	Temporarily stores metrics of slow patterns such as execution counts and times.
Slow Patterns To Scripts Mapping ST [sn_app_insights_st_slow_patterns_to_scripts_mapping]	Temporarily stores additional calculated relationship mapping of slow patterns to scripts with execution count and time.
Scheduled Jobs [sn_app_insights_st_sysauto]	Temporarily stores metrics of scheduled jobs including processing duration, error, and run counts.
Conditional threshold [sn_app_insights_trigger_condition]	Stores the configuration of metric thresholds, which are a logical combination of the individual triggers in the sn_app_insights_metric_trigger table.

Application Insights properties

These system properties control the behavior of the Application Insights application.

Application Insights properties

To change Application Insights properties, navigate to **System Properties** > **All Properties**.

These properties can be used to control the threshold trigger behavior and the p1 prediction model feature.

System Properties

Property	Type	Description
sn_app_insights.minutes_between_triggers	integer	How many minutes to wait before sending out a notification for an identical metric trigger. Default: **30**
sn_app_insights.p1_predict_factor.semaphores	integer	The normalizing factor for data for the semaphores metric in the p1 prediction model. Default: **21**
sn_app_insights.p1_predict_factor.sys_load	integer	The normalizing factor for data for the sys_load metric in the p1 prediction model. Default: **923**
sn_app_insights.p1_predict_max_cooldown	string	The number of consecutive "no p1 predicted" minutes required to exit the p1 alert state. Default: **5**
sn_app_insights.p1_predict_threshold	integer	The minimum confidence required for the p1 predict model to predict a p1 alert state. Default: **90**
sn_app_insights.default_load_duration_in_days	integer	The default time duration for metric graphs which include overview

System Properties (continued)

Property	Type	Description
		graphs, individual metric graphs, and drill-down slow pattern graphs. Default: **1**

Application Insights roles

Application Insights is installed with these roles.

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

Application Insights administrator [sn_app_insights.admin]

Visualize and monitor system health using Application Insights.

Contains Roles

Application Insights administrator [sn_app_insights.admin] contains the clotho_admin role.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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

Table rotation

Table rotation preserves instance performance and averts risk associated with querying growing data sets. This feature uses the sys_created_on field to separate data sets into multiple tables based on date.

Functionality

The administrator specifies the time parameter (duration) of the process and the number of tables (rotations) within. After the rotation writes the last table in a rotation, the rotation overwrites the first table in the rotation. Please contact ServiceNow Technical Support before applying table rotation to a custom table.

Examples:

- The query `Records created between 2015/12/10 08:49 and 2015/12/09 07/34 where topic=SystemCommand` is translated to a SQL query on a single table, because the clause on **sys_created_on** targets a single shard.
- The query `Records updated between 2015/12/10 08:49 and 2015/12/09 07/34 where topic=SystemCommand`, or without a date range, needs to target all shards and therefore is translated as a union query on all shards.

Advantages

- Allows deletion of old data without affecting current data (for example, to drop or truncate a table).
- Ensures tables only grow to a reasonable size.
- Reduces working set of data when date is known for query.

Disadvantages

- Queries that do not use the table rotation date (for example, by using the **sys_created_on** field), force an inefficient union query to query time ranges that span multiple tables and can be extremely slow if the number of sub-tables is large.

  To improve performance, it is recommended that the query includes a window of created dates.

- Users can't dot-walk to a rotated table.

You can use table rotation for sequentially-written tables or for insert-only tables. You cannot use table rotation for sys_import tables or tables that extend the Task [task] table.

Apply table rotation

You can apply a table rotation to preserve instance performance.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Table Rotations**.
2. Click **New**, or select the table rotation group to modify.
3. Fill in the fields:

   Table rotation group fields

Field	Description
Name	Enter the name of the table.
Type	Choose Rotation, Extension, or Shard. Rotation and Extension are time-based, Shard is based on a document ID. See Database rotation for more information.
Rotations	Sets the number of tables to maintain through the duration. Applies to `Rotation` type.
Duration	Set the overall duration for the action. Applies to `Rotation` and `Extension` type.
Clean base rotation	Sets the date to clean (truncate) the base table. Applies to `Rotation` type.

4. Click **Submit** or **Update**.

Note:
Deleting a rotation deletes the additional tables and all the data. Do not delete the rotation if you still need the data.

When you define a new rotation, a schedule creates and new data writes to one of the tables in the rotation group. The group includes the original table plus several additional tables.

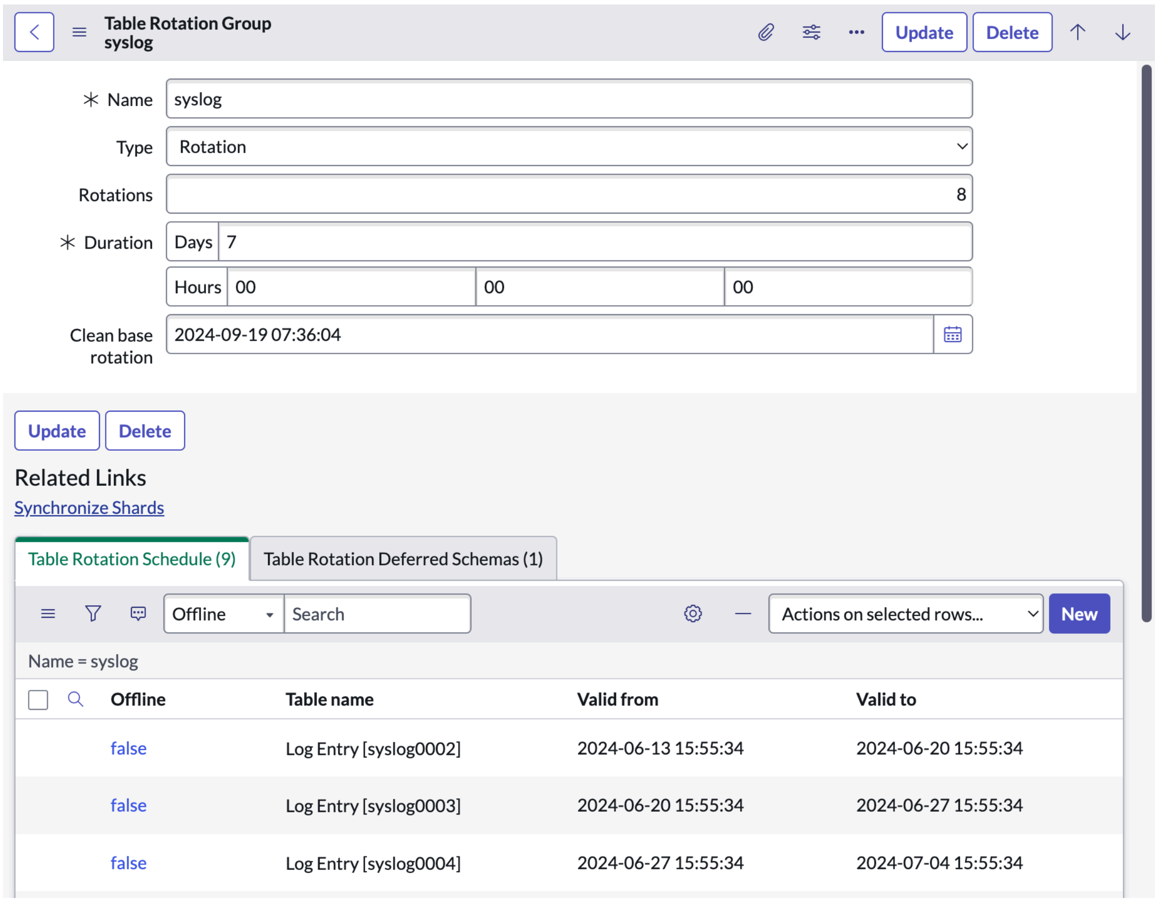

Table flattening

Table flattening stores a hierarchy of related tables as one table in a relational database.

Extension models

The system offers these extension models to store a table hierarchy on a relational database.

Available extension models

Extension model	Flattens tables?
Table per class	No
Table per hierarchy	Yes
Table per partition	Yes

Table per class

The Table per class extension model stores each table of the hierarchy in its own physical table on the relational database. Each physical table uses the table prefix of the source table

each stores a different class of records. An example of the Table per class extension model is the Asset [alm_asset] table, and its child tables: Hardware [alm_hardware], Consumable [alm_consumable], Facility [alm_facility], and Software License [alm_license]. The parent table of the hierarchy, Asset, stores a copy of every record in its descendant tables.

To find records in the Table per class extension model, the system queries records from multiple tables and joins the results. For example, when searching for hardware in a related facility, the system must join results from the Hardware, Facility, and Asset tables.

Table joins cause a performance bottleneck on relational databases. The more classes a query includes, the worse the query performance. Therefore any query for records from the top of the table hierarchy has the worst performance because it requires joining all descendant tables.

The system uses the Table per class extension model by default when creating tables. Most system tables also use the Table per class extension model as there is no performance benefit from flattening them.

Table per hierarchy

The Table per hierarchy extension model stores an entire table hierarchy in a single flat physical table on the relational database. The physical table is named after the parent table of the hierarchy, such as Task. The physical table contains all records of the table hierarchy and assigns a class name column value to each descendant table of the hierarchy. The system uses the name of the source table as the class name value. For example, Task records can have class names such as Change, Incident, or Problem.

To find records in a table hierarchy, the system queries the physical table and uses the class name column to constrain the results. Since such queries do not require joining results from multiple tables, the system provides better search performance.

The system uses the Table per hierarchy extension model for the Task table hierarchy on MySQL databases. Other tables use the Table per class extension model because there is no performance benefit to flattening them. To use Table per hierarchy on an Oracle database, contact Technical Support.

Table per partition

The Table per partition extension model stores an entire table hierarchy in a single flat logical table on the relational database. Each logical table can have multiple physical storage tables called partitions supporting it. Each partition optimizes the database resources available to a physical table such as the column count, index count, and row size. The system adds a partition whenever the logical table needs additional relational database resources.

Each logical table is named after the parent table of the hierarchy, and each supporting physical partition consists of the logical name plus a partition name. For example, the Base Configuration Item [cmdb] table starts as a logical table with no partitions. Suppose your hardware configuration items consume enough database resources that the system creates a partition called **cmdb$par1** to store them. Later, computer configuration items could consume enough database resources to warrant the system creating a second partition called **cmdb$par2** to store these records.

Within each logical table, the system assigns a class name column value to each descendant table of the hierarchy. For example, within the Base Configuration Item logical table there are records with class names for Application, Computer, and IP Router. The system also assigns a two-digit class path value to each descendant table of the hierarchy. The class path is based on the table location in the hierarchy. For example, the parent class Hardware might have a class path such as **/!!/!D** and the child class Computer might have a class path such as **/!!/!D/!!**.

To find records in the Table per partition extension model, the system queries the logical table and its partitions and uses the class path column to constrain the results. Since these queries do not require joining results from multiple tables, the system provides better search performance. In addition, the class path reduces the total number of records to search, which further improves search performance.

The system uses the Table per partition extension model for the Base Configuration Item [cmdb] table hierarchy on MySQL databases. To use Table per partition on an Oracle database, contact Technical Support.

View a table hierarchy and the extension model

The extension model used by a table is not immediately apparent. While a hierarchy can use a single physical table, the platform displays tables as if each logical table has a unique physical table. For example, when specifying a table for a workflow, you can select Change [change_request] or Incident [incident] even though the parent table, Task [task], uses a single physical table.

Before you begin

Role required: admin

About this task

Administrators can view the status of flattened table hierarchies, but cannot flatten additional hierarchies. You can configure the form to add the **Extension model** field if necessary.

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Select a table record.
3. Review the **Extension model** field value, which indicates whether the table hierarchy uses multiple unique tables or a single flat table.
   - **None**: Indicates that the table uses the table per class model. Defines a unique physical table per logical table in a hierarchy. For example, there are separate physical tables for the Asset [alm_asset] table, and its child tables: Hardware [alm_hardware], Consumable [alm_consumable], Facility [alm_facility], and Software License [alm_license].

     Note:
     If a table has an empty Extension Model field, it is not necessarily a Table Per Class table. If the table extends another table, it actually inherits the model from any parent tables it is extending, and only when none of the parents have a defined extension model is the table's extension model Table Per Class. This can be seen on many tables that extend the Task table. Incident is a logical table that extends Task and is physically located on Task, but by default has a blank Extension Model value.

     When a table doesn't extend anything at all, and has a blank Extension Model field, it is a Table Per Class table.

   - **Table per hierarchy**: Defines a single physical table per logical table hierarchy. For example, there is a single Task [task] table on the MySQL database that contains all the Task, Incident, Problem and Change records. This single physical table is represented as separate logical tables.

Add a module to test connection speed

A connection test can indicate the connection speed between your computer and your instance. A connection speed test (/connection_test.do) is available as a UI page.

Before you begin

Role required: admin

About this task

Connection Test page

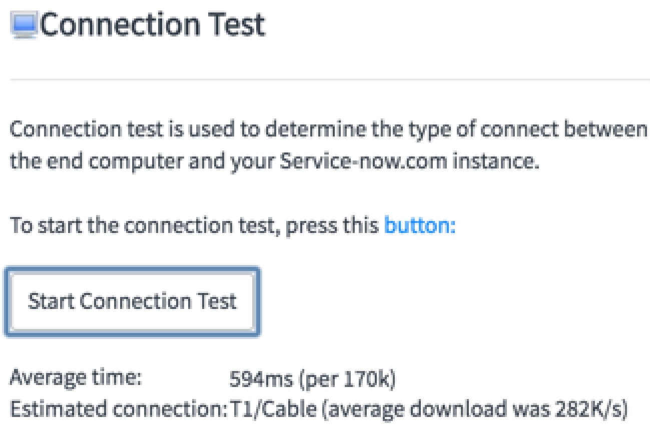

Procedure

1. Perform the appropriate action for your version of the UI:

Core UI	**a.** Navigate to **System Definition** > **Application Menus**. **b.** Open the application menu to which you want to add the module, for example, **System Diagnostics**.
UI15	Right-click an application menu, such as **System Diagnostics**, and select **Edit Application Menu** or click the pencil icon.

2. Click **New** in the list of modules.
3. Complete the Module form using the following values:
   - **Title**: Unique name such as `Connection Test`
   - **Link type**: **URL (from Arguments)**
   - **Arguments**: `connection_test.do`

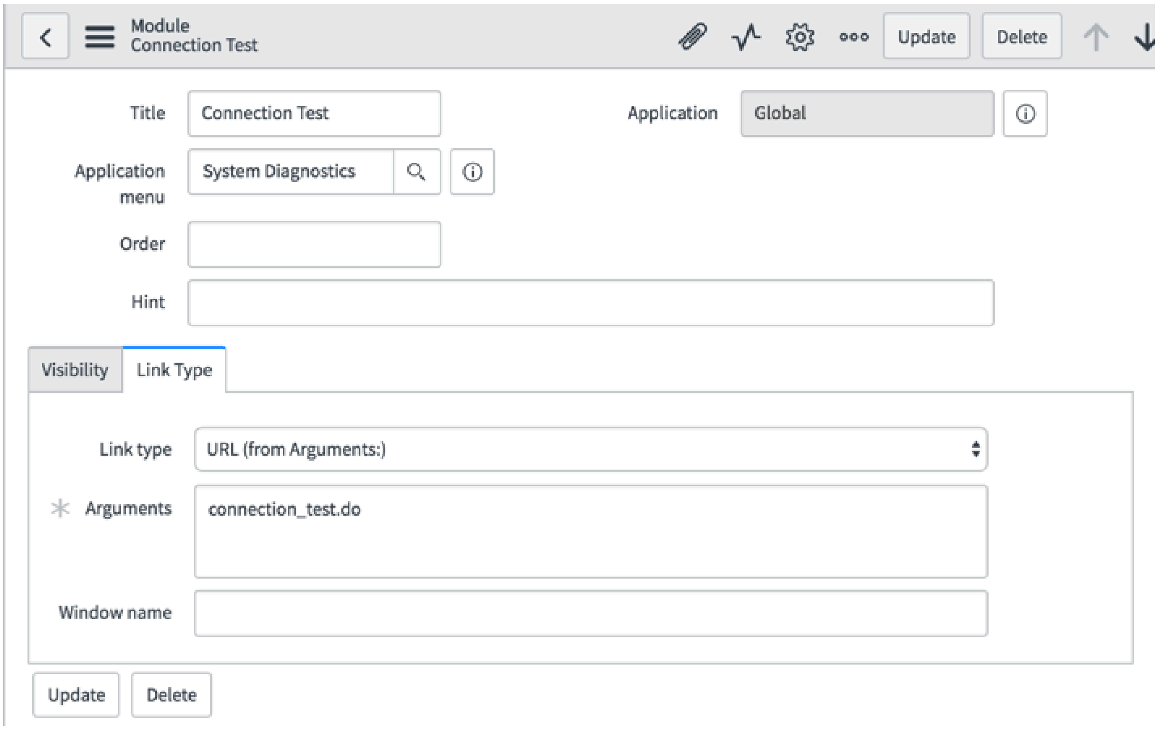

4. Click **Submit**.

Transaction quotas

Transaction quotas allow you to define a quota policy for different types of transactions. A transaction quota cancels any transaction in violation of the policy and notifies the user of the cancellation.

The Transaction Quotas plugin is active by default on all new and upgraded instances.

Transaction canceled

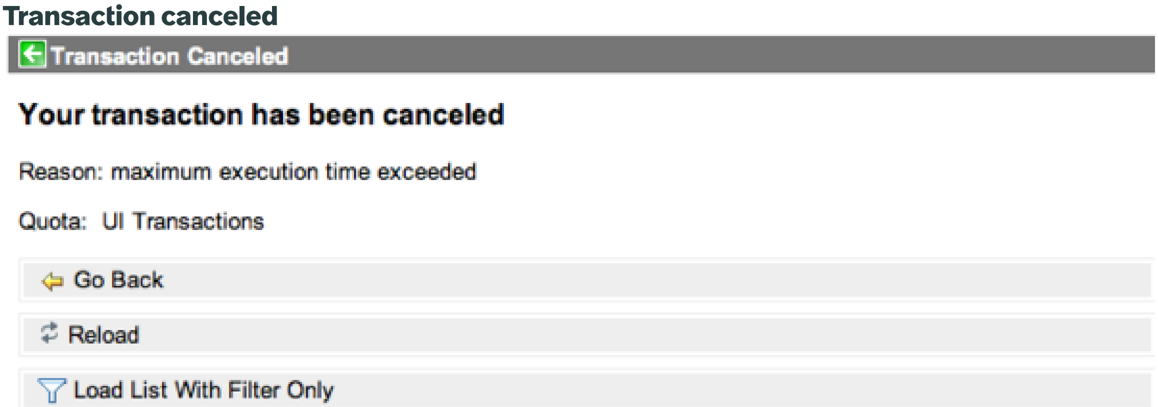

The transaction quota also writes the cancellation message to the log file as a warning.

Transaction canceled log

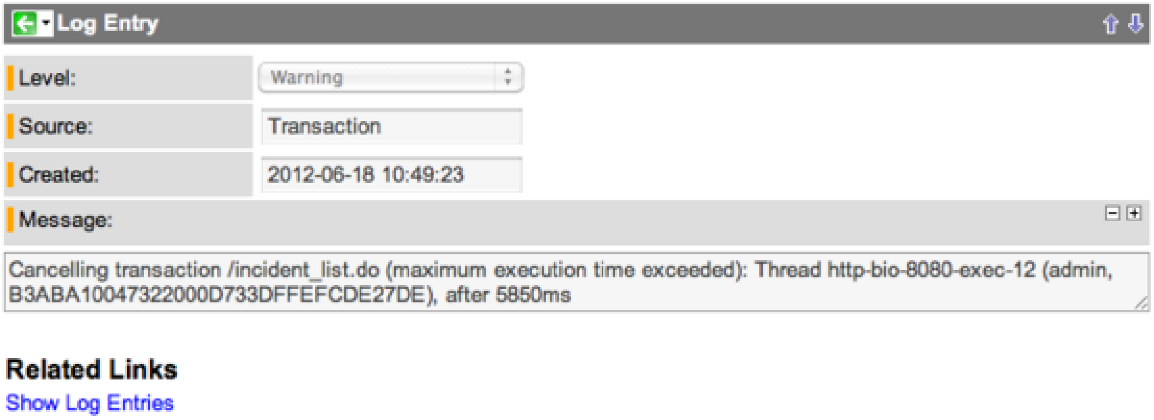

Typically, administrators set transaction quotas to prevent poorly performing queries and scripts from consuming system resources. This ensures that no transaction consumes enough resources to prevent other transactions from running. Administrators can also view cancellation log messages to identify transactions that might consume excessive resources.

How transaction quotas work

Transactions use the Quota Manager, which is a background thread that cancels transactions. The Quota Manager performs the following actions.

1. Obtains a list of active transactions, similar to the list under **User Administration** > **Active Transactions**.
2. Cycles through each transaction and checks all quota rules with conditions matching the transaction. The first occurrence that exceeds any of the quota limitations triggers a transaction cancellation. If a transaction changes or there is a new quota rule, the Quota Manager re-evaluates the transaction.

   Note:
   The **Order** field on a quota rule affects the order in which the quota rules are checked. The Quota Manager checks lower-order rules first (for example, order 80 before order 90), but ultimately checks all rules. This might have marginal performance implications, depending on the conditions of the rules involved.

3. Cancels the transaction if it has been running longer than the specified quota maximum.
4. Logs the running transactions.
5. Sleeps until the next heartbeat, which is controlled by the `glide.quota.manager.heartbeat` system property.

Transaction cancellation examples

Transactions can be canceled for more than one reason. Look for the following indicators.

- **maximum execution time exceeded**: This message appears when the glide.quota.manager threads cancels the transaction.
- **canceled by other transaction**: This message appears when the transaction was canceled by the session thread that initially issued it, and not by the glide.quota.manager thread.
- **canceled by user request**: This message appears when the user clicked the red X button to cancel the running transaction.

If you are looking for transactions that were canceled by glide.quota.manager, search for the **[Message] [Contains] [maximum execution time exceeded]** combination. For example:

```
Cancelling transaction /home.do - Default-thread-11.0
 (maximum execution time exceeded): Thread Default-thread-11.0
 (Default-thread-11.0, F530DD111B11111111FC031767DA158E), after
 30000ms
```

.

Add variable information to the cancellation message

You can use the following methods of the *jvar_transaction* variable to add information to the cancellation message.

Adding Variable Information to the Cancellation Message

Method	Description
getURL()	Returns the URL for the transaction.
getReason()	Returns the reason for page cancellation, localized to the user's language.
getQuotaId()	Returns the sys_id of the quota rule.
getQuotaName()	Returns the name of the quota rule.
getRunTime()	Returns the total run time for the page up to this point.
getType()	Returns the type of transaction, such as form, list, report, other.
getUser()	Returns the sys_id of the user.
getHomepage()	Returns the sys_id of the homepage.
getHomepageWidget()	Returns the name of the homepage widget.

Related topics

Modify the transaction cancellation page

Transaction quota properties

An administrator can add the following system properties to manage transaction quotas.

Configuring Transaction Quota Properties

Property	Description
glide.quota.manager.heartbeat	The number of seconds between the start of each Quota Manager heartbeat. This value determines how often the Quota Manager checks for transactions that exceed a quota and how often it writes status in the log file. Type: Integer Default value: 1 Location: Add to the [sys_properties] table
glide.quota.manager.minimum_transaction_time	The minimum number of seconds a transaction must run before the Quota

Configuring Transaction Quota Properties (continued)

Property	Description
	Manager matches it to a transaction quota. ServiceNow recommends setting this value to at least 1 second to avoid performance issues. For optimal performance, set this value to the value of your most restrictive quota. For example, if your most restrictive quota cancels transactions longer than 1 minute, set the minimum transaction time to 60 seconds. Type: Integer Default value: 1 Location: Add to the [sys_properties] table
glide.quota.manager.debug	Controls whether to display (true) or hide (false) additional debugging information related to the Quota Manager. Debugging information includes running transactions, canceled transactions, and what quotas are matched to transactions. Type: True | False Default value: True Location: Add to the [sys_properties] table

Related topics

Add a system property

Enable transaction quota debugging

You can enable the logging of transaction quota debugging information by enabling the system property.

Before you begin

Role required: admin

Procedure

1. Add the `glide.quota.manager.debug` system property, and set the value to true.
2. Navigate to **System Diagnostics** > **Session Debug** > **Debug Quotas**.
   The Debug Quota script runs to enable debugging.

What to do next

After setting the above system property to **true**, go to your system log to see messages from the Quota Manager.

Related topics

Add a system property

Default quota rules

Various transaction quota rules are available in the base system.

Fix Script Processor

Allows the fix script processor to run for four hours.

Presence

Cancels presence requests quickly when the system is busy.

REST Import Set API request timeout

Prevents inbound REST Import Set API transactions from running for longer than 60 seconds.

REST Table API request timeout

Prevents inbound REST Table API transactions from running for longer than 60 seconds.

AMB Transactions

Cancels AMB transactions lasting longer than the specified maximum duration. Applies to all AMB transaction types: message send and message receive.

Reference Completer

Stops the reference completer transaction after five seconds.

REST Aggregate API request timeout

Prevents inbound REST Aggregate API transactions from running for longer than 60 seconds.

Homepage Widgets

Prevents all homepage widgets from running longer than 30 seconds.

REST Attachment API request timeout

Prevents inbound REST Import Set API transactions from running for longer than 60 seconds.

UA Count Persistor Quota

Prevents UA count persistor scheduled jobs from running for more than one hour.

UI Transactions

Cancels UI transactions that are two seconds away from the 5-minute server disconnect. ServiceNow datacenter load balancers display an internal server error (HTTP Error 500) to users after five minutes. This quota rule returns a transaction canceled page instead of an error. This rule includes an exception to prevent the cancellation of background scripts.

PA Diagnostics All

Prevents any Diagnostics transaction that originates outside the scope of the Diagnostics application from running longer than 30 minutes. Diagnostics transactions refer to executing a diagnostic, opening a Diagnostics form, or opening a Diagnostics list.

Normal transaction activities

Before setting transaction quotas, navigate to **User Administration** > **Active Transactions** to review tthe normal transaction activities for your instance. You can view and terminate long-running transactions if necessary. Over time, you can determine the normal transaction load for your instance and set your quotas to match these norms.

Configure a transaction quota rule

Transaction quota rules allow you to define a quota policy for different types of transactions.

Before you begin

Role required: admin

About this task

Transaction quota rules allow you to specify the following conditions.

- The conditions under which the policy applies.
- The order in which transaction quotas apply.
- The maximum duration of a transaction before the quota manager cancels it.

Warning:
Setting transaction quotas too low can severely impact your users and prevent normal instance operations. Test your transaction quotas rigorously before implementing them in production.

Procedure

1. If necessary, activate the Transaction Quotas plugin.
2. Navigate to **System Definition** > **Quota Rules**.
3. Click **New**.
4. Complete the form.

Transaction quota rule fields

Field	Description
Name	A unique name to identify the rule.
Active	A check box that determines if this rule is active (selected).
Maximum Duration	The number of seconds a transaction has to complete before the quota cancels the transaction.
Maximum Business Rules	The number of business rules executions allowed.
Maximum Database Time	The total number of seconds for all SQL requests.
Maximum SQL Statement Time	The number of seconds a SQL statement can run.
Maximum Outbound Requests	The number of outbound HTTPs requests allowed.
Order	A number that represents the priority of the quota transaction in relation to other quotas. The transaction quota with the lowest timeout, the lowest order value, and matching conditions determines the applied quota policy.

Field	Description
Maximum Events	The number of sysevent inserts allowed.
Maximum Jobs	The number of sys_trigger inserts allowed.
Maximum SQL Queries	The number of SQL queries allowed.
Maximum Outbound Request Duration	The number of seconds for an outbound HTTP request.
Description	A description of the transaction quota rule.

5. Set a condition to specify when the transaction quota should take effect.
   The transaction quota rule condition builder displays conditions that are only applicable to transactions as listed in the following table.

Field	Description
URL	The URL of the page.
Thread Name	The name of the execution thread that is running the transaction. Foreground threads have names like http-bio-8080-exec-1. Background threads can have various names, and are useful where the URL is not sufficiently descriptive.
Foreground	Whether the transaction was completed in the foreground or background: ◦ **True**: a foreground transaction. ◦ **False**: a background transaction.
Type	The type of transaction: ◦ **List**: Any list transaction, such as incident_list.do. ◦ **Form**: All forms, UI pages, CMS pages, and so on. ◦ **XMLHttp**: Transactions that run through GlideAjax, where the URL is xmlhttp.do. ◦ **Report**: The page `sys_report_template.do`. ◦ **SOAP**: SOAP transactions. ◦ **Export**: When a list is exported in a format such as XML or Excel. ◦ **Scheduler**: When a scheduled job is performed. ◦ **Text Search**: The text search transaction or any of its related operations. ◦ **Other**: Any type that does not meet the qualities of another type. **Note:** Match the transaction quota **Type** to the transaction type listed in the Quota Manager log entry for the transaction.
User	The user performing the transaction.
Homepage	The homepage that is running. This condition is populated only if the URL is `home.do`.

