onestream

SensibleAI Agents Admin Guide

PV9.3.0 SV1.3.0

Copyright © 2026 OneStream Software LLC. All rights reserved.

All trademarks, logos, and brand names used on this website are the property of their respective owners. This document and its contents are the exclusive property of OneStream Software LLC and are protected under international intellectual property laws. Any reproduction, modification, distribution or public display of this documentation, in whole or part, without written prior consent from OneStream Software LLC is strictly prohibited.

Table of Contents

Overview 1

Components Covered in This Guide 1

Administrative Model (High-Level) 2

Prerequisites (All Agents) 2

SensibleAI Infrastructure Setup 3

Installation 3

Setup 4

Documents Marketplace Solution Setup 5

Creating the Wernicke Service Account 5

Finance Analyst Configuration 7

Finance Analyst Agent Overview 7

Finance Analyst Setup 7

Scope 7

Accessing Finance Analyst Configuration 7

Required Setup Steps (Initial Enablement) 9

Optional (Recommended) Setup Steps 10

Finance Analyst Configuration Portal 11

Cube Connections 12

Purpose 12

Permissions 12

Initial Setup: Connect Cubes 13

Ongoing Management: Resync and Disconnect 13

Personas (Finance Analyst Personas) 14

Company Persona 14

Creating Personas 15

UD Types 16

Overview 16

Configure UD Types 16

Manage UD Types 18

Users 18

Overview 18
Permissions 19
Users Table 19
Assign a Default Persona 19
Assign a Default Set (Advanced) 20
Recommendations 20
Accounting Period 20
Overview 20
Cube Context 22
Overview 22
Cube Configuration 22
Dimension Defaults 23
Overview 23
Company & Personas 25
Default Sets (Advanced) 25
Preview Inheritance 26

Assign Users 26
Edit or Delete a Default Set 27
My Defaults 27
Dimension Members 28
Overview 28
Member Details 29
Bulk Actions 30
Business Context 31
Overview 31
Guidelines: What to Include 32
Guidelines: What Not to Include 32
Report Row Templates 34
Overview 34
Creating a Template 34
Recommendations 35

Finance Analyst Analysis Plans 36
Accessing Analysis Plans 36
Analysis Plans Page 37
Plan Detail 38
Creating a Plan 39
Editing a Plan 41
Scheduled Analyses 41
Activities 42
Example Questions (Finance Analyst) 44
Rules 44
Narrative Summarizer 45
End User Behavior (Admin-Relevant Summary) 45
Chat Settings 45
Query Results 46
SensibleAI Agents Documents Overview 47
Home Page and Global Options 47

Global Settings 48
Uninstall 49
Documents 49
Knowledge Clusters 50
Files 53
Activity 55
Data Connectors 58
Key Concepts 58
Source Configurations (Admin Only) 59
Connectors 61
Connector Activity 63
Reconnect 67
Sync 68
SharePoint Configuration 69
Confluence Configuration 69
Search Agent Overview 71

Cluster Configuration 72
Access Configuration 72
Analytics 72
Deep Analysis Agent Overview 72
AI Control Tower Monitoring 72
AI Control Tower Overview 74
Chat Explorer 75
Credits 75
Credit History 76
Credit Limits 77
Analytics 78
AI Platform Connections 82
Claude 82
ChatGPT 83
Microsoft Copilot 83
Gemini 84

Document Connector Sources 85
Evaluations 85
Overview 86
Purpose 86
Configurations 87
Runs 88
System 89
Upgrade Portal 89
Application Logs 90
Audit Logs 91
Document Storage 92
Settings 93
Application Settings 93
Roles 94
Agents Settings 95
ROI & Analytics 96

Support 96
Data Manipulator 97
Support Task Execution 97
Traces Explorer 97
OneStream MCP Servers in Gemini 98
Complete the Initial Steps 98
Gemini Enterprise Business 101
Gemini Enterprise Standard / Plus 111
Environment Upgrade 127
Step 1: Review Release Notes After the Upgrade Is Released 127
Step 2: Submit RITM 128
Step 3: Set up the Environment 129
Step 4: Run the Port 131
Step 5: Confirm the Environment is Working 132

Overview

SensibleAI Agents provide governed AI capabilities within OneStream. Agents are configured and secured using OneStream roles and security groups, and are governed through administrative tooling (AI Control Tower) and documentation infrastructure (SensibleAI Agents Documents).

Components Covered in This Guide

- **Finance Analyst Agent**: Finance Analyst enables finance users to pull and analyze data from their OneStream cubes using natural language. The agent interprets the users request and can pull existing cube views or create its own dynamic reports to best fulfill the user's intent.
- **Finance Analyst Configuration Portal**: Provides a dedicated administrative interface for configuring the Finance Analyst Agent, including cube connections, personas, user assignments, dimension defaults, business context, and report row templates.
- **Search Agent**: Retrieves and synthesizes content from curated documentation stored in Knowledge Clusters.
- **Deep Analysis Agent**: Enables workflow-based automation for document-heavy analysis.
- **SensibleAI Agents Documents**: Supports cluster-based document ingestion, approval workflows, security access, personas, and external connectors.
- **AI Control Tower**: Serves as the central governance portal for analytics, credits, access, and auditing.

Administrative Model (High-Level)

Administration is performed across two primary surfaces:

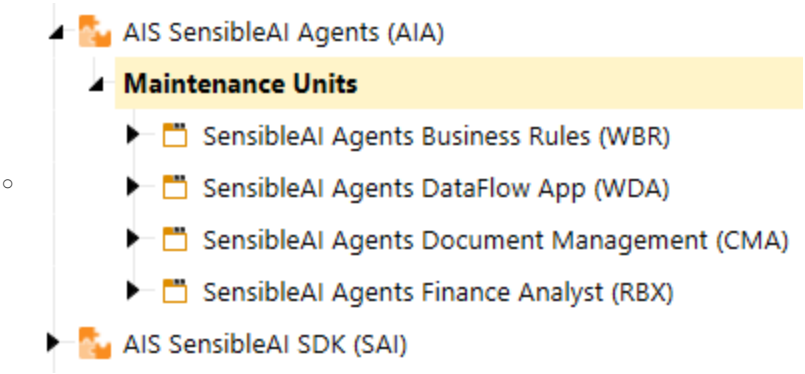
- **AI Control Tower**: Manages governance, access mapping, usage analytics, credit limits, auditing, and platform settings.
- **SensibleAI Agents Documents**: Supports document ingestion and curation for Search and Deep Analysis Agents.
- **Finance Analyst**: Inherits metadata access from existing OneStream Security settings.

Prerequisites (All Agents)

Ensure the following requirements are met before configuring agents:

- AI Agents have been successfully deployed to the environment.
- The administrator meets one of the following criteria:
    - Is a OneStream Administrator, or
    - Is assigned the SensibleAI Admin role in the AI Control Tower
- The following maintenance units are installed:
    - AIS SDK (SAI) Workspace
    - AIS SensibleAI Agents (AIA)
        - SensibleAI Agents Business Rules (WBR)
        - SensibleAI Agents DataFlow App (WDA)

- SensibleAI Agents Documents Configuration (CMA)
- SensibleAI Agents Finance Analyst Configuration (RBX)



- SensibleAI Agents Documents has been opened and the initial solution setup steps have been completed.

> **IMPORTANT:** Incomplete prerequisites may prevent configuration pages from appearing or block setup tasks from completing.

SensibleAI Infrastructure Setup

Installation

The *Getting Started with Solutions Guide* provides general instructions to install and configure OneStream Solutions. It also includes the recommended display settings, package file name information, advice for solution database migration, and considerations for modifications to solutions.

See the *Getting Started with Solutions Guide*.

Setup

Once the solution infrastructure has been deployed and the required Application Workspaces have been installed, Administrators will see the SensibleAI logo in the top right-hand corner of the OneStream interface (Windows App or Browser UX).

1. Click **Build Infrastructure** and wait for the Infrastructure setup succeeded message.
2. Select or type in the default application name. This is also the application the marketplace solutions were loaded to.

   > **NOTE:** This only captures the Groups in an environment. After setup, SensibleAI Agents are available in any application in the environment. To use Finance Analyst in other applications, the marketplace solutions must be loaded and Finance Analyst reconfigured.

3. Fill out the Roles and Agent Access settings, then save and continue:
   a. SAI Administrator: All Viewer permissions; can manage credit limits; full access to Settings (Application Settings, Roles, Agents Settings, and ROI & Analytics).
   b. SAI Support: All Administrator permissions; elevated access to Data Explorer (Explore); full access to the System section, including Document Storage, Audit Logs, and Upgrade Portal, as well as the Support section (Data Manipulator, Support Task Execution, and Traces Explorer).
   c. SAI Control Tower View Only: Can view Analytics, Credit Overview, Credit History, Explore (Chat Reviews and Interaction Explorer), and the Control Tower dashboard.
4. Sync users, then click Split Evenly to distribute credits.

> **NOTE:** Split Evenly distributes the allotted credits evenly among all OneStream users in the environment. Per-user credit allocation can be edited later in Control Tower after environment setup.

5. Click **Complete Setup**.

After completing setup, the agents are available in the slide-out panel.

Documents Marketplace Solution Setup

Navigate to the SensibleAI Agents Documents solution from OnePlace and open the library page to confirm it loads correctly. If the solution hasn't been set up yet, a setup process will guide you through the following steps:

1. Create the tables required for the solution to function.
2. Configure the backend and chat endpoints:
    a. Backend endpoint URL format: *https://hostname-agents-envtype.onestreamcloud.com/wernickebackend*
    b. Chatbot endpoint URL format: *https://hostname-agents-envtype.onestreamcloud.com*
3. Save the endpoint configuration, then launch the solution to complete setup.

Creating the Wernicke Service Account

Before creating this account, explain the need to the customer and obtain their explicit approval to proceed.

Function: wernicke.server is a background service account that Wernicke (the SensibleAI Agents engine) uses to run scheduled, non-interactive jobs — such as syncing OneStream user tables to Control Tower — inside SensibleAI Agents. It requires Administrator group membership so it can access cubes, data adapters, and admin-only endpoints without hitting permission errors.

1. Navigate to Security.
2. Create the user account and fill in the following details:
    a. **Name**: wernicke.server
    b. **Description**: Wernicke Service Account
    c. **User Type**: Interactive
    d. **External Authentication Provider**: OIDC Compliant Sign On
    e. **External Provided User Name**: wernicke.server
    f. Do not touch or change the password.
3. Click **Add Groups**, select Administrators, and move it over with the arrow. Click **OK**.
4. Verify all fields are correct.

Finance Analyst Configuration

Finance Analyst Agent Overview

Finance Analyst enables finance users to pull and analyze data from their OneStream cubes using natural language. The agent interprets the users request and can pull existing cube views or create its own dynamic reports to best fulfill the users intent.

This section covers:

- Initial setup required before users can access Finance Analyst.
- Ongoing administration using the Finance Analyst Configuration Portal.

Administrative controls that affect end-user behavior (personas, user assignments, defaults, accounting period, metadata context).

Finance Analyst Setup

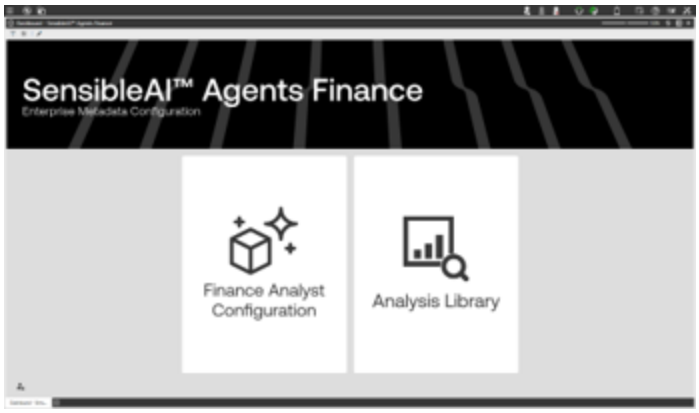
Scope

Finance Analyst can be utilized on any application within an environment. However, you must complete the entire the setup process for each desired application. To render Quick Views through Finance Analyst, adjust the selected Application in User Settings.

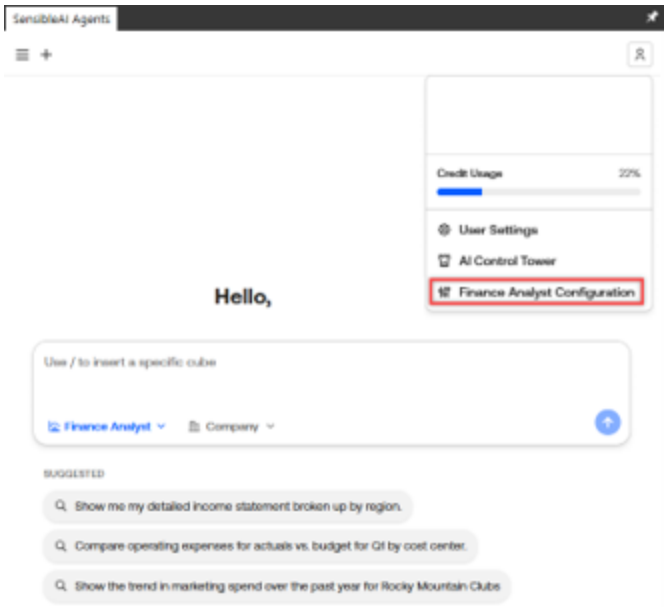
Accessing Finance Analyst Configuration

Administrators configure Finance Analyst using the Finance Analyst Configuration application, accessible through:

- SensibleAI Agents Finance marketplace solution, or

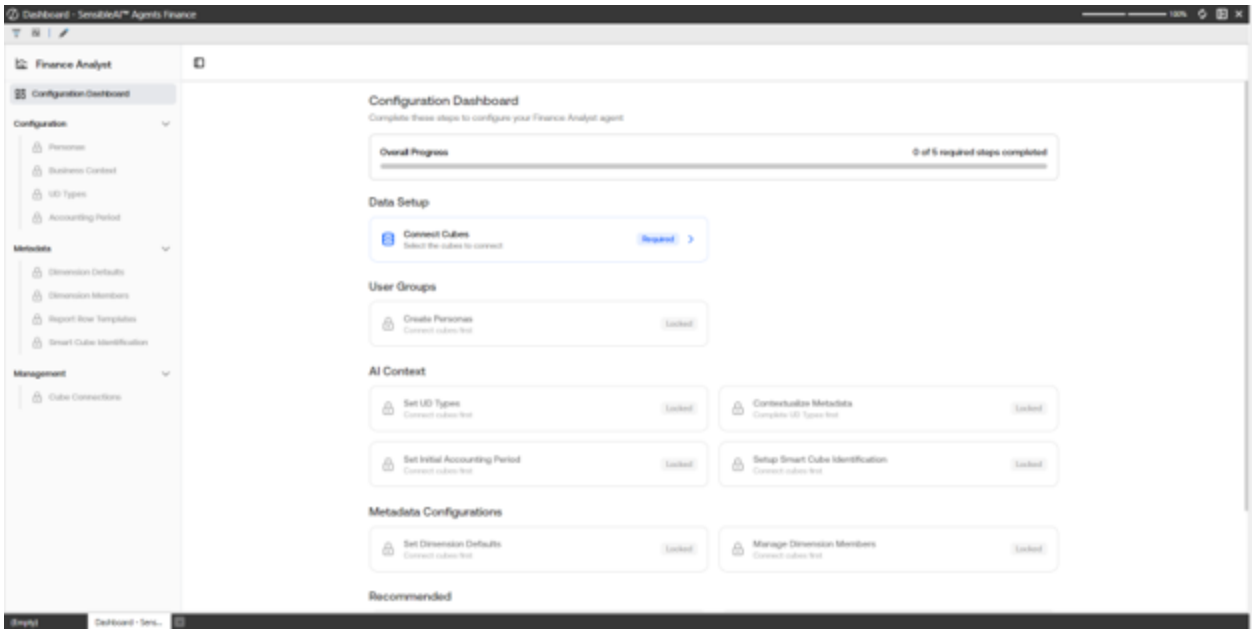


- Chat settings (where enabled)



The initial landing experience is the Configuration Dashboard, which lists the required setup steps. Steps unlock sequentially based on dependencies.

> **IMPORTANT:** Finance Analyst is not available to end users until all required steps are completed.



Required Setup Steps (Initial Enablement)

Complete these steps in the required order:

1. **Connect Cubes (required first)**: Select which cubes Finance Analyst can access. All other steps remain locked until cubes are connected.
2. **Set UD Types**: Assign user-defined (UD) dimension types to provide semantic meaning.
3. **Create Personas**: Define user groups and their cube access and behavior.
4. **Set Accounting Period**: Define the most recently closed accounting period used when answering questions about actuals.

> **NOTE:** After cubes are connected, Steps 2, 3, and 4 can be completed in any order.

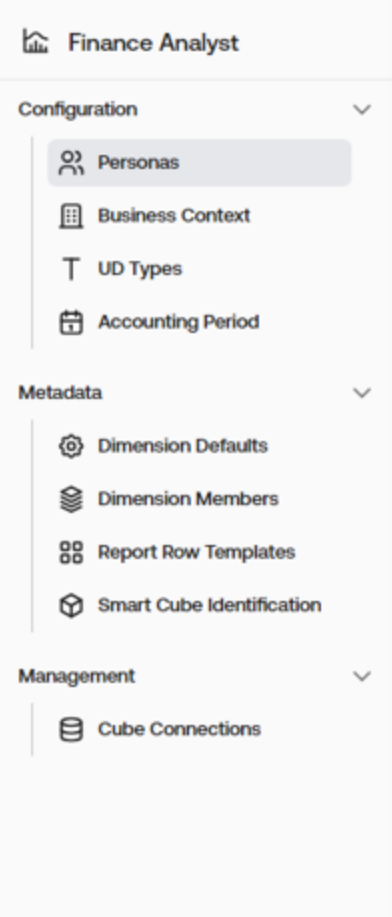
Optional (Recommended) Setup Steps

After cubes are connected, the following configuration areas become available and are recommended to improve accuracy and usability:

- Dimension Defaults (consistent behavior for unspecified dimensions)
- Users (assign each user's default persona and, where needed, a default set)
- Dimension Members (disable irrelevant members and provide disambiguating descriptions)
- Business Context (organization-specific rules not captured in metadata)
- Cube Context (provide descriptions of your cubes)
- Report Row Templates (predefined account groupings for structured reporting)
- Cube Views (provide descriptions for your existing OneStream Cube Views)

Finance Analyst Configuration Portal

After all required setup steps are complete, Finance Analyst is enabled, and the Configuration Dashboard transitions to the Configuration Portal for ongoing management.



The portal sidebar organizes configuration pages by area. The following sections describe each page.

Cube Connections

Purpose

Connect and synchronize cubes used by Finance Analyst. Cube synchronization must be completed before Finance Analyst can query data from a cube.

Permissions

- Only OneStream Administrators can manage cube connections
- Non-admin users can view the page but cannot make changes



Initial Setup: Connect Cubes

1. Locate the cube on the **Cube Connections** page.
2. Select **Connect** for each cube to enable.
3. Select **Preview Sync** to review queued actions.
4. Select **Sync Now** to begin synchronization.

> **IMPORTANT:** Finance Analyst is disabled for all users while synchronization is in progress.

After synchronization completes, continue to the Configuration Dashboard to configure all steps.

Ongoing Management: Resync and Disconnect

After initial setup:

- **Resync**: Use after structural cube changes (for example, new dimensions, renamed members, or hierarchy changes)
- **Disconnect**: Removes the cube from Finance Analyst and prevents user queries against that cube

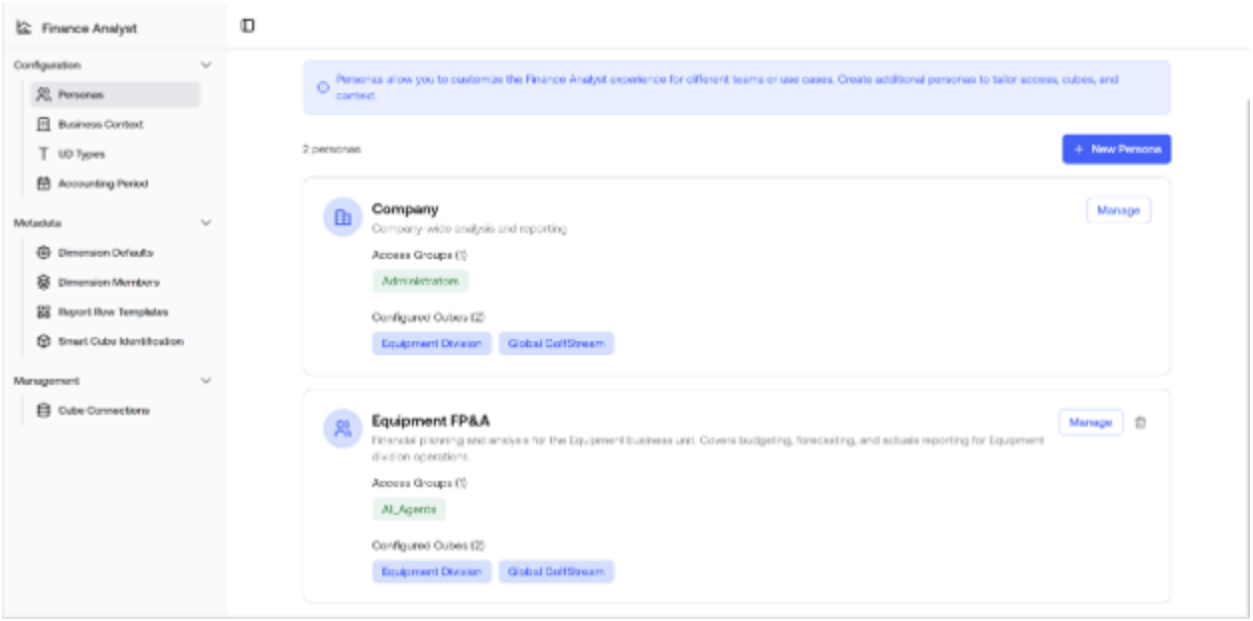
Actions are staged and executed together using Preview Sync.

Recommendations

- Connect only cubes that users will query. Additional cubes increase the search space and may reduce accuracy
- Resync after structural changes in OneStream

Personas (Finance Analyst Personas)

Personas represent distinct use cases within Finance Analyst. A user may belong to multiple personas if their responsibilities span different use cases.



Company Persona

The first time the Personas page is opened, administrators must configure the Company Persona.

The Company Persona:

- Serves as the default persona when no persona is selected
- Provides company-level defaults that other personas can inherit

Fixed properties:

- Name
- Description
- Access (available to all users when no persona is selected)

Configurable properties:

- Available Cubes
- Most Recent Closed Period (set through the Accounting Period, not edited directly in the Company Persona form)

> **IMPORTANT:** Configure the Company Persona before creating additional personas.

Creating Personas

Select New Persona to create a persona.

Access & Permissions

- **Name (required)**: Unique persona name
- **Description (required)**: Use-case description
- **Persona Access (required)**: OneStream groups allowed to use the persona
- **Additional Context Editors**: Users or groups allowed to edit persona-specific context

Configurations

- Available Cubes (required): At least one cube must be selected
- Most Recent Closed Period: Optional override of the company default

Recommendations

- Define personas by use case, not only by organizational structure (for example, Corporate FP&A vs. Regional Reporting).
- Use persona scope to reduce noise and improve accuracy.
- Delegate persona-level configuration using Additional Context Editors.

UD Types

Overview

UD Types provide semantic meaning for user-defined dimensions. Assigning UD Types enables the agent to interpret what members represent.

This page supports:

- Configure UD Types
- Manage UD Types

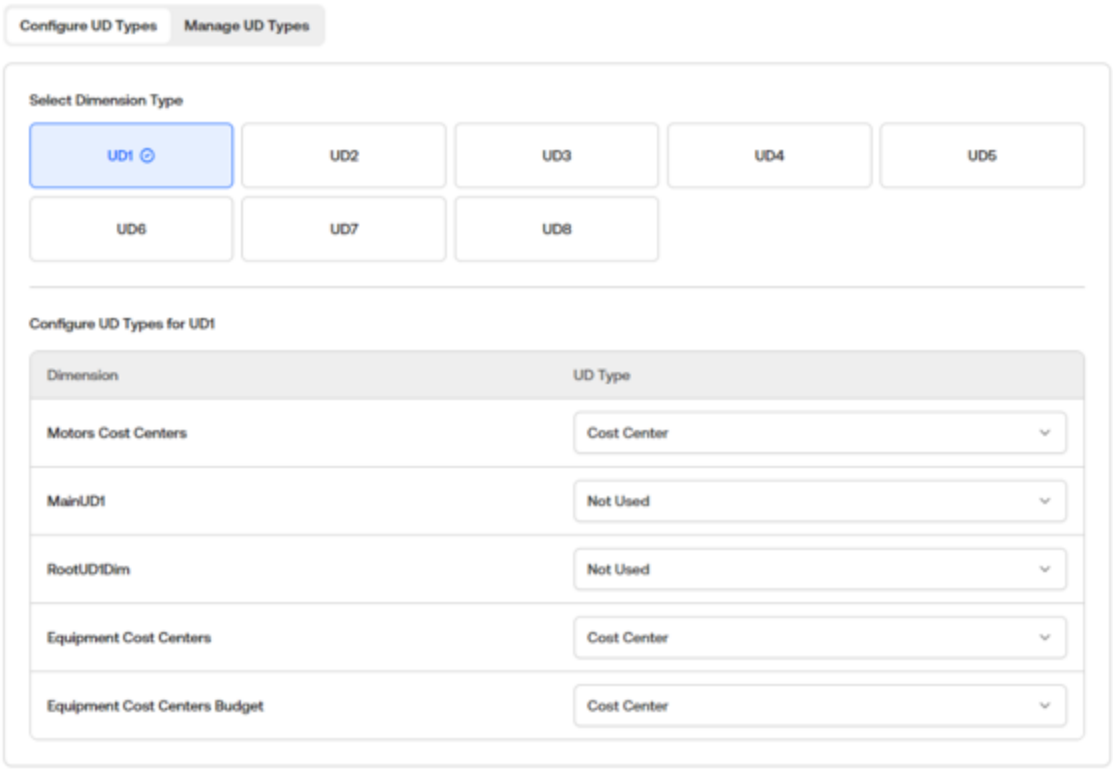
Configure UD Types

Assign a UD Type to each dimension.

1. Select a dimension type (UD1, UD2, and so on).
2. For each dimension, select the appropriate UD Type.

Rules

- Every user-defined dimension must have a UD Type assigned.
- If a dimension is not used, select Not Used:
    - Members are not visible to Finance Analyst.
- The POV value is set to the dimension default.
- A checkmark indicates completion for each dimension type.
- A UD Type can only be assigned once within a single UD dimension type. If already used, it appears disabled.



Manage UD Types

UD Types define what dimensions represent. The system includes a built-in Not Used type.

Actions:

- New UD Type
- Edit (pencil icon)
- Delete (trash icon)

Recommendations

- Create UD Types before assigning them.
- Write descriptions that clearly explain business meaning:
  - **Better**: “Organizational cost centers used for departmental expense tracking.”
  - **Avoid**: “UD1 dimension hierarchy.”