Field	Description
Homepage Widget	The widget that is rendered in the transaction. When a homepage or non-responsive dashboard renders, each widget renders in parallel. When a responsive dashboard renders, each widget renders in sequence. When the quota is exceeded, users see the message: `Widget cancelled - maximum execution time exceeded`. Increase the quota value if you don't want to show this message.
Attributes	Miscellaneous attributes that are related to the transaction.

6. Click **Submit**.

Related topics

Optimize widget rendering time on responsive dashboards

View a canceled transaction

The Quota Manager logs each canceled transaction as a warning message in the system log.

Before you begin

Role required: admin

About this task

Use the following procedure to search for transaction warnings in the system log.

Procedure

1. Navigate to **All** > **System Logs** > **System Log** > **Warnings**.
2. Edit the filter to add the condition **[Message] [starts with] [Cancel]**.
3. Click **Run**.

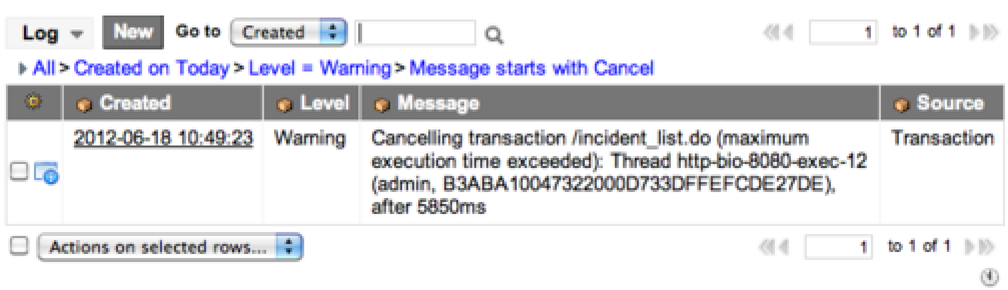

Example system log messages

An example of system log messages for transactions.

- At every heartbeat interval, which is one second by default, the Quota Manager prints the running transactions:

```
2012-02-13 12:34:08 (096) glide.quota.manager SYSTEM
 URL= /incident_list.do?
sysparm_userpref_module=b55fbec4c0a800090088e83d7ff500de&active
=true&sysparm_query=active=true^EQ,
THREAD= http-bio-8080-exec-3, FG= true, TYPE= 1, STATE= 2,
 USER= null, TIME= 8,807, MEM= 0, ATTRIBUTES= {}
```

- Every time the Quota Manager matches a quota to a transaction, it prints a message similar to the following example:

```
2012-02-13 13:25:31 (900) glide.quota.manager SYSTEM
 QuotaFinder: Assigning quota
"TEST PROBLEM FORM" with filter:
 type=form^urlLIKEsys_id=46fb9e31a9fe198101492060c2a4f8cb^EQ to
transaction:
 URL= /problem.do?sys_id=46fb9e31a9fe198101492060c2a4f8cb,
 THREAD= http-bio-8080-exec-1,
FG= true, TYPE= 1, STATE= 4, USER= null, TIME= 1,121, MEM= 0,
 ATTRIBUTES= {}
```

- Every time the Quota Manager cancels a transaction, it prints a message similar to the following example:

```
2012-02-13 13:25:33 (930) glide.quota.manager SYSTEM WARNING
 *** WARNING ***
Transaction: Cancelling transaction /problem.do (maximum
 execution time exceeded):
Thread http-bio-8080-exec-1
```

Modify the transaction cancellation page

The Quota Manager uses a UI page to control the contents of the transaction cancellation message.

Before you begin

Role required: admin

About this task

Knowledge of Apache Jelly is highly recommended when modifying the UI page. See Extensions to Jelly syntax for more information.

Procedure

1. Navigate to **All** > **System UI** > **UI Pages**.
2. Open the UI page with the name **transaction_canceled_quota**.
3. In the **HTML** field, add or modify the new cancellation message.
   The following table describes available variables.

Cancellation message variables

Method	Description
getURL()	Returns the URL for the transaction.
getReason()	Returns the reason for page cancellation, localized to the user's language.
getQuotaId()	Returns the sys_id of the quota rule.
getQuotaName()	Returns the name of the quota rule.
getRunTime()	Returns the total run time for the page up to this point.
getType()	Returns the type of transaction, such as form, list, report, other.
getUser()	Returns the sys_id of the user.

Method	Description
getHomepage()	Returns the sys_id of the homepage.
getHomepageWidget()	Returns the name of the homepage widget.

4. Click **Update**.

Application quotas

You can set a limit on the number of events or jobs that can run in an application's scope within a specified time.

Application quota rules set application quotas. Each scoped application can have one application quota rule. When an application exceeds the quota limit, all transactions running in the application scope are canceled. New transactions are canceled when they start until the next update period. These actions effectively block transactions from an application scope from running for the update period.

An application quota rule applies to transactions created by all instances of an application and all transactions created by methods of the application called by other applications.

If two users are running the same application, the application quota rule does not distinguish between transactions for each instance of the application. If, together, they violate the quota, all transactions in the application's scope are canceled.

If you check the **Log Only** option, transactions are not canceled by a quota violation. Instead, entries are added to the local host log that indicate the transactions are running under violation.

Transaction and application quota rules are evaluated separately, by defining an application quota rule, you simply introduce another restriction. A transaction is canceled if it violates a transaction quota rule, or if collectively with other application transactions, it violates its application quota rule.

You cannot define an application quota rule for the global application.

Application-quota property

An administrator can add a system property to specify how often application quota rules are evaluated.

Application quota property

Property	Description
glide.quota.manager.cluster.update.seconds	Indicates the application quota update period in seconds. If this value is less than the glide.quota.manager.heartbeat property, the glide.quota.manager.heartbeat value is used. Type: Number Default value: 300 (5 minutes) Location:Add to the System Properties `[sys_properties]` table

Configure an application-quota rule

You can configure an application-quota rule to limit the number of events or jobs that can run in an application's scope within a specified time.

Before you begin

Role required: admin

About this task

You must be in the application scope of the application for which you want to create a rule.

Procedure

1. Navigate to **All** > **System Definition** > **Application Quota Rules**.
2. Click **New**.
   The Application Quota Rule form is shown.
3. Fill in the fields.

   #### Configuring application quota rules

Field	Description
Name	A unique name to identify the rule.
Application	The scoped application's name. To change this field, use the application picker to change the application. This field cannot be Global.
Order	A number that represents the priority of the quota transaction in relation to other quotas.
Active	A check box that determines if this rule is active (selected).
Log only	When checked, transactions are not canceled by a quota violation, but entries are added to the local host log that indicate the transactions are violating the quota rule. An entry is added to the Transaction Cancellation Log where the log only field is set.
Maximum Events	The maximum number of events allowed during the update period.
Maximum Jobs	The maximum number of jobs allowed during the update period.
Condition	The conditions that determine when the transaction quota applies. The condition builder displays conditions that are only applicable to transactions. For example, Foreground is true.
Description	A description of the transaction quota.

4. Set a condition to specify when the application quota rule should take effect.
   The application quota rule condition builder displays conditions that are only applicable to transactions as listed in the following table.

Field	Description
Attributes	Miscellaneous attributes that are related to the transaction.
Created	When the transaction was created.

Field	Description
Created by	Who created the transaction.
Foreground	Whether the transaction was completed in the foreground or background: ◦ True: a foreground transaction. ◦ False: a background transaction.
Homepage	The homepage that is running. This value is populated only if the URL is `home.do`.
Homepage Widget	The homepage widget that is rendered in the transaction. When a homepage is rendered initially, each homepage widget is rendered in parallel.
Initial Scope Name	Scope where the method that created the transaction is defined. Setting this name to anything other than the selected application means that the rule is never used. The rule only looks at transactions of the selected application, and this value has the rule look for a different application, so the condition is never satisfied.
Scope Name	Scope of the application that called the method that created the transaction.
Sys ID	A sys_id of the transaction.
Thread Name	The name of the execution thread that is running the transaction. Foreground threads have names like `http-bio-8080-exec-1`. Background threads may have various names, and may be useful where the URL is not sufficiently descriptive.
Type	The type of transaction: ◦ **List**: Any list transaction, such as incident_list.do. ◦ **Form**: All forms, UI pages, CMS pages, and so on. ◦ **XMLHttp**: Transactions that run through GlideAjax, where the URL is xmlhttp.do. ◦ **Report**: The page `sys_report_template.do`. ◦ **SOAP**: SOAP transactions. ◦ **Export**: When a list is exported in a format such as XML, Excel. ◦ **Scheduler**: When a scheduled job is performed. ◦ **Text Search**: The text search transaction or any of its related operations. ◦ **Other**: Any type that does not meet the qualities of another type. **Note:** Match the transaction quota **Type** to the transaction type listed in the Quota Manager log entry for the transaction.
URL	The URL of the page.
Updated	When the transaction was updated.
Updated by	Who updated the transaction.
Updates	What was updated.

Field	Description
User	The user performing the transaction.

5. Click **Submit**.

Use table extension

This example shows you how to set the Table Extension parameter to a table rotation group.

Before you begin
Role required: admin

About this task

> **Note:**
> Deploy this plugin in partnership with a ServiceNow representative.

Administrators typically specify the time parameter (duration) of the process, but may alternatively specify the number of tables (rotations). After the last table in a rotation is written, new tables are added and old tables are archived. Using table extension, tables are never overwritten.

An advantage of table extension is to partition data across tables. It also allows you to archive data while ensuring that tables stay reasonably-sized. The working set of data is reduced when a date is known for the query.

The disadvantage is that table extension requires a union query when you query for a time range that spans multiple tables. Union queries are less efficient than queries against a single table.

A good practice is to use table extension when you have sequentially-written tables or insert-only type tables (there are exceptions to this parameter). Table extension is also useful in tables where data is needed for long periods of time.

The following example describes how to set table extension for a custom table.

Procedure

1. Navigate to **All** > **System Definition** > **Table Rotations**.
2. Click **New**.
3. Enter the following information.

Field	Action
Name	Enter the name of the table to apply the action.
Duration	Enter the overall duration for the action. 30 days is the duration in this example.
Type	Select **Extension**.

4. Click **Submit**.

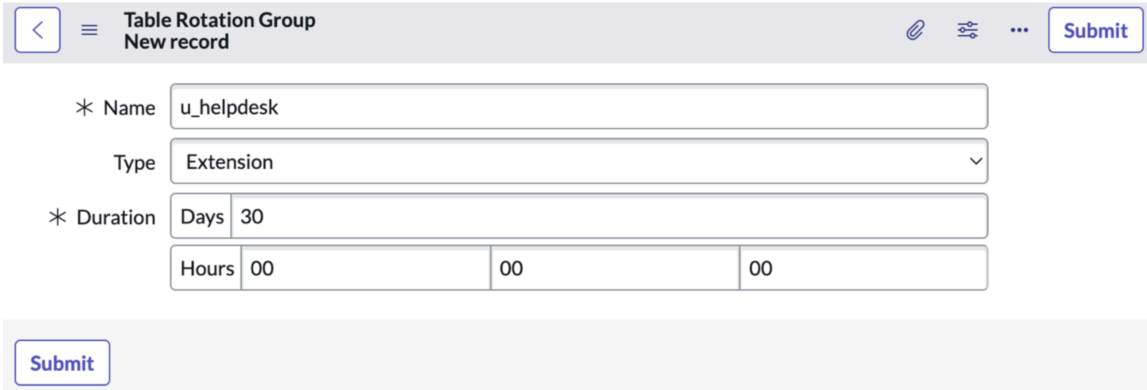

Operational toggles

You can control how much bandwidth an application uses on your system by mapping system run levels to operational toggles. An excessive bandwidth could impact your system performance and your users' experience.

Operational toggles overview

You can use system run levels to map descriptive names to operational toggle levels. The system run level is the level that you have defined the system to operate within, using your toggles and toggle levels. You may want to throttle the performance of your application when your instance is experiencing high volumes of traffic. An example of an operational toggle that you can create is a type-ahead search polling toggle that helps the user complete searches faster. You can set operations to enable users to perform a type-ahead search that is based on keystroke values.

Operational toggle levels

You can create different levels of your desired toggle. For example, you may want to perform a type-ahead search every keystroke or every fifth keystroke. For more information about toggle levels, see Create an operational toggle level.

Run level toggle mapping

You can associate the system run level to the toggle level with toggle mapping. Depending on your instance health, the API returns the appropriate toggle level value and exception roles. For more information about run level toggle mapping, see. Create run level toggle mapping.

Create an operational toggle

Create an operational toggle to maintain system run levels that throttle the application operation. Operational toggles enhance your users' experience.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Run Level** > **Operational Toggles**.
2. Click **New**.
3. On the form, fill in the fields.

Operational toggles form

Field	Description
Toggle Name	Unique name for the operational toggle.
Toggle ID	Toggle ID field that is inactive.
Short Description	Description of what function the toggle performs. **Note:** An example of this is a type-ahead search polling toggle that performs a type-ahead search every 10th keystroke.
Application	Scope that is to a specific application or is a global scope.

4. Click **Submit**.

Create an operational toggle level to define thresholds

Create operation toggle levels so that you can define the threshold limits that your application can function within.

Before you begin
Role required: admin

About this task

Operational toggle levels example

Toggle level	Description
Fully operational	The toggle performs every keystroke.
Slightly degraded	The toggle performs every fifth keystroke.
Moderately degraded	The toggle performs every 10th keystroke.
Severely degraded	The toggle is inactive.

Procedure

1. Navigate to **All** > **System Run Level** > **Operational Toggle Levels**.
2. Click **New**.
3. On the form, fill in the fields.

Operational toggle level form

Field	Description
Toggle	Toggle that you are creating levels for.

Field	Description
	Note: An example of a toggle is a type-ahead search poll.
Short Description	Function of the toggle level that you create. **Note:** An example of a level is a type-ahead search poll every 10th keystroke.
Application	Scope that is to a specific application or is a global scope.
Value	Behavior of the toggle level.

4. Click **Submit**.

Create run level toggle mapping

Create run level toggle mapping so that you can associate the system run level to the toggle level.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **System Run Level** > **Run Level Toggle Mapping**.
2. Click **New**.
3. On the form, fill in the fields.

Run level toggle mapping form

Field	Description
System Run Level	Toggle level to associate the system run level with.
Toggle	Behavior of the toggle.
Toggle Level	Degree that the toggle operates.
Application	Scope of the toggle level to specific application or global.
Value	Behavior of the run level. **Note:** An example of a value is as follows: 1 or every keystroke that the system performs a type-ahead search.
Exception Roles	Exception roles that the restrictions don't apply to.

Field	Description
	Note: An example of these exception roles is as follows: ◦ activity_admin ◦ activity_creator ◦ admin

4. Click **Submit**.

Using ServiceNow AI Platform upgrade tools

Get the latest in ServiceNow® products and services on the ServiceNow AI Platform®. The system helps you plan for updates, and monitor and work with your applications before, during, and after upgrading.

Upgrade Center	Upgrades and Conversions
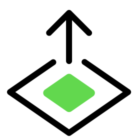	
Plan and manage all your upgrades in one place, past and present, with data visualizations, upgrade histories, and more.	Take advantage of tools to bring you the latest releases for your instance on the ServiceNow AI Platform

Upgrade Center

Plan for all your upgrades on the ServiceNow AI Platform, and get clear oversight of the process at every step. The Upgrade Center includes multiple options for each aspect of your upgrading plan. Explore various updates available to your infrastructure with Upgrade Preview, including detailed insights before actually upgrading. Visual task boards (VTB) in the Upgrade Center give you all the necessary info to manage your workflow post-upgrade, and past updates are detailed with the Upgrade History module. Tie it all together with the Upgrade Monitor to plan, schedule, observe, and debug upgrades.

See Upgrade Center.

Upgrades and Conversions

Keep your instance running smoothly during and after upgrading. ServiceNow AI Platform® Upgrades and Conversions tools make the process easy, with information on upgrades, patches, and releases, as well as options for backups, post-upgrade instance testing, and migration from ServiceNow® Express List to the Service Management platform.

See Upgrades and conversions.

Upgrade Center

Use ServiceNow® Upgrade Center to plan and manage your upgrades. Use the Upgrade Preview module to explore different release versions available to you and gain insights about the experience of an upgrade without actually upgrading your instance. You can also monitor the status of any ongoing upgrade and view the summary of your latest upgrade with the new Upgrade Monitor module. The Upgrade Center also features an Upgrade visual task board (VTB) to manage post-upgrade skips. A separate Upgrade History module can be used to view all your past upgrades.

Get started

Explore Learn about Upgrade Center	**Preview** Preview your upgrade	**Operate** Schedule and monitor your upgrades
Use Review your upgrade history	**Reference** Learn about miscellaneous information about Upgrade Center	

Exploring Upgrade Center

Get an overview of Upgrade Center so you can plan your upgrade ahead.

Key features

One-stop experience

You can see all the details that are related to an upgrade.

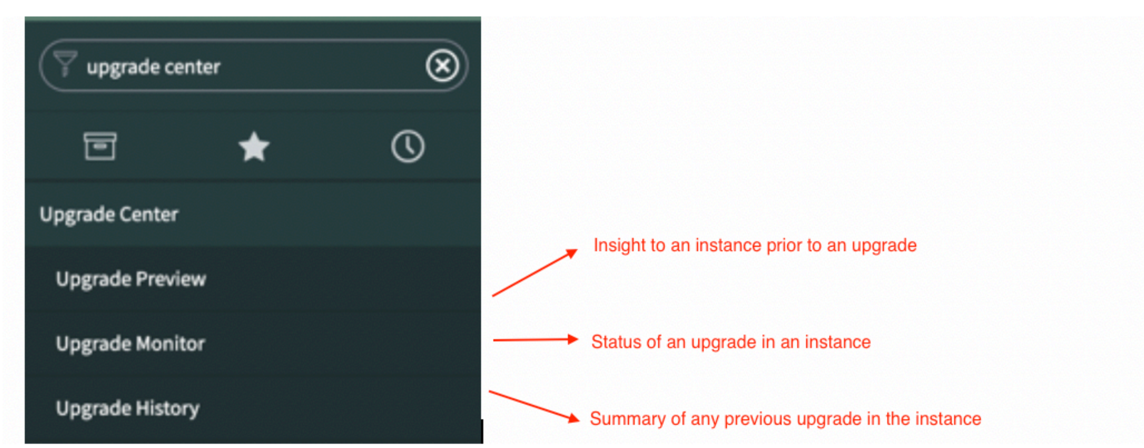

Preview your upgrade

Use the Upgrade Preview module to have an unprecedented insight to an instance prior to an actual or a scheduled upgrade.

You can see the estimated time required to complete an upgrade. The estimated time is calculated from previous upgrades.

Note:
If this is your first upgrade, **Estimated upgrade duration** doesn't show up.

Reduce post-upgrade review time by previewing the records that have been customized and are different from the upgrade version being previewed and might get skipped due to your customizations in the instance. You can then prioritize your work depending on the records listed in the skip list.

Monitor an upgrade

Use the Upgrade Monitor module to schedule and view the status of an ongoing upgrade. You can also reschedule the upgrade by clicking **Reschedule upgrade**. To check for an immediately available upgrade, select **Check for upgrade**.

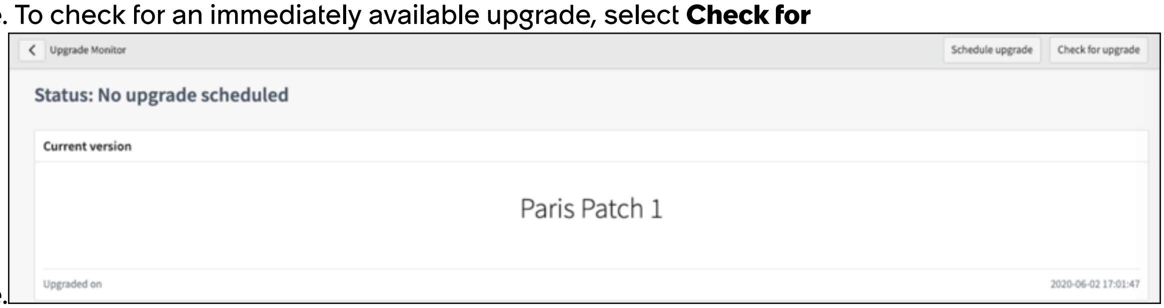

If an upgrade is in progress, the user is automatically directed to the upgrade in progress page.

Review previous upgrade detail

Use the Upgrade History module to view the complete report of any upgrade made to an instance.

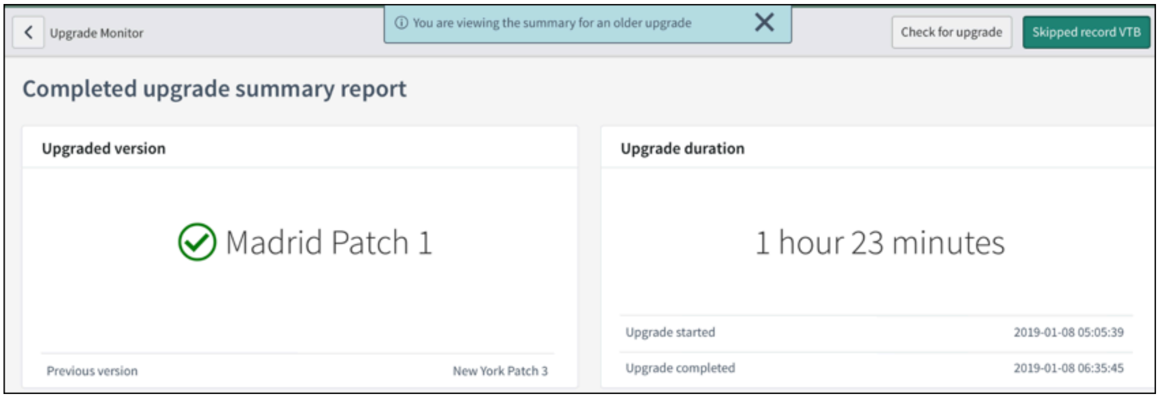

You can also view the resolution status of any previous upgrade with skipped records using the visual task board (VTB) view.

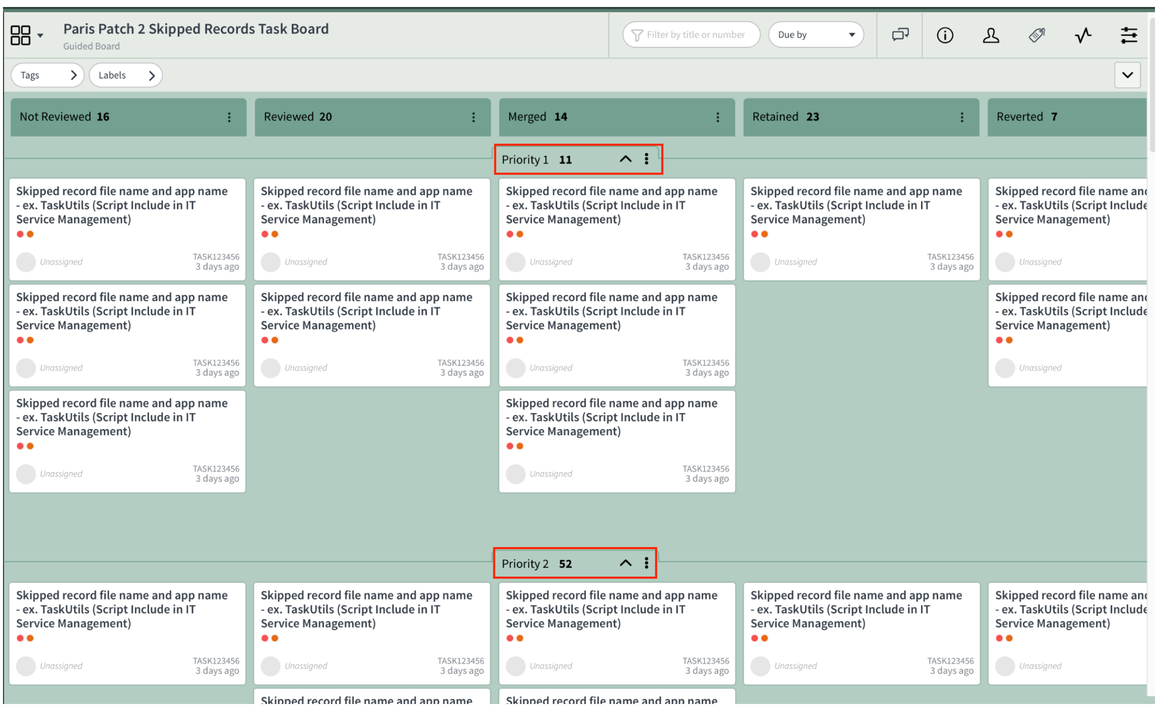

Upgrade Preview module

The Upgrade Preview module enables you to have an unprecedented insight to an instance prior to an actual upgrade. You can explore and preview upgrades to different ServiceNow release versions and see how your instance might be impacted with your current configurations. The Upgrade Preview utility helps you to plan, schedule, and prepare for an upgrade.

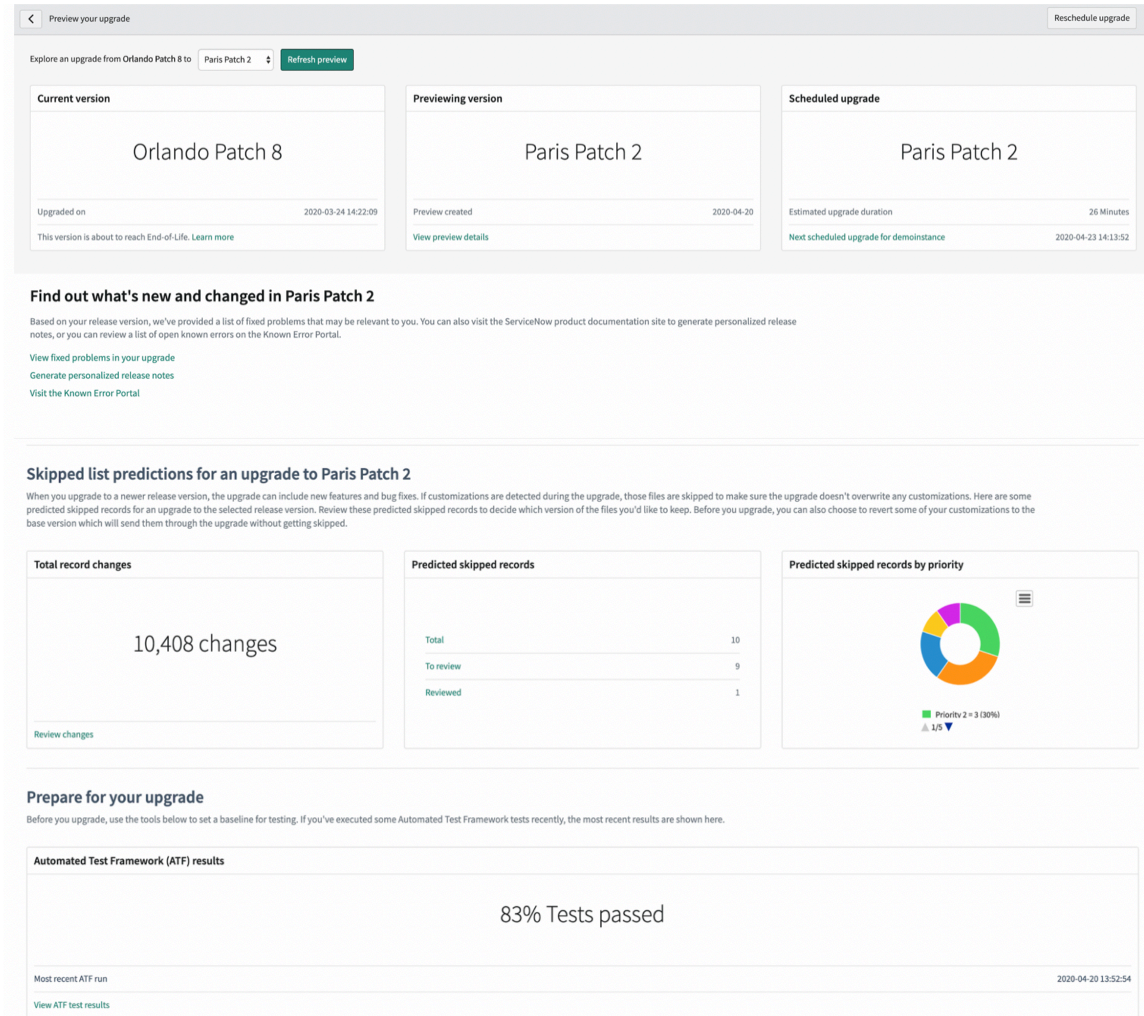

Note:
Depending on the eligibility of your instance, the list of available target versions for preview varies. Only versions that are allowed for a particular instance to be upgraded to, show up on the list. If your instance is not eligible to be upgraded to any version, the drop-down menu is empty.

The following cards show up with more information.

Card names	Description
Current version	Name of the current version. • Upgraded on: Date on which the instance got upgraded to the current version • End of life: Link to KB0610454 for more information about end of life
Previewing version	Information about the previewing version shows up. If you haven't previewed any version, you are expected to select a version for preview and click **Go**. If you have already previewed a version, the following information shows up.

Card names	Description
	• Preview created: Date on which the preview was created • View preview details: Link to the Upgrade Preview form **Note:** Click **Refresh preview** if you want to regenerate a new preview of the same target version. Although clicking **Refresh preview** regenerates the skipped list, the reverted records won't get skipped. If you have reviewed and retained some records, they get replaced as soon as you click **Refresh preview**.
Scheduled upgrade	Name of the next upgrade that has been scheduled to start. If there is no currently scheduled upgrade, the card shows No upgrade scheduled. **Note:** This card can have different messages depending on the schedule of the upgrade. If there is an upgrade scheduled, the following information shows up. • Estimated upgrade duration: Estimated duration needed to complete the scheduled upgrade **Note:** If this is your first upgrade, there will be no information about estimated upgrade duration. • Next scheduled upgrade: Date and time on which the next upgrade has been scheduled to start.
Find out what's new and changed	Links to view the new and changed features in the current upgrade version. The following three links show up when you click **Go** to preview the upgrade version. • Problem fixes: Fixes that have been made since the last upgrade version • Personalized release notes: Release notes summary for the previewed version • Known error articles: Errors that have been identified but not yet resolved

Card names	Description
	Note: These links don't show up if you haven't previewed any upgrades.
Skipped list prediction **Note:** If you are using Upgrade Plan, the card name changes to Skipped list prediction with plans.	Information about the predicted skipped records. See Preview predicted changes for more details. • Total record changes: Total number of records that are predicted to change when the upgrade occurs. Total record changes also include possible predicted skipped files known as Predicted skipped records. Review changes: Link to the list of records that have changed and can be reviewed • Predicted skipped records: Total number of records that have been predicted to be skipped. ◦ Total: Total predicted skipped records ◦ To review: Predicted skipped records to be reviewed ◦ Reviewed: Predicted skipped records that have been reviewed ◦ Resolved: Predicted skipped records that have been resolved by implementing the rules **Note:** This entry shows up on the card only when rules are executed on the skipped records. You will also see a link to create a skipped record rule if either there are no existing rules or the existing rules aren't executed. **Note:** Skipped record rules and upgrade plans can't be implemented together. See Execute a skipped record rule automatically for more information. • Predicted skipped records by priority: Pie chart to represent the predicted skipped records by priority. **Note:** Click the pie chart to see the list of predicted skipped records. ServiceNow prioritizes the skipped records based on the importance of the file types. The prioritization is done as follows:

<table>
<tr><th>Card names</th><th>Description</th></tr>
<tr><td></td><td>

Priority 1 (highest priority): UI pages, UI macros, and more
Priority 2: Business Rules, Security ACLs, and more
Priority 3: Reports and more
Priority 4: Form Sections, Choice Sets, and more
Priority 5 (lowest priority): everything else

</td></tr>
<tr><td>Predicted skipped records by product</td><td>
Records that have been predicted to be skipped and are sorted as per their product families.


Note:
The products sorted under the Other category don't have any specific product family.
</td></tr>
<tr><td>Automated Test Framework (ATF) results</td><td>
Percentage of passing ATF tests that ran in the last 30 days.

Most recent ATF run: Date and time on which the recent ATF tests ran
View ATF results: Link to show more information about the recent ATF results

Note:
Only the tests which are finished and have passed are considered for the ATF results. If one test runs more than once, only the recent execution is considered in the results.
</td></tr>
</table>

Upgrade Monitor module

Schedule an upgrade and monitor the status of an ongoing upgrade using the Upgrade Monitor module. You can also view the upgrade summary and the list of records causing conflicts in your instance once the upgrade is complete.

Access Upgrade Monitor

If you log in with the admin role while an upgrade is underway, the system automatically displays the Upgrade Progress screen. If no upgrade is in progress, you can navigate to **All** > **Upgrade Center** > **Upgrade Monitor**.

How the Upgrade Monitor fits into the upgrade process

Note:
For detailed information about the upgrade process, see Upgrade your instance.

The Upgrade Monitor concerns only part of the larger upgrade process:

1. Clone the production instance to a test instance and a non-production instance.
2. Apply the upgrade to the non-production instance.
3. On the upgraded non-production instance, process the skipped list.
4. Test the non-production instance to confirm that the instance still works and performs adequately. Compare to benchmark data from pre-upgrade production instance.
5. Apply the upgrade to the test instance. Import the update sets created on the non-production instance when you processed the skipped list. Repeat the testing to make sure that the process is working.
6. Apply the upgrade to the production instance. Import the update sets created on the non-production instance when you processed the skipped list. Test to confirm that the instance works and performs adequately.

Within this larger process, the Upgrade Monitor helps you upgrade individual instances:

- During the upgrade, it shows where in the process the system is
- After the upgrade, it reports what the upgrade did and for how long
- As you upgrade the first non-production instance, it helps you resolve conflicts between customizations and changes that are part of the upgrade
- On non-production instances, it provides information that can help you estimate how long the upgrade might take on the production instance

Note:
Debug Upgrade provides detailed debugging output for transactions containing artifacts affected by the most recent upgrade, and is designed to assist in upgrade error resolution.

Monitoring an individual instance as it upgrades

While the upgrade is in progress the Upgrade Progress shows what the upgrade process has done, what it is doing, and what remains to be done.

When the upgrade completes, the system displays the Upgrade Summary Report. The Upgrade Summary Report provides information about conflicts between customizations versus changes in the upgrade and provides a link to reconcile these conflicts. For information about understanding and resolving these conflicts, see Process the skipped records list.

When you upgrade a non-production instance, the Upgrade Summary Report can help you estimate how long the same upgrade might take on a production instance. For details about the elements on this report and how to use this information, see Upgrade Details form.

Resolving conflicts

To prevent losing customizations, the system skips upgrading records you have customized and provides you with a list of these skipped records.

As you upgrade your first non-production instance, go through the Skipped Changes to Review related list and resolve these conflicts. The system records the changes you make during this process in update sets. See Review skipped records using related lists, for more information.

You do not need to reconcile the skipped list on any instances you later upgrade. Instead, you can apply the upgrade then import the update sets containing your changes.

For details on reconciling conflicts, see Process the skipped records list.

Factors affecting upgrade duration

Various factors affect how long the system takes to perform an upgrade. The Upgrade Monitor on non-production instance can help you understand those factors and estimate how long the upgrade to your production instance takes.

Many factors can affect the duration of the upgrade process:

- The number of records in the database
- The number of customizations in the database
- The number of nodes in the instance
- The size of tables in the instance that require a schema-change in the instance
- The number of fix scripts required and the size of the tables those fix scripts manipulate

Upgrading a non-production instance can help you estimate how long the upgrade takes on production, but differences between the instances can significantly affect the duration:

- When you cloned the production instance to the non-production instance, you may have clone-excluded some tables. This reduces the size of the database and makes the non-production instance upgrade faster than the production instance.
- The production instance may have more memory and processing power.
- The production instance may have more nodes than the non-production instance.

After upgrading the first non-production instance, examine the Upgrade Summary Report for data to help estimate the impact of these factors.

Upgrade Monitor schedule states

When an upgrade is scheduled, the Upgrade Monitor displays information about the next scheduled upgrade. It displays the date and time when the next scheduled upgrade will start. If there is no upgrade scheduled, the Upgrade Monitor displays **Status: No upgrade detected**.

No upgrade scheduled

There is no upgrade that has been scheduled currently.

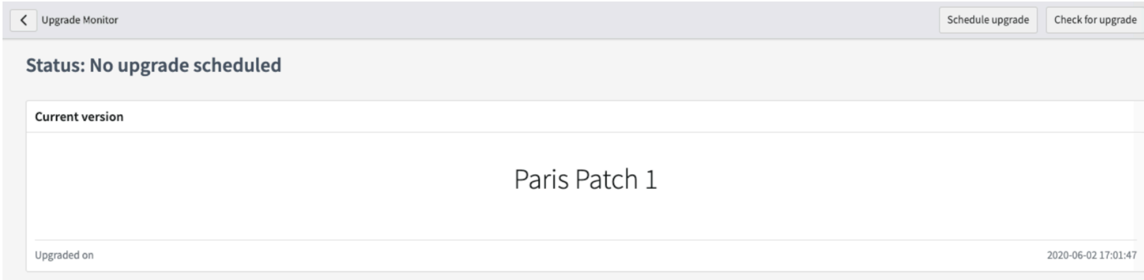

You can also schedule an upgrade by selecting **Schedule upgrade**. To check for an immediately available upgrade, select **Check for upgrade**. It also shows the current version of your upgrade and the date it was upgraded on.

Upgrade scheduled

There is an upgrade that has been scheduled to start at a given

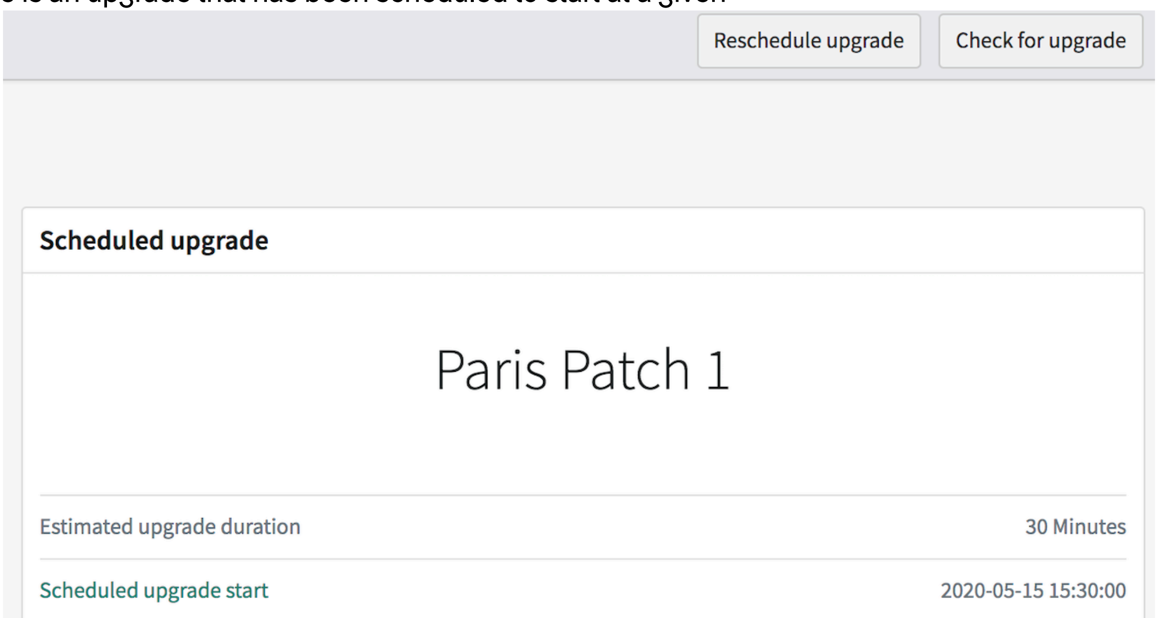

time.

You can also reschedule the upgrade by selecting **Reschedule upgrade**. To check for an immediately available upgrade, select **Check for upgrade**. By selecting either **Reschedule upgrade** or **Schedule upgrade**, you will be directed to HI Upgrade Wizard.

Upgrade monitor with issue detected

If one or both of the triggers for upgrading the system ('Check distribution for possible upgrade' and 'Check database for possible upgrade') have been customized or are missing, the Upgrade Monitor displays a warning and provides a button for resolving the issues.

Note:
If your instance is self-hosted (not hosted by ServiceNow) this message may not necessarily indicate a problem. If you have customized or disabled the upgrade job and want to keep that customization or disabled state, do not select the button to fix the upgrade issue.

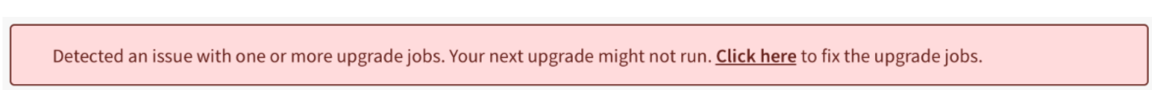

To resolve the issues with the upgrade jobs, select the link in the message. This action reverts both upgrade triggers ('Check distribution for possible upgrade' and 'Check database for possible upgrade') to their base versions.

Note:
'Upgrade' job has been renamed to 'Check distribution for possible upgrade' starting Paris. 'Check Upgrade Script' job has been renamed to 'Check database for possible upgrade' starting Paris.

Upgrade Plan overview

Accelerate your upgrades by packaging your skipped records and customizations into the relevant application container (global apps or app customizations). Upgrade Plans help you define applications and target versions to be installed in your instance. It automates the installation of these applications during upgrades, giving you a seamless upgrade experience.

Navigate to **All** > **Upgrade Center** > **Upgrade Plan**.

You will need the following instances to use Upgrade Plan.

- Builder instance: Build your upgrade plan
- Consumer instance: Implement your created upgrade plan

Note:
Ensure that you have upgraded your instance to the latest available version to build the upgrade plan. It is recommended to configure your dev instance as your builder instance. You are required to install the upgrade plan before upgrading your consumer instance. See Upgrade Plan Properties for more details.

When you upgrade an instance, resolutions are skipped, and customizations are often loaded post upgrade. This causes temporary breakdown of features until fully loaded. Tasks like committing update sets, installing new plugins and applications, and multiple update are also time consuming. Use the Upgrade Plan feature to automate these post-upgrade tasks by tracking your actions and replay the steps on all the required instances. You don’t have to manually apply post-upgrade tasks which helps in reducing downtime.

Note:
After upgrading to the latest version, you will see some skipped records. Some of those records will already be marked as reviewed and some of them will need to be reviewed.

Advantages of Upgrade Plan over Update Sets

You can achieve the following using Upgrade Plans:

- Along with the skipped record resolution, you can also track app and plugin installation

  Note:
  You won’t be allowed to choose the skipped records that are captured by the upgrade plan. Upgrade plan captures all the skipped records regardless of whether they are reviewed or not and modified or not. If the customizations are coming from different instances, then the skipped records are required to be reviewed.

- Optimizing the table alters using batch bootstrapping
- Manage your customizations using the App Repository

Note:
It is recommended to build an Upgrade Plan for each instance upgrade. When you build an upgrade plan on a builder instance, it gets created to the exactly same as the builder instance version (including patches and hot fixes). You can't use the previously created Upgrade Plan for the new instance upgrade. In case of consumer instances, when you install the upgrade plan, its version should match exactly the consumer instance version.

Persona

If you are using the app repository for active application development, use Upgrade Plan to accelerate your upgrades.

Note:
If you are currently using the Update Sets and want to catapult your upgrade process, use the Upgrade Plan feature.

Design considerations

The following are the important considerations while working with Upgrade Plan:

- Each new instance upgrade requires its own upgrade plan. It can't be shared across upgrades.
- Only one builder instance is supported to build the Upgrade Plan
- Upgrade Plan can't be uninstalled on a consumer instance. You can rollback the entire upgrade but not partially.
- The scope of the files moved to **Global Customizations - Upgrade Plan** application by upgrade plan is still global
- During an upgrade, only the upgrade plan items with **State=Ready** and **Active=true** are installed on the consumer instances. The rest of the items are skipped.
- Upgrade sets can't be included in the Upgrade Plan
- Maint only plugins are not allowed in Upgrade Plan
- Configure your instance as a builder instance before installing new applications and plugins or during skip resolutions. Otherwise, the actions are not captured by the Upgrade Plan.
- You can't view the list of customizations in the Upgrade Plan items view. You can navigate to the respective tables to ensure if the customization has been captured.

See KB1271313 for more information.

Upgrade Plan background operations

During the building of the upgrade plan, the following operations are done at the background:

- Skipped records are packaged and uploaded to the App Repository in the form of Global Application and App customization. The following are the 3 types of apps created by Upgrade Plan on your instance and App Repo.

  Note:
  After the packaging is done, you can choose the items from the Upgrade Plan. An Upgrade Plan works at the scope level, so it captures everything and publishes it to the repository.

  - The global records are packaged into the global customization upgrade plan app and published to the app repo
  - The scoped skipped records are packaged into the respective app customization packages
  - When the upgrade plan is published, a global application is created, for example, **Upgrade Plan - release name**
- Existing ServiceNow features like Global Application and App customization are used as application containers that can be installed in all the required instances. See Global application file management and Application scope for more information.
- If you install any applications or plugins, they are also captured in the Upgrade plan. But, since they are application life cycle items, they are never pushed to the App Repository.

During the consumption of the upgrade plan, the following operations are done at the background:

- At first the upgrade plan is validated and then the source from the app repo is downloaded
- Once the source is downloaded, the app is moved to the ready state

  Note:
  Only the upgrade plan items with **State=Ready** and **Active=true** are installed on the consumer instances. The rest of the items are skipped.

- Auto-generation of the preview for the upgrade plan

Building your Upgrade Plan

Build your upgrade plan to experience a seamless and an accelerated upgrade by packaging the applications in an upgrade plan.

Before you begin

Note:
Ensure that you have upgraded your instance to the latest available version to build the upgrade plan. It is recommended to build an Upgrade Plan for each instance upgrade. The target version of the Upgrade Plan must be exactly the same as the instance version (including patches and hot fixes). You can't use the previously created Upgrade Plan for the new instance upgrade.

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Center** > **Administration** > **Properties** > **Upgrade Plan Properties**.

   Note:
   It is recommended to set up your dev instance as a builder instance.

2. Navigate to **All** > **Upgrade Center** > **Upgrade Plan**.

   Note:
   If you are on the Upgrade Monitor page, click **Manage Upgrade Plan**.

   The Upgrade Plan list view shows up.
3. Click **Build** to start the upgrade plan build.

Upgrade Plan takes up all the skipped records and publish it on to the app repository. See Upgrade Plan background operations section in Upgrade Plan overview for more information.

Note:
Ensure all the applications in the instance are in ready state to be published on to the app repository.

You are allowed to have only one Upgrade Plan at any point of time.
A final confirmation message shows up to start building the upgrade plan. If any application has changes in-progress, the Upgrade Plan must be built only after the changes are done. warning message shows up to complete the review of the remaining skipped records.

4. Click **Build Upgrade Plan**.
   The upgrade plan starts to build.
5. Click **Go to Upgrade Plan** once the build is done.

   Note:
   Click **View Details** if you want to view the execution details of the Upgrade Plan. If you see a warning, it is recommended to review the execution details to know the errors.

   The upgrade plan opens and a list of upgrade plan items shows up.
6. **Optional:** Click **New** if you want to manually add any upgrade plan items.

   Note:
   The applications or plugins that you want to add manually to the upgrade plan must be already installed and active on the instance.

   If you want to skip any of the upgrade plan items from being a part of the Upgrade Plan, set the selected upgrade plan item to false.
7. **Optional:** Click **Refresh** to capture all the recent changes to the Upgrade Plan.

   Note:
   This step is applicable only if you have any update on the existing application or plugin in the Upgrade Plan. If the upgrade plan has already been published, you are required to republish it after every refresh.

8. Click **Publish** to publish the Upgrade Plan.
   See Upgrade Plan background operations section in Upgrade Plan overview for more information.

   Note:
   Install all the required applications and plugins and resolve all the skipped records before publishing the upgrade plan. If you want to delete the Upgrade Plan, click **Delete**. The **Delete** option is available only for the Upgrade Plans that are not yet published.

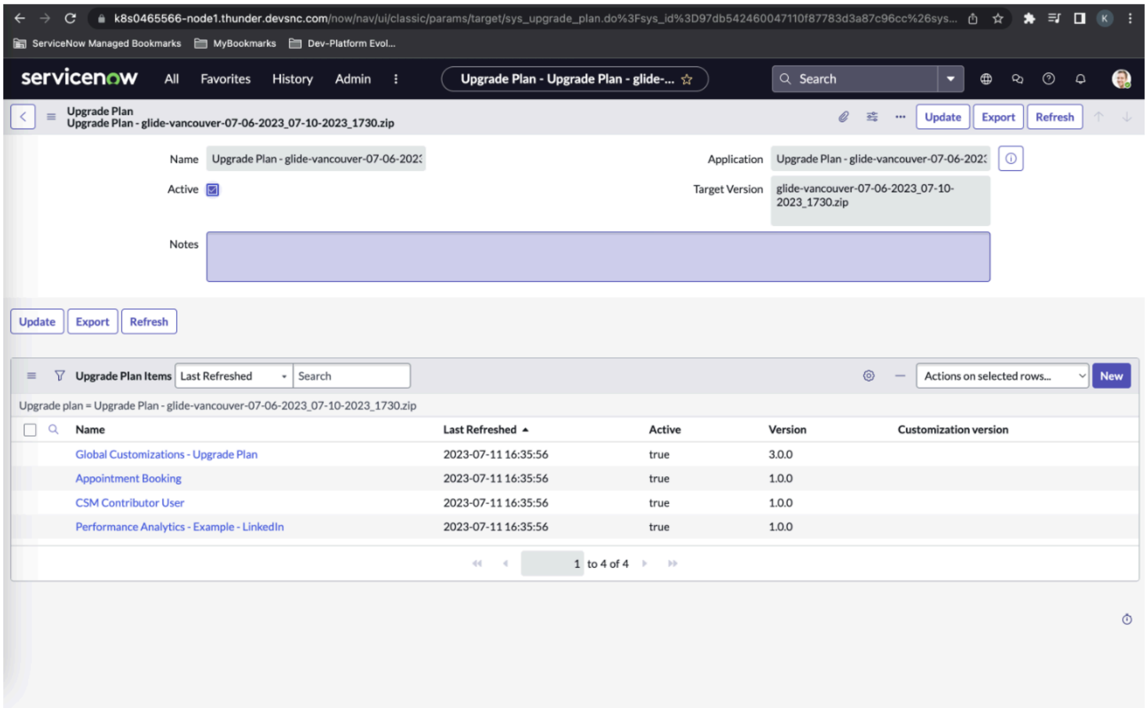

The upgrade plan is now ready to be installed in the consumer instance.

Refreshing your Upgrade Plan

Refresh your upgrade plan to package all the recently installed plugins and applications into your upgrade plan.

Before you begin

Role required: admin

Procedure

1. Go to the applications and plugins you need and select **Activate** to install them.

   Note:
   This can be done even after the Upgrade Plan is published.

   The plugin activation starts and installs successfully.

2. Go to the upgrade plan and refresh it by selecting

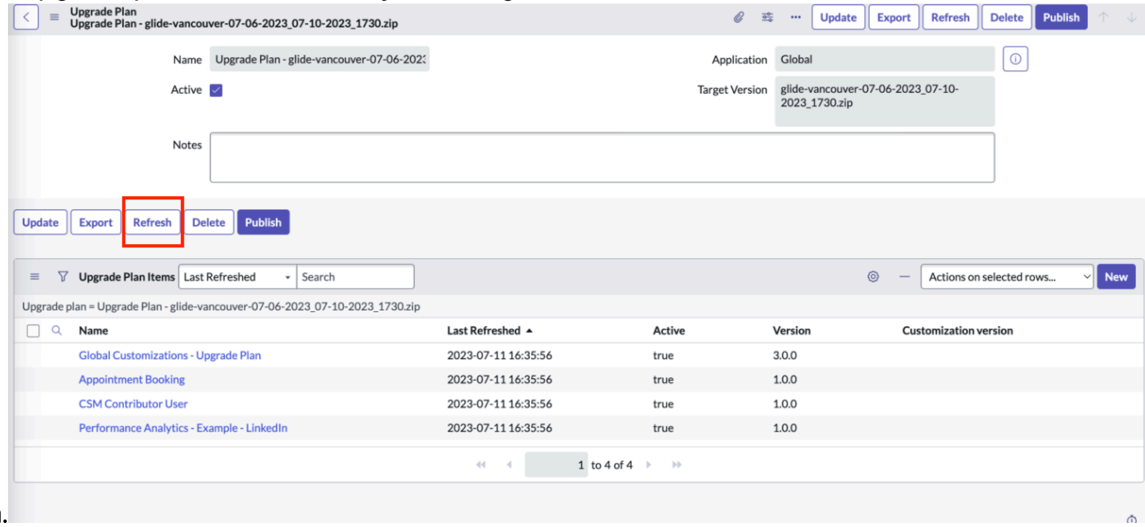

Refresh.

The Refresh Upgrade Plan confirmation message shows up.

3. After the upgrade plan refreshes, select **Go to Upgrade Plan**.

Note:
The **Publish** button is still visible because the updated Upgrade Plan needs to be published.

The recently installed plugins and applications show up under Upgrade Plan Items related list.

Installing your Upgrade Plan

Install your upgrade plan in the consumer instance to implement it in your upgrades.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Plan**.
   The list view of the upgrade plans shows up.
2. Select Retrieve to find your upgrade plan.

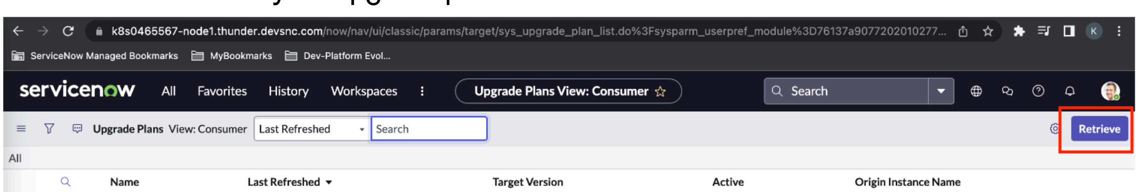

3. Search your upgrade plan by name.
4. **Optional:** Navigate to **All** > **My company applications** and refresh it.

Note:
This step is applicable only if you can't find the required version of your upgrade plan.

5. Repeat the previous steps to see the latest version of the upgrade plan.
6. Select **Install** to install the upgrade plan to your instance.

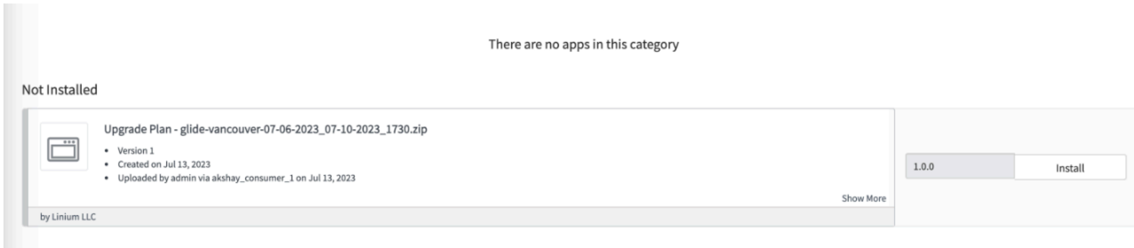

Note:
If you made any changes to the upgrade plan on the builder instance after installing it on the consumer instance, then update the upgrade plan to the latest version.

Prepare to upgrade with Upgrade Plan

Prepare your instance upgrade with Upgrade Plan by determining all the applications and plugins are ready to be implemented in the upgrade.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Plan**.
   The upgrade plan that is installed shows up on the list.
2. Select the upgrade plan.
   The upgrade plan form view shows up with all its configurations.
3. To ensure that all the applications and plugins are ready to be implemented in the upgrade, select **Refresh** in the context menu of the Upgrade Plan.
   If there are errors for any of the items, then fix them and follow the Reprocess step.

   Note:
   All the applications and plugins under the Upgrade Plan Items related list must be in the **Ready** state to proceed with the upgrade. The application or plugin in the **Already Installed** state, error or its latest version won't be considered during the upgrade process.

   The upgrade plan is now ready to be implemented.
4. Apply an existing upgrade plan to an instance anytime, without being limited to during upgrades only.

   You can now also export the upgrade plan in a batch install format by using the **Export** button. See CI/CD - POST /sn_cicd/app/batch/install for more information.
5. **Optional:** Select **Reinstall** to reinstall the failed items in an upgrade plan, after the upgrade completes. You can also review the upgrade plan history to know the reasons of failure.

   Note:
   You are required to use this option only if you have failed items in an upgrade plan.
6. **Optional:** Click **Reprocess** to reprocess the upgrade plan after the upgrade plan item errors have been resolved.

   Note:
   This step is applicable only if you have any error items in the upgrade plan.

Preview Upgrade Plan

Preview your upgrade plan before being implemented in the upgrades. Once the upgrade plan is installed, it auto generates the preview of the upgrade plan.