Users

Overview

The Users page lets administrators set Finance Analyst defaults for each user, so users do not have to configure them themselves. This is most useful during implementation, so users start in the correct persona from their first conversation.

Permissions

- Only OneStream Administrators and users assigned the SAI Administrator or SAI Support role can view the Users page.
- Assignments apply to the current application only.

Users Table

The table lists each user with access to SensibleAI Agents:

- **User**: The user's name and description.
- **Default Persona**: The persona the user starts in for new Finance Analyst conversations. Company means the Company Persona. This is also the persona used for MCP.
- **Dimension Defaults**: The default set assigned to the user. Inherited means the user receives their persona and company defaults with no user-level overrides.

Use the search bar to find users by name or description. Use the filter on the Default Persona and Dimension Defaults column headers to show only users with specific assignments.

Assign a Default Persona

1. Locate the user.
2. In the Default Persona column, select a persona.

Changes save immediately.

> **NOTE:** This is the same setting users see as Default Persona in Chat Settings. Users can change it after it is assigned.

> **IMPORTANT:** Assigning a persona does not grant access to it. The user must also belong to one of the persona's Persona Access groups. If the user does not, or if the persona is later deleted, the user starts in Company and sees a notice.

Assign a Default Set (Advanced)

1. Locate the user.
2. In the Dimension Defaults column, select a default set. To remove an assignment, select **Inherited from persona (default)**.
3. To view the assigned set's defaults, select the arrow next to the set name. To create or edit sets, select **Manage default sets…**.

See Default Sets (Advanced) under Dimension Defaults.

Recommendations

- Assign default personas during rollout so users start with the correct cubes and defaults.
- Leave Dimension Defaults as Inherited unless a user needs defaults their persona does not provide.

Accounting Period

Overview

Defines the most recently closed accounting period used for actuals reporting.

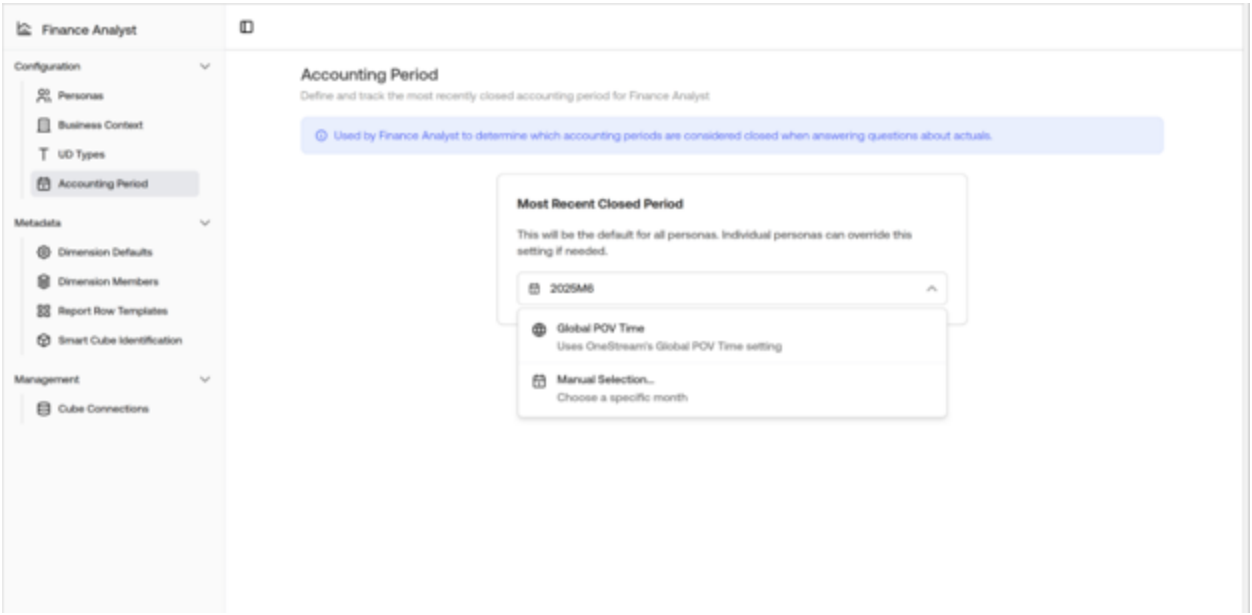
Options:

- **Global POV Time**: Uses OneStream Global POV Time.
- **Manual Selection**: Choose a specific month.

This page sets the Company Persona default. Personas may override this using Most Recent Closed Period.

Recommendations

- If Global POV aligns to the current close month (not the last closed period), use Manual Selection.
- Include Manual Selection updates in the close checklist.
- Use persona overrides when close schedules vary.



Cube Context

Overview

Cube Context helps the agent determine the correct cube when personas have access to multiple cubes. Without it, users must manually select a cube.

This feature includes:

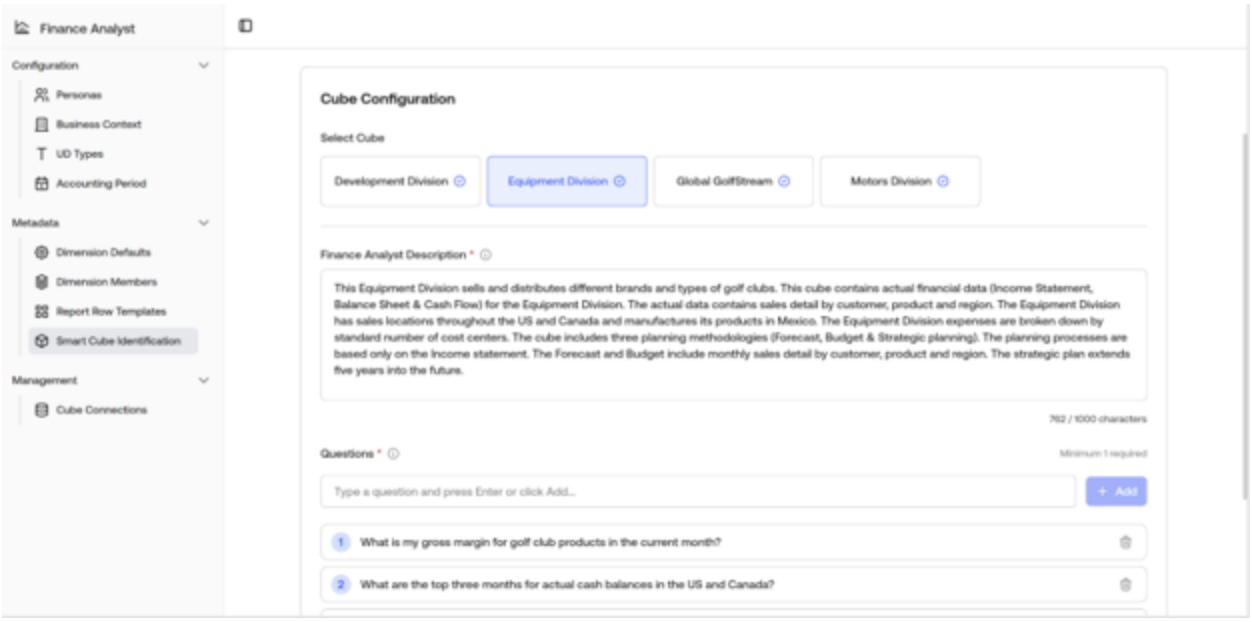
1. Cube Configuration (this page)
2. Persona-level mode selection (configured in Persona settings)

Cube Configuration

For each synced cube, provide the following:

- **Description**: Define the cube's domain and purpose.

> **IMPORTANT:** All synced cubes must be configured before Cube Context can be enabled for a persona.



Dimension Defaults

Overview

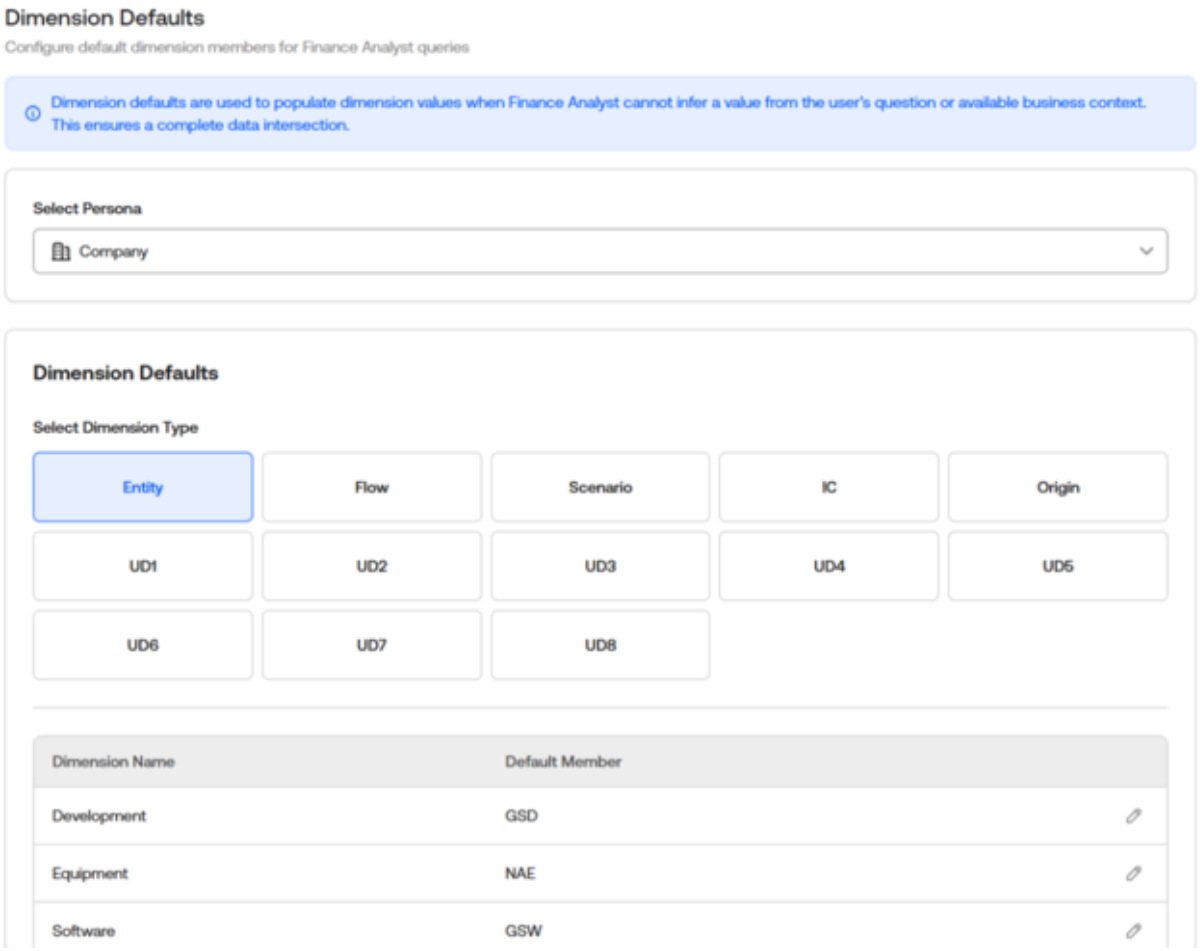
When users submit a question, unspecified dimensions are populated using defaults. Defaults are configured at the following levels:

- Company level: Applies to all users
- Persona level: Overrides company defaults for users working in that persona
- User level (Advanced): Default sets override persona and company defaults for the users assigned to them

For each dimension, Finance Analyst uses the first default it finds in this order: default set, persona, company. Each level overrides only the dimensions it sets.

The page has the following tabs:

- Company & Personas: Set company and persona defaults
- Users (Advanced): Create and manage default sets (SAI Administrators only)
- My Defaults: View the defaults that apply to your user



Company & Personas

1. Select Company or a Persona.
2. Select a dimension type (Entity, Flow, Scenario, IC, Origin, UD1–UD8).
3. Select a dimension.
4. Select the default member cell to open the member picker.
5. Choose a member and save.

Default Sets (Advanced)

Default sets let administrators set user-level defaults on top of company and persona defaults. For example, two users in the same persona could have North America and Europe as their default region.

> **NOTE:** Default sets are not needed in most implementations. Use them only when users who share a persona need different defaults.

How default sets work:

- A default set overrides only the dimensions defined in it. All other dimensions fall back to the persona, then the company.
- A default set applies in every persona the user works in. It is not tied to a specific persona.
- Each user can be assigned one default set. A set can be assigned to many users.
- Only SAI Administrators can create default sets and assign them to users.

Create a Default Set

1. On the **Users (Advanced)** tab, select **New Default Set**.
2. Enter a **Name** (required).
3. (Optional) Enter a Description.
4. Select **Create**.

Set Defaults in a Default Set

1. Select the default set.
2. Select a dimension type and dimension.
3. Select the default member cell, choose a member, and save.

Dimensions without an override show the inherited member, labeled (**Use Persona Default**) or (**Use Company Default**). To remove an override, select **Use Inherited Default**.

Preview Inheritance

Because a default set applies across personas, the inherited values depend on the persona a user is working in. Use **Preview inheritance as** to see how the set resolves for a specific persona. The preview does not change the set.

Assign Users

Assign default sets to users on the Users page. Select **Manage users** to open it.

Edit or Delete a Default Set

Use the edit and delete actions on the default set card. When you delete a set, its users fall back to their persona and company defaults.

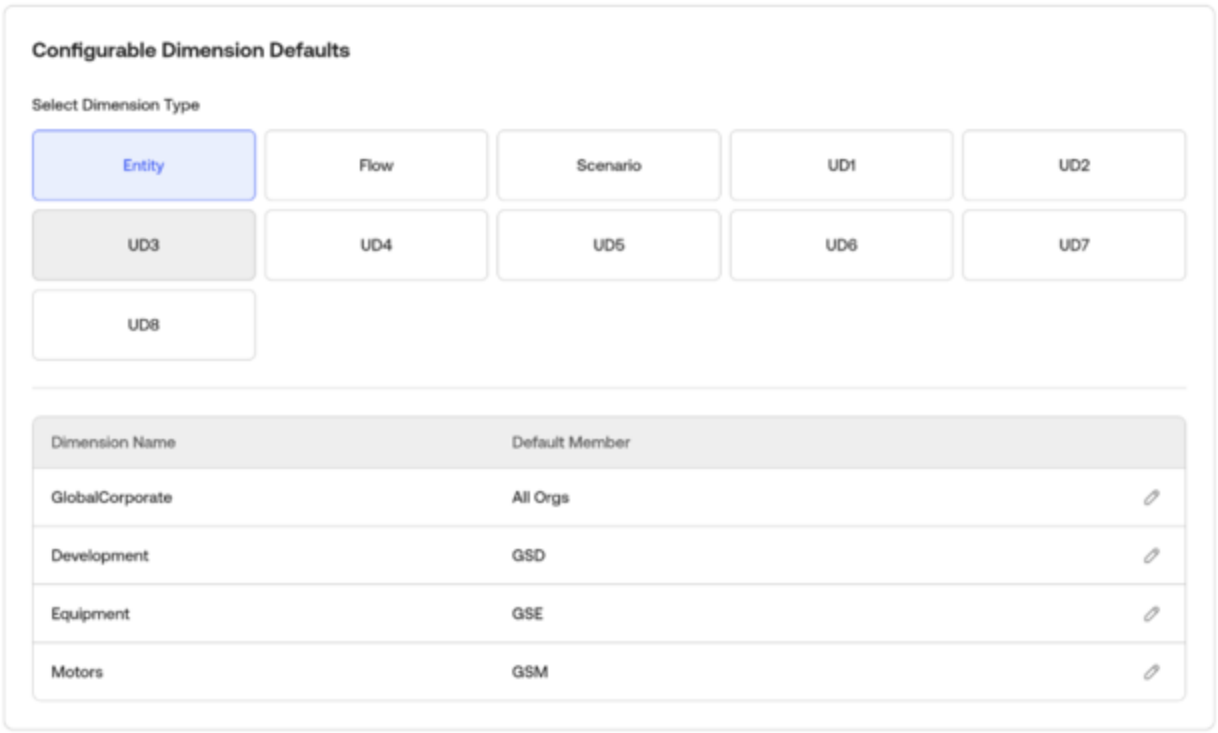
My Defaults

My Defaults shows the defaults that apply to you for each dimension. This tab is read-only. Select a persona to see how your defaults resolve in that persona. The **Source** column shows where each default comes from:

- **User**: Your assigned default set
- **Persona**: The selected persona
- **Company**: The company defaults

Recommendations

- Set company defaults first, and override only where needed.
- Defaults typically align with “Top” or “None.”
- Prefer personas over default sets. Use a default set only when users in the same persona need different defaults.
- Keep default sets small. Override only the dimensions that differ.
- Revalidate defaults after cube ingestion.



Dimension Members

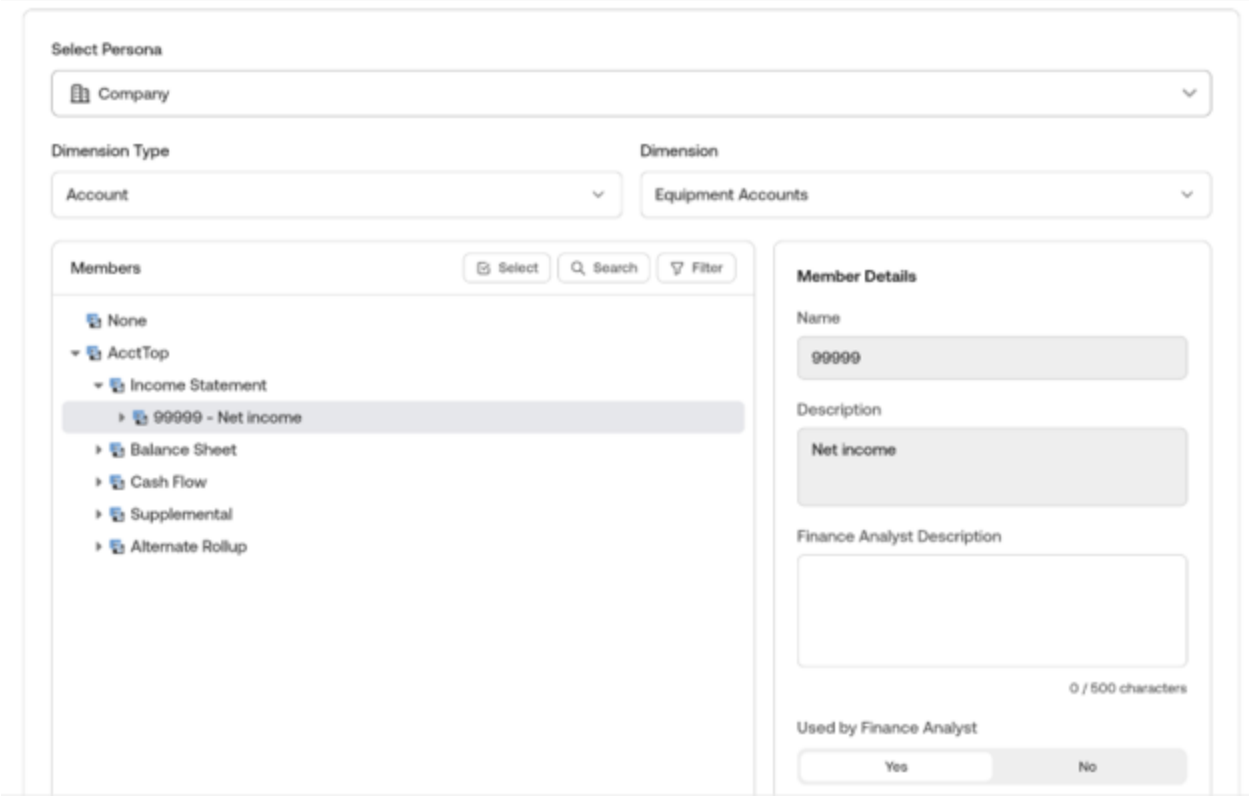
Overview

Controls which members are visible to Finance Analyst and allows custom descriptions to improve accuracy. Configuration is available at the company or persona level.

Workflow

1. Select a persona.
2. Select a dimension type and dimension.

3. Browse the hierarchy.
4. Select a member to edit.



Member Details

- **Name**: Read-only
- **OneStream Description**: Read-only
- **Finance Analyst Description**: Editable at the company level and inherited by personas
- **Used by Finance Analyst**: Yes/No toggle

Bulk Actions

Use Select above the tree:

- Hierarchy selection (applies to descendants).
- Individual selection (up to 50 members).
- Use the action bar to set selected members as Used or Not Used.

Recommendations

- Disable purely structural members to reduce noise.
- Use Finance Analyst Description to disambiguate similar members.
- Use bulk actions to manage large hierarchies efficiently.

Business Context

Overview

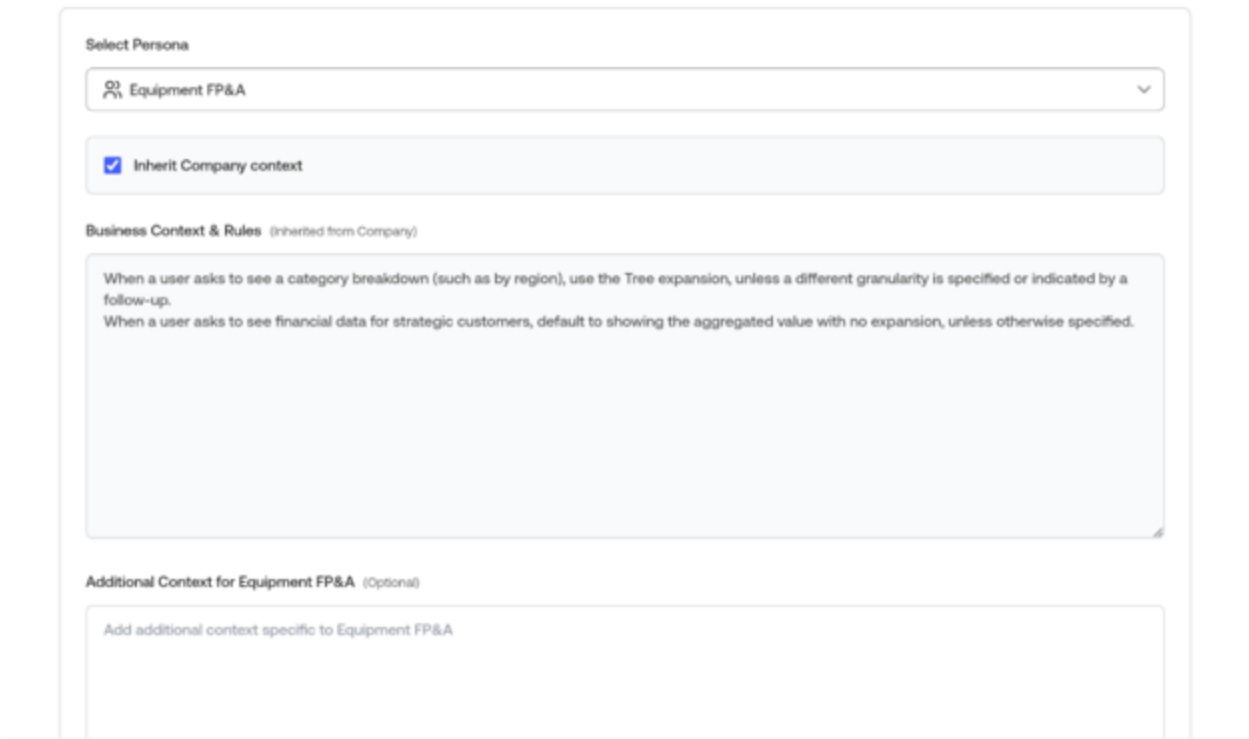
Business Context provides organization-specific rules not captured in metadata and guides interpretation.

Context Levels:

- Company
- Persona

Persona Inheritance Options:

- **If enabled**: Company context is read-only, and persona context is applied as additional context.
- **If disabled**: Persona context replaces company context.



Guidelines: What to Include

Include rules and meanings not reflected in metadata:

- Organizational terminology (for example, what “locations” means)
- Use-case definitions (for example, what “breakdown” means)
- Business rules (forecast locking, fiscal conventions)
- Account conventions (ranges, groupings)
- Organization-specific jargon mappings (scenario meanings)

Guidelines: What Not to Include

Avoid restating built-in OneStream concepts:

- Dimension definitions
- Cube definitions (handled by Cube context)
- Time formats or fiscal logic (built-in)
- Member search logic (built-in)

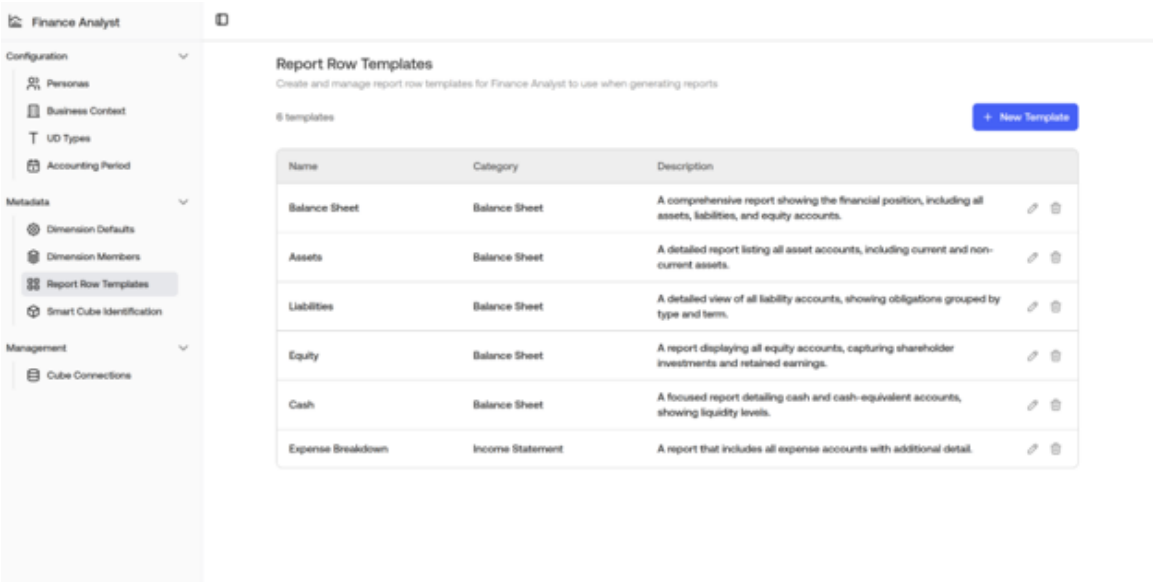
Recommendations

- Encode rules that experienced analysts know implicitly.
- Use company context for broadly applicable rules and persona context for use-case-specific rules.

Report Row Templates

Overview

Report Row Templates define predefined account groupings for structured reports (for example, Income Statement and Balance Sheet).



Creating a Template

The following fields are required:

- **Name**: Enter a template name (for example, “Detailed Income Statement”).
- **Category**: Select either Income Statement or Balance Sheet.

- **Description**: Describe what the accounts represent. This is used for matching.
- **Members**: Provide a member filter string. Expansion functions are supported.

> **Example:** A#19999.ChildrenInclusive

Recommendations

- Reuse existing cube view row templates where applicable.
- Use expansion functions to preserve structure while grouping.