Before you begin

Role required: admin

Procedure

1. Go to the Upgrade Preview page to preview the Upgrade Plan.
2. Check the Upgrade Plan mentioned on the Previewing version 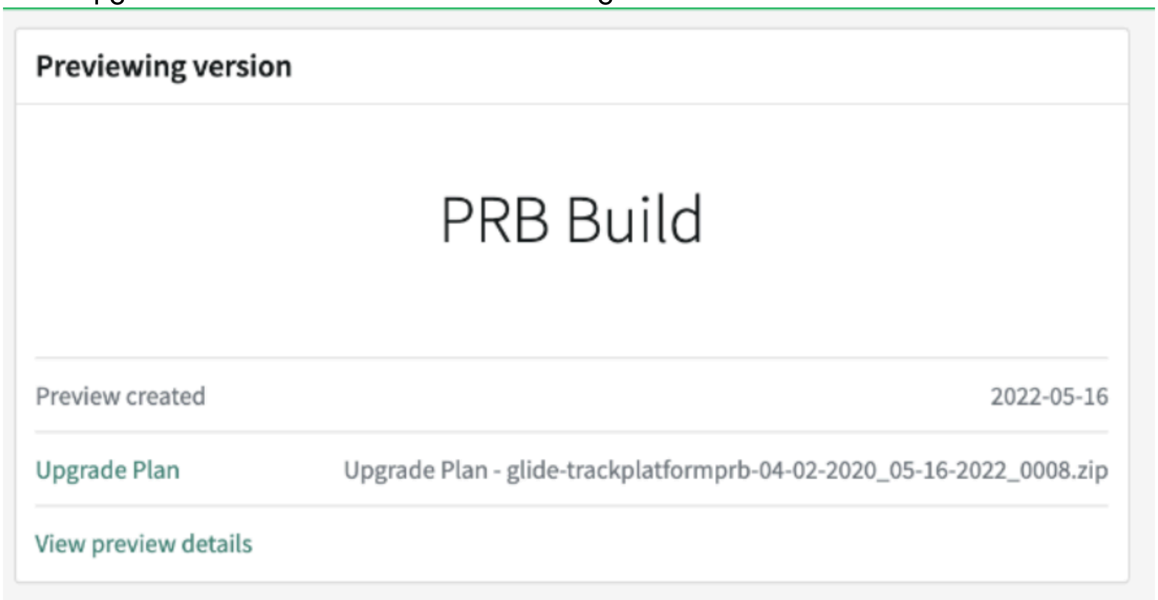

card.
3. Check the Predicted skipped records with upgrade plan card to preview the skipped 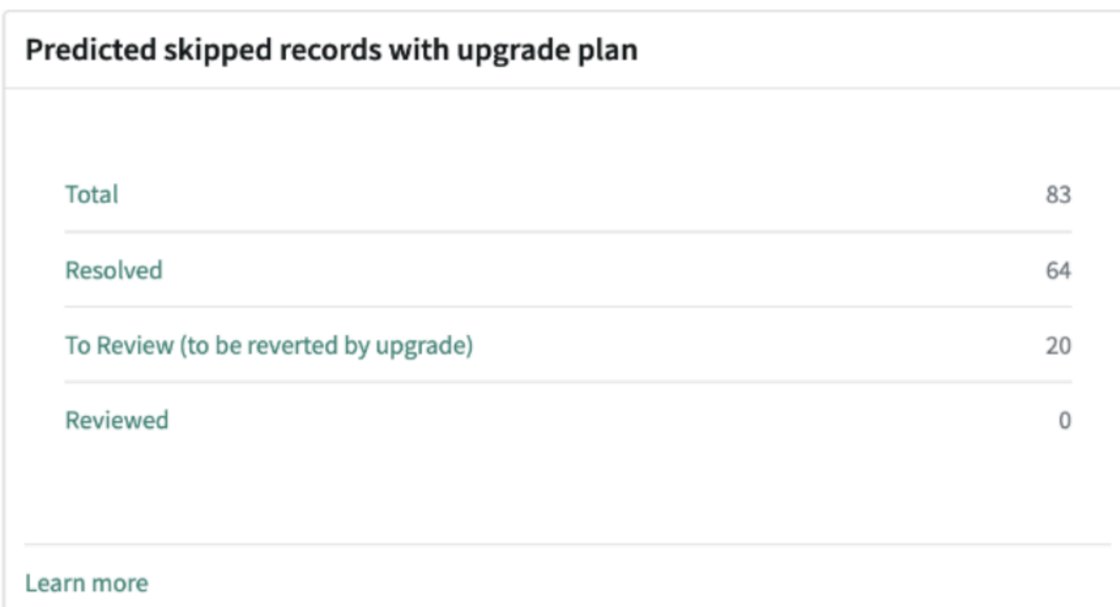

records.
    - Total: total number of skipped records
    - Resolved: skipped records included in the upgrade plan
    - Reviewed: skipped records that have been reviewed

Apply Upgrade Plan on your upgrade

Apply the selected upgrade plan to your instance upgrade.

Before you begin

Role required: admin

Procedure

1. Go to the Upgrade Monitor page to ensure that the selected upgrade plan is being implemented in your upgrade.
   If the upgrade plan has been implemented on the ongoing upgrade, the Review upgrade plan link shows up on the Target version card.

   See Upgrade Plan background operations section in Upgrade Plan overview for more information.

2. Check on the date and time of the implemented upgrade plan in the Upgrade duration card once the upgrade completes.
   You can select the Review upgrade plan link on the Upgrade duration card to know the history of the upgrade plan. It shows all the installed applications and plugins and the failed items.
3. Scroll down to see the skipped records with upgrade plan.
   - Total: Total number of skipped records
   - Resolved: Total records from the upgrade plan that have been auto-resolved
   - To review: Total records left to be reviewed manually
4. **Optional:** Click **Review upgrade plan** on the Upgrade duration card if you want to check the summary of your upgrade plan history.

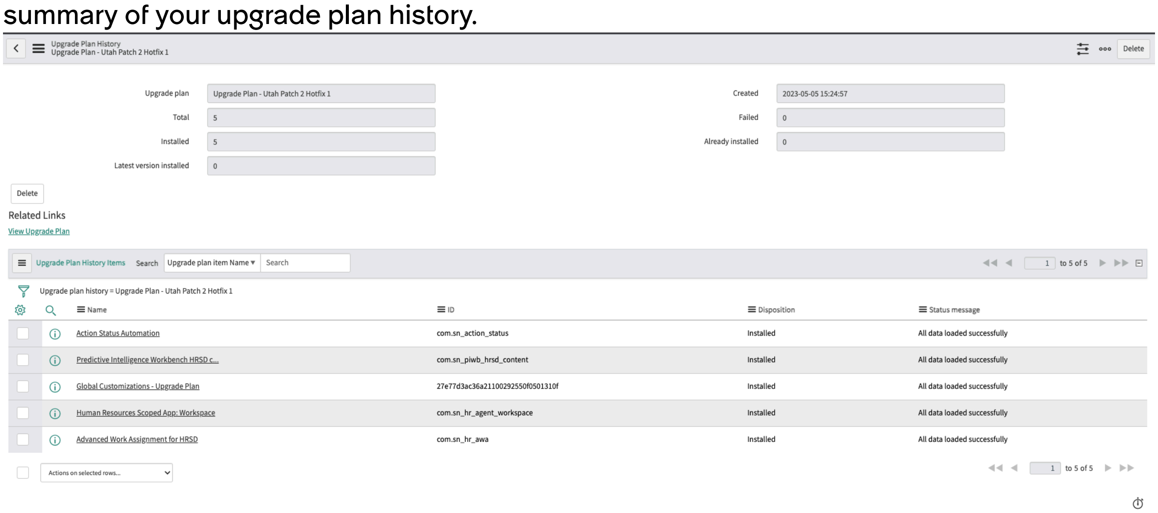

Review skipped records with upgrade plan

Review the skipped records after the completion of the upgrade.

Before you begin

Role required: admin

Procedure

1. Once the upgrade completes, scroll down to see the skipped records with upgrade plan.
   - Total: Total number of skipped records
   - Resolved: Total records from the upgrade plan that have been auto-resolved
   - To review: Total records left to be reviewed manually

   Note:
   After using the upgrade plan, the Global Application is created and all the records move to the Global Application. As soon as you create the Global Application, the new update sets show up.

   You won’t be allowed to choose the skipped records that are captured by the upgrade plan, regardless of whether they are being reviewed. If the customizations are coming from different instances, then the skipped records are required to be reviewed.

2. Select the To review link to look at the records that need to be reviewed manually.

Upgrade History module

The Upgrade History module tracks every upgrade made to an instance. You can also view the complete report of an old upgrade or a recently completed upgrade version using this module.

To view an upgrade history record, navigate to **All** > **Upgrade Center** > **Upgrade History** and select an upgrade from the list. On selecting an upgrade from the list, the System Upgrades form appears.

System Upgrades form details

Field	Description
From	Name of the previous .war file (version).
To	Name of the applied .war file (version).
Upgrade started	Time stamp when the upgrade process began.
Upgrade finished	Time stamp when the upgrade process was completed.

Upgrade History Details form section

Field	Description
Changes skipped	Total number of records that are different from the previous upgrade and the upgrade component was not applied, mostly due to customization.
Changes applied	Total number of changes that are applied as a part of this upgrade **Changes applied** is sum of updated and different records, added to the number of deleted records (where the value of changed is true) added to the number of inserted records (where the value of changed is true).

Upgrade History Details form section (continued)

Field	Description
Changes processed	Total number of records that were processed as a part of this upgrade. Changes processed is the sum of **Changes skipped** and **Changes applied**.
Copies to review	Total number of copied records to review whose base records have been upgraded

Select Skipped Record VTB related link to view and manage skipped records over VTB. See Skipped Records visual task board (VTB) for more information.

Select View Upgrade Plan History related link to view the details of the upgrade plan configuration. See Upgrade Plan overview for more information.

Select View Upgrade Summary Report related link to view the summary of an older upgrade over Upgrade Monitor module.

See Review skipped records using related lists for more information about related lists.

Previewing an upgrade

Use the Upgrade Preview module to have an insight about the experience of an upgrade prior to the actual upgrade. You can explore and preview upgrades to different ServiceNow release versions and see how your instance might be impacted with your current configurations.

- Preview predicted changes: Preview the changes that have been predicted to occur after an upgrade.
- Upgrade Preview form: Alternate view of a previewed upgrade.

View previewed upgrade

Use the Upgrade Preview form to have an alternate view of a previewed upgrade.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Center** > **Preview module**.
2. Select the target version to be previewed.
3. Click **Go** or **Refresh Preview** to run the selected preview.
4. Click **View preview details** on the **Previewing version** card to view the following related lists.
   - **Predicted Skips to Review related list**: The Upgrade Preview process informs you about the customizations that are predicted to be skipped during an upgrade. The upgrade preview process the skip files that have been customized. Predicted Skips to Review lists all the skipped files that haven't been reviewed yet.
   - **Predicted Skips Reviewed related list**: Predicted Skips Reviewed related list displays the records which have previously appeared on the Predicted Skips to Review related list and have been reviewed. When you select a skipped record to review and the **Resolution**

Status has been set to a value other than **Not Reviewed**, the updated record moves to the Predicted Skips Reviewed related list.

- **Previewed Changes related list**: Previewed Changes related list gives the total number of records that are predicted to change when the upgrade occurs. Total record changes also includes possible predicted skip files known as Predicted skipped records.

See Previewed changes for more details about the related lists.

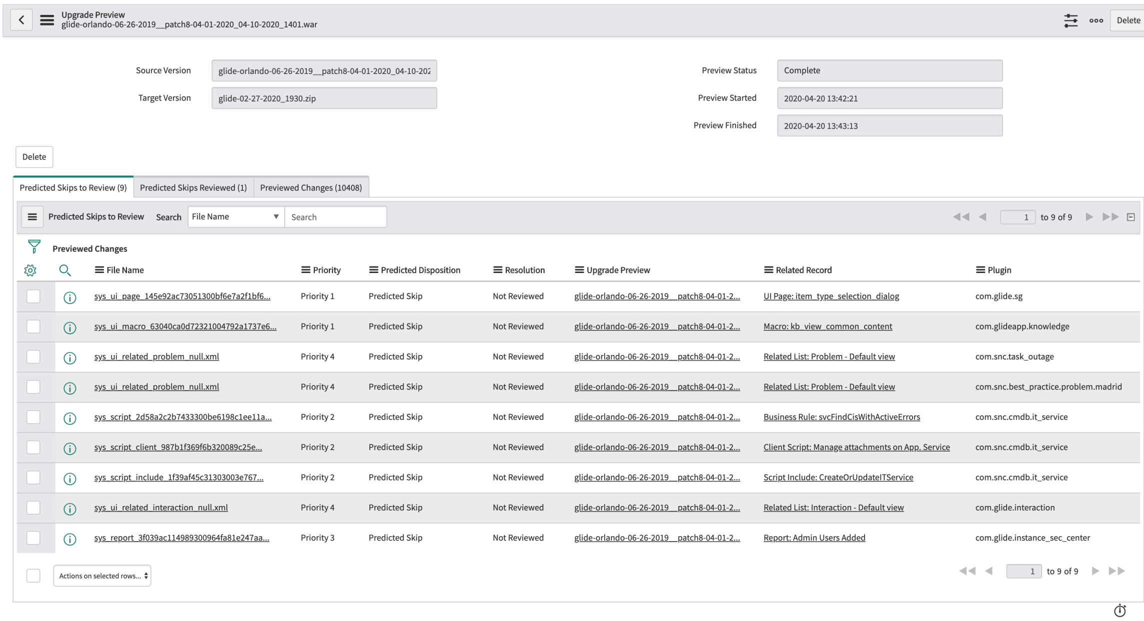

Field	Input Value
Source version	Current version of the instance
Target version	Target version to which the instance is upgraded to be previewed
Preview status	Status of the preview. ○ Pending ○ In Progress ○ Complete ○ Cancelled ○ Failed
Preview Started	Date and time on which the preview started
Preview Finished	Date and time on which the preview completed

Preview predicted changes

Use the previewed changes table to view the list of total records that are predicted to change when the upgrade occurs.

Preview skipped list prediction

The **Total record changes** also includes records that are predicted to be skipped due to various reasons. The most common reason is due to customization. The reason an upgrade

skips customizations is to avoid overwriting your customizations. Upgrade Preview attempts to predict which records will be skipped when the actual upgrade occurs. It predicts skips due to customizations, and some of the other more complex scenarios as well. It is possible to use Upgrade Preview to address these predicted customization skips before the upgrade occurs, which can prevent that skip from happening.

Click **View preview details** link on the **Previewing version** card. The Upgrade Preview form appears along with the following related lists details.

Related lists details

Field	Description
File name	Record that has been flagged for reconciliation.
Predicted Disposition	Action predicted to be performed on this file during the selected upgrade: • Predicted Insert: The system is predicted to insert a new record. • Predicted Update: The system is predicted to update this record. • Reverted: This record was reverted to the base version, and won't be skipped on the next upgrade. • Predicted Skip: The system won't change this record in order to preserve customizations. • Predicted Skip (Manual Merge): The system won't change this record because updating it requires manual intervention. • Predicted Skip (Apply Once): The system is predicted to skip this record because it had already applied an update from an xml file in the apply once folder.
Priority	Priority that has been assigned to resolve the conflict. Values range from one to five, with one representing the highest priority.

Related lists details (continued)

Field	Description
	Note: ServiceNow prioritizes the skipped records based on the importance of the file types. The prioritization is done as follows: • Priority 1 (highest priority): UI pages, UI macros, and more • Priority 2: Business Rules, Security ACLs, and more • Priority 3: Reports and more • Priority 4: Form Sections, Choice Sets, and more • Priority 5 (lowest priority): other
Resolution	Ways to resolve a conflict. • **Not reviewed**: Conflicted files that have not been reviewed yet • **Reviewed**: Reviewed but no action has been taken yet • **Review not needed**: Review of the skipped files are not required • **Reviewed and Retained**: Left customizations in place without update from upgrade • **Reviewed and Reverted**: Customizations discarded, record updated according to the upgrade version
Plugin	Plugin containing the record
Related record	Record that the changelist entry applies to

Scheduling and monitoring an upgrade

Use the Upgrade Monitor module to schedule and monitor the status of an ongoing upgrade in your instance. You can also view the upgrade summary and the list of records causing conflicts in your instance once the upgrade is complete.

- Monitoring an upgrade: Monitor the progress of an ongoing upgrade in your instance
- Process the skipped records list: Resolve the differences between the upgraded and customized versions of a record by processing the skipped records list
- Resolve conflicts of a record: Resolve conflicts between your customized record and the changes associated with the upgrade

Monitor an upgrade to an instance

Monitor the progress of an ongoing upgrade in an instance with the Upgrade Monitor. When the upgrade is done, you can view a summary of the results on the Upgrade Summary Report.

Before you begin
Role required: admin

Procedure

1. Navigate to **Upgrade Center** > **Upgrade Monitor**.

   Note:
   If an upgrade is in progress, the user is automatically directed to **Upgrade Center** > **Upgrade Monitor** page.

2. Monitor the progress of the upgrade.
   When the system finishes the upgrade, it displays the Complete upgrade summary report.

   Note:
   You can now view the status of every stage in the upgrade process. See Upgrade Progress for more information.

What to do next
From the Upgrade Summary Report, resolve any conflicts that prevented the system from upgrading records.

Upgrade Progress

When an upgrade is underway, Upgrade Progress displays progress bars and other information to help you monitor the process.

You can view the status of each of the stages during an upgrade process. It helps you to know exactly which stage you are in at any point of the upgrade process.

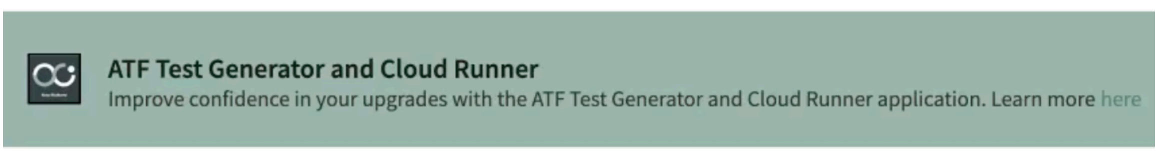

If you have the ATF Test Generator and Cloud Runner store application installed, you can schedule test generation runs from the Upgrade Monitor page. The above message shows up during an upgrade process when the following conditions are met.

- The SN_ATF_TG plugin is not active.
- The instance is not a fed instance.
- The instance is not a self-hosted instance.

Note:
You can now confidently upgrade your instance by using the ATF Test Generator and Cloud Runner store app. See for more information.

Once the database upgrade starts, the following details show up.

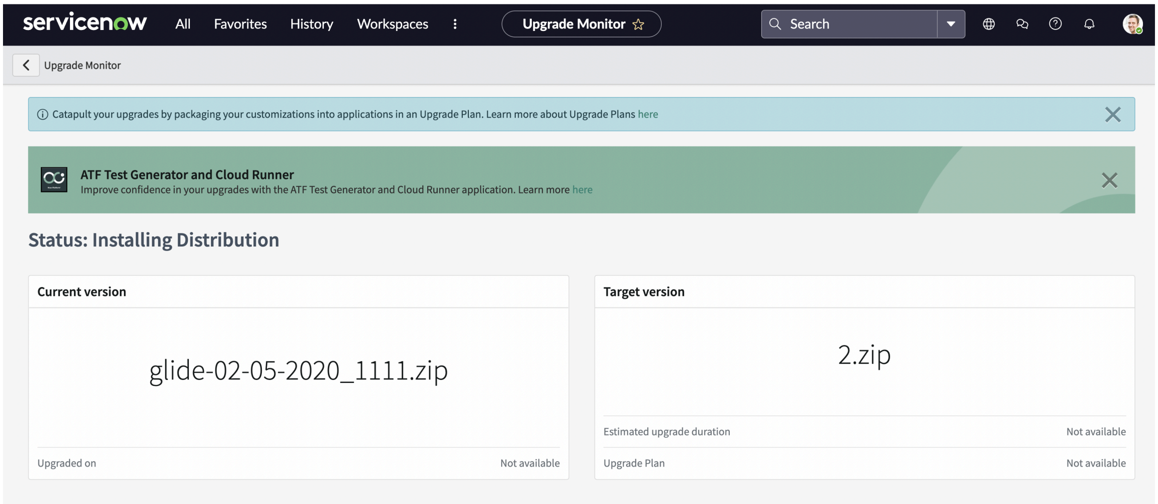

Upgrade Progress details

Field	Input Value
Current version	Current version of the instance
Target version	Target upgrade version of the instance
Upgrade Progress	A progress bar depicting where the upgrade is in the overall process. The length of a section in the bar does not indicate the relative duration of that process.
Upgrade Progress: Upgrading Platform	The system is applying upgrades to elements that form the foundation of the platform.
Upgrade Progress: Updating Schema	The system is scanning the plugins to create a list of tables that require upgrading. This prevents the system from upgrading the same table multiple times.
Upgrade Progress: Loading Plugins	The system is loading both core and optional plugins. Some features require more than one plugin, so the number of plugins listed may not match the number of optional features installed.
Upgrade Progress: Completing	The system is upgrading components that need to be completed after the previous three stages are done.
Details	Shows the current activity, a progress bar for the current activity, and the file currently being updated.
Node Upgrades	The color of the icons represents the status of each node during this upgrade: Pending, Running, Successful, Failed, or Down.
Node	The selected node indicated by the arrow. To change the selection, position the mouse cursor over the icon for the node to select.

Upgrade Progress details (continued)

Field	Input Value
Running time	How long the selected node (indicated by the arrow) has been running. If the selected node is offline, this value stops updating and shows how long a node was online before going offline.
Version	The current build for the selected node (indicated by the arrow).
Successful upgrade	When the selected node is online, shows how long the node has been online. If the selected node is offline, shows how long the node has been offline.

Upgrade Summary Report

The Upgrade Summary report summarizes the actions taken, provides tools to resolve conflicts between customizations and changes that are part of the upgrade, and provides information to help estimate time for upgrades to other instances.

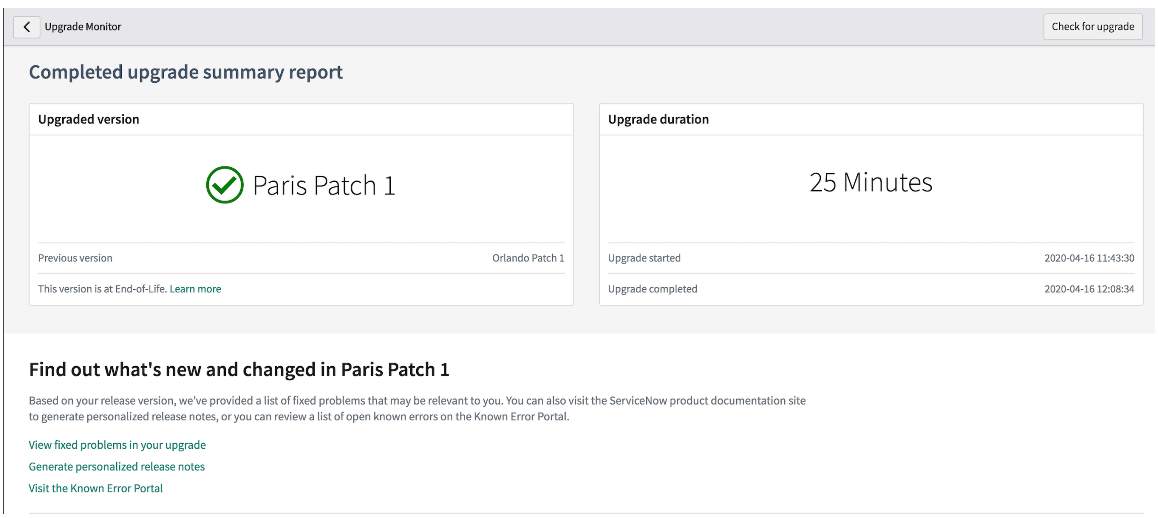

Upgrade Summary details

Field	Input Value
Upgraded version	Version to which the instance has been upgraded
Upgrade duration	Total time taken to make the upgrade
Find out what's new and changed	Links to view the new and changed features in the current upgrade version. The following three links show up.

Upgrade Summary details (continued)

Field	Input Value
	• Problem fixes: Fixes that have been made since the last upgrade version • Personalized release notes: Release notes summary for the previewed version • Known error articles: Errors that have been identified but not yet resolved

Skipped Records

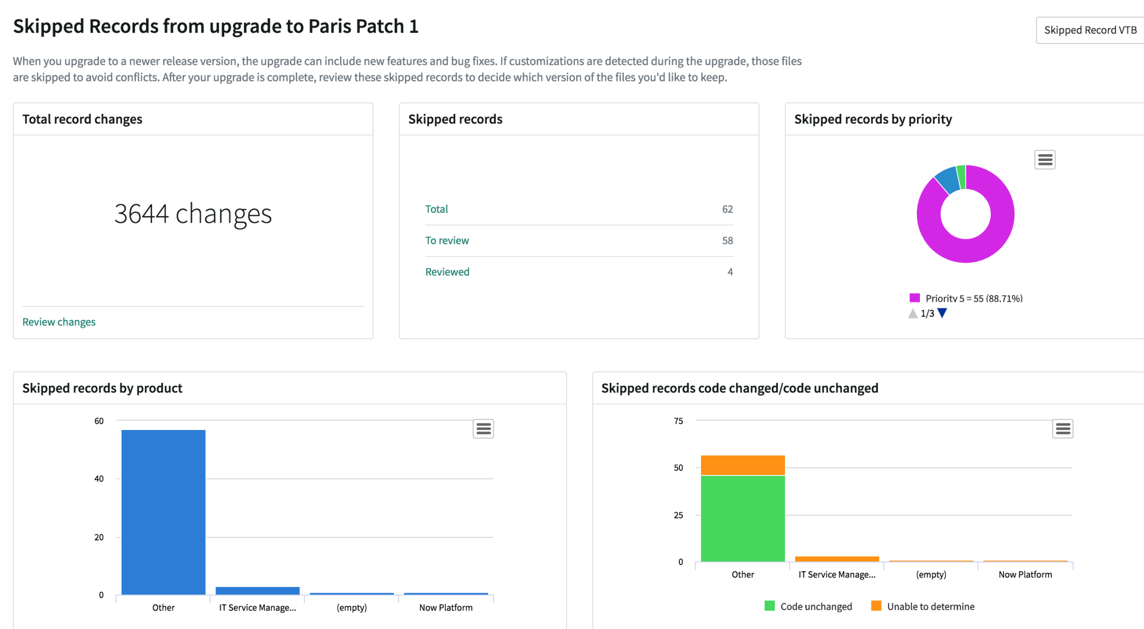

Skipped Records details

Field	Input Value
Total record changes	Total number of records that have changed since the previous upgrade version. **Review changes**: List of records that have changed can be reviewed. See System Upgrade form for more details.
Skipped records	Skipped records to prevent your customizations from being overwritten. If customizations are found in some records during the upgrade, those records are skipped. As an admin, you can resolve each update that was skipped due to a customization. To resolve a skipped update, you can review the reason for each skipped record and then

Skipped Records details (continued)

Field	Input Value
	either retain the customization or revert the customization to the base system. If you choose to revert some of the customizations to the base version, those records are sent through the upgrade without getting skipped. • **Total**: Total skipped records • **To review**: Total skipped records that are yet to be reviewed • **Reviewed**: Total skipped records that have been reviewed
Skipped records by priority	ServiceNow prioritizes the skipped records based on the importance of the file types. The prioritization is done as follows: • Priority 1 (highest priority): UI pages, UI macros, and more • Priority 2: Business Rules, Security ACLs, and more • Priority 3: Reports and more • Priority 4: Form Sections, Choice Sets, and more • Priority 5 (lowest priority): other
Skipped records by product	Skipped records are sorted by product names
Skipped records code changed/code unchanged	Skipped records are sorted depending on if there are changes in the code or not.

Note:
You can click **Skipped Record VTB** to view the resolution status of the current upgrade with skipped records using the visual task board (VTB) view. See Skipped Records visual task board (VTB) for more information.

Node Upgrades

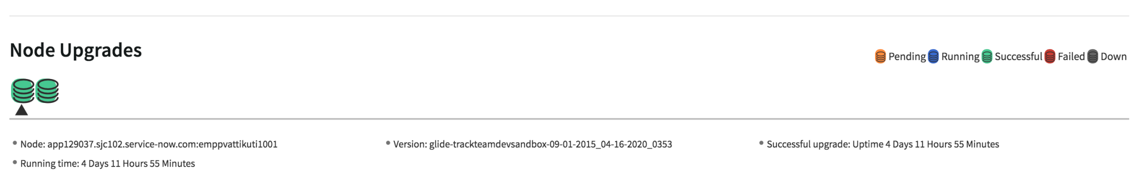

The Node Upgrades section shows the status of the upgrade for each node in the instance. The color of the icon denotes the status, as illustrated by the legend (key) and to the right of the node icons. To see details about a node, position the cursor above the icon for that node. An arrow points to the node selected, and the information below the icons pertains to that node.

Application Upgrade Status

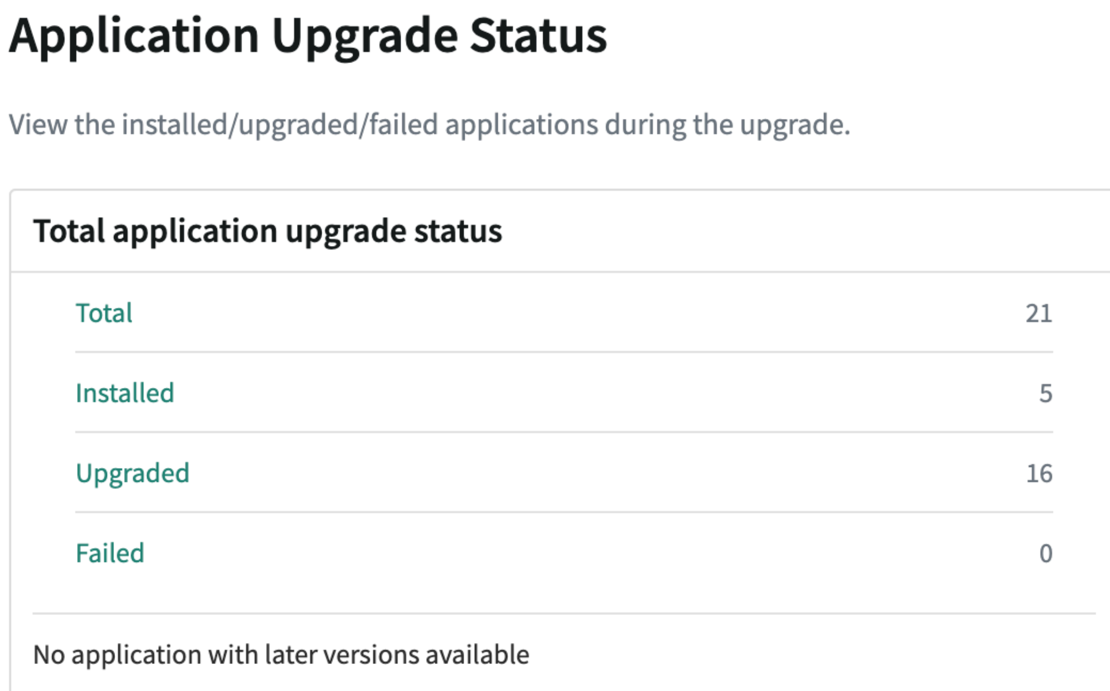

You can view the list of installed and upgraded applications during an upgrade. It also lists the applications that failed to install during the process.

Note:
The **All applications are on the latest versions** message shows up only when all the applications are on their latest versions.

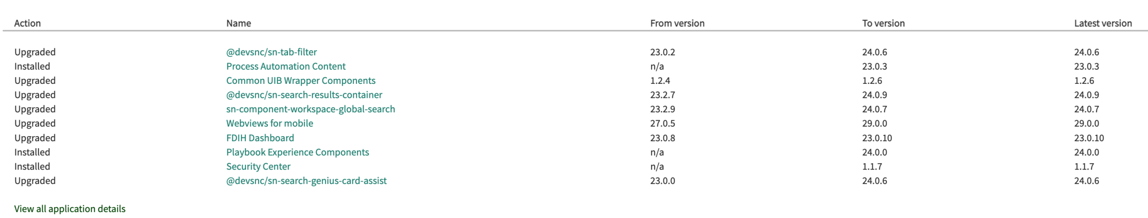

Action	Name	From version	To version	Latest version
Upgraded	@devsnc/sn-tab-filter	23.0.2	24.0.6	24.0.6
Installed	Process Automation Content	n/a	23.0.3	23.0.3
Upgraded	Common UIB Wrapper Components	1.2.4	1.2.6	1.2.6
Upgraded	@devsnc/sn-search-results-container	23.2.7	24.0.9	24.0.9
Upgraded	sn-component-workspace-global-search	23.2.9	24.0.7	24.0.7
Upgraded	Webviews for mobile	27.0.5	29.0.0	29.0.0
Upgraded	FDIH Dashboard	23.0.8	23.0.10	23.0.10
Installed	Playbook Experience Components	n/a	24.0.0	24.0.0
Installed	Security Center	n/a	1.1.7	1.1.7
Upgraded	@devsnc/sn-search-genius-card-assist	23.0.0	24.0.6	24.0.6

View all application details

You can view the first 10 applications that were installed during the upgrade. Select **View all application details** to view the information about all the installed, upgraded, and failed applications on the Upgrade App Version Histories table.

Note:
You can select any of the application links to go to its form view. Select **Go to application** on the form view to go to the application sys record.

The following message shows up if you have selected any of the failed applications

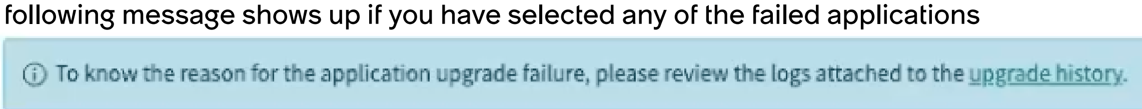

link.

Schema Changes to Clone-excluded Tables

Schema Changes to Clone-excluded Tables

- ua_sp_logouts
- ua_sp_bot_session
- ua_stats_defn
- sys_connection
- discovery_credentials
- sys_import_set
- sys_upgrade_history_log
- ua_ih_usage
- sys_email_account
- ua_sp_usage

View all clone-excluded tables

The Schema Changes to Clone-excluded Tables section shows a list of tables affected by the upgrade that were clone-excluded when you cloned the production instance to this instance. Because clone-excluded tables are empty, upgrading them takes less time than upgrading those same tables on the production instance. To estimate how much longer the production upgrade takes, note the size of the clone-excluded tables on the production instance.

Top 10 Fix Scripts by Duration

Top 10 Fix Scripts by Duration

Duration	Name/Description	Duration	Name/Description
1 Second	fix_outage_number.xml Conditionally loads outage number plugin only if the customer does not already have a custom prefix	0 Seconds	init_vncc.xml Fix script on upgrade to populate semaphore sets
0 Seconds	fix_enable_sqa_ui_plugin_on_upgrade.xml For Social QA upgrade customers Social QA UI plugin will be activated by default in case we want to ship any critical UI fixes	0 Seconds	fix_asmt_metric_type_extendable.xml Update metric types to active
0 Seconds	fix_business_calendar_span_schema.xml Correct the start and end columns to datetime for business_calendar_span	0 Seconds	sys_script_fix_95412ef70f0300105605539ac4767e9b.xml
0 Seconds	fix_sys_hub_flow_base_description.xml Fix job for setting the mandatory attribute of description field to false on sys_hub_flow_base table	0 Seconds	zzzzz_zzzzz_set_upgrade_complete.xml Set the upgrade impending property to false now that upgrade is finished.
0 Seconds	sys_script_fix_6a982fa13b4300105a18802b13efc487.xml	0 Seconds	fix_upgrade_history_log_priority_choice.xml Add 'sys_upgrade_history_log' to property 'glide.ui.permitted_tables' for Upgrade Center Skipped Records Chart UI.

View all fix script durations

The Top 10 Fix Scripts by Duration helps you understand which fix scripts required the most time.

Top 10 Schema Changes by Duration

Top 10 Schema Changes by Duration

Duration	Table name	Alter Type(s)	Element Names	Row Count
8 Seconds	sys_upgrade_history_log	add_columns	[claim_status]	83473
5 Seconds	sys_upgrade_history_log	batch_alter	sys_upgrade_history_log	85607
1 Second	sys_cache_flush	batch_alter	sys_cache_flush	43994
0 Seconds	sys_aw_form_header	batch_alter	sys_aw_form_header	142
0 Seconds	sys_cluster_state	batch_alter	sys_cluster_state	2
0 Seconds	upgrade_center_metadata	create_table	upgrade_center_metadata	0
0 Seconds	sys_declarative_action_group	batch_alter	sys_declarative_action_group	1
0 Seconds	pa_breakdowns	batch_alter	pa_breakdowns	31
0 Seconds	cxs_table_config	batch_alter	cxs_table_config	9
0 Seconds	sys_sg_ui_parameter	batch_alter	sys_sg_ui_parameter	17

View all schema change durations

The Top 10 Schema Changes by Duration helps you understand which schema changes required the most time.

Top 10 Plugins by Duration

The Top 10 Plugins by Duration helps you see the plugins that required the most time. Click **View all plugin duration** to see the **System Upgrade Metrics** list filtered by current sys upgrade history log and sorted by duration. See View loaded files for a plugin for more information.

View loaded files for a plugin

Get a related list view of all the files loaded for a plugin by clicking **View all plugin duration**.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **Upgrade Center** > **Upgrade Monitor**.
   When the system finishes the upgrade, it displays the Upgrade Summary Report.
2. Scroll down to **Top 10 Plugins by Duration**.
3. Click **View all plugin duration**.
   The **System Upgrade Metrics** list of all the plugins and their durations shows up.
4. Click one of the plugins from the list to open the **System Upgrade Metrics** form view of that plugin.

   The Plugin Files related list of all the loaded files for that plugin is displayed.

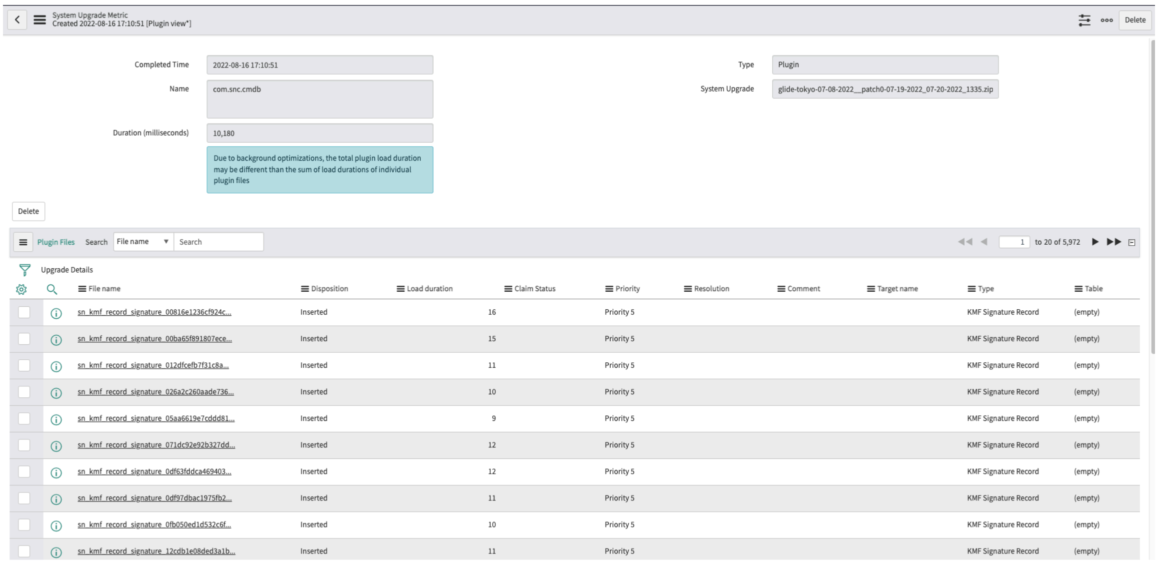

> **Note:**
> The total plugin load duration is not the sum of the load duration of all the loaded files for that plugin.

Process the skipped records list

Process the skipped records list to resolve the differences between the upgraded and customized versions of a record. If you have customized or altered a record affected by this upgrade, such as a business rule or script, the upgrade generates a skip log record.

Before you begin

Role required: admin

About this task

Review the changes you made to baseline records, such as business rules and scripts, that appear on the skipped records list and revert to the baseline version if appropriate. Post-upgrade, thoroughly test all changes you made to these records.

Procedure

1. Navigate to **All** > **Upgrade Center** > **Upgrade Monitor**.

   If the upgrade is still in progress, the system displays the Upgrade Progress screen. When the upgrade finishes, the system displays the Upgrade Summary Report.

2. After the system displays the Upgrade Summary Report, click the Review changes link in the Skipped box.

The system displays the System Upgrades form.

3. Navigate to Review Skipped Records section and, if necessary, scroll to the Skipped Changes to Review related list.
   See Review skipped records using related lists for more information.
4. Click the row for the first record you want to reconcile.

   The system displays the Upgrade details form for that record.

5. Evaluate how you want to resolve the conflict for this record and take the appropriate action:

Action	Description
Retain the customized record as is and do not update it	After reviewing the changes, set **Resolution Status** to **Reviewed and Retained**. The record moves from the Skipped Changes to Review to Skipped Changes Reviewed related list. See Review skipped records using related lists.
Retain the customization by merging changes from the updated object	**a.** Click **Resolve Conflicts** to navigate to the Resolve Conflicts form. **b.** Review the differences. **c.** To merge a field: ▪ Click the right-arrow button for the field. ▪ Click a text box to view and edit the detailed differences. ▪ When you have merged all appropriate fields, click **Merge**. After merging the customization changes: ○ The **Disposition** changes from **Skipped** to **Merged**. ○ The **Resolution Status** changes to **Reviewed and Merged**. ○ The record moves from the Skipped Changes to Review to the Skipped Changes Reviewed related list.
Discard the customization and update the record to match the base system for this upgrade	After reviewing the changes, click **Revert to Base System**. ○ The **Disposition** changes from **Skipped** to **Reverted**. ○ The **Resolution Status** changes to **Reviewed and Reverted**. ○ The system creates a Customer Update record.

Action	Description
	◦ The record moves from the Skipped Changes to Review to the Skipped Changes Reviewed related list. **Note:** At any time after you revert a customization, you can click **Reapply Changes** to reapply the customization (undo the revert).
Review the skip and perform no action on the object	After reviewing the changes, set **Resolution Status** to **Reviewed**. The record moves from the Skipped Changes to Review to the Skipped Changes Reviewed related list.
Leave on the skipped list for a later decision, and note that you have not reviewed the record	From the **Resolution** list, choose **Not Reviewed** to defer the decision on how to handle this conflict. The record stays on the Skipped Changes to Review related list.

Note:
The system tracks changes to records in an update set so you can apply these changes to another instance later. However, the system does not migrate the upgrade details records from one instance to the next. These records apply to a specific upgrade of a specific instance. If you want to preserve the Comments, Resolutions, or other information from the skipped list, export it from this instance.

6. In the **Comment** field, write the reasons for making your decision and other information you want to document.
7. Click **Update**.
   Post-upgrade, thoroughly test all changes you made to the records on the skipped record list.

System Upgrade form

When an upgrade is complete, the System Upgrades form displays key statistics about the upgrade and a related list of skipped records (the skipped list).

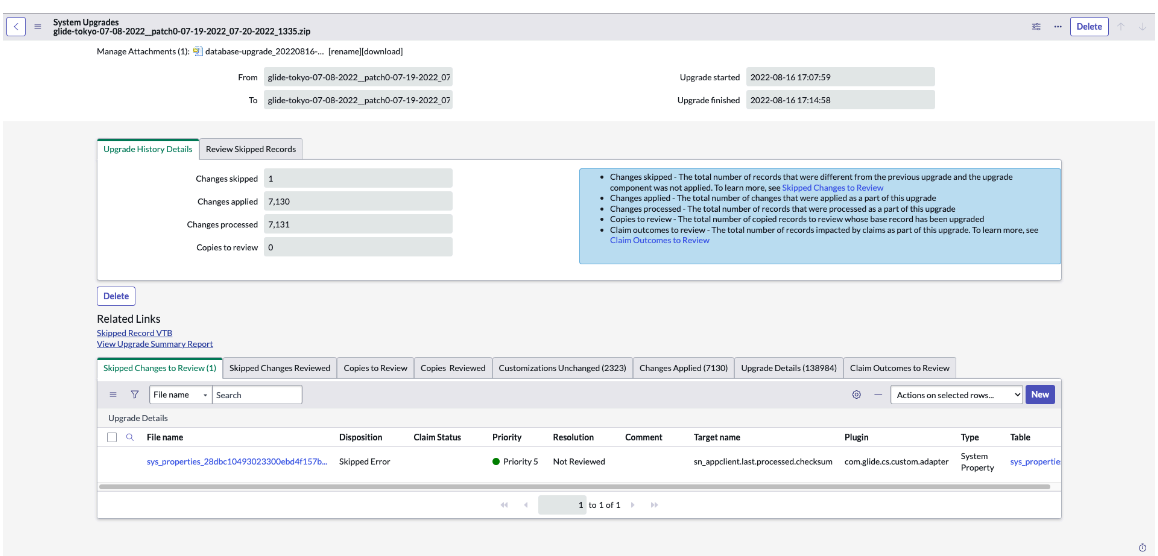

System Upgrade details

Field	Input Value
From	Version of the instance before upgrading.
To	Version of the instance after upgrading.
Upgrade started	Time of the start of the upgrade.
Upgrade finished	Time at which the upgrade completed.
Changes skipped	Total number of records that are different from the previous upgrade and the upgrade component was not applied, mostly due to customization.
Changes applied	Total number of changes that are applied as a part of this upgrade **Changes applied** is sum of updated, different records, deleted records (where the value of changed is true) and inserted records (where the value of changed is true).
Changes processed	Total number of records that were processed as a part of this upgrade. Changes processed is the sum of **Changes skipped** and **Changes applied**.
Copies to review	Total number of copied records to review whose base records have been upgraded

Resolve conflicts for an individual record

Reconcile differences between your customized record and the changes associated with the upgrade.

Before you begin

Role required: admin

Procedure

1. From the Upgrade Details form for the record you are reconciling, click **Resolve Conflicts**.

   The system displays the Resolve Conflicts form, which highlights differences between the two versions of the record. The form displays information about the base system record on the left and the customized record on the right.

   Note:
   The system creates a new customer update record when you click **Resolve Conflicts**.

2. Compare the base system with the customized record for each field on this form.
   For non-script fields, edit the customized record on the right-hand side to include what you want from the base system and the customization.
3. If this record contains a script, check it for conflicts and resolve.

a. Click inside the Script field.

The system displays the Resolve Conflicts - Script form highlighting areas where the two versions of the script differ.

b. Edit the right-hand side so that the script contains whichever content you want. To move a block of code from the left to right side, click the small arrows corresponding to that block in the middle column.

c. Click **OK**.

The system returns to the Resolve Conflicts form.

4. To save your changes to the record, click **Save Merge**.
The system sets the **Resolution** for this record to **Reviewed and Merged**.

Upgrade details form

From the Upgrade Details form, you can review an individual record affected by the upgrade and reconcile conflicts between the upgrade and customizations.

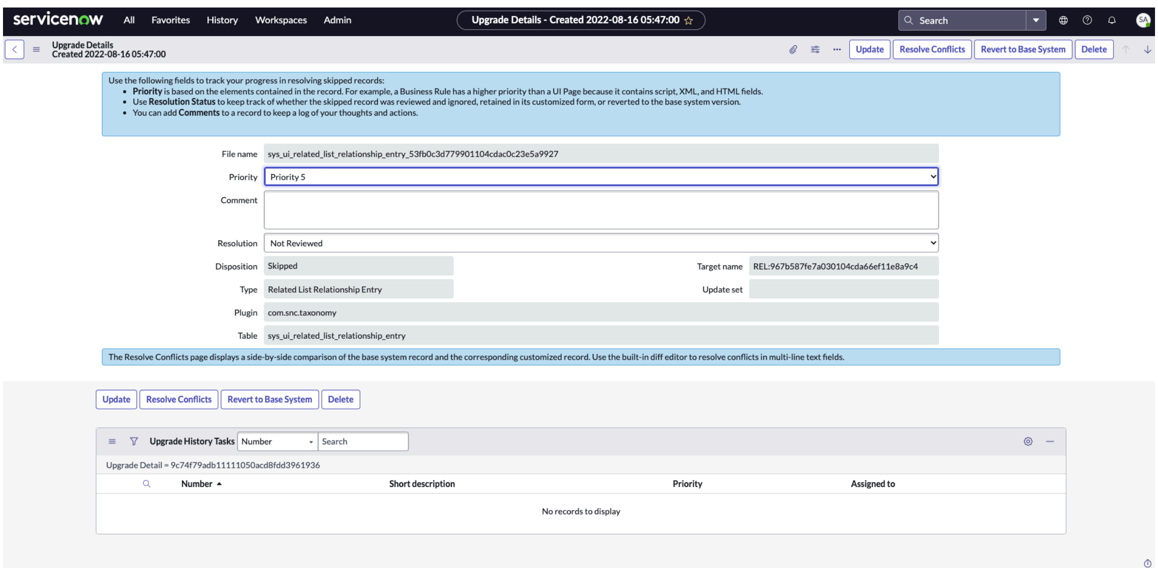

Upgrade Details form details

Field	Input value
File name	The record the system has flagged as needing to be reconciled.
Priority	The priority the system has assigned to resolving this conflict. Values range from one to five, with one representing the highest priority.
Comment	Comments to document your decisions about reconciling this record.
Resolution	How you elected to resolve this conflict: • **Not reviewed**: Not yet reviewed • **Reviewed**: Reviewed but no action yet taken

Upgrade Details form details (continued)

Field	Input value
	• **Reviewed and Merged**: Made changes to the record to reconcile the customized and upgraded versions • **Reviewed and Retained**: Left customizations in place without update from upgrade • **Reviewed and Reverted**: Customizations discarded, record updated according to upgrade For more information, see Process the skipped records list.
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record • Updated: The system updated the record • Deleted: The system deleted the record • Skipped: The system did not change this record in order to preserve customizations • Reverted: This record was reverted to the base version • Table not found: The system could not find the table that contains this record • Unchanged: The system did not change this record because the baseline component has not changed since the last release • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Type	Type of the record, for example, Script include.
Target name	Name of the skipped record, if applicable.
Update set	Unused.
Plugin	The plugin containing this record.
Table	The table containing this record.

Resolve Conflicts form

The Resolve Conflicts form you compare to the base system version with the customized version of a record and reconcile the differences.

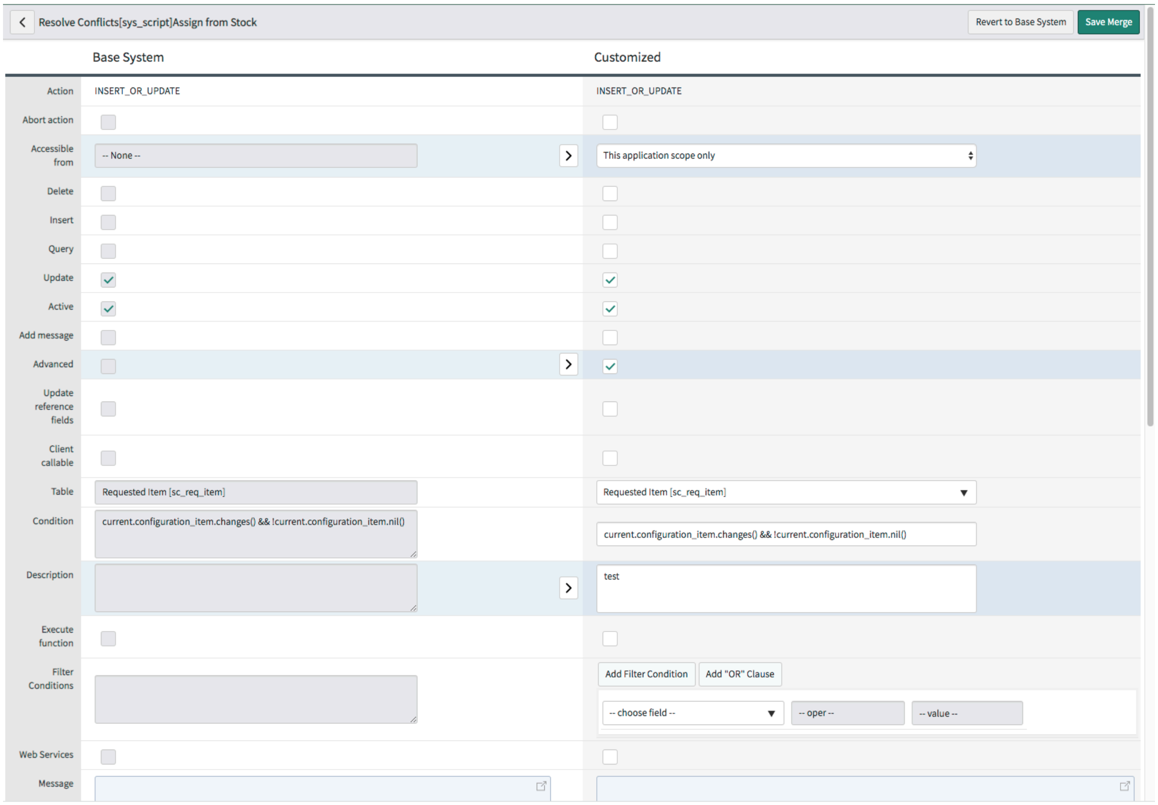

Fields

The fields this form displays depend on the type of record you are reconciling. The left column shows the records fields in the base system, including the proposed changes that are part of the upgrade. The right column shows the fields for your customized record.

Reviewing upgrade history

The Upgrade History module tracks every upgrade made to an instance. You can also view the complete report of an old upgrade or a recently completed upgrade version using this module.

- Upgrade history related lists: Use the related lists to resolve, track and review the skipped records in an upgrade
- Visual task board (VTB) view of skipped records: Use the VTB view to see the resolution status of any previous upgrade with skipped records

Review skipped records using related lists

Use different related lists to resolve, track and review the skipped records in an upgrade.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Center** > **Upgrade History**.
   A list of upgrades is displayed.

2. Select an upgrade from the list.
   System Upgrades form is displayed.
3. Select any of the following related lists to review the skipped records.

   Skipped Changes to Review related list

   Resolve the skipped update records by tracking and reviewing them in the Skipped Changes to Review related list. To prevent your customizations from being overwritten during system upgrades, the upgrade process skips (does not apply the update to) objects that have been customized. See Resolve a skipped update and set a resolution status for more information.

   Skipped Changes Reviewed related list

   View the list of the reviewed skipped records in the Skipped Changes Reviewed related list. When you select a skipped record to review and set a **Resolution Status** to a value other than Not Reviewed, the update record moves from Skipped Changes to Review related list to the Skipped Changes Reviewed related list.

   Copies to Review related list

   Copies to Review lists all the records that need to be reviewed.

   Copies Reviewed related list

   Copies Reviewed lists all the records that have been reviewed.

   Customization Unchanged related list

   Customizations Unchanged lists all records that were skipped (due to a customization), but there have been no changes from the last upgrade. See Revert a customization for more information.

   Changes Applied related list

   Changes Applied lists all changes that were applied in this upgrade.

   Claim Outcomes to Review related list

   During system upgrades, the Claim Status tab displays outcomes to review and resolve.

   Upgrade Details related list

   Upgrade Details lists all Upgrade Details records for this upgrade.

   Note:
   The Copies to Review and Copies Reviewed related lists are Platform specific related lists.

Upgrade history related lists details

The following table gives the description of each field in all the related lists within the upgrade history module.

Related lists details

Field	Input Value
File name	File name has different description for different related lists

Related lists details (continued)

Field	Input Value
	• Skipped Changes to Review related list: Name of skipped Upgrade Detail record • Skipped Changes Reviewed related list: Name of the reviewed Upgrade Detail record • Copies to Review related list: Name of the record that needs to be reviewed • Copies Reviewed related list: Name of the record that has been reviewed • Customization Unchanged related list: Name of unchanged sys_upgrade_history_log record. Select it if you want to access Upgrade Details and add comment text, or set a resolution status for it • Changes Applied related list: Current upgrade file name • Claim Outcomes to Review related list: Name of skipped Upgrade Detail record • Upgrade Details related list: Current upgrade file name
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations. • Skipped Error: The system didn't have the upgrade on this file but skipped due to reasons like the record already exist in the system in a different scope or a system property exists. **Note:** This is editable only by a maint. • Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release.

Related lists details (continued)

Field	Input Value
	• Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Priority	Relative importance of the conflict that caused the skip. **Note:** ServiceNow prioritizes the skipped records based on the importance of the file types. The prioritization is done as follows: • Priority 1 (highest priority): UI pages, UI macros, and more • Priority 2: Business Rules, Security ACLs, and more • Priority 3: Reports and more • Priority 4: Form Sections, Choice Sets, and more • Priority 5 (lowest priority): other
Reason	States the reason why a particular record has been moved to the Skipped Error disposition.
Resolution	• Reviewed • Retained • Reverted
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.
Type	Current file type (such as Business Rule or UI Policy).

Related lists details (continued)

Field	Input Value
Table	Table that contains the record.

Revert a customization

To prevent customizations from being overwritten by system upgrades, the upgrade process automatically skips changes to objects that have been customized. You may want to overwrite your customizations when a software upgrade contains a feature that you would like to implement.

Before you begin

Role required: admin

About this task

To identify customized objects, the system adds a corresponding record in the Customer Updates [sys_update_xml] table. The table maintains the current version information for all objects that have been customized. The upgrade process skips changes to objects that have entries in the table. The upgrade process does not skip objects if only excluded fields have changed.

Procedure

1. Navigate to **All** > **Upgrade Center** > **Upgrade History**.
2. Select the desired software version.
3. Filter the Upgrade Details related list by **Disposition is Skipped**.
4. **Optional:** Add another filter condition for **Changed is True** to return only the objects that have changed since the last upgrade.
5. Select the update record to implement.
   The **File differences** field displays a side-by-side comparison of the customization and the default version. Deletions are highlighted in red, additions in green, and modifications in yellow.
6. Click **Revert to base system** to overwrite your customized object with the system default version.
   - The **Disposition** changes from **Skipped** to **Reverted**.
   - After you revert a customization, you have the option to click **Reapply Changes** to reapply your customizations (undo the revert).

Resolve a skipped update and set a resolution status

Review the reason for each skipped record to resolve each skipped update after an upgrade.

Before you begin

Role required: admin

About this task

You resolve an update by either retaining the customization or by merging or overwriting the customization with the base system update.

Note:
Objects that are customized and that did not change in the base system since the last upgrade require no action on your part.