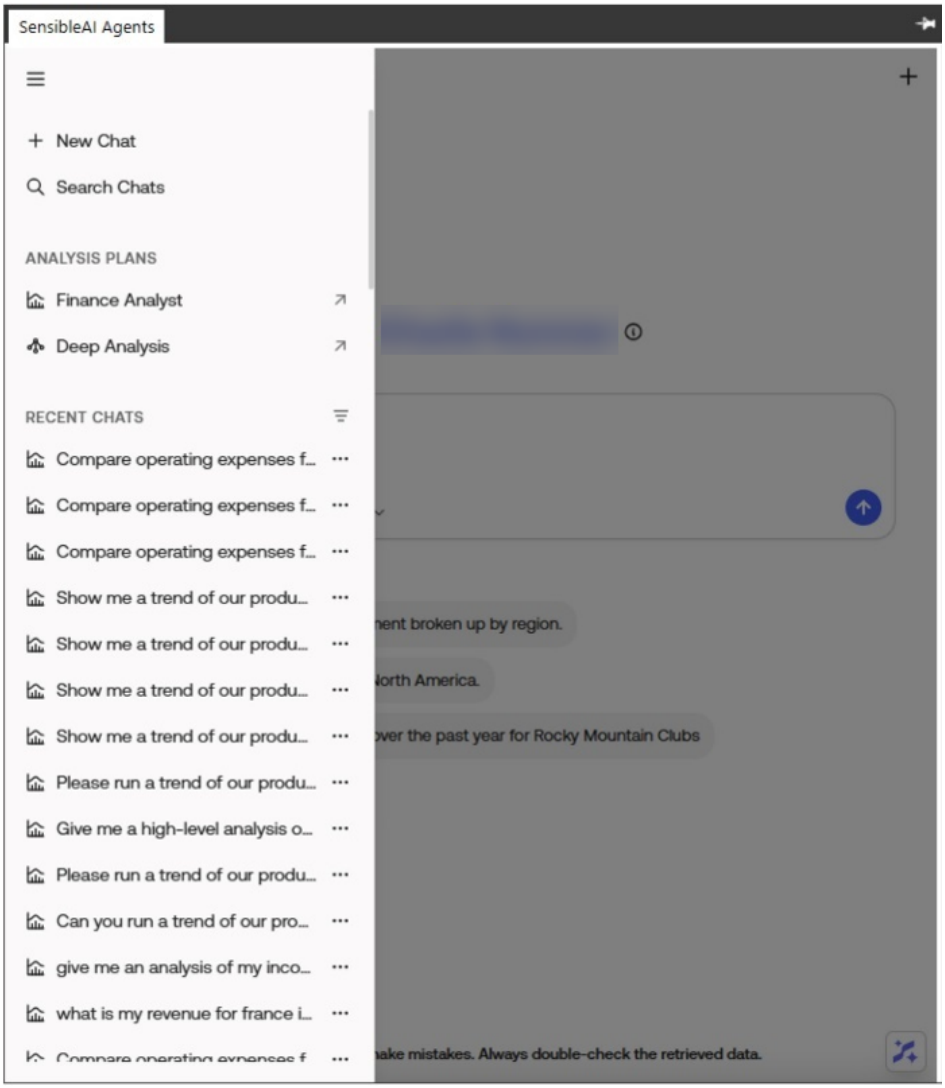
Finance Analyst Analysis Plans

Analysis Plans enable finance teams to define, save, and automate recurring analysis workflows. Common use cases include variance analysis, flux analysis, monthly product reviews, cost center performance reviews, and other recurring reporting processes that run against updated OneStream data each time. Users configure a plan once and run it on demand or on a schedule. Each user sees only their own plans and results. Plans are intended to be created by power users or others with knowledge of both the analysis process and the relevant OneStream metadata.

Accessing Analysis Plans

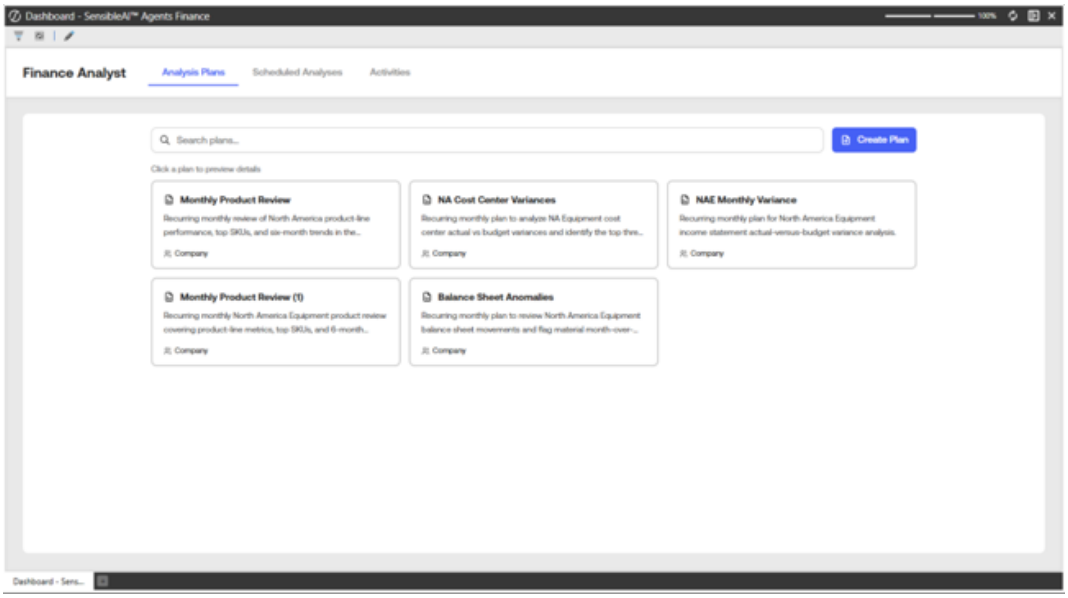
Analysis Plans are accessible from:

- The SensibleAI Agents Finance marketplace solution
- The Analysis Plans section in the SensibleAI Agents chat sidebar



Analysis Plans Page

The Analysis Plans tab displays all plans created by the current user. Click on an existing plan to view its details or select Create Plan to start a new one.



Plan Detail

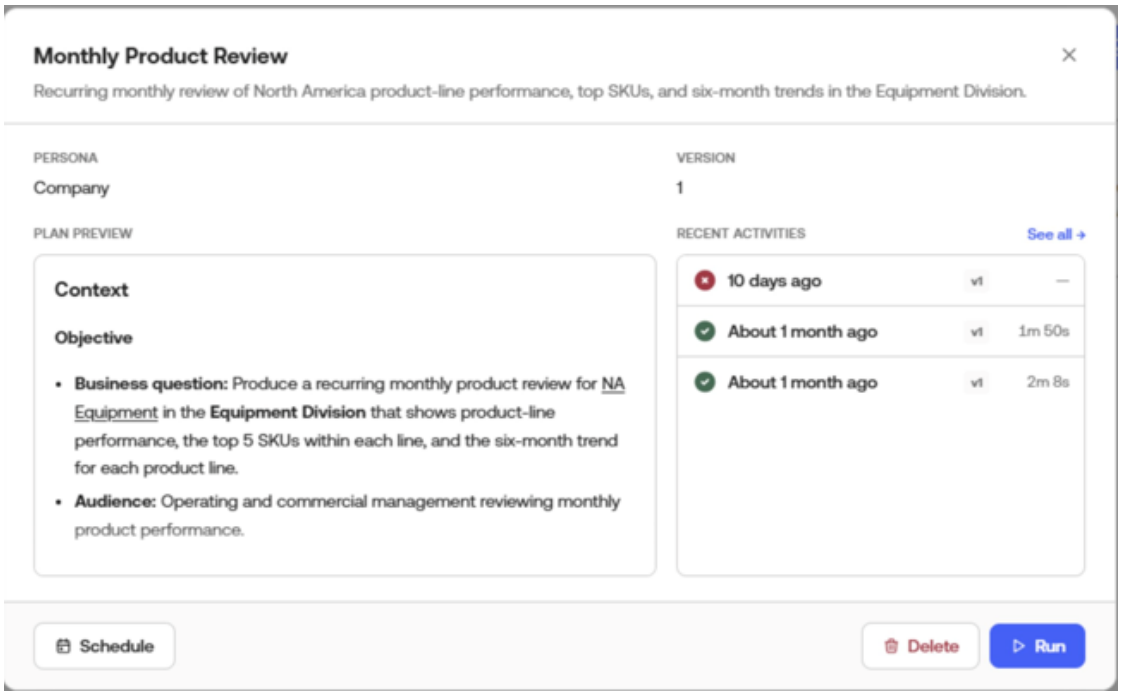
The Analysis Plans tab displays all plans created by the current user. Click on an existing plan to view its details or select Create Plan to start a new one.

Select a plan to open its detail view. Review the following fields:

- **Persona**: The Finance Analyst persona used when the plan runs
- **Version**: The current version of the plan
- **Plan Preview**: The structured plan document, including context, objective, analysis type, and assumptions
- **Recent Activities**: A log of past executions showing timestamps, versions, run durations, and success or failure status

The following actions are available from the detail view:

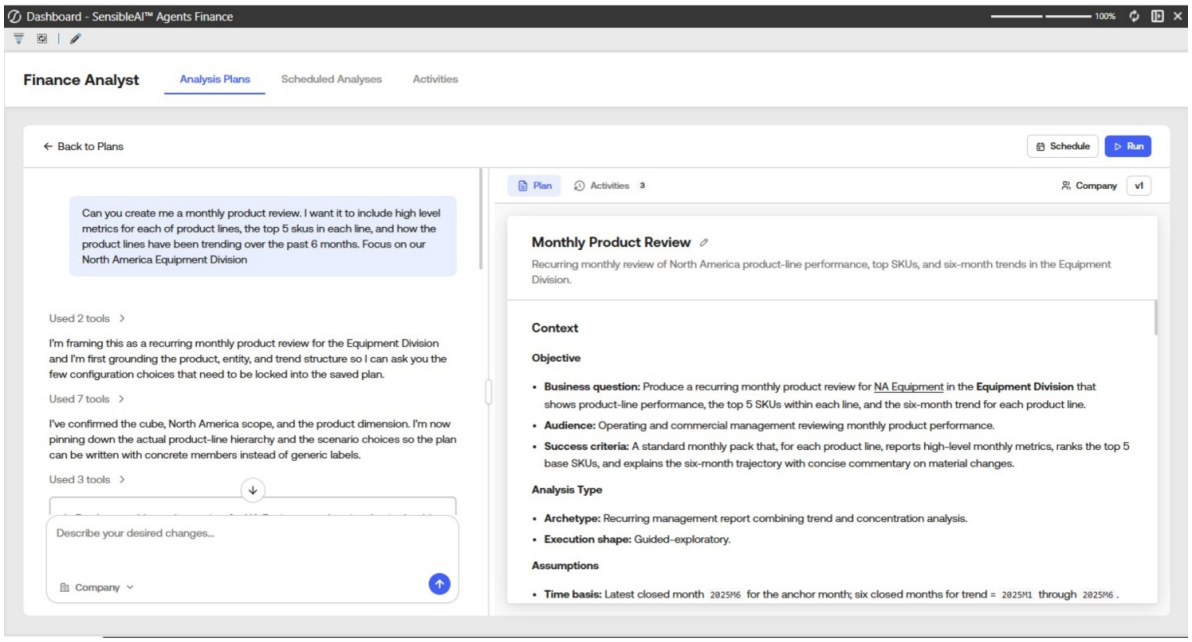
- **Run**: Executes the plan immediately
- **Schedule**: Opens the scheduling configuration
- **Delete**: Permanently removes the plan



Creating a Plan

Plans are created through an iterative planning agent. To create a plan, complete the following steps:

1. Go to the Analysis Plans tab and select Create Plan.
2. Describe the desired analysis in the input field.
3. The planning agent asks follow-up questions to resolve data scope, time basis, persona, and output structure.
4. Continue iterating with the agent via natural language until the plan captures all necessary details.
5. Once all details are confirmed, the plan is automatically saved and available to run or schedule.



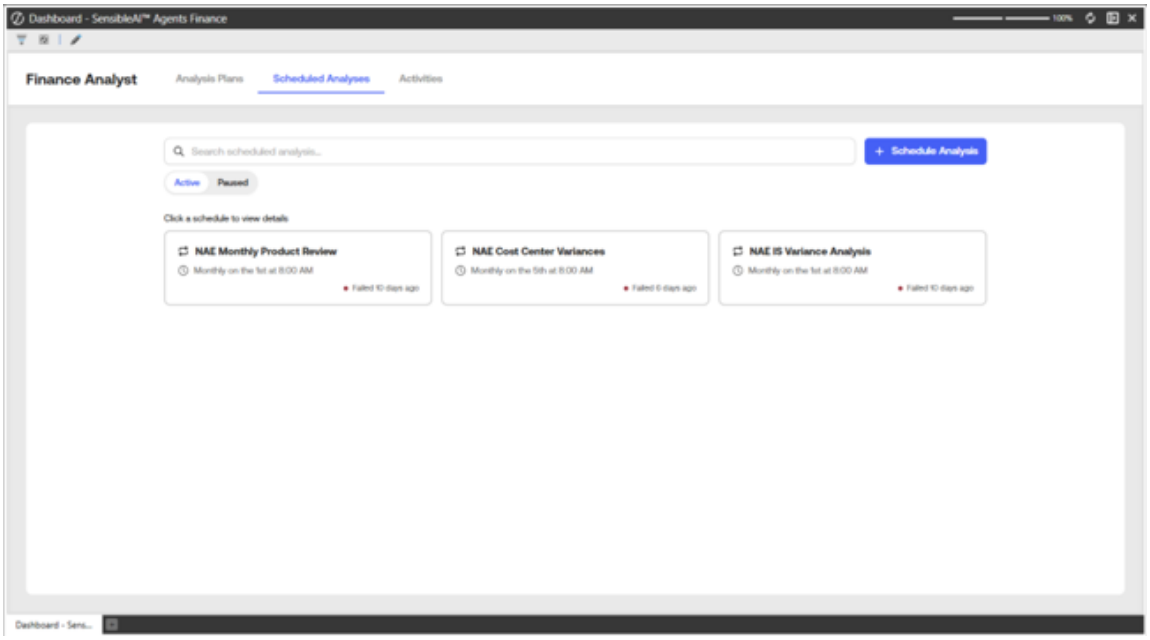
Editing a Plan

To edit an existing plan, open it and submit your desired change in natural language to the planning agent. Each edit automatically creates a new version, preserving the prior version history.

Scheduled Analyses

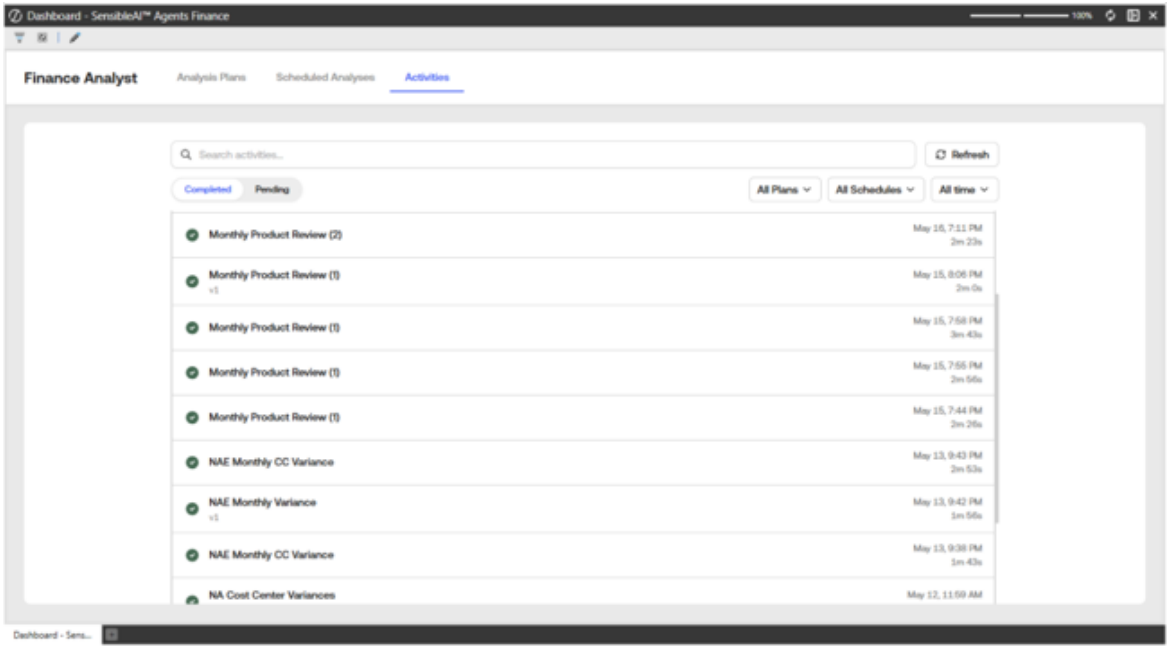
The Scheduled Analyses tab displays all active and paused schedules configured by the current user. Each schedule card shows the plan name, recurrence cadence, and the status of the last run. Use the search bar to find a specific schedule, or toggle between Active and Paused views.

To create a new schedule, select Schedule Analysis, choose an existing plan, configure the time-based recurrence, and assign the schedule a name. From this tab, users can also pause or delete any existing schedule.



Activities

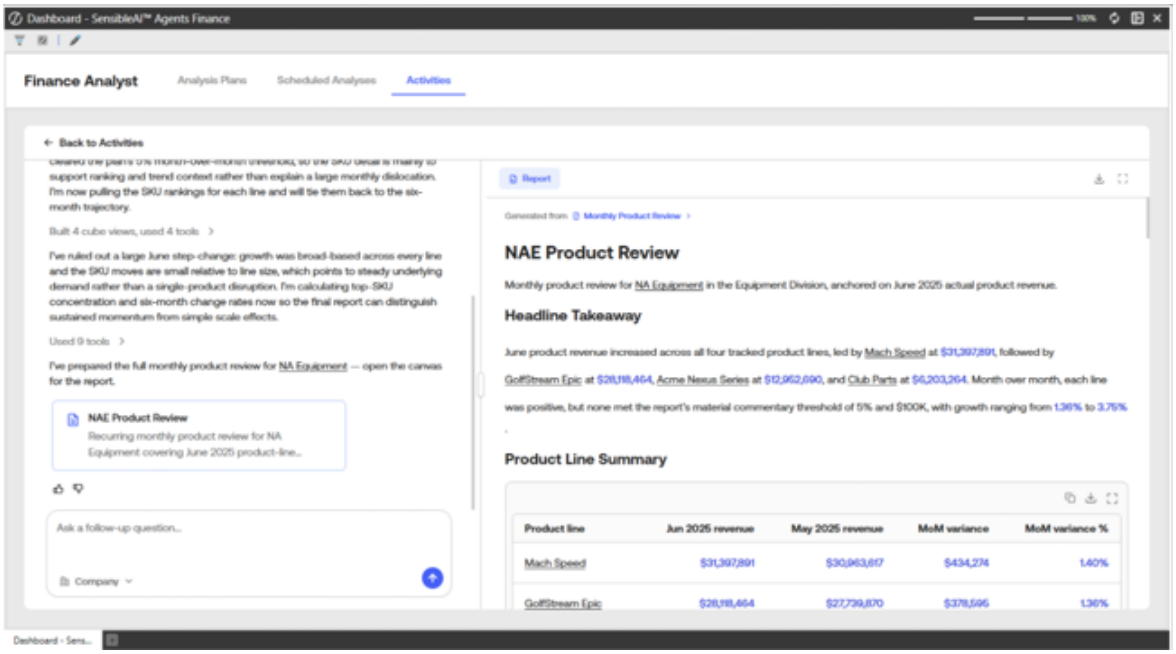
The Activities tab provides a complete log of all analysis executions for the current user. Toggle between Completed and Pending views to review past runs or monitor what is queued. Use the plan, schedule, and time range filters to narrow the list.



Select any completed activity to open the full output report. The report view has two panels:

- **Agent Process (left panel)**: A step-by-step log of how the agent built the analysis, including data pulls and reasoning. Users can ask follow-up questions or continue drilling into specific insights from this panel.
- **Report (right panel)**: The structured analysis output, including headline takeaways, commentary, and data tables. Key figures include source citations that link directly to the underlying OneStream data pull used to generate the number.

Currently reports cannot be natively shared with other users. However, users can download the reports directly to PDF with each data pull presented in an accompanying excel spreadsheet.



Example Questions (Finance Analyst)

Example questions are configured in the AI Control Tower under Finance Analyst settings.

Rules

- A minimum of 3 and a maximum of 10 questions must be configured.
- Three questions are randomly selected for each new conversation.
- Default questions are used if none are configured.

Recommendations

- Use real user phrasing.
- Cover multiple use cases and reporting needs.
- Example questions are currently shown for all personas.

Narrative Summarizer

Narrative Summarizer enables Finance Analyst to read and summarize the commentary your team has written in OneStream, alongside the underlying financial data, rolling child account notes up into a concise parent narrative.

This feature works across all five OneStream text-based view members:

- Annotation
- VarianceExplanation
- Assumptions
- AuditComment
- Footnote

For best results, scope each summarization to a single parent rollup, and specify your desired materiality threshold, summary style, and commentary location.

Narrative Summarizer is available both as an ad hoc request in Finance Analyst chat and as a scheduled Analysis Plan, allowing users to automatically generate parent-level narratives every close cycle.

End User Behavior (Admin-Relevant Summary)

Chat Settings

Users can configure the following:

- **Active Persona**: Determines cubes and defaults. Changes apply to new conversations.

Query Results

Finance Analyst returns interactive Dynamic Reports, accessible in the following ways:

- **Windows Application**: Cube View or Quick View
- **Modern Browser**: Cube View
- **Excel Add-in**: Quick View

SensibleAI Agents Documents Overview

SensibleAI Agents Documents manages documentation ingestion and grouping for:

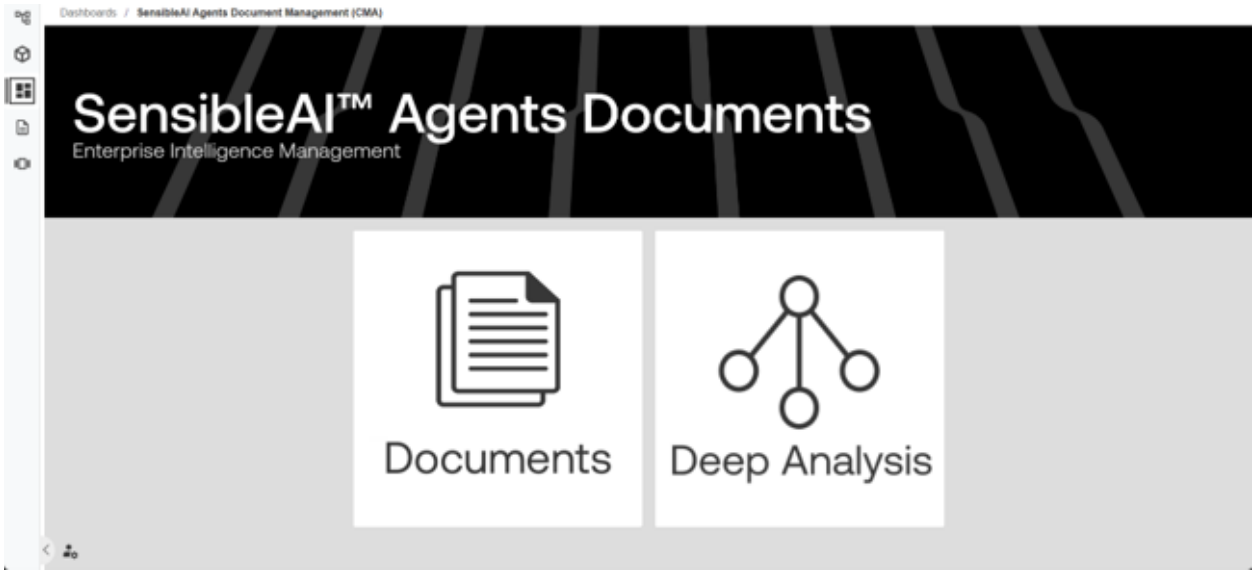
- Search Agent
- Deep Analysis Agent

Documents are uploaded into Knowledge Clusters, which define access boundaries and retrieval scope by utilizing existing OneStream Security Groups.

Home Page and Global Options

The Home page provides access to the following sections:

- Documents
- Deep Analysis



To access Global Options, select the **Settings** icon in the left toolbar.

Global Settings

Configure the following settings:

- (Security) Power User Role
- Backend Endpoint Parameter
- Chatbot Endpoint Parameter



Uninstall

- **Uninstall UI**: removes the app but retains data tables. This is upgrade-friendly where supported.
- **Uninstall Full**: removes the app and all related data tables.

> **IMPORTANT:** Use Uninstall Full when a full reset is required.

Documents

SensibleAI Agents Documents opens on the **Overview** page. The left sidebar has the following tabs: Overview, Knowledge Clusters, Files, Connectors, Activity, and Personas. You manage connector source configurations in AI Control Tower > Integrations, not in Documents.

The Overview page summarizes the content for the current user: the number of Knowledge Clusters they can access, the number of files, and the number of personas. The cluster count breaks down into clusters they own, clusters they are a member of, and, for SAI Administrators, clusters visible only through the administrator role. The page also includes an Activities to review table where users can approve, reject, or cancel pending items without opening the **Activity** page.

The Documents forms for clusters, uploads, file edits, and connectors are single scrolling pages. Each section unlocks after you complete the previous section.

Knowledge Clusters

Knowledge Clusters group documents and control access. Review the following guidance:

- Files can be added to multiple clusters.
- Users must have membership/ownership to access content within a given cluster.

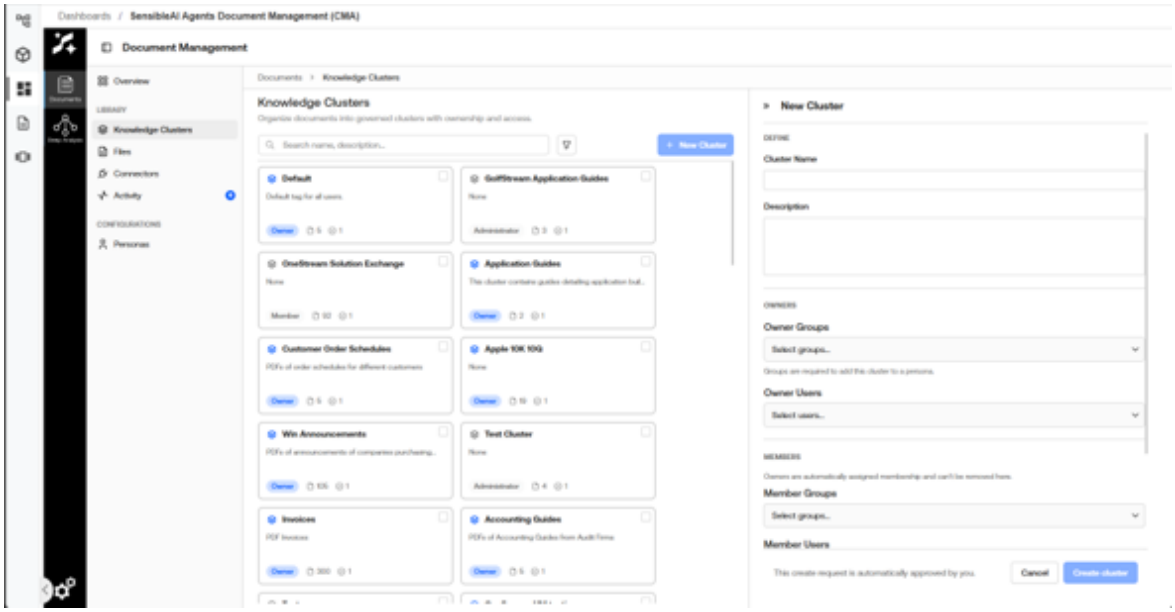
There are two cluster roles:

- **Member**: A user who can access cluster information.
- **Owner**: A user with member permissions who can also approve or reject activities for the cluster. A cluster can have multiple owners.

Create a Cluster

To create a cluster, complete the following steps. The cluster form is a single scrolling page, and each section unlocks after the previous section is complete.

1. Go to **KnowledgeClusters**.
2. Select **New Cluster**.



3. Enter a name (required) and description (optional). See Cluster Naming Rules.
4. Select **Owners**.
5. Set Required Approvals: The number of approvals needed before an activity on the cluster is approved. Approvals can come from cluster Owners or SAI Administrators. The default is 1.
6. Select **Members**.
7. Select **Create cluster**. If you are an Owner of the new cluster or an SAI Administrator, your approval is applied automaticall, when Required Approvals is 1, the cluster is created without manual review. Otherwise, the creation activity appears in **Activity** > **Review**.

> **NOTE:** Version 0.47.0 and later: cluster creation requires approval by a SensibleAI Agents Administrator role.

Cluster Naming Rules

Cluster names must be unique with no leading or trailing spaces. Names have a maximum of 50 characters. The following names are reserved:

- OneStream
- All Clusters

The following characters are restricted: ' " { } [ ] ( ) & * ! \ / | @ , ? % : < >

Cluster Detail

Select a cluster card on the **Knowledge Clusters** page to open the **Cluster Detail** modal. From there, administrators can do the following tasks:

- Edit Owners and Members
- View files and properties
- Preview documents in OneStream File Viewer

File attributes include:

- File Name, Staged by, Clusters
- Creation Time, Last Updated
- Traffic Count
- Positive/Negative Reviews (when cited)

Delete a Cluster

Only cluster owners and SAI Administrators can delete a cluster. Deletion requires approval. When deleting a cluster, choose whether to delete its files or keep them. Kept files that do not belong to any other cluster become orphaned. Only SAI Administrators and the users who uploaded orphaned files can view or edit them. To make orphaned files available again, edit them on the **Files** page and add them to another cluster.

Files

The **Files** page lists every file you can access. From here, you can upload files to one or more clusters and edit files to change which clusters they belong to. Review the following supported file types:

- **PDF**: Text + images + tables (advanced processing)
- **DOCX**: Text + images + tables (advanced processing)
- **PPTX**: Text + images + tables (advanced processing)
- **XLSX**: Text only
- **CSV/TXT**: Text only
- **MD**: Text + basic tables

Upload Files

Upload files or folders from your local machine, either from the Files page or to a specific cluster from the Knowledge Clusters page. You cannot upload zip files. Extract them first and upload the files or folders.

In the web version (opened from the user menu in chat) and in the Modern Browser Experience solution, you can drag and drop files or folders onto the page or the upload form. Drag and drop is not available in the Windows application. Use the file or folder picker instead.

Review the following limits:

- 85MB per file (including files within an uploaded files)
- 1000 files per upload request

To upload files, complete the following steps:

1. Go to **Files**.
2. Select **Upload**.
3. **Details**: Enter a request name (required) and description (optional).
4. **Files**: Add files or folders.
5. **Clusters**: Select clusters for access.
6. Review and submit.

After submission, the upload displays in **Activity**. Every file undergoes a security scan, and then cluster owners or SAI Administrators approve or reject the upload in **Activity** > **Review** before indexing. If you are an owner of a target cluster or an SAI Administrator, your approval applies automatically. When your approval meets the **Required Approvals** for the cluster, no manual review is needed.

Duplicate File Handling

Duplicate names degrade response quality and retrieval diversity. File names are checked when you submit an upload. If any names match existing files, a modal opens where you can resolve all conflicts at once or choose an option for each file:

- **Replace**: Replaces the existing file after approval. This option is available only when the existing file belongs to the clusters you are uploading to.
- **Rename**: Uploads the file under a suggested new name.
- **Exclude**: Removes the conflicting file from the upload.

If the existing file belongs to a cluster outside the upload, only Rename and Exclude are available. To add the existing file to more clusters instead, use Edit Files.

Edit Files

On the **Files** page, select one or more files and select **Edit** to change which clusters they belong to:

- Add cluster
- Remove cluster

Review the following approval rules:

- **Add cluster**: Approved by the target cluster owner.
- **Remove cluster**: Approved by the owner of the cluster being removed.

Delete Files

Deleting a file removes it from all clusters and from the system. At least one owner from each associated cluster must approve the deletion.

Activity

Activity is the single queue for the following document and connector workflows:

- File uploads, edits, and deletions
- Cluster creation, edits, and deletions
- Connector creation, reconnection, sync, and deletion

It replaces the former **Library Requests** and **Connectors Events** tabs. The **All** tab lists every activity with its current status and history. Select an activity to view file-level details. The **Review** tab lists activities awaiting your approval, which also display in the Activities to review table on the **Overview** page.

Activities can have the following statuses:

- Pending
- Scanning
- Staging Files
- Pending Review
- Approved
- Uploading
- Rejected
- Failed
- Partially Failed
- Completed

Review

Resource owners (cluster owners or the connector owner) and SAI Administrators can approve or reject activities from the **Review** tab or from the Activities to review table on the **Overview** page. Connector sync previews are approved or canceled instead. Cancel discards the sync without applying any changes.

A cluster can have multiple owners and sets its own number of Required Approvals. A connector has a single owner. SAI Administrators can approve any activity, and their approvals count toward the Required Approvals for a cluster.

Connector sync previews and activities with security scan failures always require review. For all other activities, when the submitter is an owner of the affected resource or an SAI Administrator, their approval applies automatically. If only one approval is required, the activity is approved without manual review.

If any files fail the security scan, the activity goes to the **Review** tab, where each owner or SAI Administrator reviewing it can choose to bypass individual failed files. A failed file is bypassed and included in the upload only if every reviewer who approves the activity bypasses it. All other failed files are excluded.

Approved items proceed to Uploading and then Completed or Failed. After approval or rejection, items leave the **Review** tab but remain visible in All.

Troubleshooting Failed Activities

Do the following tasks if the file processing fails:

1. Review OneStream Error Log for error details.
2. Validate that the file is not corrupted. Re-save/re-export.

3. Retry the upload (transient errors may occur).
4. Contact support if failures continue.

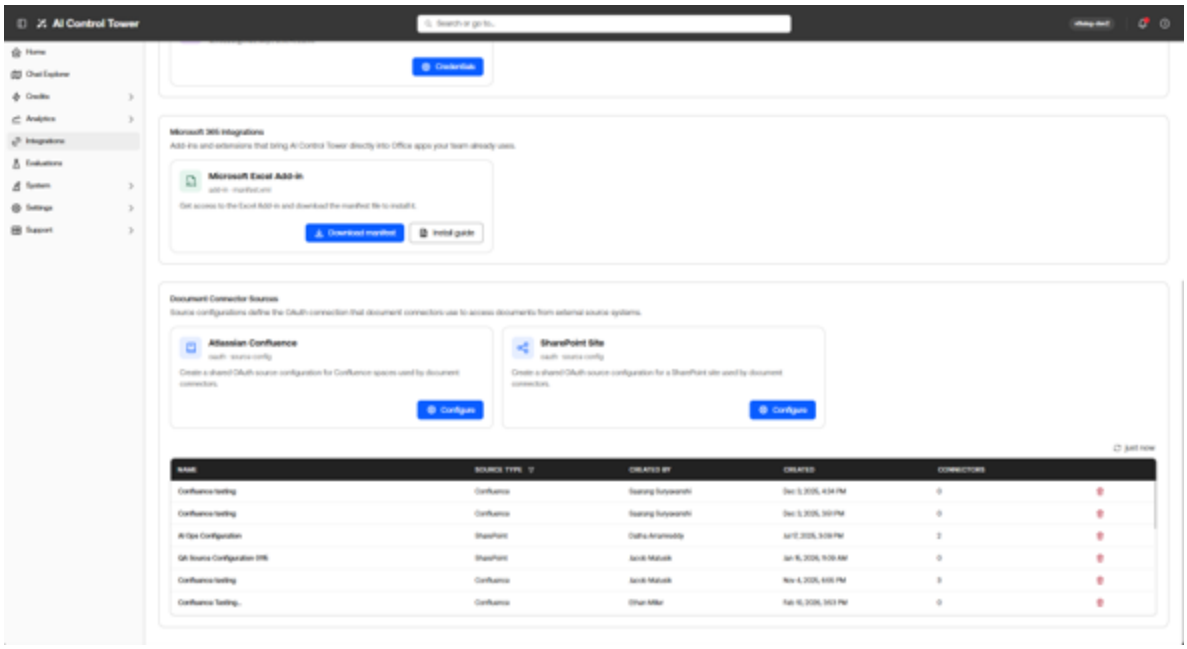
Data Connectors

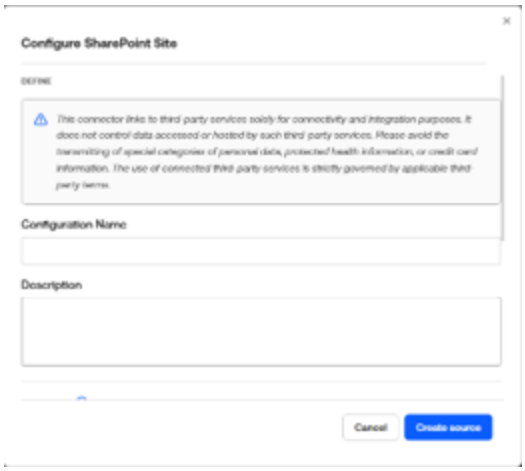
Data Connectors ingest documentation from external systems into clusters.

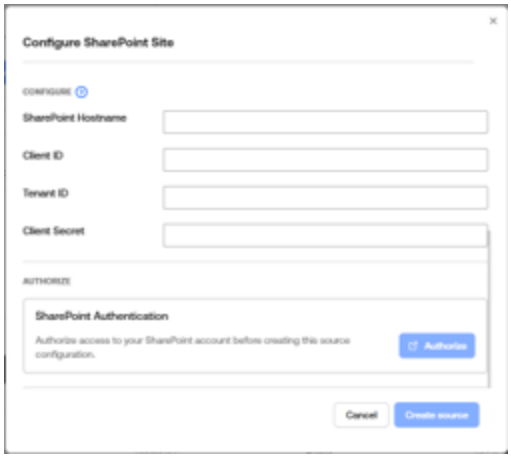
Key Concepts

- **Source Configuration**: Credentials and connection settings that SAI Administrators create in **AI Control Tower** > **Integrations**. Many connectors can use one source configuration.
- **Connector**: Created on the **Documents Connectors** page. A connector uses an existing source configuration to retrieve files from a specific location and sync them into clusters.

Source Configurations (Admin Only)

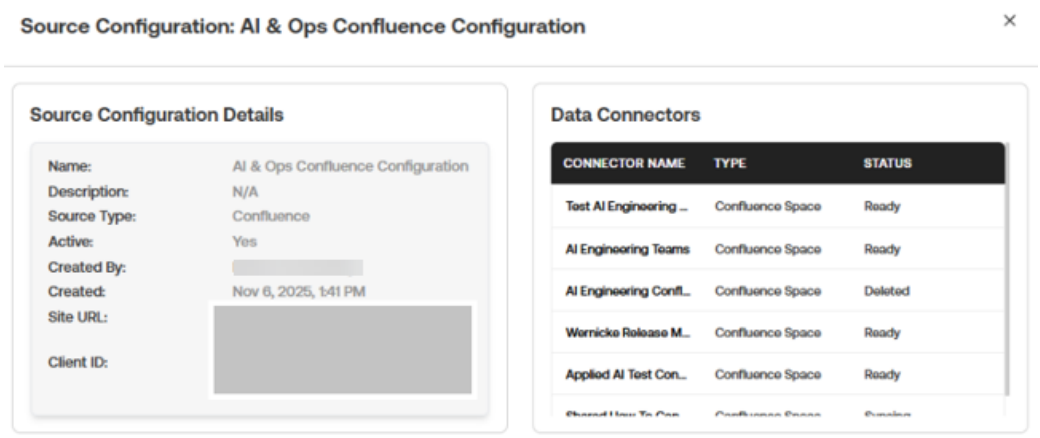






SAI Administrators configure and manage connector source configurations in AI Control Tower > Integrations. See Document Connector Sources. At least one source must exist before users can create connectors. For each source, administrators can do the following:

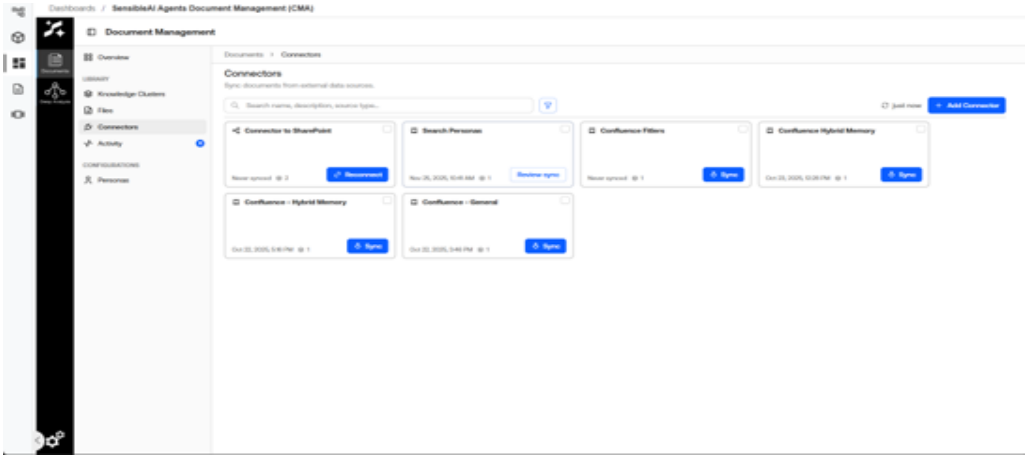
- View details and associated connectors
- Edit the source configuration
- This will pop up an informational dialog that displays information about the selected Source Configuration, including the Source Configuration's details and associated Data Connectors.



- Delete (deletes all associated connectors)

> **IMPORTANT:** Review associated connectors before deleting a source configuration.

Connectors



The **Connectors** page shows connectors as a card grid. Non-administrators see only the connectors they own. SAI Administrators can switch between All connectors and My connectors. Each connector card shows the following details: Non-administrators see a message asking them to contact an administrator, and administrators see a link to add a source in **AI Control Tower** > **Integrations**. The table on this page contains the following details for each connector:

Field	Description
Connector Name	Name of the connector
Source	Type of data source
Clusters	Clusters linked to the Connector

Field	Description
Status	Ready, Creating, Validating, Syncing, Deleting, or Disconnecting
Owner	Creator of the Connector
Created Time	When the connector was created.
Last Sync Time	Time of last ingestion

Connector fields:

- Name, Source Type, Clusters, Status, Owner, Created Time, Last Sync Time

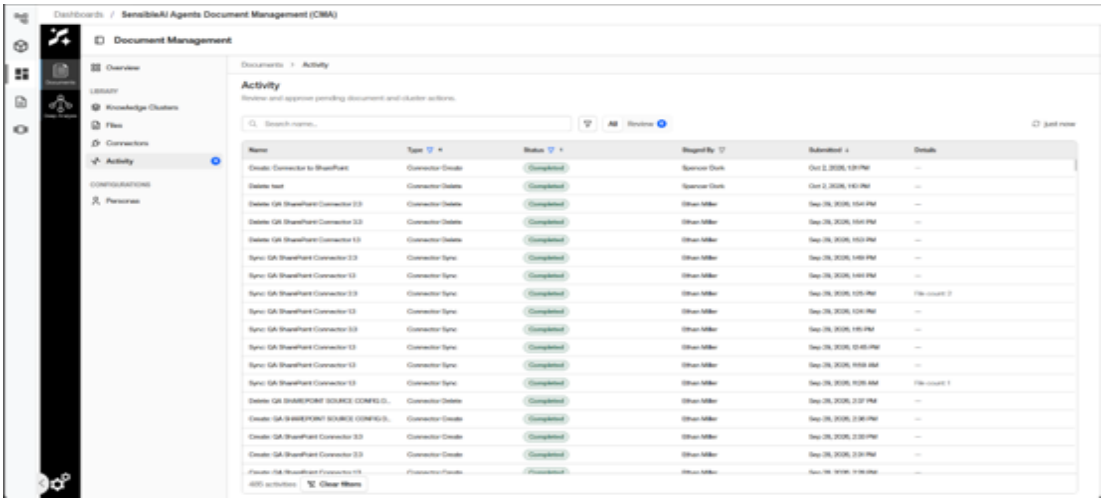
Statuses include:

Connector Statuses
Ready
Creating
Validating
Syncing
Deleting
Disconnected

Connector Activity

Connector creation, reconnection, sync, and deletion display in Activity. See Activity under Documents. Each sync creates a sync preview that lists the files to add, update, or remove. The preview displays in the Review tab, where the connector owner or an SAI Administrator can approve it to apply the changes or cancel it to discard them. You cannot reject sync previews. Use Cancel instead. Sync previews are never approved automatically. Even when the connector owner starts the sync, the preview must be reviewed and approved before any changes are applied. Only the connector owner can start a sync. SAI Administrators cannot sync a connector that another user owns, but they can approve or cancel its sync previews. Because a connector has a single owner, one approval completes the review. Reconnect results and completed syncs display in the All tab. Each connector activity shows the following details:

Activity Type	Details Shown
Create	Activity Details, Connector Details, and Connector Access Breakdown
Reconnect	Activity Details and Connector Details
Sync	Activity Details, Connector Details, and Individual File Statuses
Delete	Activity Details, Connector Details, and Connector Event History



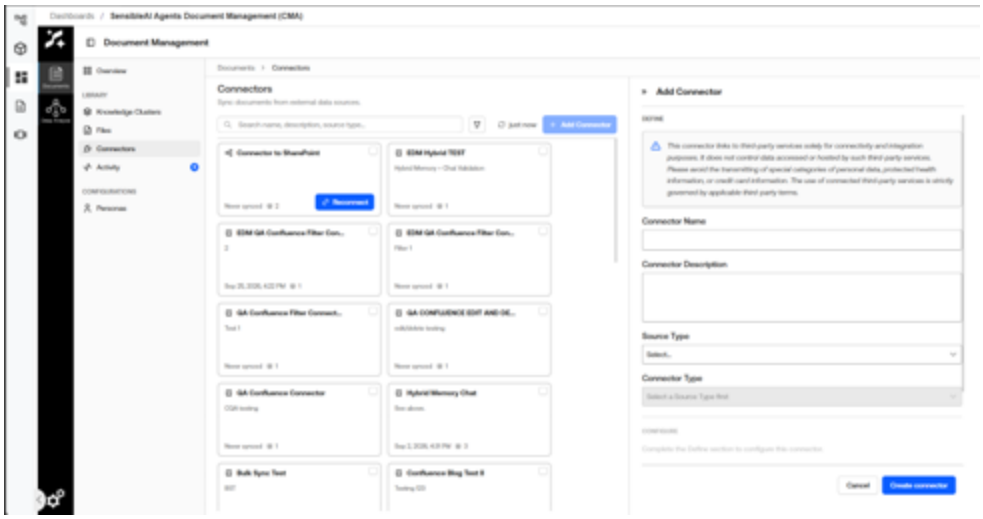
Connector activity statuses include:

Connector Activity Statuses
Staging
Scanning
Pending Review
Syncing
Completed
Failed
Partially Failed

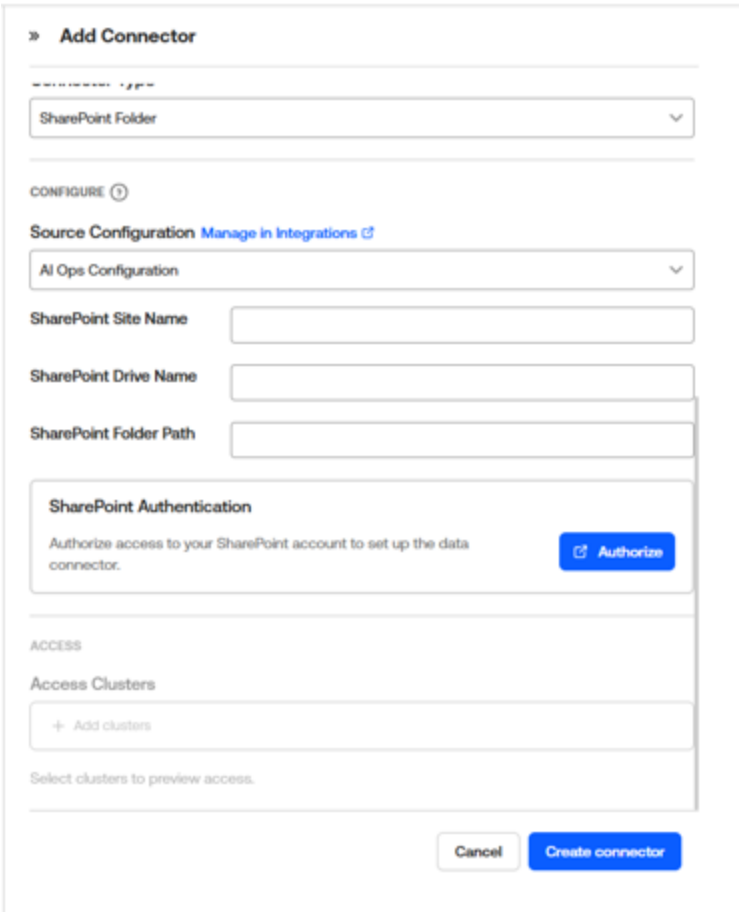
Connector Creation

On the **Connectors** page, select **Add Connector**. At least one source configuration must exist in **AI Control Tower** > **Integrations**.

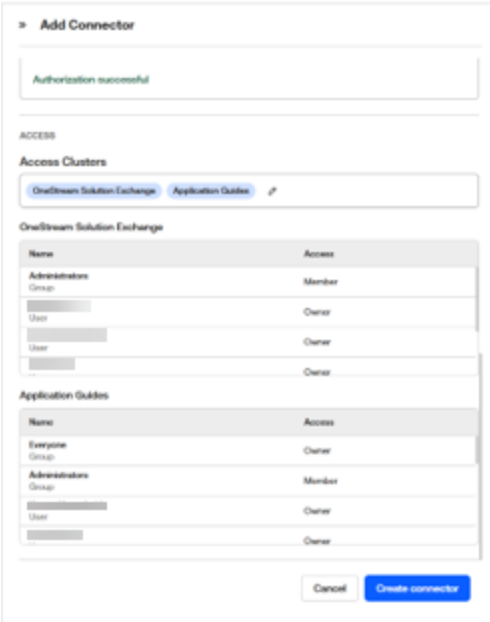
1. **Define**: Enter a name and an optional description.



2. Select the **Source Type** and **Connector Type**.
3. **Configure**: Select an existing source configuration and complete the source-specific fields, such as folder path or space key. You do not enter credentials here. They come from the source configuration in **AI Control Tower** > **Integrations**. The connector owner can edit any connector setting until its first sync. After the first sync, the owner can edit only metadata, such as name and description. The configuration is locked.



4. **Access**: Select clusters for access.



5. Review and submit.

> **NOTE:** If the connector immediately shows Disconnected, the configuration is incorrect. Delete and recreate.

Reconnect

If a connector is Disconnected, its owner can reconnect it to revalidate the connection. There is no separate Validate action.

1. Select the connector.
2. Select **Reconnect**. If you have not synced the connector yet, you can edit its connection settings before reconnecting.
3. Review the result in **Activity** > **All**.

Sync

A sync starts manually and applies changes to the document index. Sync previews display in **Activity** > **Review**, where the connector owner or an SAI Administrator can approve or cancel them.

Rules

- Only the connector owner can sync to avoid documents being ingested by non owners as it bypasses the request process.
- Connector must be in Ready status.

Sync Stages

Pending → Staging Files → Scanning → Pending Review → Syncing → Completed (or Failed / Partially Failed / Cancelled)

Pending Review Preview Categories

- New Files
- Updated
- Overwritten, same process as the duplicate file name handling
- Deleted
- Edited
- Scan Failures

Limits

- Maximum 1,000 files per sync

Connector Deletion

- The connector record is permanently deleted from the system. Files already synced into Knowledge Clusters are not removed.
- Any data snapshots associated with the connector are removed.
- The connector is disassociated from its Source Configuration.
- If the deleted connector was the only connector belonging to that owner, any stored authorization secrets for that owner are also removed.

Rules

- Status must be Ready or Disconnected
- Only the owner or an administrator can delete

SharePoint Configuration

SharePoint can be connected as a data source using a source configuration and one or more connectors. Source configurations are created and managed by SAI Administrators in **AI Control Tower** > **Integrations**. Each SharePoint site must be configured separately for security.

Confluence Configuration

Confluence can be connected as a data source using a source configuration and one or more connectors. Source configurations are created and managed by SAI Administrators in **AI Control Tower** > **Integrations**. Each Confluence site must be configured separately for security.

How It Works

Once a source configuration is established, additional connectors can be created from it by specifying a different folder path, without requiring a new source configuration. This allows content hosted across different SharePoint sites or folders to be connected by reusing existing credentials, provided the target content is accessible under the same configuration.

Creating a New SharePoint Source Configuration

The following fields are required when setting up a new SharePoint source configuration in **AI Control Tower** > **Integrations** (administrators only):

Field	Description
Configuration Name	The display name for this configuration
SharePoint Hostname	The SharePoint hostname **Example:** contoso.sharepoint.com
Client ID	Azure App Registration Client ID
Client Secret	Azure App Registration Client Secret
Tenant ID	Azure Active Directory Tenant ID

Creating Additional Connectors from an Existing Source Configuration

When adding a new connector using an existing source configuration, only the SharePoint folder details need to be specified:

Field	Description
SharePoint Site Name	The name of the SharePoint site to connect to
SharePoint Drive Name	The name of the document library / drive
SharePoint Folder Path	The folder path within the drive to sync

> **NOTE:** This allows multiple connectors to reuse the same source configuration while targeting different folders or sites, reducing the administrative overhead of managing credentials.

Search Agent Overview

The Search Agent provides retrieval and synthesis across curated documentation stored in Knowledge Clusters. Users submit natural language questions, and the agent returns synthesized answers with citations to source documents.

Cluster Configuration

Search Agent uses Knowledge Clusters managed in SensibleAI Agents Documents. To set up a cluster, see Documents.