When an object is customized, the system adds a corresponding record to the Customer Updates [sys_update_xml] table and then maintains current version information for all customized objects. The upgrade process skips changes to objects that have a current version in the Customer Updates table. When you follow the procedure, you perform one of the following actions:

- Retain a customization with no changes
- Retain a customization by merging changes from the updated object
- Revert a customized object to the updated version (that is, overwrite the customization)
- Review the skip and perform no action on the object

Procedure

1. Navigate to **All** > **Upgrade Center** > **Upgrade History**.
2. Select the desired software version.
3. In the Skipped Changes to Review related list, select the update record to resolve.
   See Review skipped records using related lists.
4. Review the list of changes.
   For text fields, you can click in the field to open the Diff/Merge tool. Review the differences. Click a text box to view and edit the detailed differences.
5. Perform one of the following actions.
   You have the option to add a **Comment** to any record, for example, to explain the action to future reviewers.

Skipped Changes to Review related list

Field	Input Value
Retain the customized record as is and do not update it.	After reviewing the changes, set **Resolution Status** to **Reviewed and Retained**. The record moves from the Skipped Changes to Review to Skipped Changes Reviewed related list.
Retain the customization by merging changes from the updated object.	a. Click**Resolve Conflicts** to navigate to the Resolve Conflicts form. b. Review the differences. c. To merge a field: ▪ Click the right-arrow button for the field. ▪ Click a text box to view and edit the detailed differences. ▪ When you have merged all appropriate fields, click **Merge**. After merging the customization changes: ○ The Resolution Status changes to Reviewed and Merged. ○ The record moves from the Skipped Changes to Review to the Skipped Changes Reviewed related list.

Field	Input Value
Discard the customization and update the record to match the base system for this upgrade.	After reviewing the changes, click **Revert to Base System**. ◦ **Disposition** changes from **Skipped** to **Reverted**. ◦ **Resolution Status** changes to **Reviewed and Reverted**. ◦ The system creates a Customer Update record. ◦ The record moves from the Skipped Changes to Review to the Skipped Changes Reviewed related list. **Note:** At any time after you revert a customization, you can click **Reapply Changes** to reapply the customization.
Review the skip and perform no action on the object.	After reviewing the changes, set **Resolution Status** to **Reviewed**. The record moves from the Skipped Changes to Review to the Skipped Changes Reviewed related list.
Leave on the skipped list for a later decision, and note that you have not reviewed the record.	From the **Resolution** list, choose **Not Reviewed** to defer the decision on how to handle this conflict. The record stays on the Skipped Changes to Review related list.

6. Click **Update**.
7. Repeat the process to resolve each update record in the list.

Result

Only skipped updates with a **Resolution Status** of **Not Reviewed** or without a resolution set appear in the Skipped Changes to Review related list. Any action you take that changes the **Resolution Status** to a value other than **Not Reviewed** or **--None--** removes the skipped update from list and moves it to the Skipped Changes Reviewed related list.

Skipped Records visual task board (VTB)

View the resolution status of any previous upgrade with skipped records using the visual task board (VTB) view. An upgrade history record is created for each upgrade that is run.

To view the history report navigate to **Upgrade Center** > **Upgrade History**. You can then select an upgrade from the list to view the upgrade history details. Click the **Skipped Record VTB** related link to view and manage the skipped records over VTB.

Note:
The **Skipped Record VTB** related link shows up only when there are skipped records for the version upgrade.

Upgrade History VTB details

Field	Description
Not Reviewed	Tasks which have not been reviewed
Reviewed	Tasks which have been reviewed
Merged	Tasks which have both the old and new changes
Retained	Tasks which have retained the updates from the latest upgrade
Reverted	Tasks which have reverted its changes to the base system

The skipped records are prioritized based on the importance of the file types. The prioritization is done as follows:

- Priority 1 (highest priority): UI pages, UI macros, and more
- Priority 2: Business Rules, Security ACLs, and more
- Priority 3: Reports and more
- Priority 4: Form Sections, Choice Sets, and more
- Priority 5 (lowest priority): other

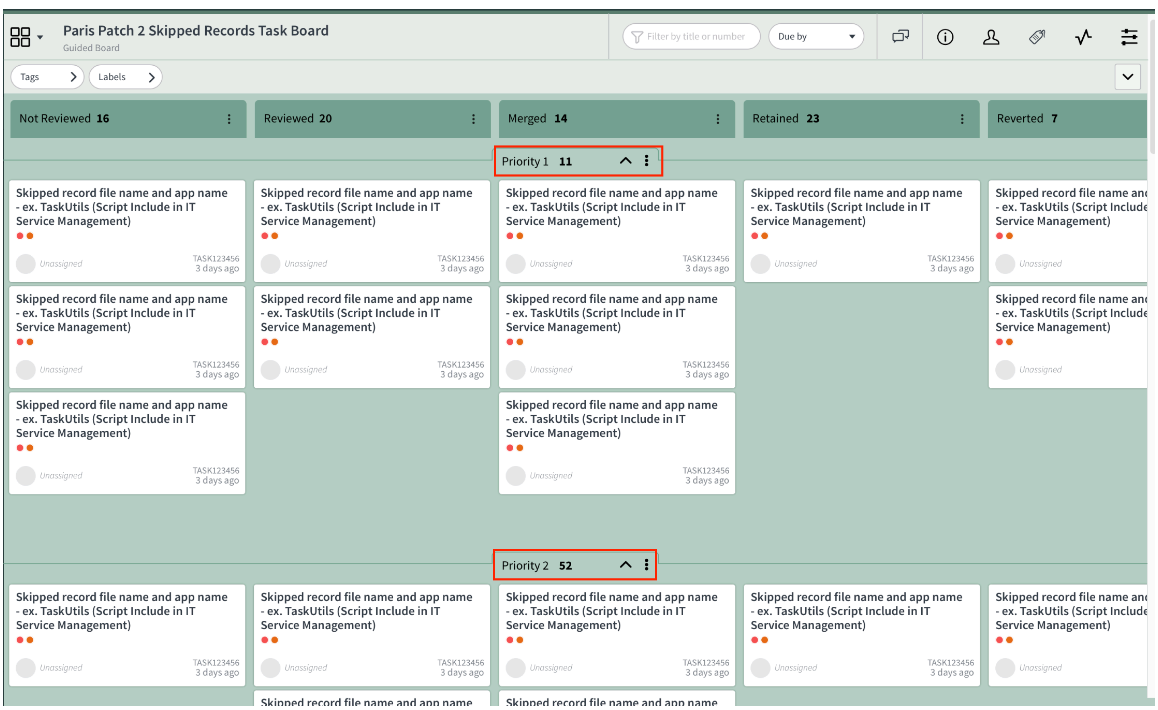

Upgrade History Task form

You can update information about a skipped record task using the Upgrade History Task form.

To make changes on a skipped record task, click a task number link on a task over VTB.

Upgrade History task form details

Field	Description
Short description	Description about the task
Priority	Prioritization of the skipped records based on the importance of the file types. The prioritization is done as follows: • Priority 1 (highest priority): UI pages, UI macros, and more • Priority 2: Business Rules, Security ACLs, and more • Priority 3: Reports and more • Priority 4: Form Sections, Choice Sets, and more • Priority 5 (lowest priority): other
Disposition	Disposition is Skipped since the task board is only for skipped records. The system did not change this record in order to preserve customizations.
Type	Type of file which determines the priority level.
State	State of the files. • Not reviewed: Tasks which have not been reviewed • Reviewed: Tasks which have been reviewed • Merged: Tasks which have both the old and new changes • Retained: Tasks which have retained the updates from the latest upgrade • Reverted: Tasks which have reverted its changes to the base system
Work notes	Notes about the skipped record task
Update	Button to update the skipped record task
Resolve Conflicts	Button to resolve conflicts by reviewing the differences
Revert to Base System	Button to revert the customizations
Delete	Button to delete the skipped record task

The Old Resolution Details related list shows the actions taken in the previous upgrade version.

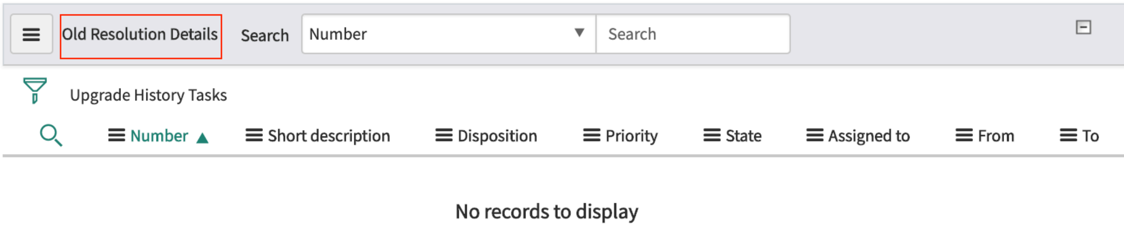

Old Resolution Details related list

Field	Description
Number	Task number of the skipped file where some actions were taken in the previous upgrade version.
Short description	Description about the task
Disposition	Action performed on this file during previous upgrade
Priority	Prioritization of the skipped records based on the importance of the file types. The prioritization is done as follows: **1.** (highest priority): xml content **2.** script or script_plain **3.** html content **4.** sys_ui_form_section, sys_ui_related_list, or sys_choice_set **5.** (lowest priority): other
State	State of the files. • Not reviewed: Tasks which have not been reviewed • Reviewed: Tasks which have been reviewed • Reviewed and Merged: Tasks which have been reviewed and have both the old and new changes • Reviewed and Retained: Tasks which have been reviewed and have retained the updates from the latest upgrade • Reviewed and Reverted: Tasks which have been reviewed and have reverted its changes to the base system
Assigned to	Name of the assigned user
From	Past previous version
To	Past current version

Previous Resolutions related list

The Previous Resolutions related list shows the history of the selected skipped record. You can see what resolutions have been done with the selected skipped record in the previous upgrades.

Upgrade History task form details

Field	Description
File name	Name of the skipped record.
Disposition	Disposition is skipped since the related list is only for skipped records.
Priority	Prioritization of the skipped records based on the importance of the file types. The prioritization is done as follows: • Priority 1 (highest priority): UI pages, UI macros, and more • Priority 2: Business Rules, Security ACLs, and more • Priority 3: Reports and more • Priority 4: Form Sections, Choice Sets, and more • Priority 5 (lowest priority): other
Resolution	The resolution made in the previous upgrades
Target name	Name of the record that you see on UI, as opposed to the sys_update_name which is the underlying name that has the table and sys_id
Type	Type of file which determines the priority level.
Table	The table where the skipped record belongs.

Update default labels in VTB view

Filter skipped records in the task board of the VTB with the implementation of color-coded labels. You can filter the skipped records by assigning a color to each of the products.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Center** > **Upgrade History**.
2. Select an upgrade from the list.
   The System Upgrades form appears.
3. Click the **Skipped Record VTB** related link to view the resolution status of the skipped records over VTB.

   Note:
   The **Skipped Record VTB** related link shows up only when there are skipped records for the version upgrade.

   The VTB view of the resolution status of the skipped records shows up.

4. Click **Labels** to show the default labels.
   The default labels consist of Include codes, Does not include code and Others categories.

   > **Note:**
   > The Include and Does not include codes categories indicate if there is any change in the code of the skipped records. All products fall under Other category of labels. The product labels are assigned with different colors. The tasks get assigned based on the label colors of the products.

5. **Optional:** If you want to add a label, click **Add Label**.
   Enter the name of the label you want to create. You can also change the color of a label by

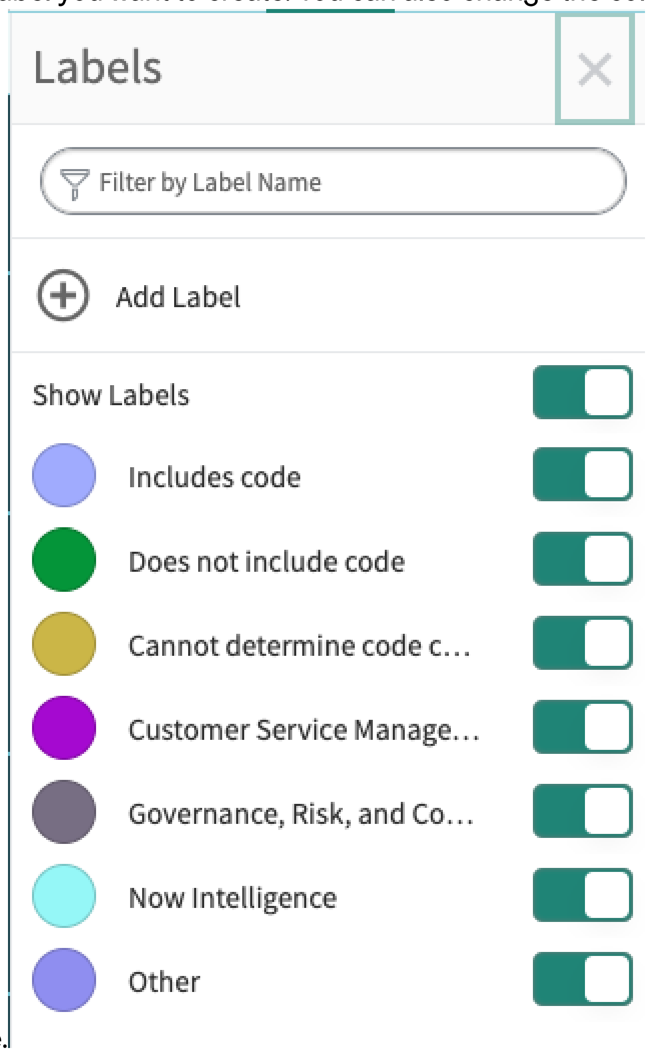

   clicking the color circle.

   > **Note:**
   > The new labels and colors are saved only for the current visual task board and don't get rolled over to the next updated board. If you want the new label updates to show up in the next updated board, navigate to **Upgrade Center** > **Administration** > **Properties** > **Upgrade Center VTB Labels**. Set the colors as required and the updated colors start showing up from the next generated board.

What to do next

Click on a skipped record in the VTB to update the task. See Upgrade History Task form for more information.

View import history

View your import history by accessing the **My Application Import History** module.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **My Application Import History**.
   The list of import history records in your instance shows up.
2. Select one of the imports from the list for more information about that import.
   The form view of sys_upgrade_history record shows up.

   Note:
   To learn more about importing from Source Control, see ServiceNow Studio. Once the importing is done, use the **My Application Import History** module to view the import history.

Quick access to plugins and history records

Reduce the navigating time by directly accessing plugins, application installation history records, and update set commits history records. The filters for the plugins and history records are automatically implemented in the plugins and history records views.

Before you begin

Role required: admin

About this task

Plugins and application installation history records

1. Navigate to **System Definition** > **Plugin Installation History**.
2. Select any of the plugins or history records to access them directly.

Update set commit history records

1. Navigate to **System Update Sets** > **Update Sets to Commit**.
2. Select any of the update sets to directly view the commit history of that update set.

Explore upgrade history log

Reduce processing time by extracting all upgrade related logs into a separate file named database-upgrade_<timestamp>.log. You can also zip the file and attach it to upgrade history.

Before you begin

Role required: admin.

About this task

When the upgrade starts, the upgrade log messages are written to both the localhost_log file and the database-upgrade_<timestamp>.log file. When the upgrade is done, the database-upgrade_<timestamp>.log file is uploaded to the upgrade history table.

Procedure

1. Navigate to **All** > **System Logs** > **Utilities**.

   Note:
   You can implement this step to view the extracted database-upgrade_<timestamp>.log file only after the upgrade is done.

2. Click **Node Log File Download**.
   The database-upgrade_<timestamp>.log file appears in the list of log files.

   Note:
   The localhost_log file consists of all the logs. The database-upgrade_<timestamp>.log file consists of logs only related to the upgrade.

   The database-upgrade_<timestamp>.log file appears only if the `glide.db.upgrade.log.save_to_db.enabled` property is set to True. By default, the `glide.db.upgrade.log.save_to_db.enabled` is set to True.

3. Go to **Upgrade Center** > **Upgrade History**.
   A list of upgrades shows up.
4. Click the required upgrade from the list.
   You can see `database-upgrade_<timestamp>.log` file is automatically attached to the upgrade.

   Note:
   The `database-upgrade_<timestamp>.log` file gets attached only if the `glide.db.upgrade.log.save_to_db.enabled` property is set to True. By default, the `glide.db.upgrade.log.save_to_db.enabled` is set to True.

5. Download the zipped database-upgrade_<timestamp>.log file.

   Note:
   The database-upgrade_<timestamp>.log file has been zipped to reduce the download time.

Administration

The Administration module contains properties for managing Upgrade Center.

Properties

On the properties form, you can set parameters that control how the system executes.

Upgrade Duration Estimation Properties

Field / Element	Description
Queries upgrade duration estimate *`glide.upgrade.preview.duration.instance.id`*	A non-production instance queries the upgrade duration estimate. If this property is absent, the system uses all remote instance records to query the upgrade duration estimate. You can also copy the sys_id of the remote instance record you wish to query for

Upgrade Duration Estimation Properties (continued)

Field / Element	Description
	upgrade duration estimate. The system uses the highest duration from the results.

Upgrade Visual Task Board Properties

Field / Element	Description
Enable an admin or user with admin role to own Upgrade VTB *glide.upgrade_center.task_board.owner*	Upgrade VTB owner must be an admin. If this property is absent, the system uses the first admin user from active admin users list. The active admin users list is sorted by username.
Enable admins or users with admin role to be members of Upgrade VTB *glide.upgrade_center.task_board.members*	Upgrade VTB members must be an admin. If this property is absent, the system adds all active admin users as members.

Upgrade Plan Properties

Field / Element	Description
Queries about the type of instance the user is configured. *glide.upgrade.plan.instance_type*	You must configure the instance either as a builder or a consumer. You must build the upgrade plan on a builder instance and implement it on a consumer instance.

Upgrade Center VTB Labels list

The Upgrade Center VTB Labels list under Administration module shows all the labels in the Upgrade Visual Task Board (VTB). You can also change the color of the existing labels.

If you want to change the color of an existing label, navigate to **Upgrade Center** > **Administration** > **Properties** > **Upgrade Center VTB Labels**

Note:
If you update the color of an existing label, the change will not be reflected in the current version upgrade related board. The change will be seen only from the next version upgrade.

Upgrade Center roles

Upgrade Center is installed with these roles.

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

System administrator [admin]

Access all tables and information within Upgrade Center on your instance.

Contains Roles

List of roles contained within the role.

- sn_templated_snip.template_snippet_admin
- sn_employee.admin
- taxonomy_admin
- sn_ace.ace_user
- sn_hr_sp.esc_admin

Groups

List of groups this role is assigned to by default.

None

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevate to a privileged role.

No.

Special considerations

None.

Managing upgrade risk

You will be notified of the possible conflicts that might result due to several customizations on a skipped metadata file, with a new UI message on the form pages of the metadata files.

One of the major conflict areas during an upgrade is resolving conflicts for skipped metadata files that are released by ServiceNow® but have been customized. These conflicts result due to multiple users updating the same files. ServiceNow might update the code in the files which you have customized over time. Whenever you install an application or a plugin, you tend to make several customizations as per your requirements. During the next upgrade, you might have several skipped records depending on the previous customizations.

> **Important:**
> Avoid making modification on the high risk and medium risk (if possible) skipped metadata files that can potentially cause conflicted skipped metadata files in your next upgrade. You can edit the files if the modifications are absolutely required.

The goal of the UI message is to make you aware of the friction that will be introduced during the upgrade process, if any high priority files have been modified. The

notification is of the form of UI messages on the form pages of the metadata

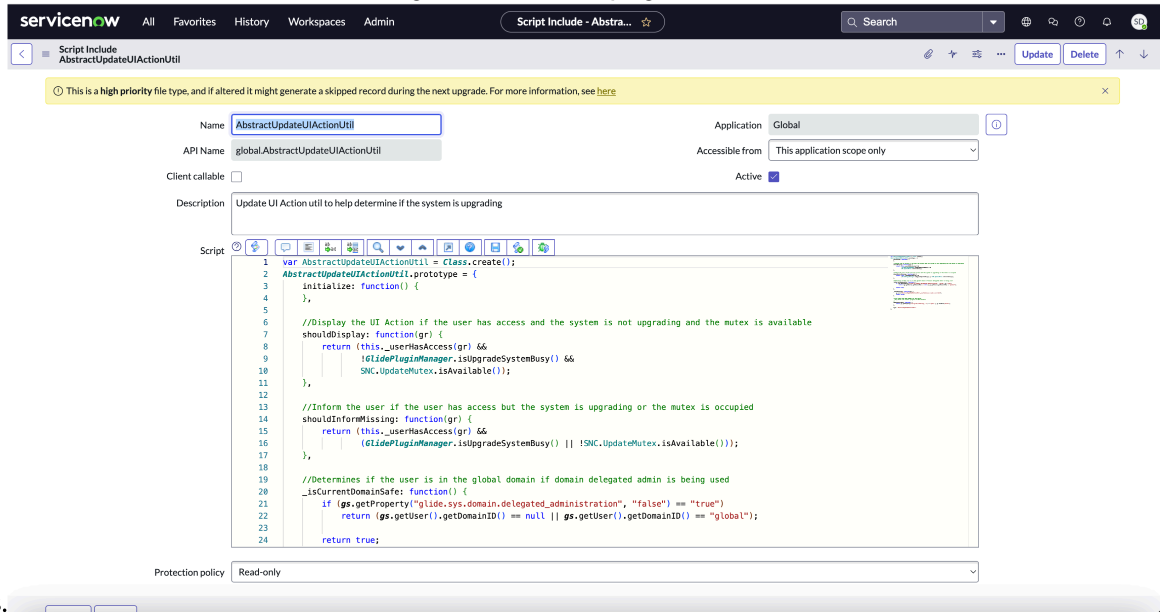

files.

Create a skipped record rule

Create skipped record rules based on your set conditions to define your customizations after an upgrade.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Center** > **Administration** > **Skipped Record Rules Editor**.
   A list of available rules shows up.

   Note:
   By default, a few inactive sample rules have been provided. You can change these rules to active to use any of them.

2. Select **New** to create a new skipped record rule.
   The Upgrade Skipped Record Rules form shows up.

   Note:
   The new rules are active by default. You can make them inactive if you don’t want to use them.

3. Fill in the fields on the form.

Field	Description
Name	Name of the rule.
Order	Order in which the rule runs. **Note:** The order number automatically increments from the last rule order created.

Field	Description
Active	Option to activate the rule.
Conditions	Conditions as to when the rules are implemented on the skipped records. **Note:** These conditions are only for skipped records that are not being reviewed. If you set the disposition condition to anything other than the skipped records, that condition doesn't work.
Action	Option to choose one of the four following actions: ◦ Keep My Modifications (Always Retain) ◦ Revert and Keep Inactive ◦ Assign Skipped Records to User ◦ Assign Tags to Skipped Records **Note:** Based on the selected action, a message shows up with more information.
Comment	Add a comment on the conditions being set.

4. Select **Submit**.
   The new rule appears on the list of available rules.

What to do next

You can now run the skipped record rules either automatically during an upgrade or can run on demand on the skipped records. See for more information.

Execute a skipped record rule

Run skipped record rules based on your set conditions to resolve skipped records in an upgrade. The rules either execute automatically during an upgrade or can run on demand after an upgrade.

Before you begin

Role required: admin

Execute a skipped record rule on demand

Execute the skipped record rules on demand on the skipped records whenever it's required.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Upgrade Center** > **Administration** > **Skipped Record Rules Editor**.
   A list of available rules shows up.

Note:
By default, a few inactive sample rules have been provided. You can set it to active if you want to use the sample rules.

2. Execute the rules on demand using either of the following ways.

Execute a skipped record rule automatically

Execute the skipped record rules automatically during the upgrade on the skipped records.

These rules are executed on the skipped records automatically during the upgrade process. The prediction of the rules execution on the skipped records can be previewed on the Skipped list prediction card in the Upgrade Preview module.

Once the execution of the rules are done, the summary can be reviewed in the Upgrade Summary of Upgrade Monitor.

Note:
Skipped record rules and upgrade plans can't be implemented together in an upgrade.

Upgrades and conversions

The ServiceNow platform includes tools to help you during and after an upgrade, and also allows you to convert a ServiceNow Express List instance to a ServiceNow Service Management enterprise platform.

Washington DC

The upgrade process moves your instance to a new ServiceNow release version. Understand the difference between upgrading and patching, release definitions, rollback and backup options, and how to test your non-production and production instance upgrades.

Upgrade History module: Track every upgrade

The Upgrade History module tracks every upgrade made to an instance. Administrators can use the module to resolve upgrade conflicts and optionally to revert customizations to base system versions to take advantage of new features.

An upgrade history record is created for each upgrade that is run. To view an upgrade history record, navigate to **System Diagnostics** > **Upgrade History** and click the upgrade.

Note:
Debug Upgrade provides detailed debugging output for transactions containing artifacts affected by the most recent upgrade, and is designed to assist in upgrade error resolution. See Debug upgrade.

Note:
The **Payload** and **Payload Hash** fields have been removed from the Upgrade History record in a previous release.

Upgrade History record

Field	Description
From	Name of the previous .war file (version).
To	Name of the applied .war file (version).

Upgrade History record (continued)

Field	Description
Upgrade started	Time stamp when the upgrade process began.
Upgrade finished	Time stamp when the upgrade process was completed.
Changes skipped	Total number of records that were different from the previous upgrade but were skipped, mostly likely due to customization. Changes skipped is the sum of the records that have disposition of skipped manual merge (where the value of changed is true), added to the number of records that have disposition of skipped error, added to the number of records that were skipped and different. **Note:** To prevent your customizations from being overwritten during system upgrades, the upgrade process skips (does not apply the update to) objects that have been customized. One of your responsibilities as the administrator is to resolve each update that was skipped due to a customization. To resolve a skipped update, you review the reason for each skipped record and then either merge the customization or revert the customization to the base system.
Changes applied	Total number of the changes that were applied in this upgrade. Changes applied is sum of updated and different records, added to the number of deleted records (where the value of changed is true) added to the number of inserted records (where the value of changed is true).
Changes processed	Total number of items processed during this upgrade. Changes processed is the sum of Changes skipped, plus Changes applied.

Upgrade History Details form section

Field	Description
Updated and different	Number of Upgrade Detail records for which the value of the disposition is updated, and the value of changed is equal to true. This does not appear on the form by default.
Updated and not different	Number of Upgrade Detail records for which the value of the disposition is updated, and the value of changed is equal to false. This does not appear on the form by default.
Skipped and different	Number of Upgrade Detail records for which the value of the disposition is skipped, and the value of changed is equal to true. This does not appear on the form by default.
Skipped and not different	Number of Upgrade Detail records for which the value of the disposition is skipped, and the value of changed is equal to false.

Upgrade History Details form section (continued)

Field	Description
	This does not appear on the form by default.

Review Skipped Records form section

The Skipped Changes to Review related list displays each record that was skipped during the upgrade process. Use the list to review the reason for each skipped record in the list and then either merge your customization or revert your customization to the base system.

Skipped Changes to Review related list

To prevent your customizations from being overwritten during system upgrades, the upgrade process skips (does not apply the update to) objects that have been customized. To assist you in tracking and resolving skipped update records that need review, Skipped Changes to Review lists all updates skipped during the upgrade process.

By default, the list is filtered by Disposition Skipped Error, Disposition Skipped Manual Merge, Disposition Skipped, or Disposition Skipped (second pass), and the resolution status is either empty or not reviewed.

Navigate to **System Diagnostics** > **Upgrade History** to view the Skipped Changes to Review related list.

Skipped Changes to Review related list

Field	Input Value
File name	Name of skipped Upgrade Detail record.
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations. • Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Claim Status	Outcomes:

Skipped Changes to Review related list (continued)

Field	Input Value
	• N/A: There are no claims on the record (no need to review) • Won conflict: The loading application loaded the file but is in conflict with the previously loaded file. Recommendation: Review the two versions of the file to confirm your intended outcome. • No conflict: There may be one or more applications loading but there is no conflict. However, it is still good to review the outcome to see if it's what you intended. • Lost conflict: The file was not loaded because another application has a better claim on the file. Recommendation: Review the two versions of the file to see which should be loaded (or if they should be merged).
Priority	Relative importance of the conflict that caused the skip based on the following criteria: • 1: UI pages, UI macros, and more • 2: Business rules, security ACLs, and more • 3: Reports and more • 4: Form sections, choice sets, and more • 5 Everything else
Resolution	• Not reviewed: Records which have not been reviewed • Reviewed: Records which have been reviewed • Reviewed and Merged: Records which have been reviewed and have both the old and new changes • Reviewed and Retained: Records which have been reviewed and have retained the updates from the latest upgrade • Reviewed and Reverted: Records which have been reviewed and have reverted changes to the base system
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.
Type	Current file type (such as Business Rule or UI Policy).
Table	Table that contains the record.

Resolve a skipped update and set a resolution status

To prevent your customizations from being overwritten during system upgrades, the upgrade process skips (does not apply the update to) objects that have been customized. One of your responsibilities as the administrator is to resolve each skipped update after an upgrade.

Before you begin

Role required: admin

About this task

You resolve an update by either retaining the customization or (when the software upgrade contains a feature that you would like to implement) by merging or overwriting the customization with the base system update.

Note:
Objects that are customized and that did not change in the base system since the last upgrade require no action on your part.