Access Configuration

Search Agent access is mapped to OneStream security groups in Control Tower > Settings > Agents.

Analytics

Search Agent usage is tracked in AI Control Tower under Analytics > Clusters and Analytics > Documents.

Deep Analysis Agent Overview

Deep Analysis automates document-heavy workflows by processing input documents in a structured fashion to provide deeper analytical insights across uploaded documentation. Deep Analysis shares the same documentation foundation as Search through Knowledge Clusters managed in SensibleAI Agents Documents. Go to **SensibleAI Agents Documents** > **Deep Analysis**.

AI Control Tower Monitoring

Deep Analysis usage metrics are available in AI Control Tower under Analytics > Analysis Plans:

- Plans saved
- Documents processed
- Most used plans by execution count
- Most used plans by documents processed

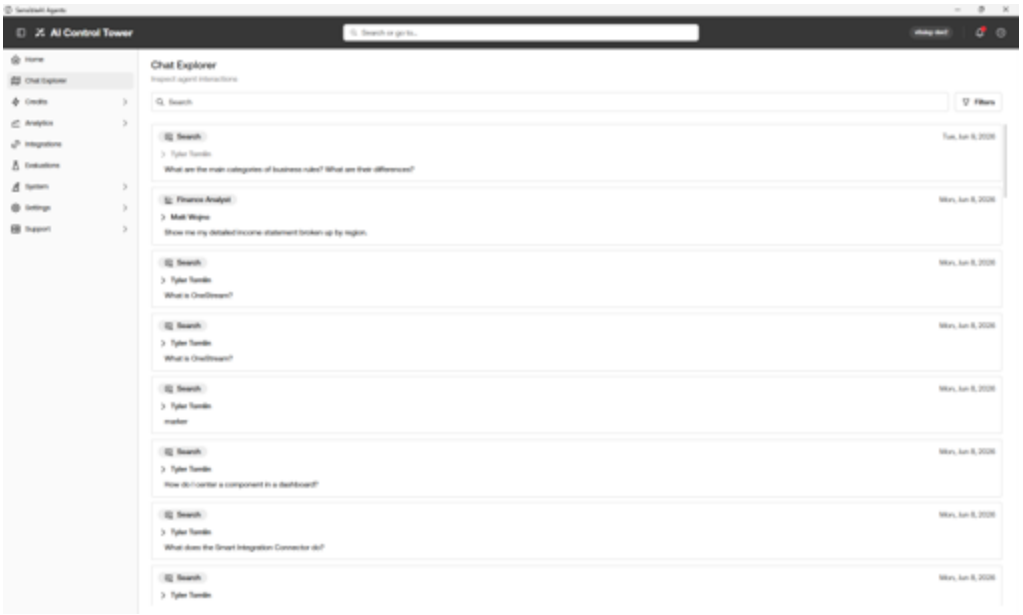
AI Control Tower Overview

AI Control Tower is the administrative portal for governance and visibility across the AI platform. It supports:

- Chat Exploration visibility
- Cost and Credit management
- Analytics
- Third Party Integrations
- Performance Evaluations
- System
- Settings
- Support

Access is permission-based and displays in the user profile menu when they are granted access through OneStream security groups.

Chat Explorer



This is a detailed log of interactions: Shows all turns (interactions) broken down by conversation. It shows full questions and system responses along with review details.

Credits

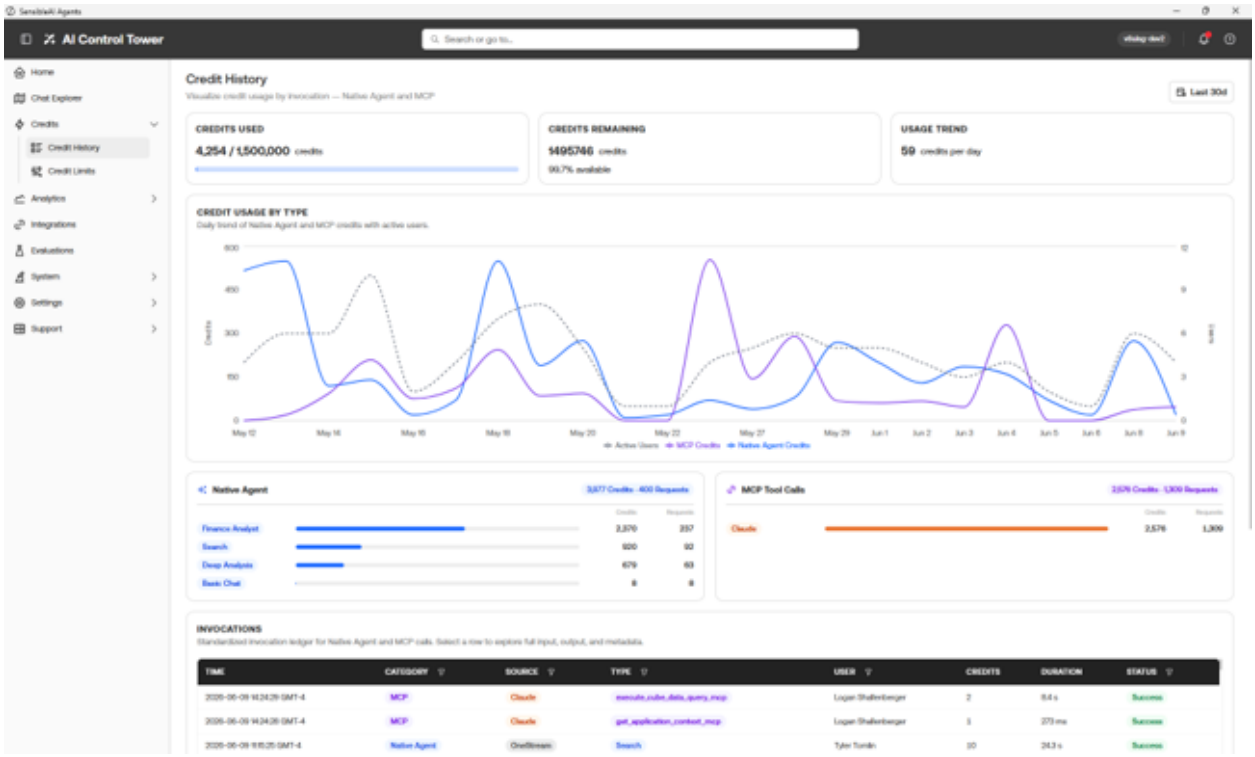
Credits provide visibility into AI credit consumption and limits. Credits refresh annually. This page displays a daily usage chart and the following KPIs:

- Credits Used
- Credits Remaining
- Days Until Refresh
- Usage Trend



Credit History

Credit History provides a breakdown by users or agents.

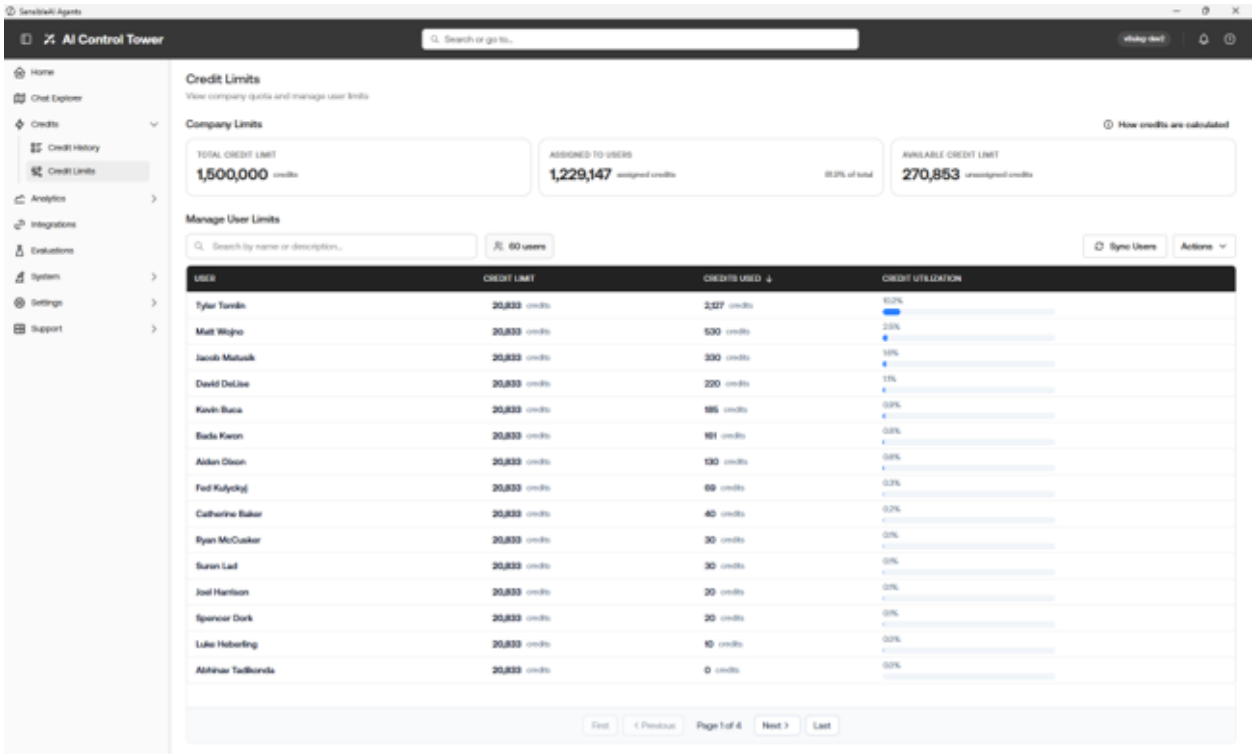


Credit Limits

> **NOTE:** This is SAI Administrator+.

This page manages company limits and user limits:

- Split evenly
- Configure details
- Sync Users
- **Configure Environment Operations Quota**: A credit limit for document uploads, connector syncs, and evaluation runs.



Review the following notes:

- Errors are not changed.
- Whole numbers only and rounded to whole numbers.

Analytics

The Analytics page provides usage and ROI views across:

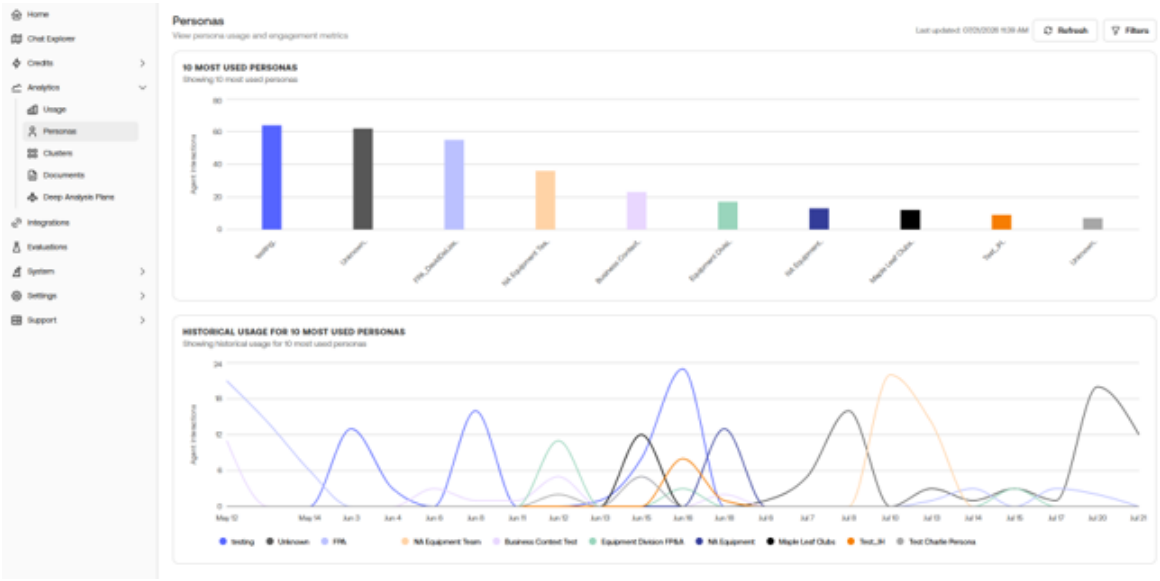
- Finance Analyst
- Search
- Deep Analysis

The dashboard provides the following KPIs:

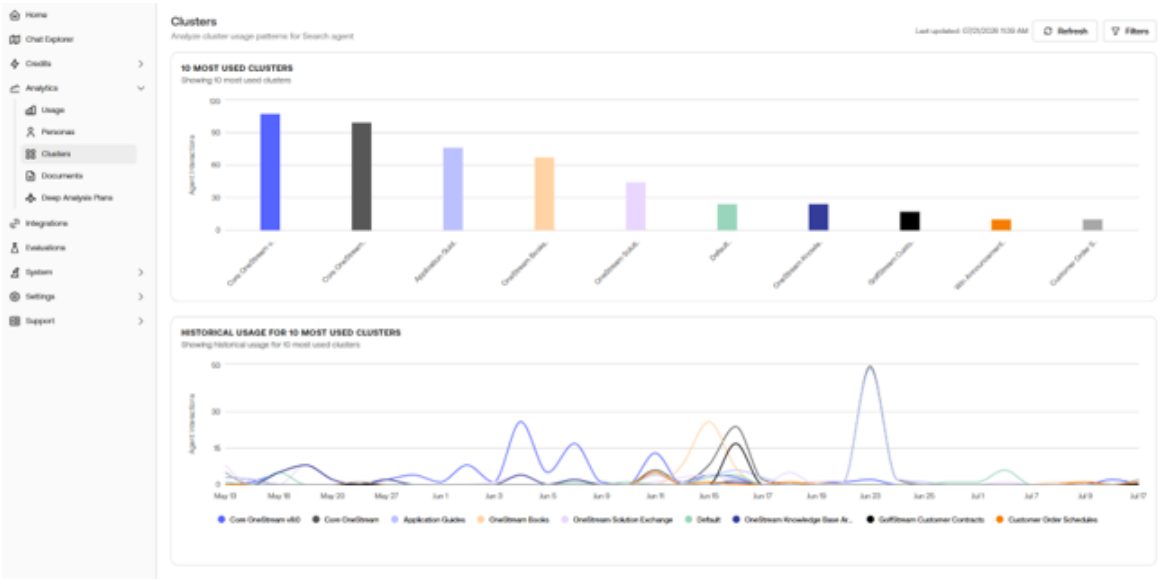
- Active Users
- Agent Interactions
- Time Savings
- Cost Savings



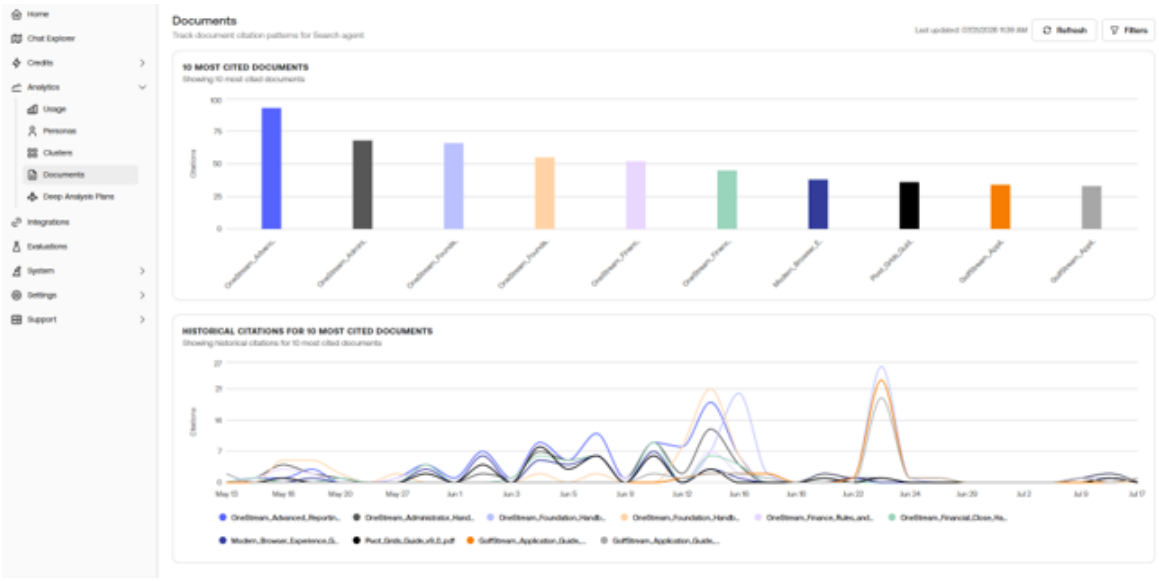
- **Personas**: Top 10 personas by interactions.



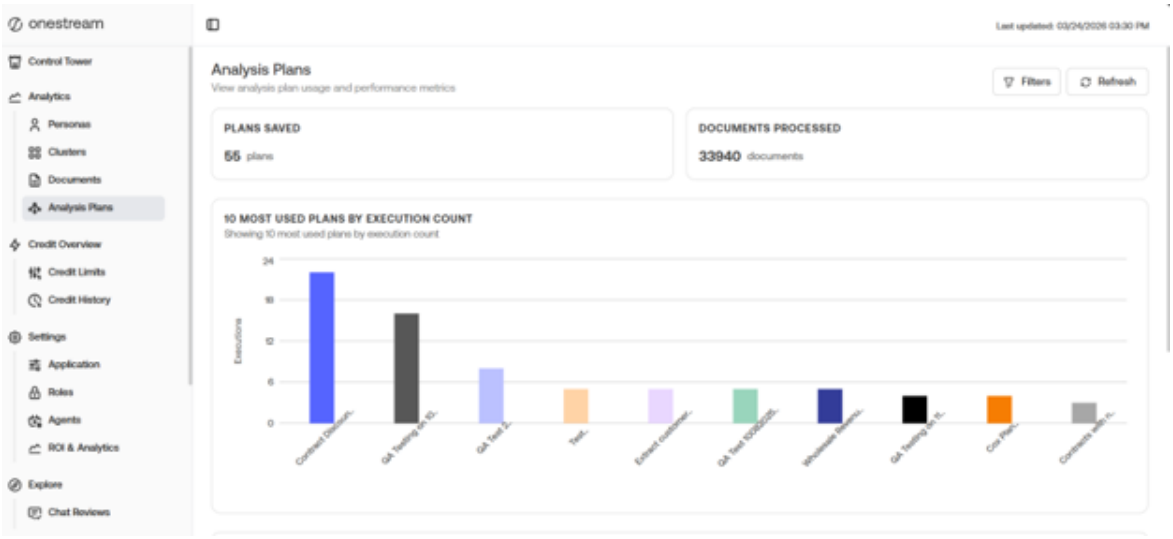
- **Clusters**: Search-specific.



- **Documents**: Search-specific.



- **Analysis Plans**: Deep Analysis-specific.



AI Platform Connections

The following guides provide step-by-step instructions for administrators to connect OneStream MCP servers to supported AI platforms and to set up document connector sources, such as SharePoint and Confluence, for SensibleAI Agents Documents. Each integration is initiated from AI Control Tower > Integrations.

Claude

Step-by-step guide for setting up OneStream MCP servers in a Claude environment.

1. Navigate to the Integrations page in the AI Control Tower and select **Configure** for the Claude option.
2. Copy the connection details required to connect your Claude environment.
3. Continue in the Control Tower to the skill download section and download the skill file.
4. Open Claude Desktop and navigate to Organization Settings.
5. Go to the **Connectors** tab and add a custom web connector: click Add > Custom > Web. Provide a name and paste the connection details from step 2 into their respective fields. Click **Add** when complete.
6. Verify the MCP server appears in the Web Connectors section.
7. Navigate to the Skills section (above Connectors) and upload a new skill. Select the markdown file downloaded in step 3, this provides Claude with context on when and how to use the OneStream tools.

Users in the Claude environment are now ready to connect to the MCP server from their device.

ChatGPT

Step-by-step guide for setting up OneStream MCP servers in a ChatGPT environment. Note: this integration requires a support ticket with your OneStream product delivery team before configuration can be completed.

1. Navigate to the Integrations page and select **Configure** for the ChatGPT option.
2. Create a customer support ticket with your OneStream product delivery team. ChatGPT generates random callback URIs, so OneStream must allowlist your specific Callback URI before the connection can be established.
3. Obtain your ChatGPT Callback URI: navigate to Organization Settings > Apps > + Create. Enter any name and server URL, then click Advanced OAuth Settings and copy the custom Callback URL.
4. Paste the Callback URI into the OneStream SensibleAI Control Tower input section and include it in the support ticket.
5. Once the ticket is processed, re-enter the ChatGPT configuration modal and copy the connection details.
6. Input the connection details into the ChatGPT MCP app connector.
7. Continue through the Control Tower connection flow and download the skills markdown file.
8. In ChatGPT, navigate to Workspace Settings > Skills and create a new organization skill. Upload the skills markdown file downloaded in step 7.

Once all steps are completed, users are ready to connect to the MCP server.

Microsoft Copilot

Step-by-step guide for setting up OneStream MCP servers in a Microsoft Copilot environment.

1. Navigate to the Integrations page and select Initialize Connection for Copilot. You will receive multiple configuration values used to set up OneStream authentication for the Copilot agent.
2. Navigate to the Microsoft Developer Portal and go to Tools > OAuth Client Configuration.
3. Create a new OAuth Client Registration using the values from step 1, completing the following fields:
    a. App Settings: Registration Name, Base URL, scope set to ‘My organization only’, App type set to ‘Any Teams App’
    b. Auth Settings: Client ID, Client Secret, Authorization Endpoint, Token Endpoint, Refresh Endpoint, Scope, enable Proof Key for Code Exchange (PKCE), Request Body Parameters (default)
4. Save the configuration, copy the Reference ID to your clipboard, and return to the OneStream Copilot connection flow to enter it.
5. Download the appPackage zip file from the Control Tower.
6. Navigate to the M365 Admin Center and go to All Agents > Add Agent.
7. Upload the manifest from the appPackage downloaded in step 5 and grant the required permissions.

Once uploaded, users can access the OneStream tools from within their Copilot tenant.

Gemini

Step-by-step guide for setting up OneStream MCP servers in a Gemini environment.

1. Navigate to the Gemini integrations connector in the AI Control Tower.
2. Copy the JSON config object displayed in the Control Tower.
3. Locate your .gemini folder (typically at C:\Users\<user>\.gemini) and open settings.json.
4. Paste the JSON config object from step 2 into settings.json and save the file.
5. Verify the MCP server is recognized by running /mcp in Gemini. Once it appears, initiate authentication with /mcp auth <server name>.
6. Once authenticated, the MCP server is ready to use.

Document Connector Sources

Document connector sources hold the credentials and connection settings for data sources such as SharePoint and Confluence. SAI Administrators and SAI Support users manage them in **AI Control Tower** > **Integrations**. To add a source, select **Configure** on the SharePoint Site or Atlassian Confluence card, complete the form, and save. Saved sources list their name, source type, creator, creation date, and number of connectors. Select a source to view its details and linked connectors, edit it, or delete it. Deleting a source also deletes every connector that uses it. Users can create connectors in Documents only after a source exists. For SharePoint field requirements, see Creating a New SharePoint Source Configuration under Data Connectors.

Evaluations

Role Access: SAI Control Tower Viewer and above

Overview

The Evaluations page in the OneStream Control Tower enables users to test and monitor the performance of configured AI agents, such as Finance Analyst, over time.

This page is intended for administrators, implementers, and delivery teams who need to verify that agent responses are:

- Accurate
- Repeatable
- Explainable
- Aligned with business context

Purpose

Evaluations provide a structured and repeatable method for validating AI agent responses against expected outcomes.

Without Evaluations, teams must manually:

- Test prompts
- Compare outputs
- Track results across iterations

This approach does not scale, particularly when changes are made to:

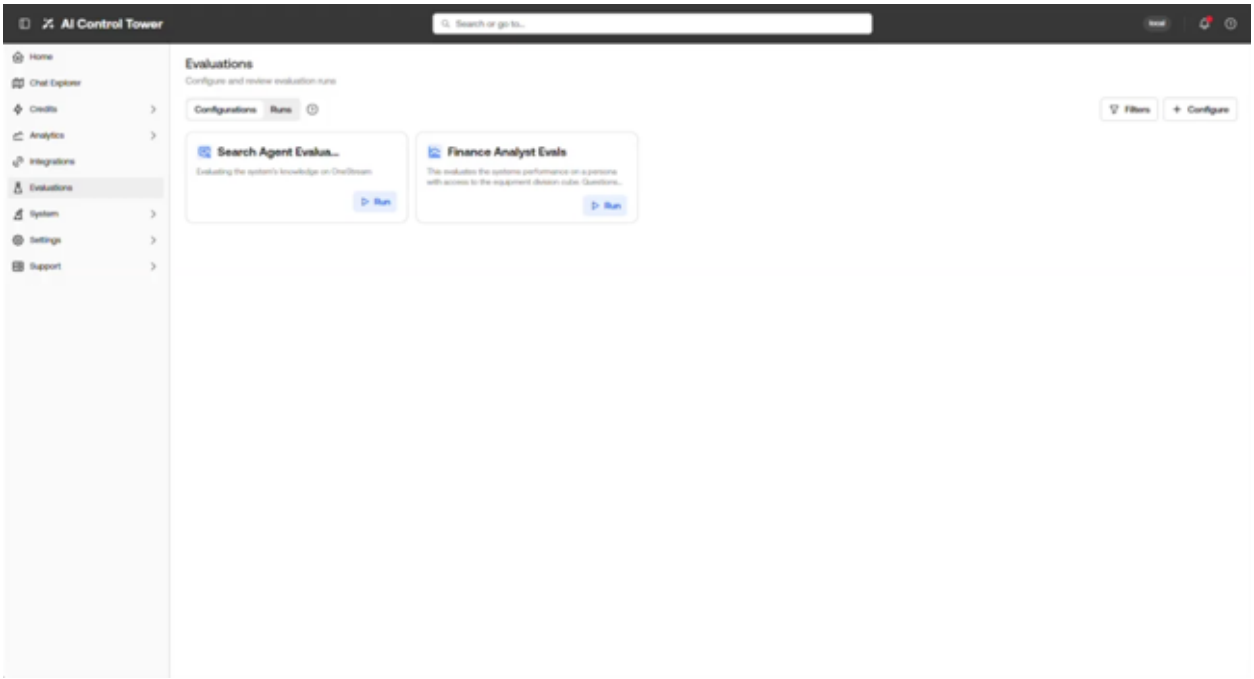
- Prompts
- Models
- Data sources

- Personas
- Agent configurations

The Evaluations engine enables:

- Automated validation against expected outputs
- Consistent testing across multiple runs
- Scoring with explanation using the LLM Judge

Configurations



The Configurations tab is used to create and manage evaluation configurations. Each configuration defines a reusable validation setup.

When creating a configuration, define the following:

- Agent Type: Specifies the agent to evaluate (for example, Finance Analyst or Search).
- Agent Configuration: Defines the parameters used by the agent during the evaluation.
    - Finance Analyst: This includes the persona (which controls data access and business context, such as access to a specific cube) and the time period (which ensures consistency across runs by fixing the evaluation to a specific point in time, for example 2025M3 or 2025M4). Configuration options vary by agent type.
    - Search: This includes the clusters that the agent searches against.
    - Deep Analysis: Includes the plan and cluster that are to be used.
- Repetitions: Specifies how many times the evaluation is executed. Running multiple repetitions helps validate consistency across responses.
- Input Dataset: Provides the set of inputs used to test the agent. This may include expected outputs for comparison.

Runs

The Runs tab displays each execution of an evaluation configuration.

Selecting a run provides detailed results, including:

- Input provided to the agent
- Expected output (if defined)
- Actual agent response
- Evaluation score (if LLM judge was used)

The LLM Judge evaluates each response and provides:

- A score
- An explanation of which elements were correct or incorrect

This enables users to understand not only the result, but also the reasoning behind the evaluation.

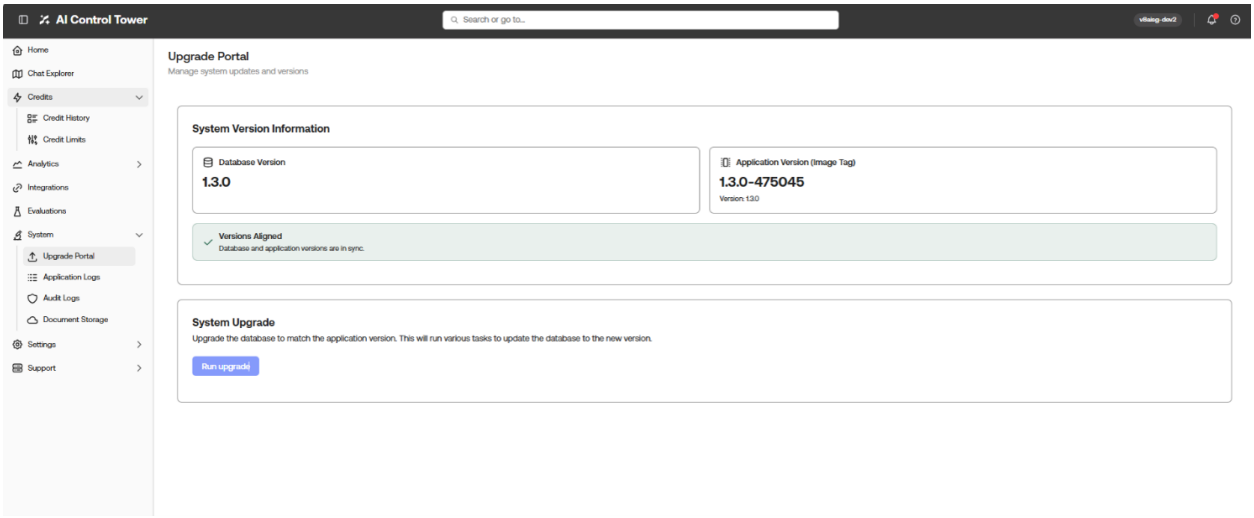
System

The System section provides SAI Support users with some of the main tools needed to monitor platform health: check the Upgrade Portal, monitor application logs, review audit logs, and manage document storage.

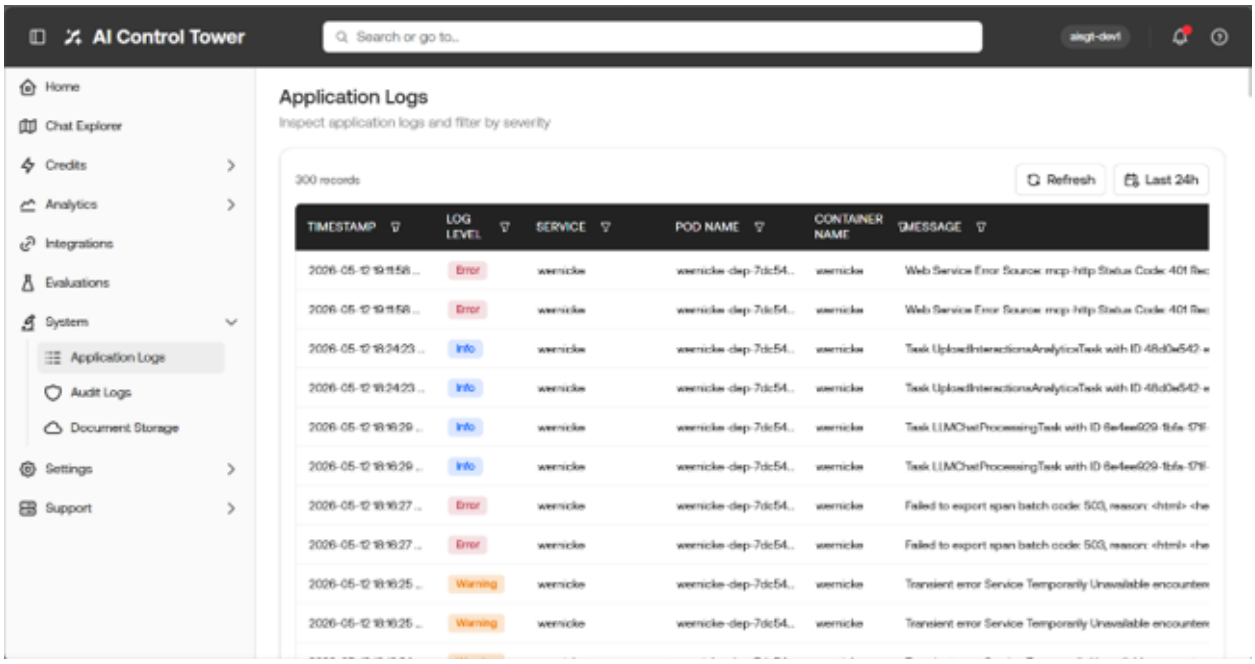
Upgrade Portal

Role Access: SAI Support and above

This page shows the database version, application version, and sync status. Upgrades may require users to be logged out.



Application Logs

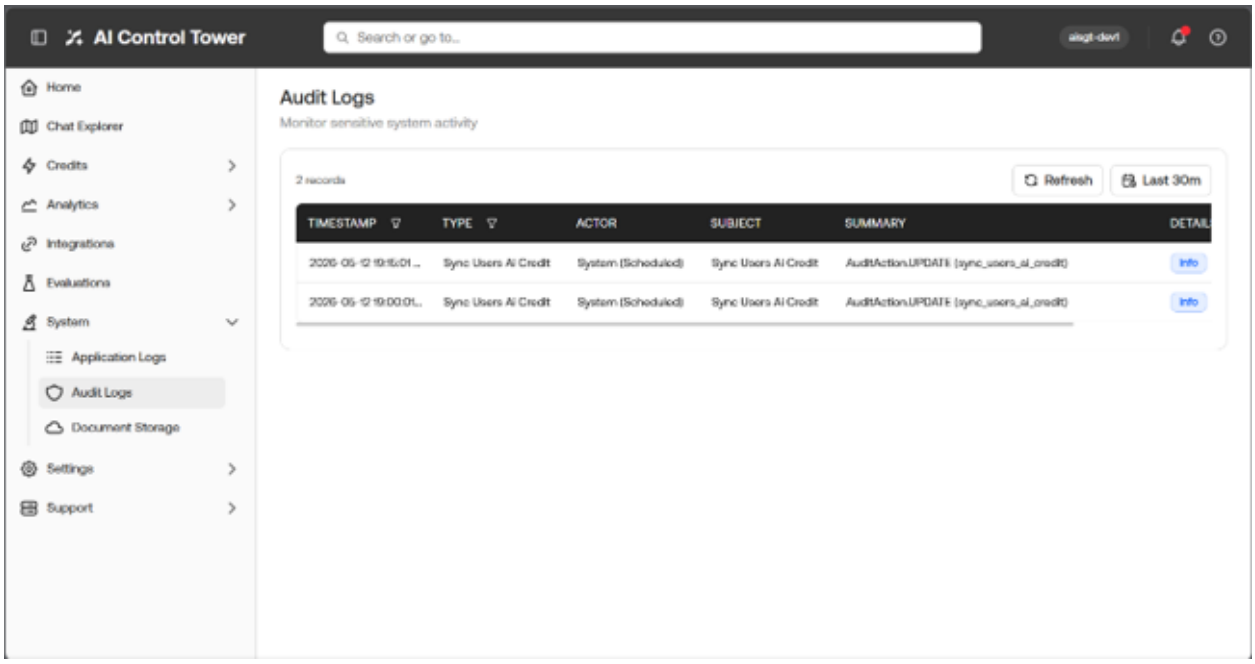


Role Access: SAI Support and above

The Application Logs page provides a full record of the application logs that occur across all services within the system. Each entry in the log captures the timestamp, log level, service, pod name, container name, and the log message. Individual records can be drilled into for a full inspection.

This page is a critical tool for troubleshooting, giving SAI Support users a complete trail of all application logs.

Audit Logs

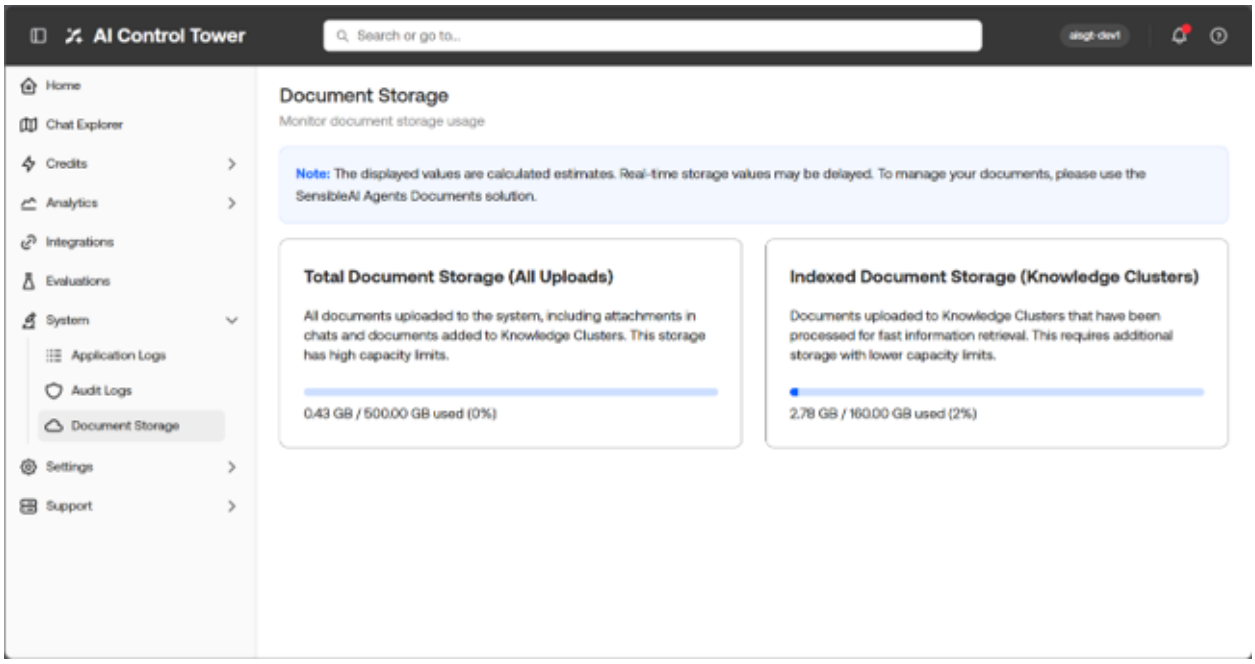


Role Access: SAI Support and above

The Audit Logs page provides a full record of sensitive system activity across the platform. Each entry in the log captures the timestamp, action performed, entity type, entity ID, user ID, and the before and after state of the record. Individual records can be drilled into for a full inspection.

This page is a critical tool for compliance, security monitoring, and troubleshooting, giving SAI Support users a complete trail of all significant system-level changes.

Document Storage



Role Access: SAI Support and above

The Document Storage page gives administrators visibility into how much storage the platform is consuming across two categories:

Storage Type	Description
Total Document Storage (All Uploads)	All documents uploaded to the system, including chat attachments and Knowledge Cluster documents. This storage has high capacity limits.

Storage Type	Description
Indexed Document Storage (Knowledge Clusters)	Documents uploaded to Knowledge Clusters that have been processed for fast information retrieval. This requires additional storage with lower capacity limits.

> **NOTE:** Displayed values are calculated estimates and real-time values may be delayed. To manage documents, use the SensibleAI Agents Documents solution.

Settings

> **NOTE:** This is SAI Administrator+.

Settings includes the following pages:

- Application Settings
- Roles
- Agents
- ROI & Analytics

Application Settings

Configure application settings and preferences. The following settings are currently included:

- **Company Default Application**: The OneStream application name used by default.
- **Feature Flags**: Turn application features on and off. Currently, the only feature flag is agent_memory.

agent_memory allows the Finance Analyst agent to save user-level memories based on a user's interactions with the agents, building up a store of user preferences over time.

Roles



The following roles are mapped to OneStream security groups:

- **SAI Control Tower Viewer**: Base read-only access. This role can view Analytics, Credit Overview, Credit History, Explore (Chat Reviews and Interaction Explorer), and the Control Tower dashboard.

- **SAI Administrator**: Viewer access plus the ability to manage Credit Limits and full Settings (Application Settings, Roles, Agents Settings, and ROI & Analytics).
- **SAI Support**: All Administrator access, plus elevated access to the Data Explorer (Explore) and the entire System section (Upgrade Portal, Application Logs, Audit Logs, Document Storage) and the Support section (Data Manipulator, Support Task Execution, Traces Explorer).
- **OS Admin**: OneStream platform administrator inherits all Control Tower access.

Agents Settings

This maps Control Tower roles to OneStream security groups and agent access to OneStream security groups (Finance Analyst, Search, and Deep Analysis).



ROI & Analytics

This page enables you to configure inputs saved in Time Savings and Cost Savings:

- Time Savings Multiplier
- Average Salary



Support

Support includes the following pages:

- Data Manipulator
- Support Task Execution
- Traces Explorer

Data Manipulator

Data Manipulator Login: Requires data manipulator access.

Data Manipulator: Restricted to support engineers. Authorized users can directly view and edit SensibleAI Agents database records.

Support Task Execution

Role Access: SAI Support and above

Support Task Execution lets you start background system tasks on demand and track them while they run. Use it to run maintenance jobs outside their normal schedule, clean up data, or check system health.

> **IMPORTANT:** These tasks run with elevated access and can affect the whole platform.

Traces Explorer

Role Access: SAI Support and above

Traces Explorer shows how an agent handled a request, step by step: the question, each step the agent took, how long each step ran, and the final response or error. Use it to find out why an answer was wrong, why a request was slow, or where a request failed.

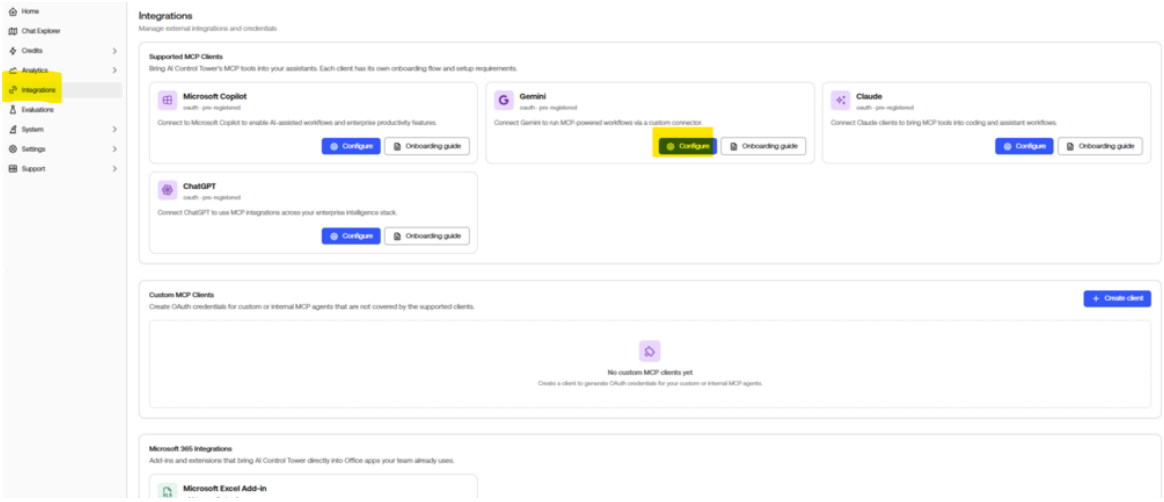
OneStream MCP Servers in Gemini

This section describes how to set up OneStream MCP servers in your Gemini environment. Gemini Business and Gemini Enterprise use different admin consoles. Use the subsection that matches your setup.

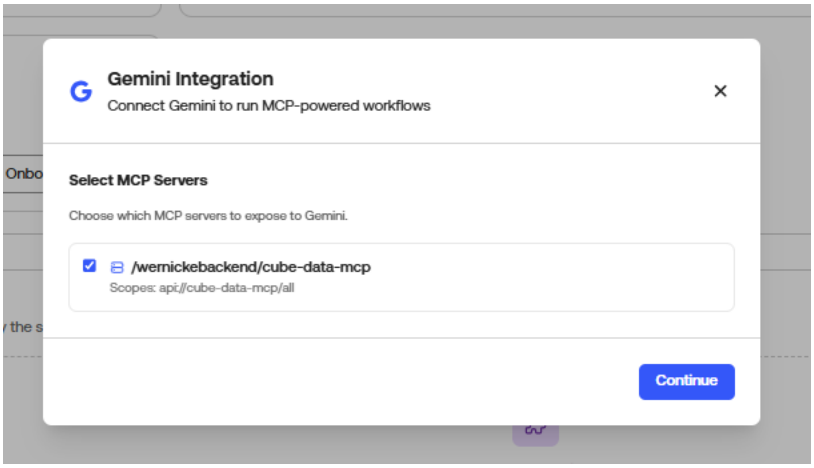
Complete the Initial Steps

Complete the following steps, and then follow the section for your Gemini edition.

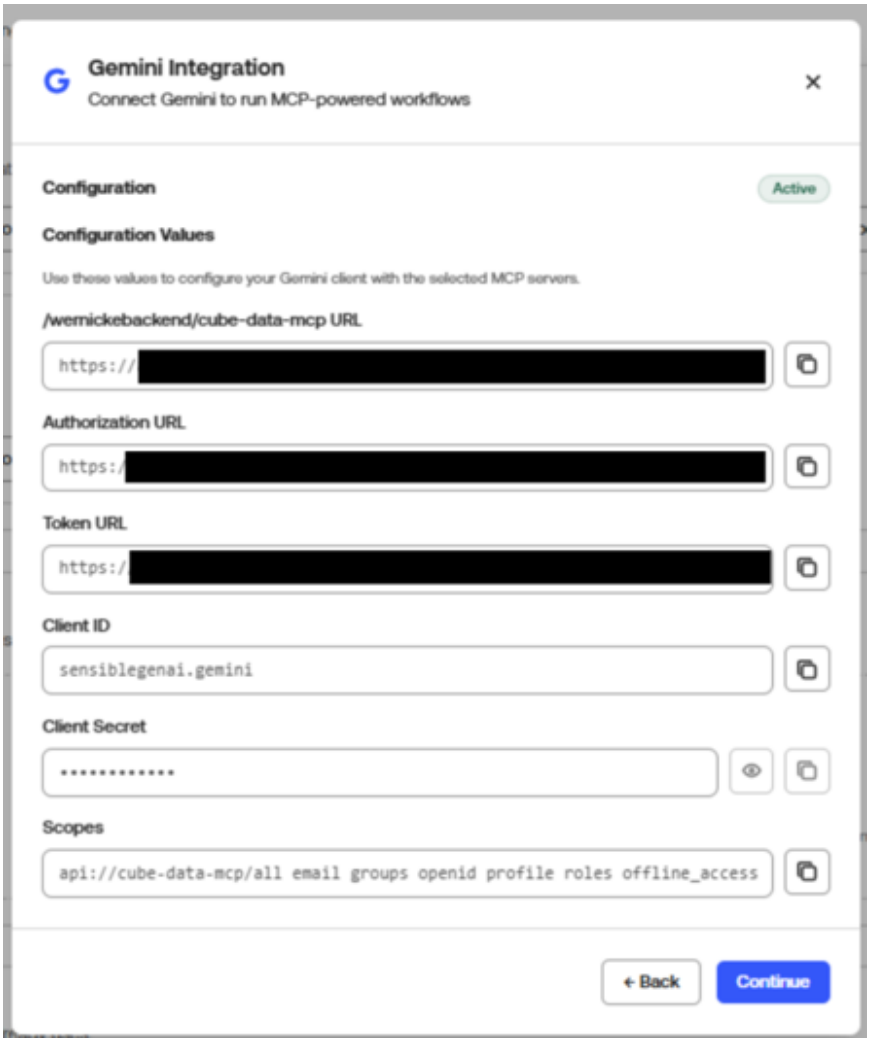
1. Go to the integration page and, for the Gemini integration, select **Configure**.



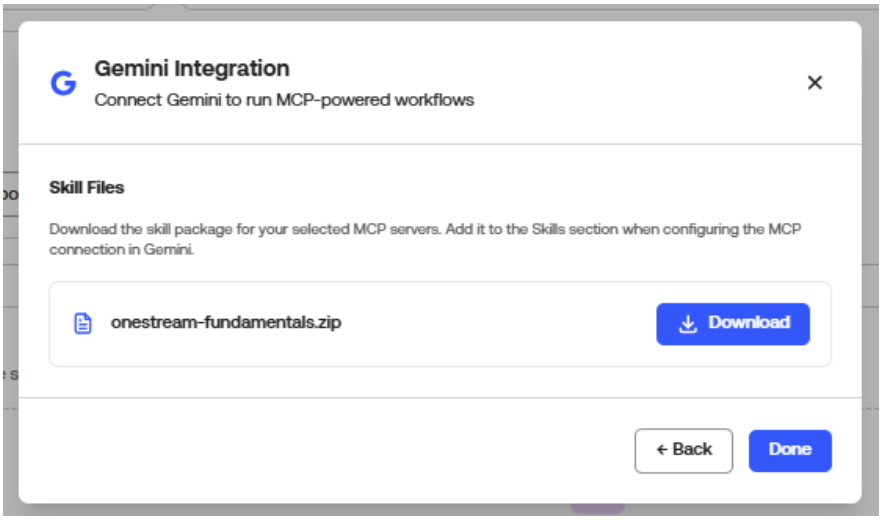
2. Select **/wernickebackend/cube-data-mcp**, and then select **Continue**.



3. Copy the details required to connect your Gemini environment. The values in the image are blocked out because your details will be different.

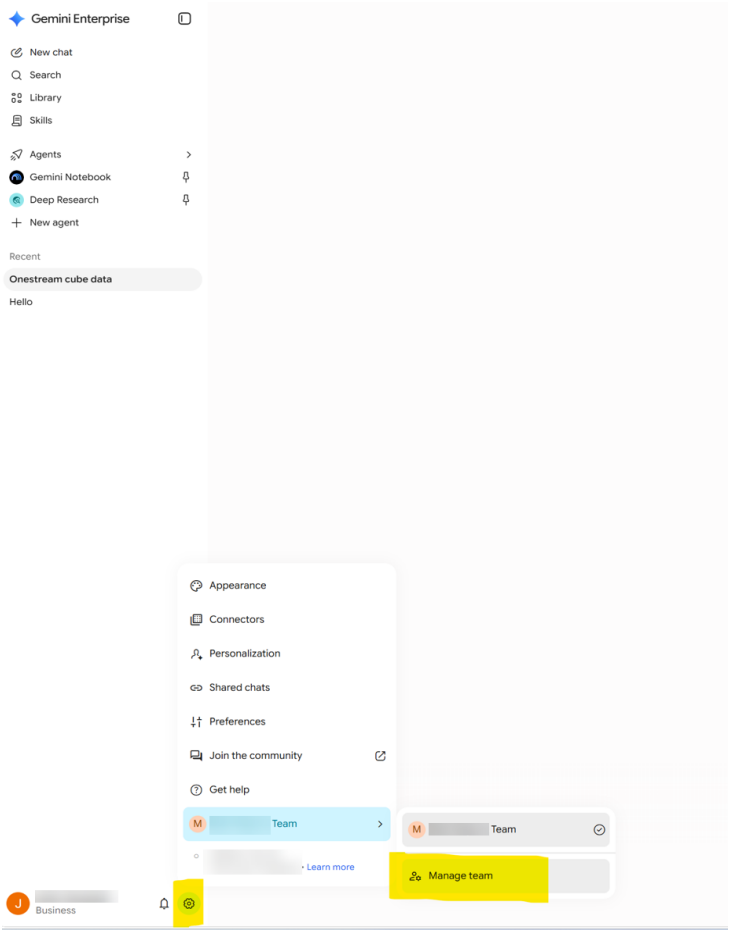


4. Download the skill file.

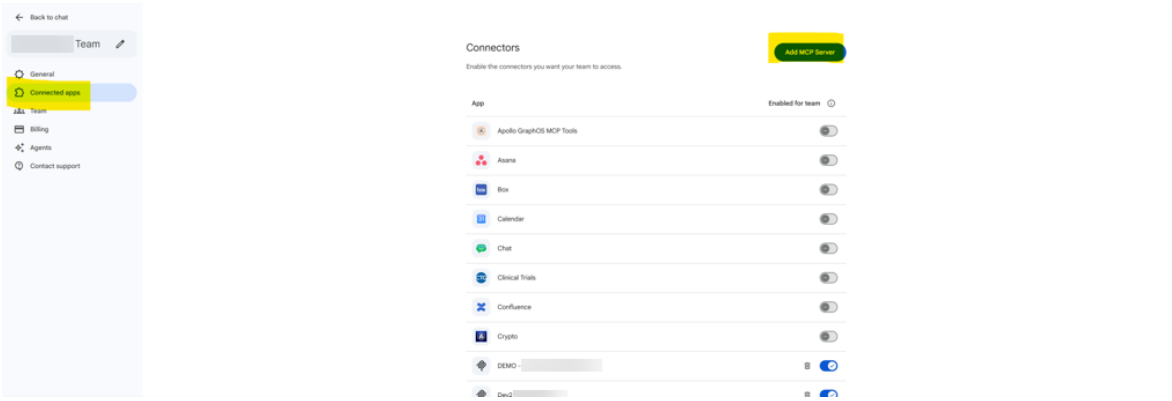


Gemini Enterprise Business

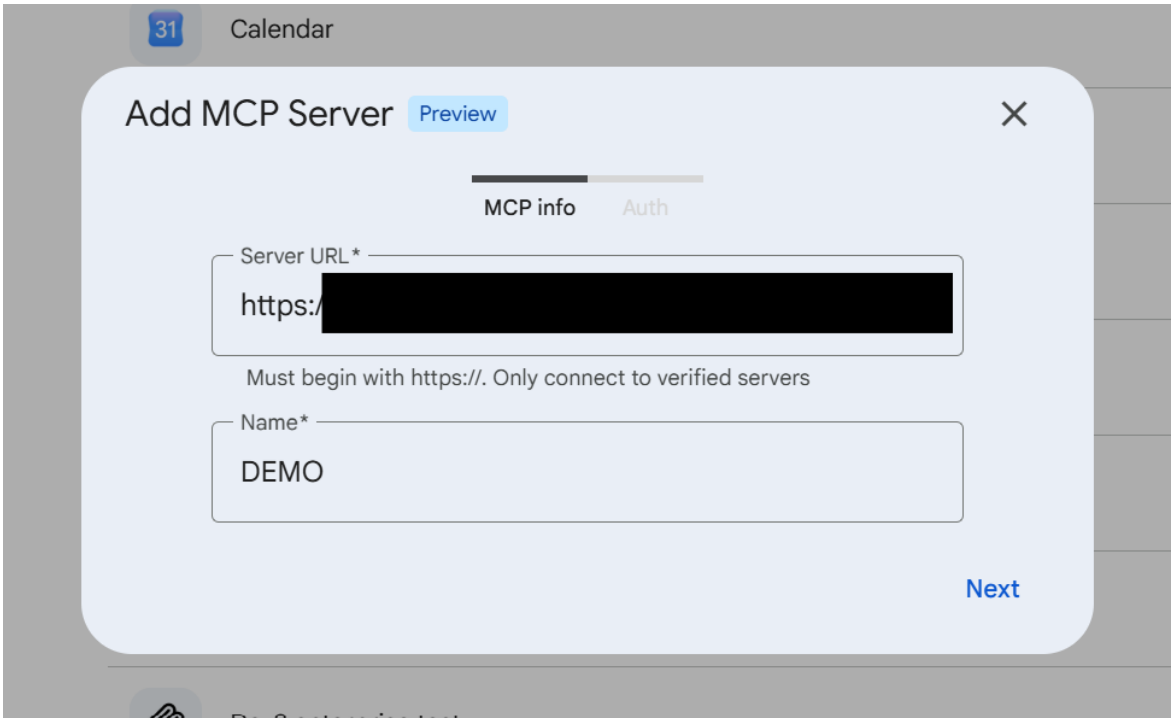
1. Sign in to your admin account at Gemini Enterprise (Business edition) at Gemini Business Sign In.
2. Select **Settings**, select your Team, and then select **Manage team**.