When an object is customized, the system adds a corresponding record to the Customer Updates [sys_update_xml] table and then maintains current version information for all customized objects. The upgrade process skips changes to objects that have a current version in the Customer Updates table. When you follow the procedure, you perform one of the following actions:

- Retain (keep) a customization with no changes
- Retain a customization by merging changes from the updated object
- Revert a customized object to the updated version (that is, overwrite the customization)
- Review the skip and perform no action on the object

Procedure

1. Navigate to **All** > **System Diagnostics** > **Upgrade History**.
2. Select the desired software version.
3. In the Skipped Changes to Review related list, select the update record to resolve.

   Note:
   By default, the list displays records with disposition of Skipped and resolution status of Not Reviewed.

4. Review the list of changes.
   For text fields, you can click in the field to open the Diff/Merge tool. Review the differences. Click a text box to view and edit the detailed differences.
5. Perform one of the following actions.
   You have the option to add a **Comment** to any record, for example, to explain the action to future reviewers.
6. Click **Update**.
7. Repeat the process to resolve each update record in the list.

Result

Only skipped updates with a **Resolution Status** of **Not Reviewed** or without a resolution set appear in the Skipped Changes to Review related list. Any action you take that changes the **Resolution Status** to a value other than **Not Reviewed** or **--None--** removes the skipped update from list and moves it to the Skipped Changes Reviewed related list.

Related topics

Upgrade Details related list

Skipped Changes to Review related list

Skipped Changes Reviewed related list

Skipped Changes Reviewed related list

Skipped Changes Reviewed lists update records that previously appeared on the **Skipped Changes to Review** related list and have been reviewed. When you select a skipped record to review and set a **Resolution Status** to a value other than Not Reviewed, the update record moves to the Skipped Changes Reviewed related list.

Navigate to **System Diagnostics** > **Upgrade History** to view the Skipped Changes Reviewed related list.

Skipped Changes Reviewed related list

Field	Input Value
File name	Name of reviewed Upgrade Details record.
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations. • Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Claim Status	Outcomes: • N/A: There are no claims on the record (no need to review) • Won conflict: The loading application loaded the file but is in conflict with the previously loaded file. Recommendation: Review the two versions of the file to confirm your intended outcome.

Skipped Changes Reviewed related list (continued)

Field	Input Value
	• No conflict: There may be one or more applications loading but there is no conflict. However, it is still good to review the outcome to see if it's what you intended. • Lost conflict: The file was not loaded because another application has a better claim on the file. Recommendation: Review the two versions of the file to see which should be loaded (or if they should be merged).
Priority	Relative importance of the conflict that caused the skip based on the following criteria: • 1 (highest priority): xml content • 2: script or script_plain • 3: html content • 4: sys_ui_form_section, sys_ui_related_list, or sys_choice_set • 5 (lowest priority): other
Resolution	• Reviewed • Retained • Reverted
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.
Type	Current file type (such as Business Rule or UI Policy).
Table	Table that contains the record.
File name	
File differences	Comparison of the file in the upgrade with the customized version.
Changed by vendor	Indicates whether the file has been changed by the vendor since the last upgrade.

Customizations Unchanged related list

Customizations Unchanged lists all records that were skipped (due to a customization), but the changes that were going to be applied in this upgrade have not changed from the last upgrade.

Navigate to **System Diagnostics** > **Upgrade History** to view the Customizations Unchanged related list.

Customizations Unchanged related list

Field	Input Value
File name	Name of unchanged sys_upgrade_history_log record. Select it if you want to access Upgrade Details and add comment text, or set a resolution status for it.
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations. • Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Priority	Relative importance of the conflict that caused the skip based on the following criteria: • 1 (highest priority): xml content • 2: script or script_plain • 3: html content • 4: sys_ui_form_section, sys_ui_related_list, or sys_choice_set • 5 (lowest priority): other
Resolution	• Reviewed • Retained • Reverted
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Changed by vendor	Indicates whether the file has been changed by the vendor since the last upgrade.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.

Customizations Unchanged related list (continued)

Field	Input Value
Type	Current file type (such as Business Rule or UI Policy).
Table	Table that contains the record.

Revert a customization

To prevent customizations from being overwritten by system upgrades, the upgrade process automatically skips changes to objects that have been customized. You may want to overwrite your customizations when a software upgrade contains a feature that you would like to implement.

Before you begin

Role required: admin

About this task

To identify customized objects, the system adds a corresponding record in the Customer Updates [sys_update_xml] table. The table maintains the current version information for all objects that have been customized. The upgrade process skips changes to objects that have entries in the table. The upgrade process does not skip objects if only excluded fields have changed.

Procedure

1. Navigate to **All** > **System Diagnostics** > **Upgrade History**.
2. Select the desired software version.
3. Filter the Upgrade Details related list by **Disposition is Skipped**.
4. **Optional:** Add another filter condition for **Changed is True** to return only the objects that have changed since the last upgrade.
5. Select the update record to implement.
   The **File differences** field displays a side-by-side comparison of the customization and the default version. Deletions are highlighted in red, additions in green, and modifications in yellow.
6. Click **Revert to base system** to overwrite your customized object with the system default version.
   - The **Disposition** changes from **Skipped** to **Reverted**.
   - After you revert a customization, you have the option to click **Reapply Changes** to reapply your customizations (undo the revert).
   - During replace on upgrade, no changes are made to the update_exempt fields like the **Active** field. The rest of the record receives the update.

     Note:
     If the attribute is not explicitly specified as an update_exempt field, the **Active** field on a tracked table is treated as update_exempt by default.

   Since we skip updating update_exempt fields, two versions of the record in the sys_update_version table are generated.

- The version record with the source '_to_platform_version_' contains the payload from ServiceNow without any modifications. (State= HISTORY)
- The version record with the previous customization source, for example, '_update_set_', that has the customer's customizations overridden by the values from the incoming ServiceNow version, except for the update_exempt fields. The customer's customizations for update_exempt fields are preserved in this version. (State= CURRENT)

Related topics

Update set administration

Changes Applied related list

Changes Applied lists all changes that were applied in this upgrade.

This list is ordered by priority and displays Upgrade Detail records for this upgrade that have a disposition of Updated, Updated (second pass), Inserted, Inserted (second pass), Deleted, Deleted (second pass), and for which the changed flag is true for all of the dispositions.

Navigate to **System Diagnostics** > **Upgrade History** to view the Changes Applied related list.

Changes Applied related list

Field	Input Value
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations. • Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Priority	Relative importance of the conflict that caused the skip based on the following criteria: • 1 (highest priority): xml content • 2: script or script_plain • 3: html content

Changes Applied related list (continued)

Field	Input Value
	• 4: sys_ui_form_section, sys_ui_related_list, or sys_choice_set • 5 (lowest priority): other
Resolution	• Reviewed • Retained • Reverted
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.
Type	Current file type (such as Business Rule or UI Policy).
Table	Table that contains the record.
File name	Current upgrade file name.
File differences	Comparison of the file in the upgrade with the customized version.
Changed by vendor	Indicates whether the file has been changed by the vendor since the last upgrade.

Claim Outcomes to Review related list

During system upgrades, the Claim Status tab displays outcomes to review and resolve. See the Claim Status field in the table below.

By default, the list is filtered by Disposition Skipped Error, Disposition Skipped Manual Merge, Disposition Skipped, or Disposition Skipped (second pass), and the resolution status is either empty or not reviewed.

Navigate to **System Diagnostics** > **Upgrade History** to view the Claimed Outcomes to Review related list.

Claimed Outcomes to Review related list

Field	Input Value
File name	Name of skipped Upgrade Detail record.
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations.

Claimed Outcomes to Review related list (continued)

Field	Input Value
	• Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Claim Status	Outcomes: • N/A: There are no claims on the record (no need to review) • Won conflict: The loading application loaded the file but is in conflict with the previously loaded file. Recommendation: Review the two versions of the file to confirm your intended outcome. • No conflict: The installing application loaded the file without any conflicts; however, there may be one or more applications loading the same file so it is still recommended to review the outcomes. • Lost conflict: The file was not loaded because another application has a better claim on the file. Recommendation: Review the two versions of the file to see which should be loaded (or if they should be merged).
Priority	Relative importance of the conflict that caused the skip based on the following criteria: • 1: UI pages, UI macros, and more • 2: Business rules, security ACLs, and more • 3: Reports and more • 4: Form sections, choice sets, and more • 5 Everything else
Resolution	• Not reviewed: Records which have not been reviewed • Reviewed: Records which have been reviewed • Reviewed and Merged: Records which have been reviewed and have both the old and new changes • Reviewed and Retained: Records which have been reviewed and have retained the updates from the latest upgrade • Reviewed and Reverted: Records which have been reviewed and have reverted changes to the base system

Claimed Outcomes to Review related list (continued)

Field	Input Value
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.
Type	Current file type (such as Business Rule or UI Policy).
Table	Table that contains the record.

Upgrade Details related list

Upgrade Details lists all Upgrade Details records for this upgrade.

Navigate to **System Diagnostics** > **Upgrade History** to view the Upgrade Details related list.

Upgrade Details related list

Field	Input Value
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations. • Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Priority	Relative importance of the conflict that caused the skip based on the following criteria: • 1 (highest priority): xml content • 2: script or script_plain • 3: html content

Upgrade Details related list (continued)

Field	Input Value
	• 4: sys_ui_form_section, sys_ui_related_list, or sys_choice_set • 5 (lowest priority): other
Claim Status	Outcomes: • N/A: There are no claims on the record (no need to review) • Won conflict: The loading application loaded the file but is in conflict with the previously loaded file. Recommendation: Review the two versions of the file to confirm your intended outcome. • No conflict: There may be one or more applications loading but there is no conflict. However, it is still good to review the outcome to see if it's what you intended. • Lost conflict: The file was not loaded because another application has a better claim on the file. Recommendation: Review the two versions of the file to see which should be loaded (or if they should be merged).
Resolution	• Reviewed • Retained • Reverted
Comment	During the process of resolving a skipped update, you have the option to add a **Comment** to any record. For example, the comment might explain the action that you took to future reviewers.
Target name	Name of the record corresponding to the current file.
Plugin	Plugin that contains the record.
Type	Current file type (such as Business Rule or UI Policy).
Table	Table that contains the record.
File name	Current upgrade file name.
Changed by vendor	Indicates whether the file has been changed by the vendor since the last upgrade.

Related topics

Upgrade History module: Track every upgrade

Resolve a skipped update and set a resolution status

Upgrade Monitor module: Upgrade an individual instance

The Upgrade Monitor helps you upgrade an individual instance. You can monitor the progress of an upgrade and resolve conflicts between the upgrade and customizations.

Set up	Use	Troubleshoot and get help
Upgrade Monitor overview	• Monitor an upgrade to an instance • Process the skipped records list • Resolve conflicts for an individual record	• Debug upgrade for last upgrade session. • Ask or answer questions in the Developer Community • Search the Known Error Portal for known error articles • Contact Customer Service and Support

Upgrade Monitor overview

The Upgrade Monitor helps you upgrade an individual instance. You can monitor the progress of an upgrade and resolve conflicts between the upgrade and customizations.

Set up for the Upgrade Monitor

In Washington DC, the Upgrade Monitor is installed by default.

If you log in with the admin role while an upgrade is underway, the system automatically displays the Upgrade Progress screen. If no upgrade is in progress, you can navigate to **System Diagnostics** > **Upgrade Monitor**.

How the Upgrade Monitor fits into the upgrade process

Note:
For detailed information about the upgrade process, see Upgrade your instance.

The Upgrade Monitor concerns only part of the larger upgrade process:

1. Clone the production instance to a test instance and a non-production instance.
2. Apply the upgrade to the non-production instance.
3. On the upgraded non-production instance, Process the skipped records list
4. Test the non-production instance to confirm that the instance still works and performs adequately. Compare to benchmark data from pre-upgrade production instance.
5. Apply the upgrade to the test instance. Import the update sets created on the non-production instance when you processed the skipped list. Repeat the testing to make sure that the process is working.
6. Apply the upgrade to the production instance. Import the update sets created on the non-production instance when you processed the skipped list. Test to confirm that the instance works and performs adequately.

Within this larger process, the Upgrade Monitor helps you upgrade individual instances:

- during the upgrade, it shows where in the process the system is
- after the upgrade, it reports what the upgrade did and for how long

- as you upgrade the first non-production instance, it helps you resolve conflicts between customizations and changes that are part of the upgrade
- on non-production instances, it provides information that can help you estimate how long the upgrade takes on the production instance.

Note:
Debug Upgrade provides detailed debugging output for transactions containing artifacts affected by the most recent upgrade, and is designed to assist in upgrade error resolution. See Debug upgrade.

Monitoring an individual instance as it upgrades

While the upgrade is in progress the Upgrade Progress shows what the upgrade process has done, what it is doing, and what remains to be done.

When the upgrade completes, the system displays the Upgrade Summary Report. The Upgrade Summary Report provides information about conflicts between customizations versus the upgrade and provides a link to reconcile these conflicts. For information about understanding and resolving these conflicts, see Resolve Conflicts form.

When you upgrade a non-production instance, the Upgrade Summary Report can help you estimate how long the same upgrade takes on a production instance. For details about the elements on this report and how to use this information, see Upgrade Progress.

Resolving conflicts

To prevent losing customizations, the system skips upgrading records you have customized and provides you with a list of these skipped records.

As you upgrade your first non-production instance, go through the Skipped Changes to Review related list and resolve these conflicts. The system records the changes you make during this process in update sets.

You do not need to reconcile the skipped list on any instances you later upgrade. Instead, you can apply the upgrade then import the update sets containing your changes.

For details on reconciling conflicts, see Process the skipped records list.

Factors affecting upgrade duration

Various factors affect how long the system takes to perform an upgrade. The Upgrade Monitor can help you understand those factors and estimate how long the upgrade to your production instance takes.

Many factors can affect the duration of the upgrade process:

- The number of records in the database
- The number of customizations in the database
- The number of nodes in the instance
- The size of tables in the instance that require a schema-change in the instance.
- The number of fix scripts required and the size of the tables those fix scripts manipulate.

Upgrading a non-production instance can help you estimate how long the upgrade takes on production, but differences between the instances can significantly affect the duration:

- When you cloned the production instance to the non-production instance, you may have clone-excluded some tables. This reduces the size of the database and makes the non-production instance upgrade faster than the production instance.
- The production instance may have more memory and processing power.
- The production instance may have more nodes than the non-production instance.

After upgrading the first non-production instance, examine the Upgrade Summary Report for data to help estimate the impact of these factors.

Note:
The performance of the upgrade engine is significantly enhanced and optimized for all types of upgrades (in-family, and family-to-family) in Madrid and later releases.

Upgrade Monitor

When an upgrade is not running, the Upgrade Monitor displays information about the next check for an available upgrade.

The Upgrade Monitor displays the next date and time when the system will check for an available upgrade.

Upgrade Monitor

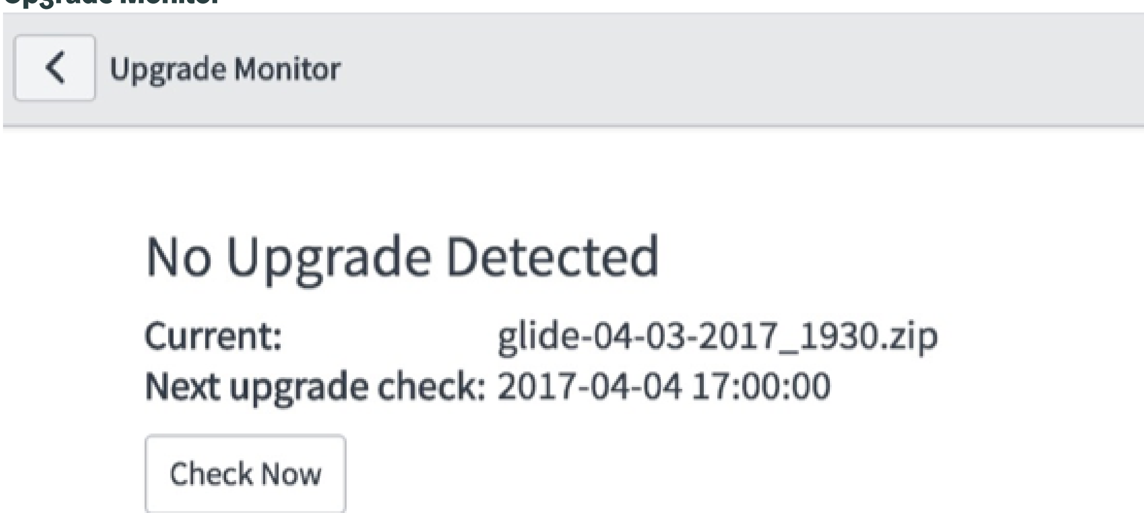

To check immediately for an available upgrade, click **Check Now**.

If one or both of the triggers for upgrading the system ('Check distribution for possible upgrade' and 'Check database for possible upgrade') have been customized or are missing, the Upgrade Monitor displays a warning and provides a button for resolving the issues.

Note:
If your instance is self-hosted (not hosted by ServiceNow) this message may not necessarily indicate a problem. If you have customized or disabled the upgrade job and want to keep that customization or disabled state, do not click the button to fix the upgrade issue.

Upgrade Monitor with Issue Detected

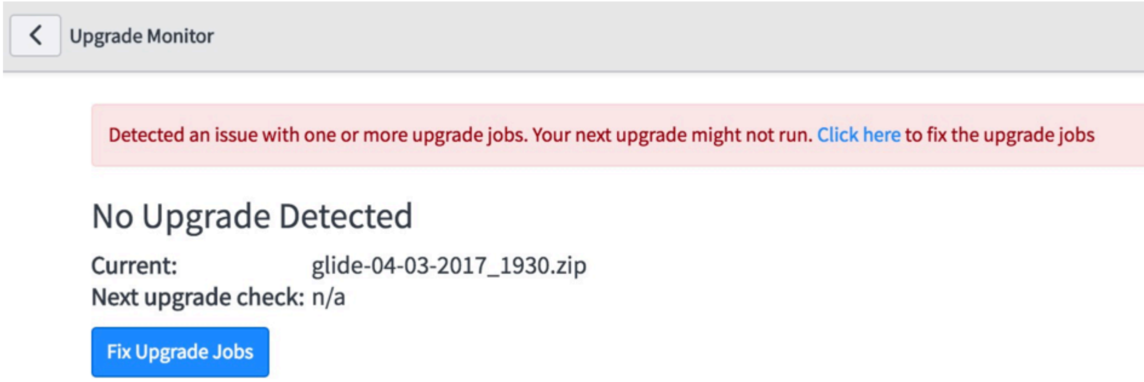

To resolve the issues with the upgrade jobs, click **Fix Upgrade Jobs**. This action reverts both upgrade triggers ('Check distribution for possible upgrade' and 'Check database for possible upgrade') to their base versions.

Note:
'Upgrade' job has been renamed to 'Check distribution for possible upgrade' starting Paris. 'Check Upgrade Script' job has been renamed to 'Check database for possible upgrade' starting Paris.

Monitor an upgrade to an instance

While the system is upgrading an instance, you can monitor its progress with the Upgrade Monitor. When the upgrade is done, you can view a summary of the results on the Upgrade Summary Report.

Before you begin

Role required: Admin.

Procedure

1. If necessary, navigate to **System Diagnostics** > **Upgrade Monitor**.
   When an upgrade is in progress, the Upgrade Progress page shows its status.
2. Monitor the progress of the upgrade.
   When the system finishes the upgrade, it displays the Upgrade Summary Report.

What to do next

From the Upgrade Summary Report, Resolve Conflicts form that prevented the system from upgrading records.

Upgrade Progress

When an upgrade is underway, Upgrade Progress displays progress bars and other information to help monitor the process.

Upgrade Progress

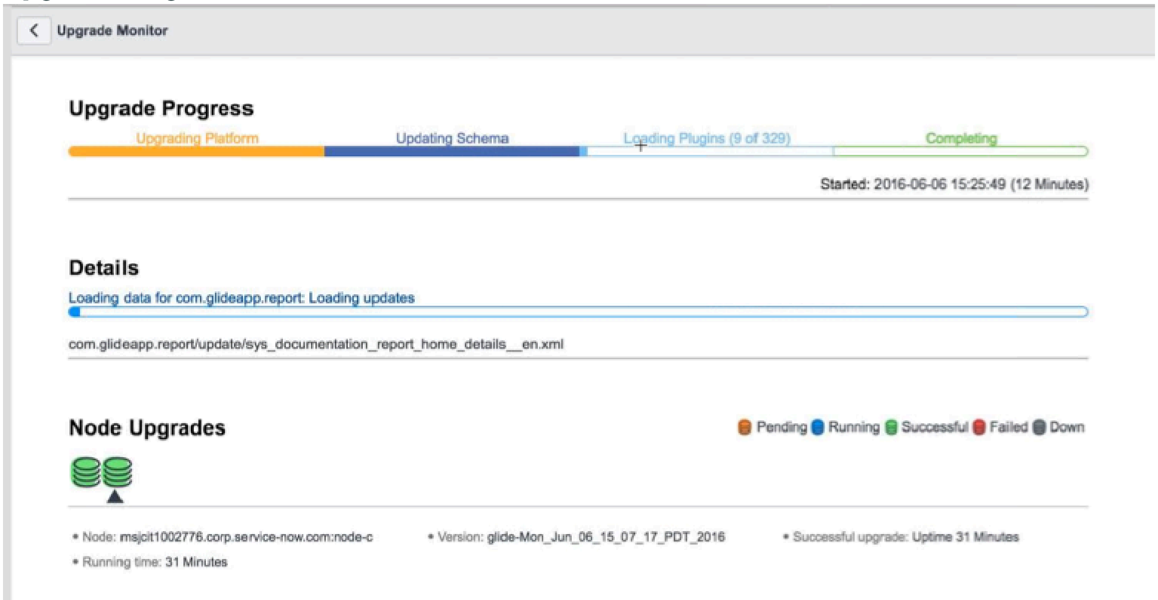

Screen elements

Screen element	Description
Upgrade Progress	A progress bar depicting where the upgrade is in the overall process. The length of a section in the bar does not indicate the relative duration of that process.
Upgrade Progress: Upgrading Platform	The system is applying upgrades to elements that form the foundation of the platform.
Upgrade Progress: Updating Schema	The system is scanning the plugins to create a list of tables that require upgrading. This prevents the system from upgrading the same table multiple times.
Upgrade Progress: Loading Plugins	The system is loading both core and optional plugins. Some features require more than one plugin, so the number of plugins listed may not match the number of optional features installed.
Upgrade Progress: Completing	The system is upgrading components that need to be completed after the previous three stages are done. **Note:** After the system finishes the Completing phase, it displays a separate screen showing Finalizing. The system tracks in update sets the changes made during the Upgrading Platform, Updating Schema, Loading Plugins, and Completing phases. It does not track changes made during Finalizing.
Details	Shows the current activity, a progress bar for the current activity, and the file currently being updated.
Node Upgrades	The color of the icons represents the status of each node during this upgrade: Pending, Running, Successful, Failed, or Down.

Screen elements (continued)

Screen element	Description
Node	The selected node indicated by the arrow. To change the selection, position the mouse cursor over the icon for the node to select.
Running time	How long the selected node (indicated by the arrow) has been running. If the selected node is offline, this value stops updating and shows how long a node was online before going offline.
Version	The current build for the selected node (indicated by the arrow).
Successful upgrade	When the selected node is online, shows how long the node has been online. If the selected node is offline, shows how long the node has been offline.

Upgrade Summary Report

This report summarizes the actions taken, provides tools to resolve conflicts between customizations and the upgrade, and provides information to help estimate time for upgrades to other instances.

Upgrade Summary Report: Database Upgrade Highlights

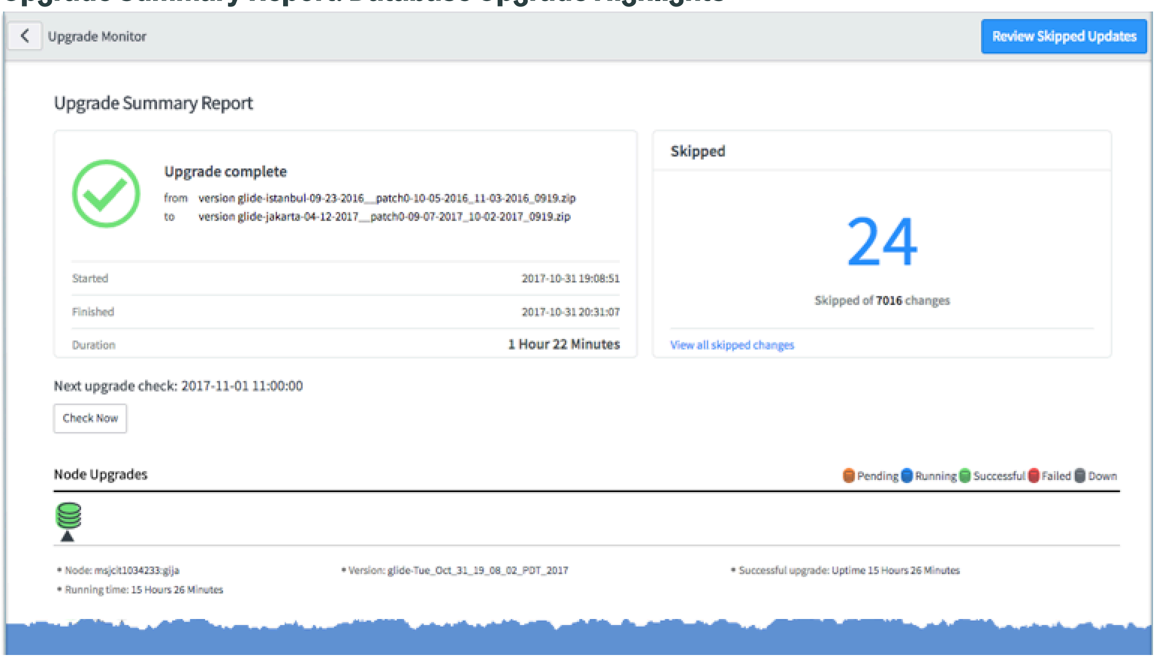

Screen elements

Screen element	Description
From version	Previous version of the instance
To version	Upgraded version of the instance
Started	When the upgrade process started
Finished	When the upgrade process finished

Screen elements (continued)

Screen element	Description
Duration	How long the upgrade process took
Skipped	In the example, `24 Skipped of 7016 changes` reads as: • **24 Skipped** - Total number of customizations that were skipped, skipped manual merge, or skipped in error, and the file changed in the distribution since the last upgrade of the instance. These are the number of records the system did not upgrade because of conflicts between customizations and the upgrade. • **7016 Changes** - Total number of file changes in the distribution since the last upgrade of the instance, plus all inserts and deletes. **Note:** The manner in which upgrade records are counted was changed in the Jakarta release, which may render different record counts than in previous releases.
Review Skipped Updates	Click to reconcile conflicts that caused the system to skip some updates

Node Upgrades

Upgrade Summary Report: Node Upgrades

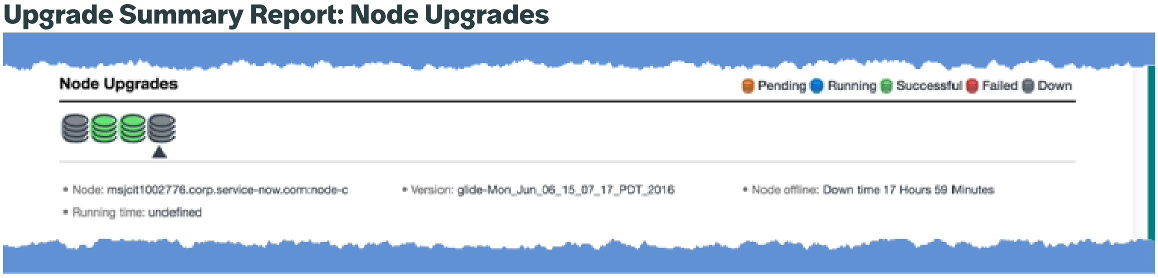

The Node Upgrades section shows the status of the upgrade for each node in the instance. The color of the icon denotes the status, as illustrated by the legend (key) and to the right of the node icons. To see details about a node, position the cursor above the icon for that node. An arrow points to the node selected, and the information below the icons pertains to that node.