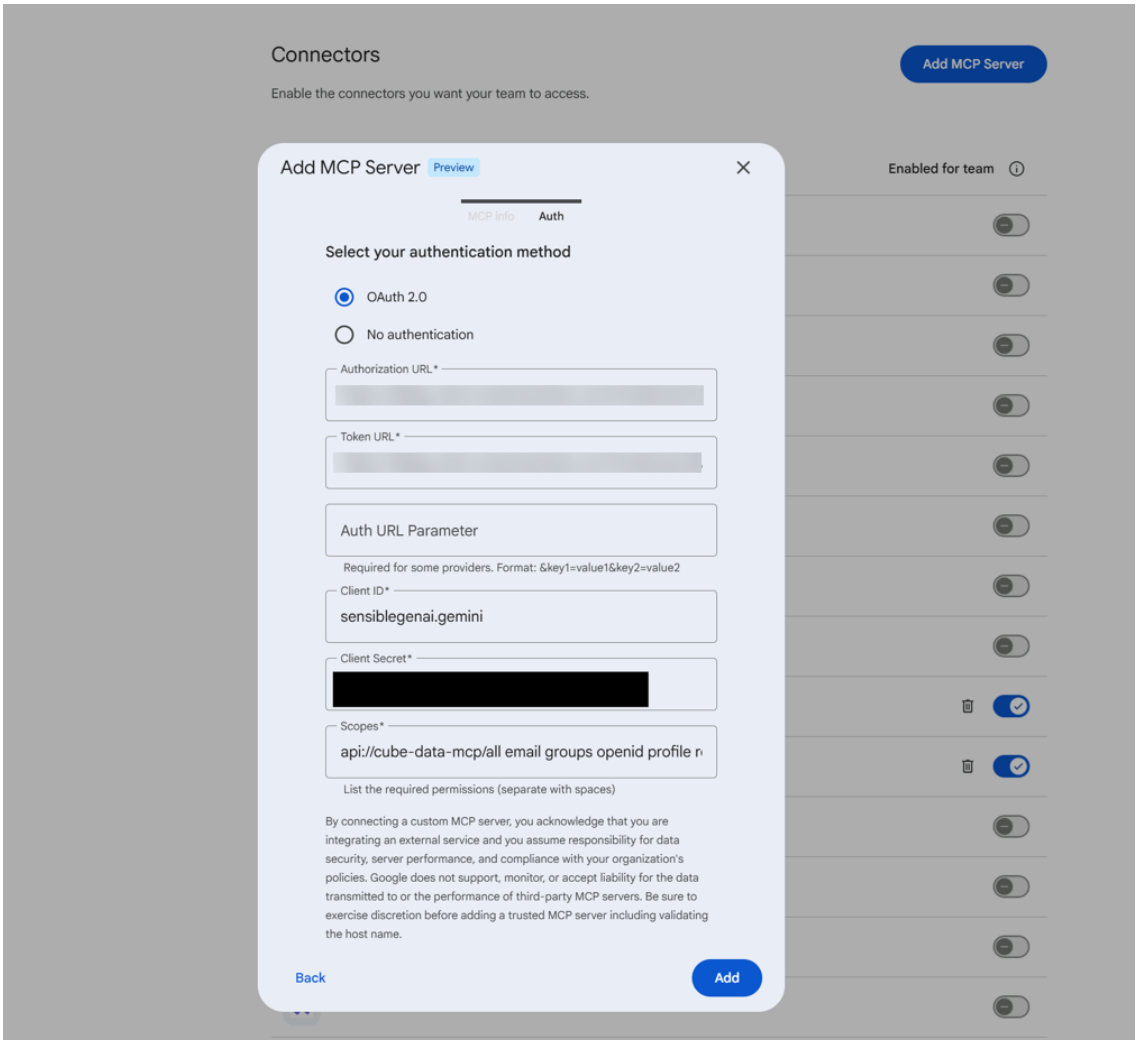
3. Select the **Connected apps** tab, and then select **Add MCP Server**.



4. In the **Server URL** field, enter the URL from OneStream. In the **Name** field, enter a name for the connector, and then select **Next**.



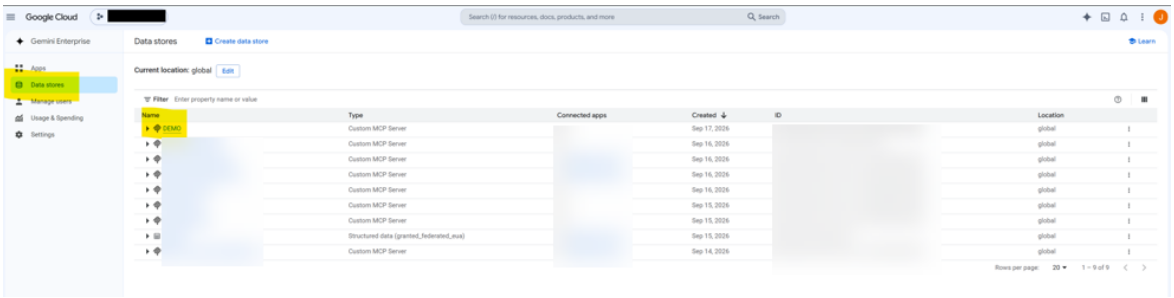
5. Under **Authentication settings**, select **OAuth 2.0.** In the **Authorization URL**, **Token URL**, **Client ID**, **Client Secret**, and **Scopes** fields, enter the values from OneStream, and then select **Add**.



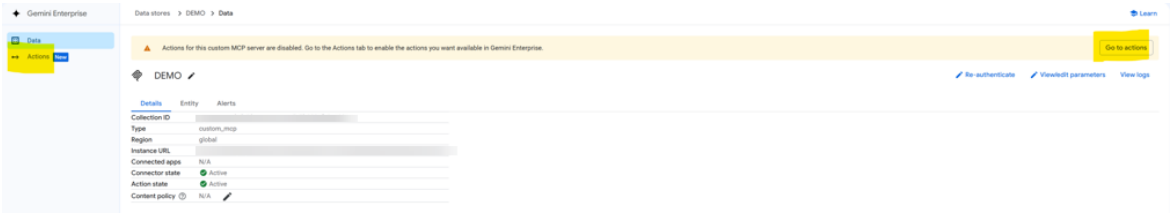
6. In **Connected apps**, turn on the new connector so your team can see it.



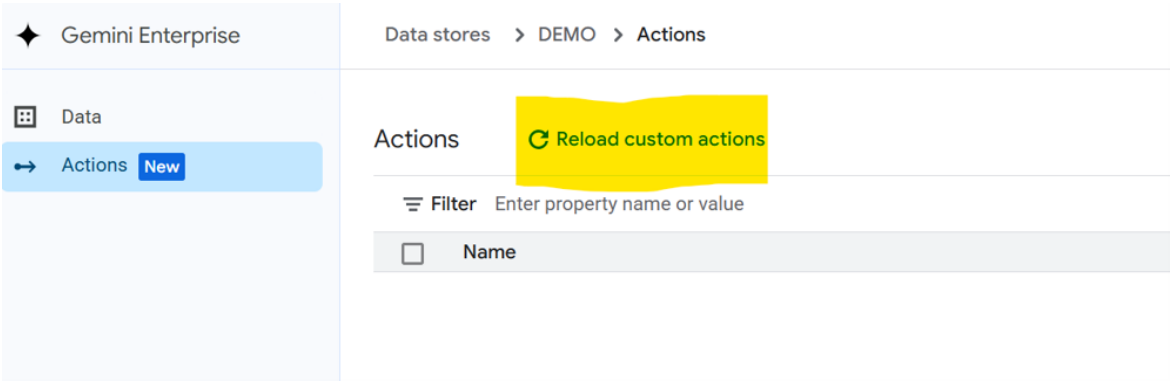
7. The connector does not return results yet. Go to https://console.cloud.google.com/gemini-enterprise/data-stores and locate the project registered to your Business account. On the **Data Stores** page, find the data store with the name you gave the connector. It displays with the type **Custom MCP Server**.



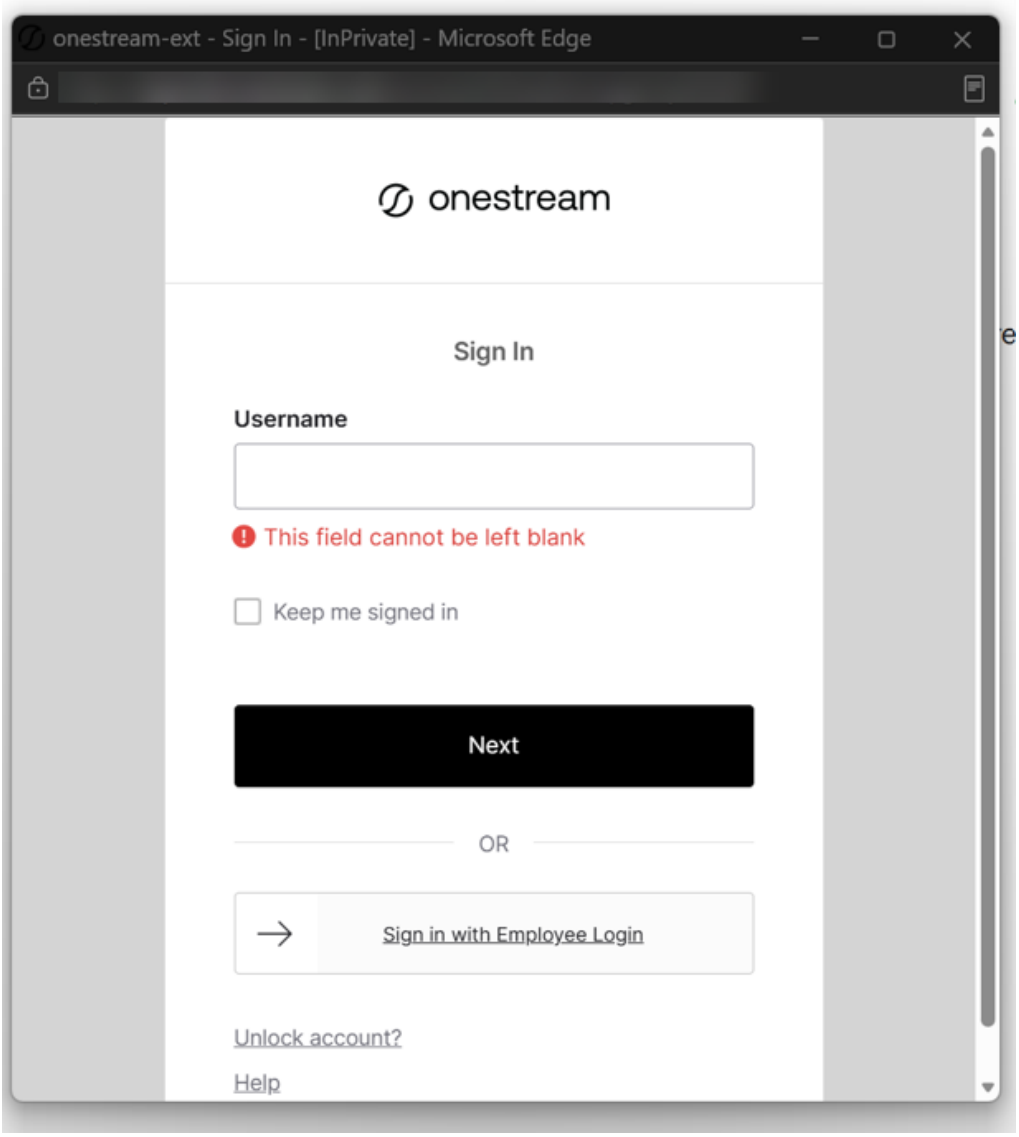
8. Select the **Custom MCP Server** you created, and then select the **Go to Actions** button on the right or **Actions** in the navigation panel on the left.



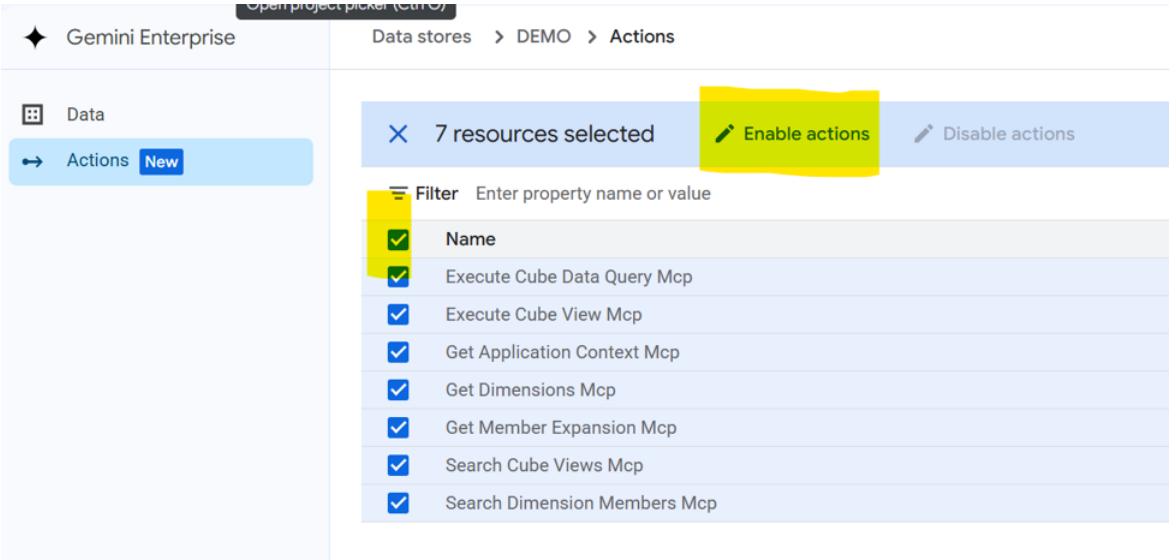
9. Select the **Reload custom actions** button.



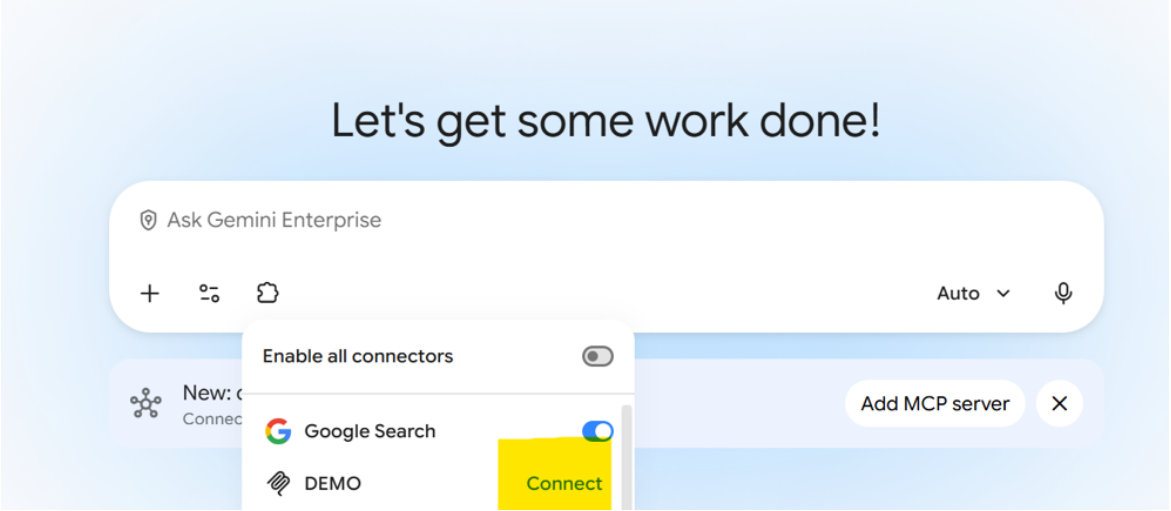
10. In the pop-up window, sign in with your OneStream credentials.



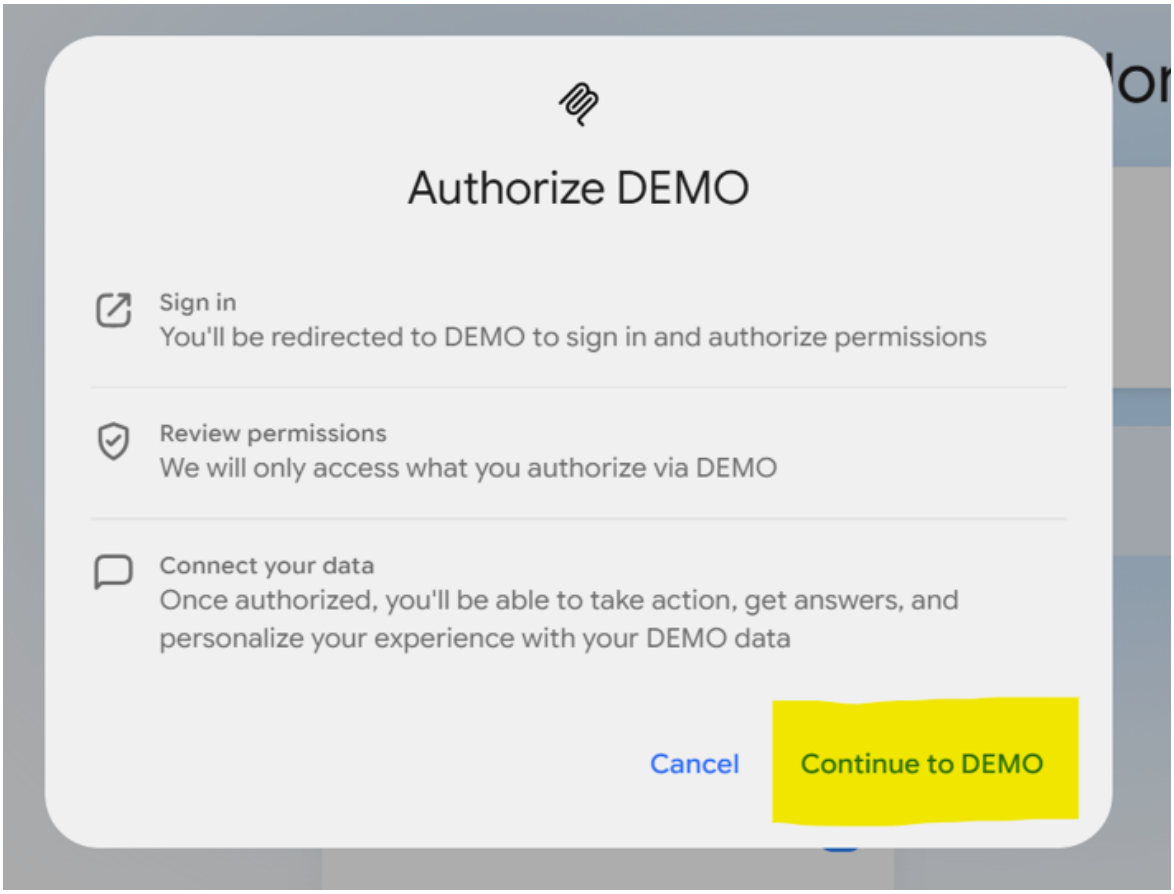
11. All MCP tools should display. Select all tools, and then select **Enable Actions**.



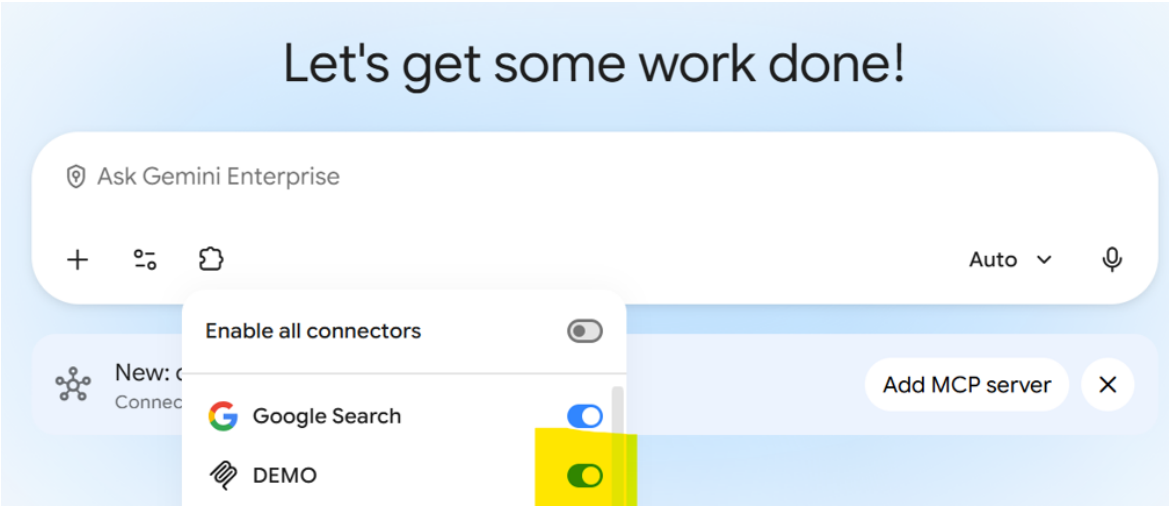
12. Return to the Gemini Enterprise chat window and select **Connectors**. Next to the name of the Custom MCP Connector you created, select **Connect**.



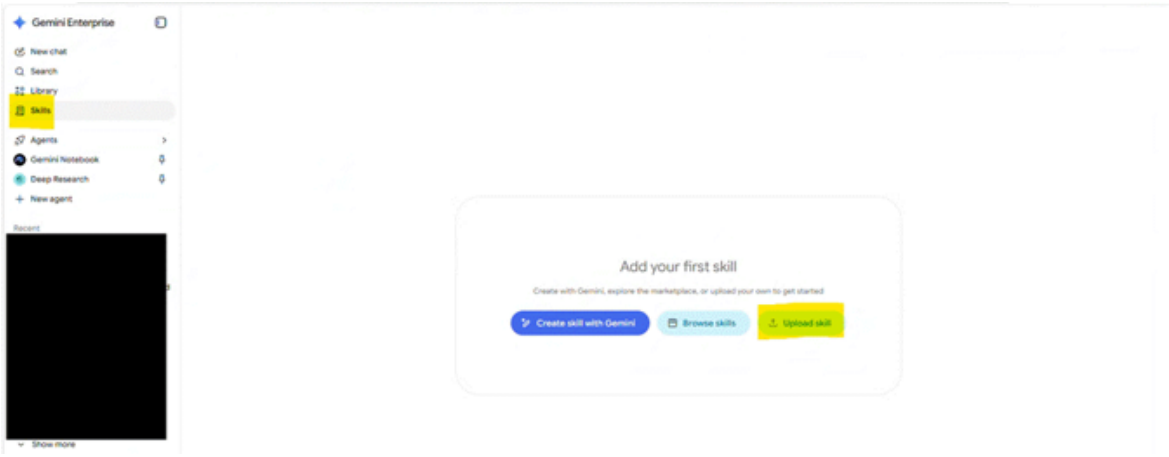
13. Select **Continue to DEMO**.