Schema Changes to Clone-excluded Tables

Upgrade Summary Report: Schema Changes to Clone-excluded tables

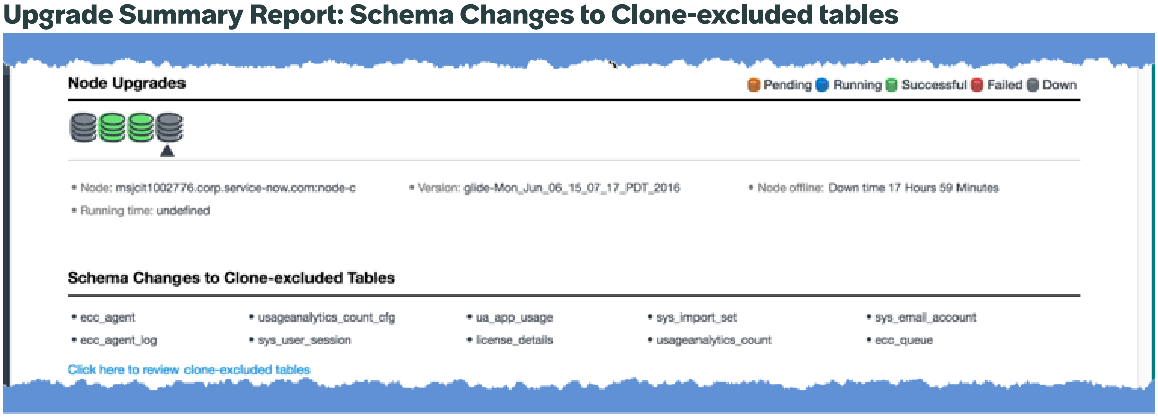

The Schema Changes to Clone-excluded Tables section shows a list of tables affected by the upgrade that were clone-excluded when you cloned the production instance to this instance. Because clone-excluded tables are empty, upgrading them takes less time than upgrading those same tables on the production instance. To estimate how much longer the production upgrade takes, note the size of the clone-excluded tables on the production instance.

Top 10 Fix Scripts by Duration

Upgrade Summary Report: Top 10 Fix Scripts by Duration

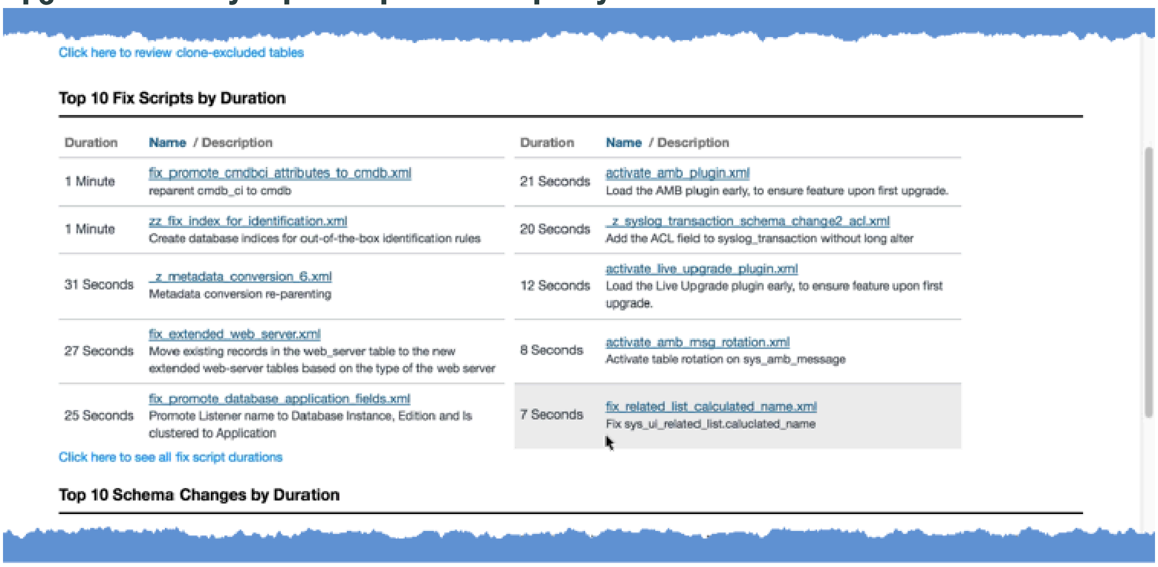

Click here to review clone-excluded tables

Top 10 Fix Scripts by Duration

Duration	Name / Description	Duration	Name / Description
1 Minute	fix_promote_cmdbci_attributes_to_cmdb.xml reparent cmdb_ci to cmdb	21 Seconds	activate_amb_plugin.xml Load the AMB plugin early, to ensure feature upon first upgrade.
1 Minute	zz_fix_index_for_identification.xml Create database indices for out-of-the-box identification rules	20 Seconds	_z_syslog_transaction_schema_change2_acl.xml Add the ACL field to syslog_transaction without long alter
31 Seconds	_z_metadata_conversion_6.xml Metadata conversion re-parenting	12 Seconds	activate_live_upgrade_plugin.xml Load the Live Upgrade plugin early, to ensure feature upon first upgrade.
27 Seconds	fix_extended_web_server.xml Move existing records in the web_server table to the new extended web-server tables based on the type of the web server	8 Seconds	activate_amb_msg_rotation.xml Activate table rotation on sys_amb_message
25 Seconds	fix_promote_database_application_fields.xml Promote Listener name to Database Instance, Edition and Is clustered to Application	7 Seconds	fix_related_list_calculated_name.xml Fix sys_ui_related_list.caluclated_name

Click here to see all fix script durations

Top 10 Schema Changes by Duration

The Top 10 Fix Scripts by Duration helps you understand which fix scripts required the most time.

Top 10 Schema Changes by Duration

Upgrade Summary Report: Top 10 Schema Changes by Duration

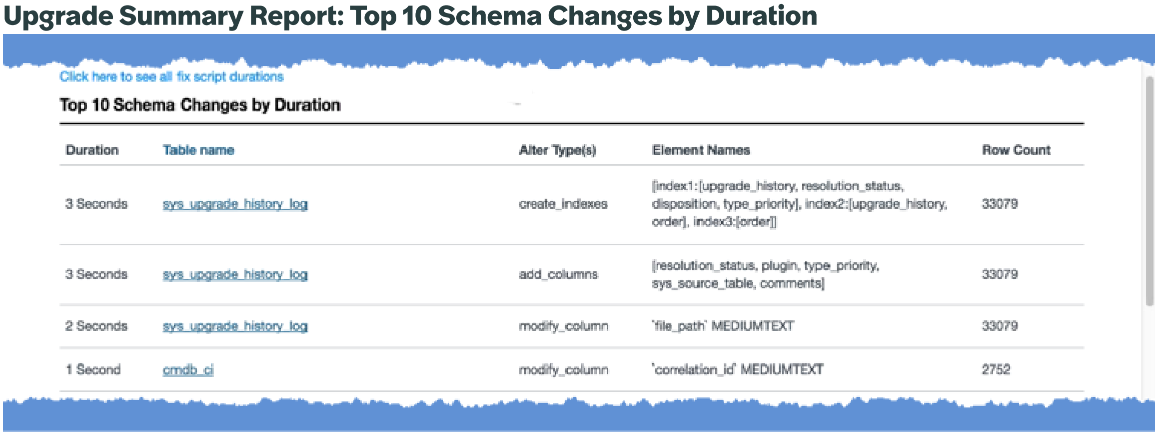

Click here to see all fix script durations

Top 10 Schema Changes by Duration

Duration	Table name	Alter Type(s)	Element Names	Row Count
3 Seconds	sys_upgrade_history_log	create_indexes	[index1:[upgrade_history, resolution_status, disposition, type_priority], index2:[upgrade_history, order], index3:[order]]	33079
3 Seconds	sys_upgrade_history_log	add_columns	[resolution_status, plugin, type_priority, sys_source_table, comments]	33079
2 Seconds	sys_upgrade_history_log	modify_column	\`file_path\` MEDIUMTEXT	33079
1 Second	cmdb_ci	modify_column	\`correlation_id\` MEDIUMTEXT	2752

The Top 10 Schema Changes by Duration helps you understand which schema changes required the most time.

Process the skipped records list

If you customized or altered a record affected by this upgrade, such as a business rule or script, the upgrade generates a skip log record. You must resolve the differences between the upgraded and customized versions of the record by processing the skipped record list, or, in other terms, just processing the skipped list.

Before you begin

Role required: admin

About this task

Review the changes you made to baseline records, such as business rules and scripts, that appear on the skipped records list and revert to the baseline version if appropriate. Post-upgrade, thoroughly test all changes you made to these records.

Procedure

1. If necessary, navigate to **System Diagnostics** > **Upgrade Monitor**.
   If the upgrade is still in progress, the system displays the Upgrade Progress screen. When the upgrade finishes, the system displays the Upgrade Summary Report.
2. After the system displays the Upgrade Summary Report, click the **Click here** link in the Skipped box.
   The system displays the System Upgrades form.
3. Navigate to Review Skipped Records section and – if necessary – scroll to the Skipped Changes to Review related list.
4. Click the row for the first record you want to reconcile.
   The system displays the Upgrade Details form for that record.
5. Evaluate how you want to resolve the conflict for this record and take the appropriate action.

   Note:
   The system tracks changes to records in an update set so you can apply these changes to another instance later. However, the system does not migrate the upgrade details records from one instance to the next. These records apply to a specific upgrade of a specific instance. If you want to preserve the Comments, Resolutions, or other information from the skipped list, export it from this instance.

6. In the **Comment** field, write the reasons for making your decision and other information you want to document.
7. Click **Update**.
   Post-upgrade, thoroughly test all changes you made to the records on the skipped record list.

Related topics

Debug upgrade

System Upgrades form

When an upgrade is complete, the System Upgrades form displays key statistics about the upgrade and a related list of skipped records (the skipped list).

System Upgrades form

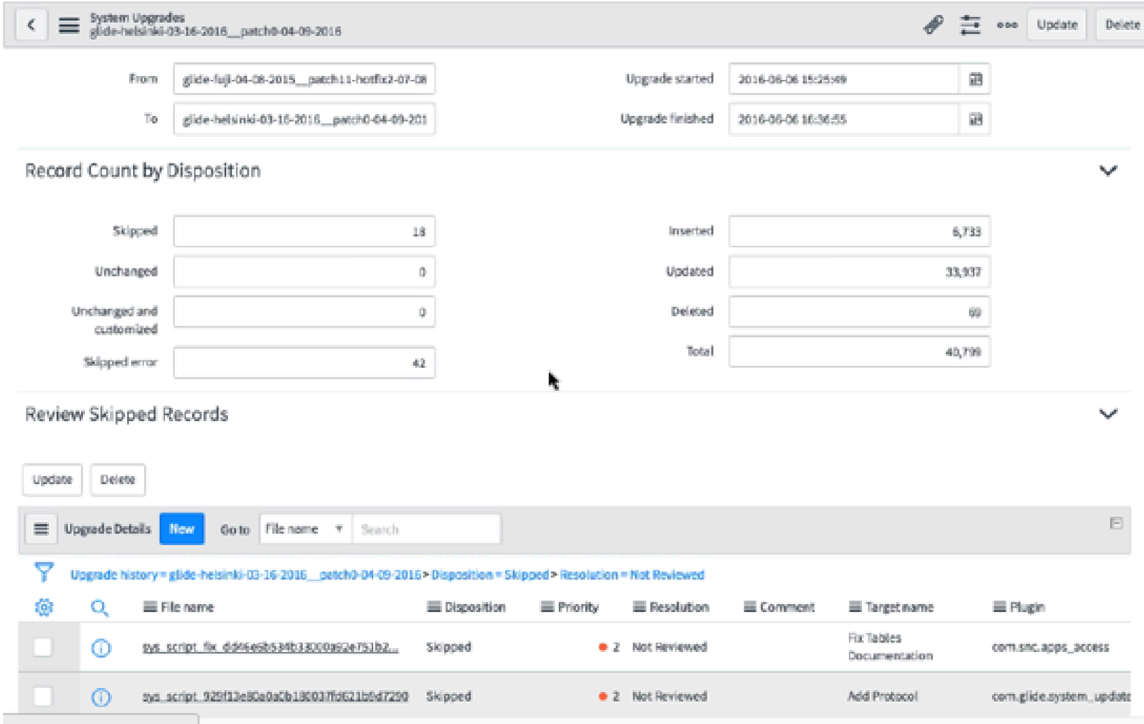

Fields

Screen element	Description
From	The version of the instance before upgrading.
To	The version of the instance after upgrading.
Upgrade started	The time upgrade started.
Upgrade finished	The time the upgrade finished.
Skipped	The number of records the system did not upgrade because of conflicts between customizations and the upgrade.
Unchanged	The number of records in the instance unchanged by the upgrade.
Unchanged and customized	The number of customized records in the instance unchanged by the upgrade.
Skipped error	The number of records the system did not upgrade because of one or more errors.
Inserted	The number of records the system inserted.
Updated	The number of records the system successfully updated.
Deleted	The number of records the system deleted.
Total	The total number of records the system inspected and/or processed for this upgrade.

Resolve conflicts for an individual record

Reconcile differences between your customized record and the changes associated with the upgrade.

Before you begin

Role required: admin

Procedure

1. From the Upgrade Details form for the record you are reconciling, click **Resolve Conflicts**.
   The system displays the Resolve Conflicts form, which highlights differences between the two versions of the record. The form displays information about the base system record on the left and the customized record on the right.
2. Compare the base system with the customized record for each field on this form, and for non-script fields, edit the customized record to include what you want from the base system and the customization.
3. If this record contains a script, check it for conflicts and resolve.
   a. Click inside the Script field.
      The system displays the Resolve Conflicts - Script form highlighting areas where the two versions of the script differ.
   b. Edit the script by clicking the small arrows that correspond to the blocks of code until script contains whichever content you want.
   c. Click **OK**.
      The system returns to the Resolve Conflicts form.
4. To save your changes to the record, click **Save Merge**.
   The system sets the **Resolution** for this record to **Reviewed and Merged**.

Related topics

- Compare to the current version
- Compare a pushed version to a local version
- Compare two versions of an article
- Resolve a collision in Team Development
- Revert a change
- View customizations and compare with current version

Upgrade Details form

From the Upgrade Details form, you can review an individual record affected by the upgrade and reconcile conflicts between the upgrade and customizations.

servicenow

Upgrade Details form

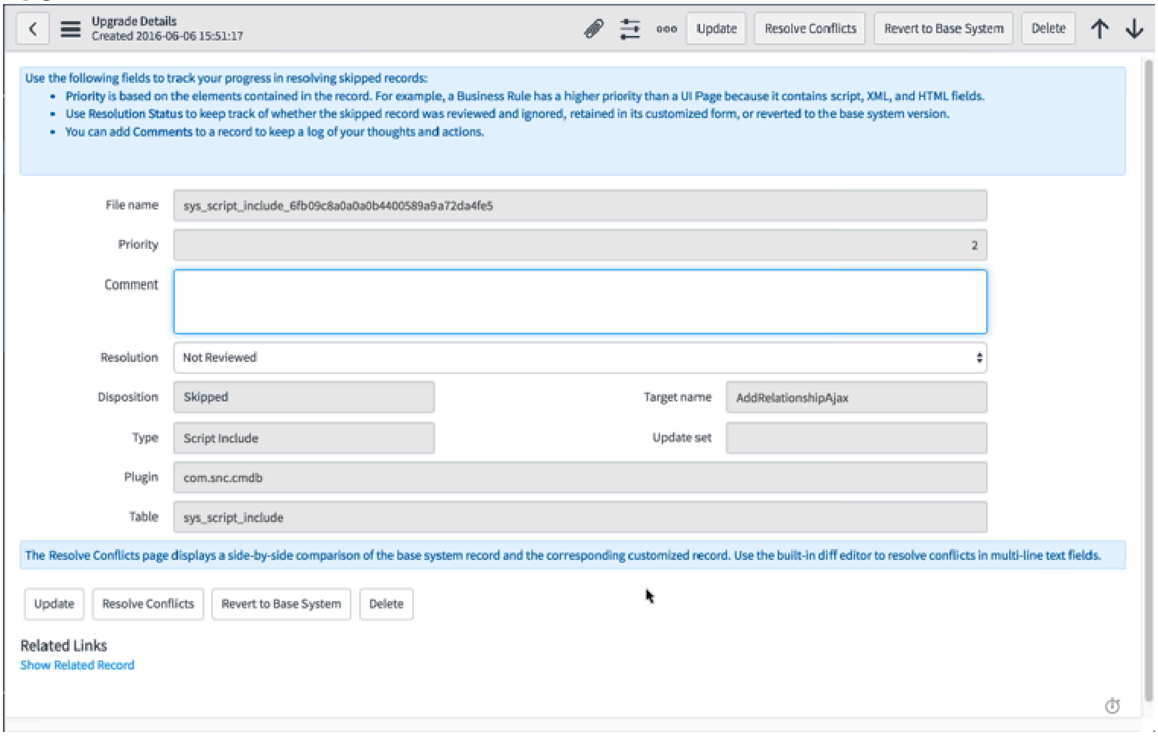

Fields

Screen element	Description
File name	The record the system has flagged as needing to be reconciled.
Priority	The priority the system has assigned to resolving this conflict. Values range from one to five, with one representing the highest priority.
Comment	Comments to document your decisions about reconciling this record.
Resolution	How you elected to resolve this conflict: • **Not Reviewed** • **Reviewed** - reviewed but no action yet taken • **Reviewed and Merged** - made changes to the record to reconcile the customized and upgraded versions • **Reviewed and Retained** - left customizations in place without update from upgrade • **Reviewed and Reverted** - customizations discarded, record updated according to upgrade For more information, see Process the skipped records list.
Disposition	Action performed on this file during the selected upgrade: • Inserted: The system inserted a new record. • Updated: The system updated this record. • Deleted: The system deleted this record. • Skipped: The system did not change this record in order to preserve customizations.

Fields (continued)

Screen element	Description
	• Reverted: This record was reverted to the base version. • Table not found: The system could not find the table that contains this record. • Unchanged: The system did not change this record because the baseline component has not changed since the last release. • Skipped Manual Merge: The system did not change this record because updating it requires manual intervention. • Skipped Apply Once: The system skipped this record because it had already applied an update from an xml file in the apply once folder. • Not Latest: The system applied a change, but this change was overwritten later during the same upgrade.
Type	The record type, for example Script include.
Target name	Name of the skipped record, if applicable.
Update set	Unused.
Plugin	The plugin containing this record.
Table	The table containing this record.

Resolve Conflicts form

The Resolve Conflicts form you compare to the base system version with the customized version of a record and reconcile the differences.

Resolve Conflicts form

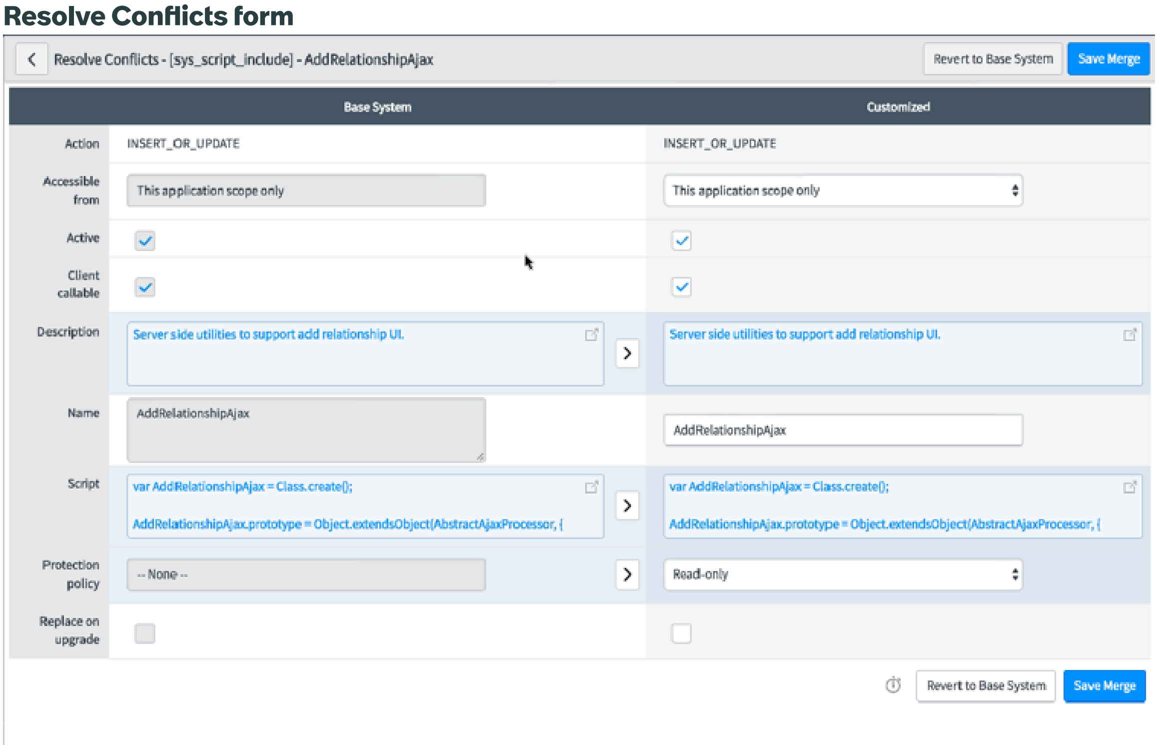

Fields

The fields this form displays depend on the type of record you are reconciling.

The left column shows the records fields in the base system, including the proposed changes that are part of the upgrade. The right column shows the fields for your customized record.

Debug upgrade

Diagnose and resolve issues you suspect may have resulted from the most recent upgrade by enabling upgrade debugging during a user session. Each transaction lists whether an artifact was skipped, customized by the customer, or modified by ServiceNow during the last family-to-family or patch version upgrade.

Administrators and technical personnel can use this output as an adjunct to the Upgrade History and Upgrade Monitor modules. Debug Upgrade is a component of standard ServiceNow session debugging. You enable Debug Upgrade, process a transaction (for example, load a form), then analyze the reported upgrade information. See .

Reported debug information

Debug Upgrade includes several expandable sections:

- Skipped During Last Upgrade
- Customer Customized
- ServiceNow Modified During Last Upgrade

These sections serve as a starting point for diagnosing when a function or artifact does not operate as expected after an upgrade. Each listing sorts by the dates the artifacts were processed, with the most recent appearing first. Expand the appropriate section, and click the name of a listed artifact to view the underlying artifact record. These records include update version history, which represents the state of a customizable object at specific times. See Version records.

The listed artifacts may require some kind of possible user intervention or action. Using underlying artifact detail, you can determine:

- what artifacts have changed?
- what artifacts have not been changed that would or should have?
- what changed artifacts were missed?
- what artifacts operate differently than expected (if referencing something that has changed during last upgrade)?

Depending on the answers to these questions, you may have to modify the selected artifact so it works in the way you expect. Conversely, you might simply accept the new updates made to the artifact, and its changed functionality.

Skipped during last upgrade

Lists artifact records that were skipped during the last upgrade that also ran during the last transaction. These records are standard objects that missed changes from ServiceNow because they were modified at some point by a user. See Resolve conflicts for an individual record.

Customer customized

Lists all artifact records that were customized at some point by users that also ran during the last transaction. This listing is not strictly related to the actual activities performed during the last

upgrade session. The listing does include the records that appear under Skipped During Last Upgrade. However, not all records that appear in the Customer Customized list are included in Skipped During Last Upgrade.

Note:
Customizations that you designate for replacement during upgrades, by setting the **Replace on Upgrade** field in the customization record to true, do not appear on the Customer Customized section list.

While this listing is not directly related to the activities performed during the last upgrade, it may be worthwhile to look at individual customization records listed on it. Some of the listed customization records may be referencing other records changed or skipped during upgrade. The activities performed during the last upgrade may have an impact on, or affect these relationships.

ServiceNow modified during last upgrade

Lists all standard (non-customized) artifact records modified by ServiceNow during the last upgrade that also ran during the last transaction.

Related topics

Upgrade History module: Track every upgrade

Upgrade Monitor module: Upgrade an individual instance

Enable and use debug upgrade

Enable Debug Upgrade to analyze post-upgrade issues that may require follow up, then disable it when finished with the upgrade debugging session.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Diagnostics** > **Session Debug**.
2. Select **Debug Upgrade** to enable the upgrade debugger.
3. Process a transaction (for example, load a form).
4. Click the **Skipped During Last Upgrade**, **Customer Customized**, or **ServiceNow Modified During Last Upgrade** header to view the appropriate list.

Upgrade Transaction List

Field	Description
Name	Name of record processed during the last upgrade session.
Type	Type of record processed during the upgrade. **Business rule** Server-side script that runs when a record is displayed, inserted, updated, or deleted, or when a table is queried. See Classic Business rules. **Client script** Client JavaScript run when client-based events occur. See Client scripts. **Data policy**

Field	Description
	Policies that enforce data consistency. See Data policy. **Script include** Script includes that store JavaScript run on the server. See Script includes. **UI action** Buttons, links, and context menu items on forms and lists. See UI actions. **UI macro** Discrete scripted components added to the user interface. See UI macros. **UI page** Pages used to create and display forms, dialogs, lists, and other UI components. See UI pages. **UI policy** Policies that change the behavior of information on a form and control custom process flows for tasks. See UI policies. **Note:** These are the only types of artifacts logged by the Upgrade Debugger.
Last Modified	Date and time the record was last modified.

5. Click **Show more** to view all artifacts.
6. Click the name of a listed artifact to view the underlying artifact record.
   For example, if you click a Script Include artifact, the underlying Script Include record appears.
7. To turn off Debug Upgrade and all session debugging after completing your session:
   a. Navigate to **System Diagnostics** > **Session Debug**.
   b. Select **Disable All** to disable the upgrade debugger.

Additional resources for ServiceNow AI Platform products and solutions

If you're looking for ServiceNow AI Platform best practices, troubleshooting, or other implementation guidelines, select a feature or resource type to discover ServiceNow resources on other relevant websites.

Note:
Many resources below require you to log in to sites like ServiceNow University, Now Create, or the ServiceNow Community. If the expected resource does not load, please log in and try to access the resource again.

Resource links

ServiceNow AI Platform features or products	Resource type	Resources
Admin Center	Getting Started	Getting started with Admin Center
Admin Center	FAQs	Admin Center FAQs
Admin Center	Troubleshooting	Troubleshooting for Admin Center
AI Search	Getting Started	Getting started with AI Search
AI Search	Best practices	Best practices for AI Search
AI Search	Release notes	AI Search release notes
AI Search	Training	AI Search: Indexing Explore AI Search
Application Manager	Getting Started	Getting started with Application Manager
Application Manager	Release notes	Application Manager release notes
Clone Admin Console	Release notes	Clone Admin Console release notes
Core Configuration	Release notes	Core ServiceNow AI Platform release notes
Data Management	Getting Started	Getting started with Data Management
Data Management	Best practices	Best practices for Data Management Recommended best practices for Data Management
Data Management	Tips and examples	Tips for Data Management
Domain Separation	Getting Started	Getting started with Domain Separation
Domain Separation	Best practices	Best practices for Domain Separation

Resource links (continued)

ServiceNow AI Platform features or products	Resource type	Resources
Domain Separation	Tips and examples	Domain Separation examples
Domain Separation	Troubleshooting	Troubleshooting for Domain Separation
Domain Separation	Training	Domain Separation Basics Domain Separation Implementation
Dynamic Translation	Tips and examples	Dynamic Translation examples
Dynamic Translation	Release notes	Dynamic Translation release notes
Instance Scan	Getting Started	Getting started with Instance Scan
Instance Scan	Training	Introduction to ServiceNow Instance Scan
Localization Framework	Release notes	Localization Framework release notes
Notifications	Best practices	Best practices for Notifications
Notifications	Tips and examples	Notifications tips
Notifications	Troubleshooting	Troubleshooting for Notifications
Notifications	Release notes	Notifications release notes
ServiceNow AI Platform	Training	System administrator Troubleshoot the ServiceNow AI Platform ServiceNow AI Platform Implementation ServiceNow Administration fundamentals on demand Performance Troubleshooting overview ServiceNow AI Platform Implementation On Demand

Resource links (continued)

ServiceNow AI Platform features or products	Resource type	Resources
		Process and Platform Owner Learning Path ServiceNow Administration Fundamentals on Demand
ServiceNow AI Platform	Getting Started	Introducing the ServiceNow AI Platform Getting started with the ServiceNow AI Platform Platform Academies overview ServiceNow System administrator overview ServiceNow Admin full course How to use ServiceNow - A comprehensive guide Setup the ServiceNow AI Platform and get upgrade ready ServiceNow Fundamentals ServiceNow Admin Complete Course|System Administrator Beginner
Platform Performance	Getting Started	Getting started with Platform Performance
Platform Performance	Best practices	Best practices for Platform performance
Subscription Management	Getting Started	Getting started with Subscription Management
Subscription Management	Release notes	Subscription Management release notes
System Clone	Getting Started	Getting started with System Clone
System Clone	Best practices	Best practices for System Clone
System Clone	FAQs	System Clone FAQs
System Clone	Tips and examples	System Clone tips

Resource links (continued)

ServiceNow AI Platform features or products	Resource type	Resources
System Clone	Release notes	System Clone release notes
System Localization	Getting Started	Getting started with System Localization
System Localization	Training	ServiceNow Localization Overview
Table Administration	Best practices	Best practices for Table Administration
Table Administration	Training	Table Administration and Security Overview
Upgrade Center	Getting Started	Getting started with Upgrade Center
Upgrade Center	Best practices	Best practices for Upgrade Center
Upgrade Center	FAQs	Upgrade Center FAQs
Upgrade Center	Tips and examples	Upgrade Center examples
Upgrade Center	Troubleshooting	Troubleshooting for Upgrade Center
Upgrade Center	Release notes	Upgrade Center release notes
Upgrade Center	Training	Upgrade Center overview Upgrade Skipped Records
User Administration	Best practices	Best practices for User Administration
Zing text indexing and search engine	Release notes	Zing text indexing and search engine release notes