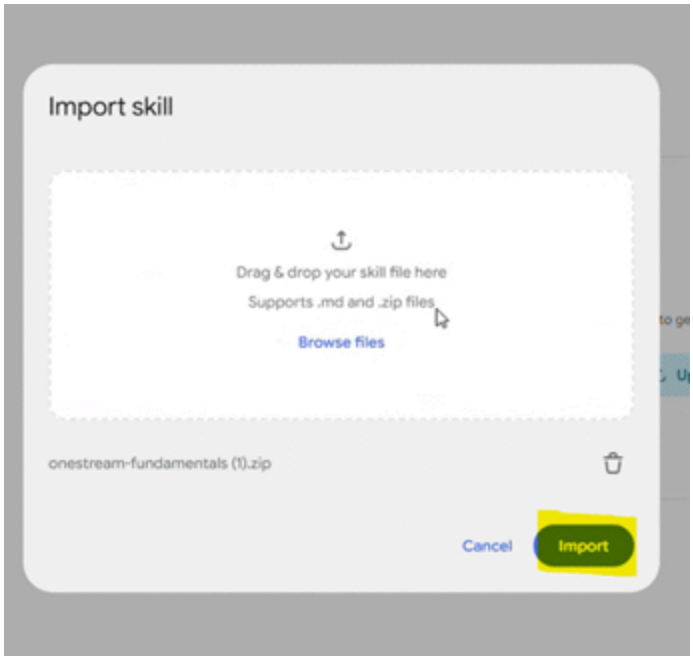


14. After that finishes, turn on the toggle for your OneStream Custom MCP.

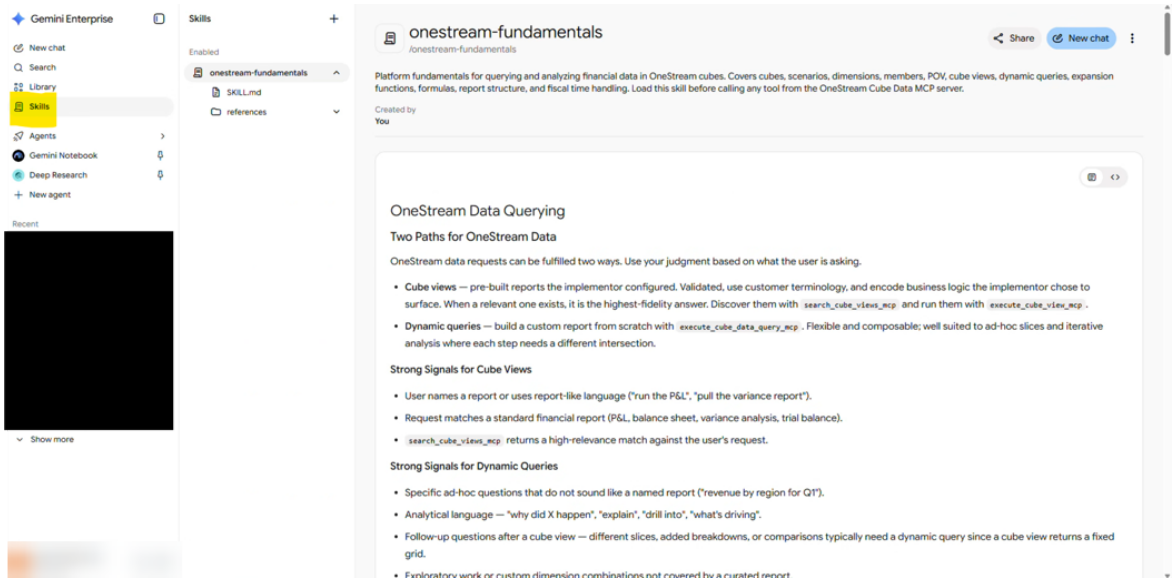


15. Under the **Skills** tab, upload the onestream-fundamentals.zip file you downloaded earlier.



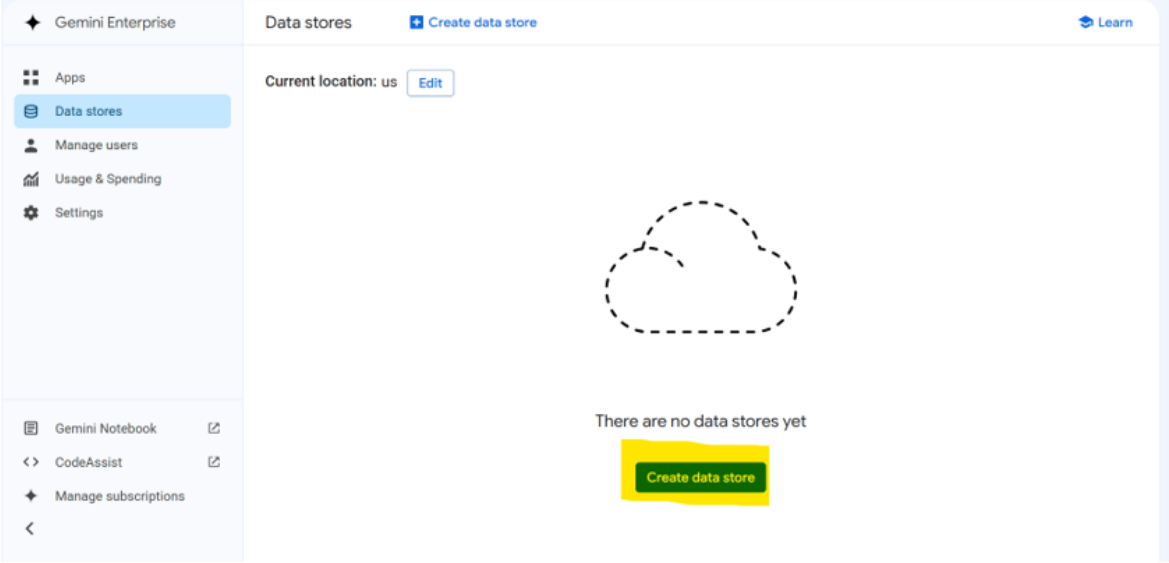


16. A successful skill upload looks like the following:

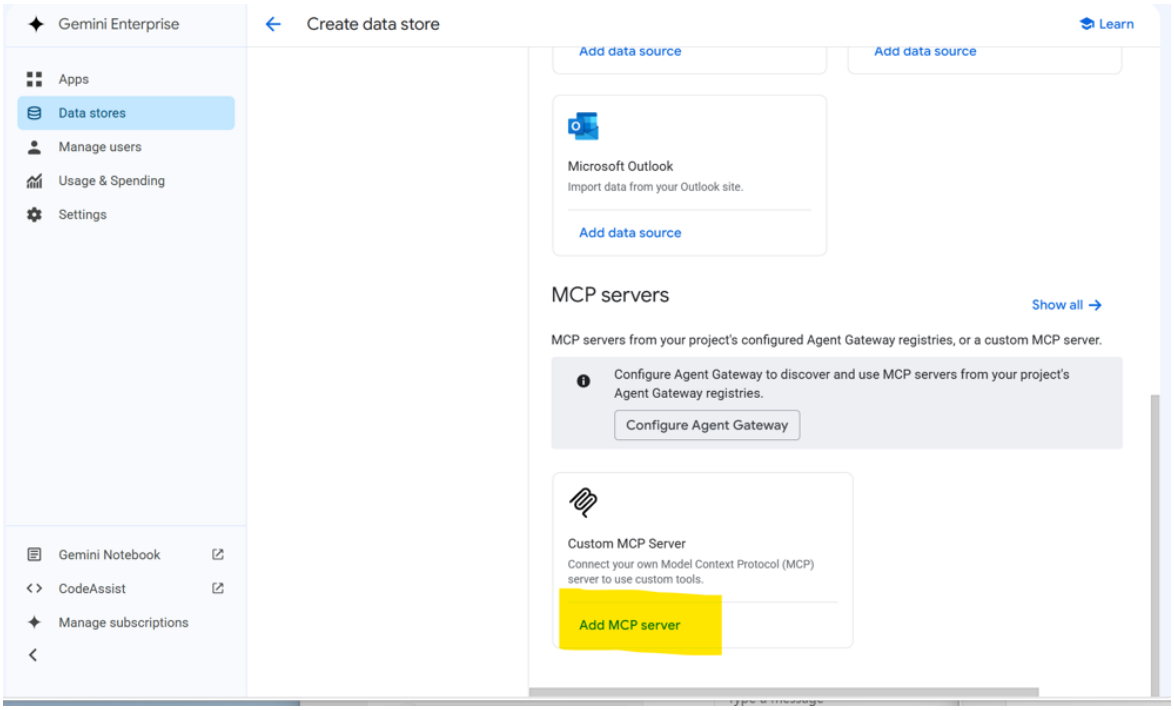


Gemini Enterprise Standard / Plus

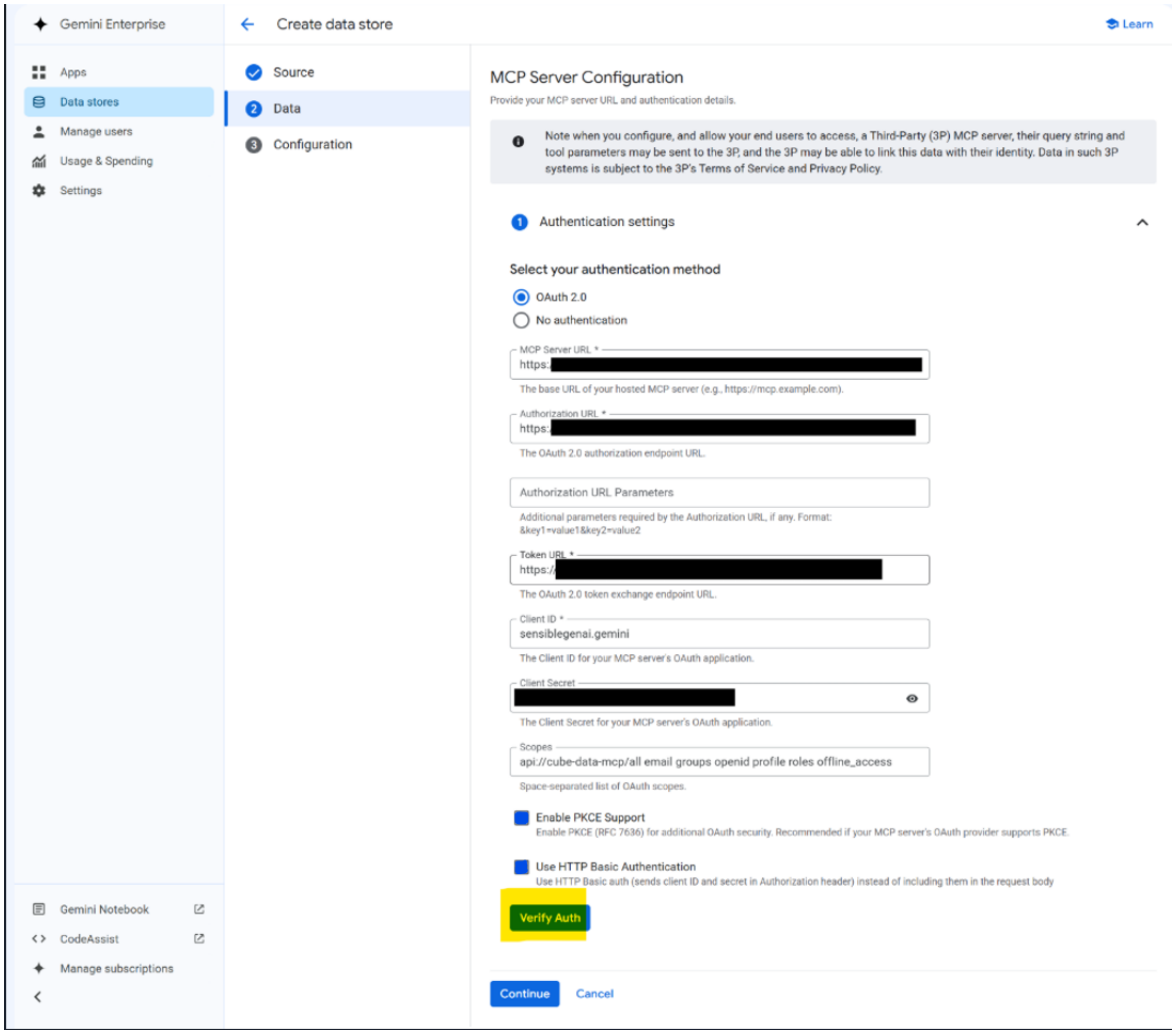
1. Go to the **Data Stores** page in Gemini Enterprise at https://console.cloud.google.com/gemini-enterprise/data-stores, and then select **Create data store**.



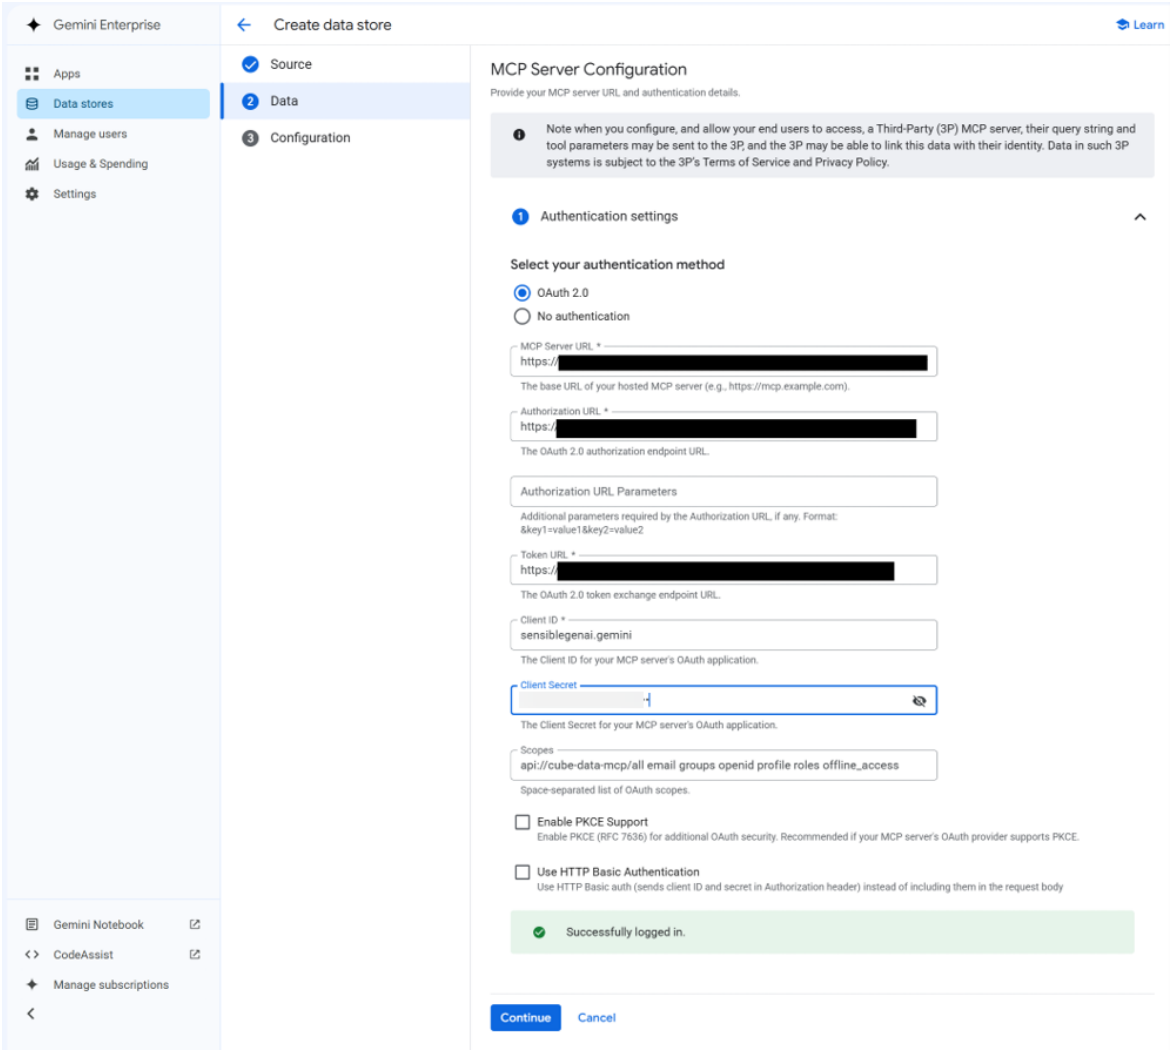
2. Scroll down and select **Add MCP server**.



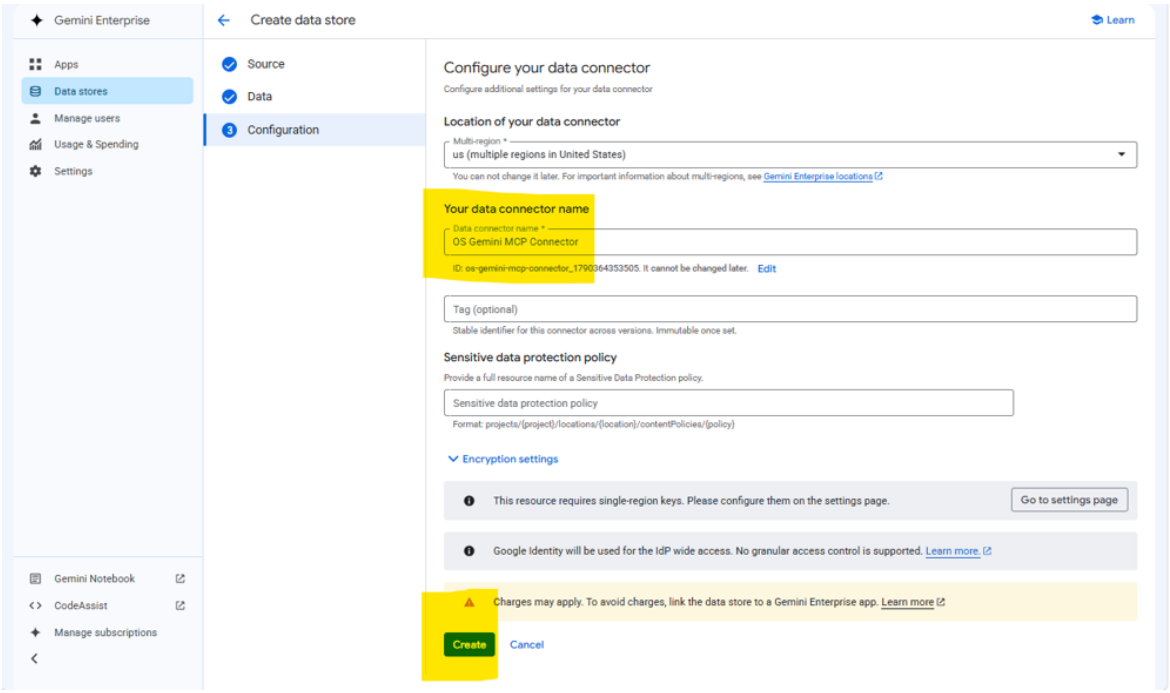
3. Enter the credentials from OneStream, and then select **Verify Auth**.



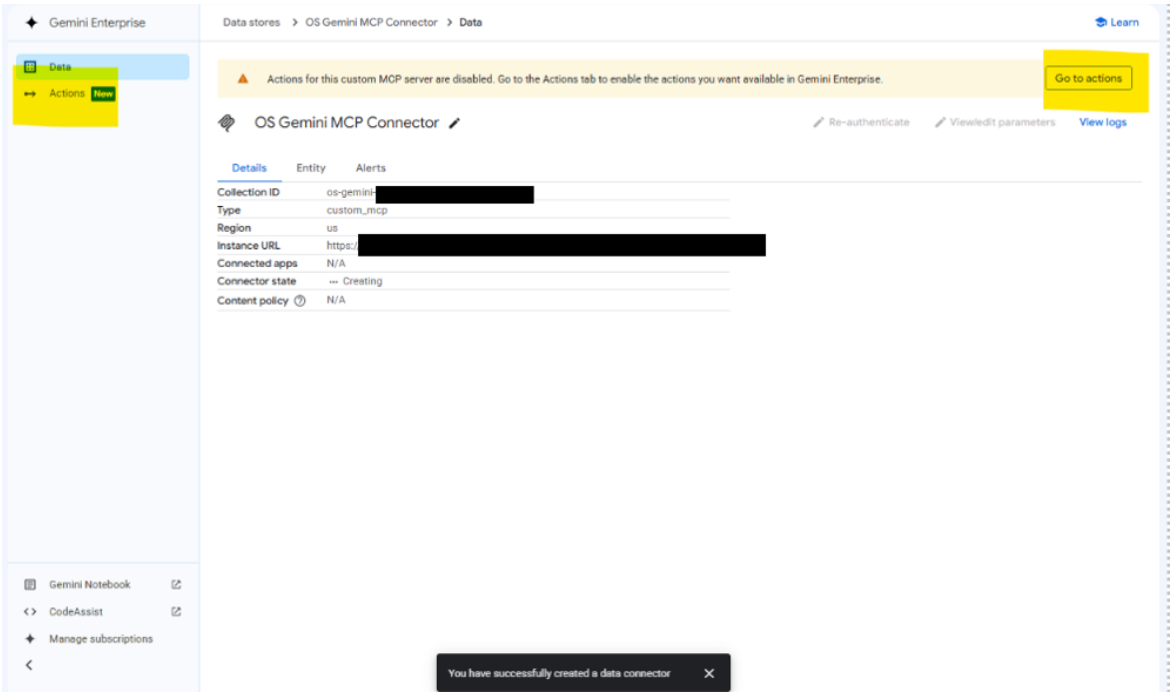
4. The **Successfully logged in** message displays. Select **Continue**.



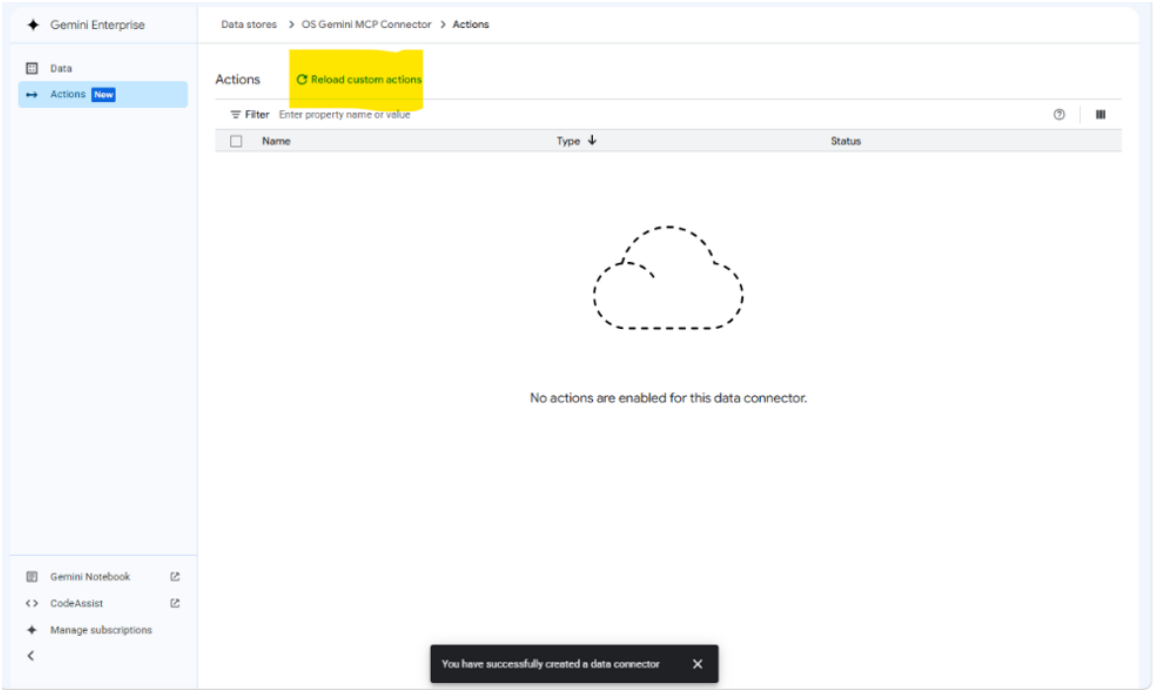
5. Enter a data connector name, and then select **Create**.



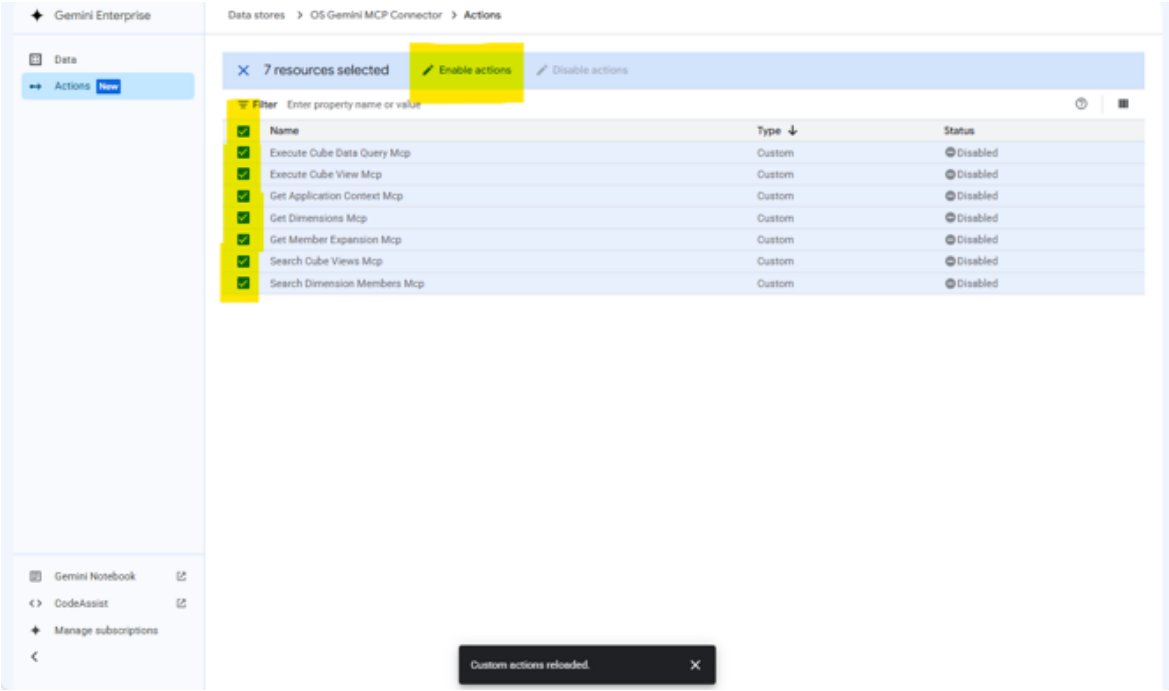
6. Select either highlighted **Actions** button to open the actions page.



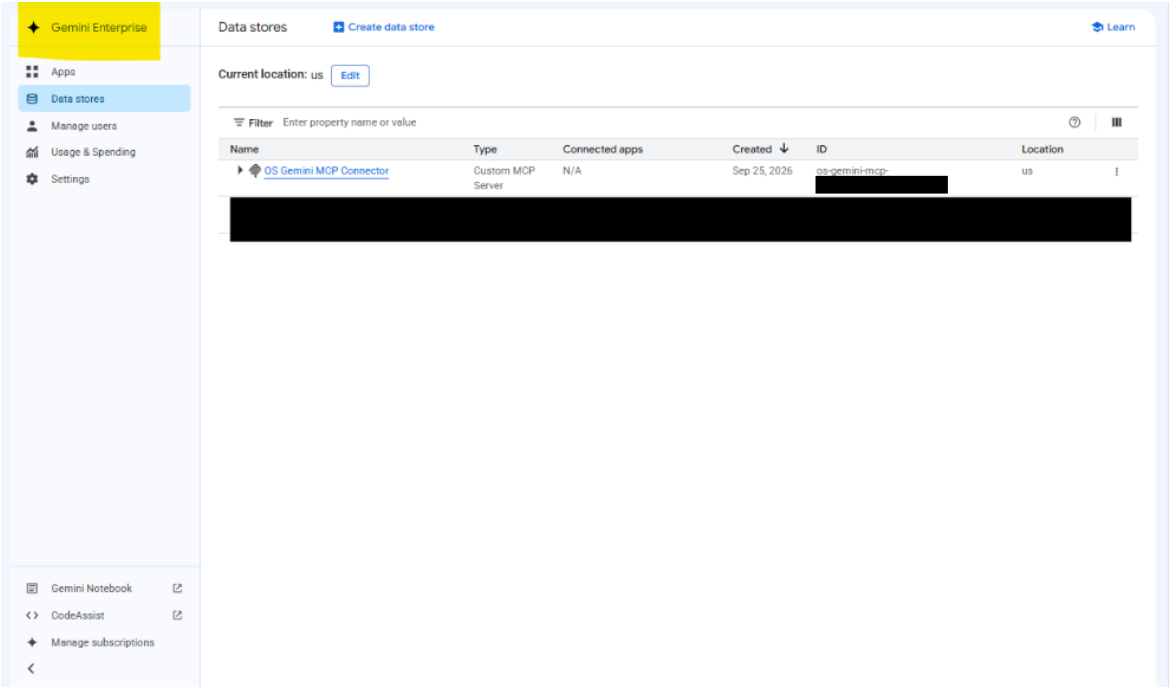
7. Select **Reload custom actions**.



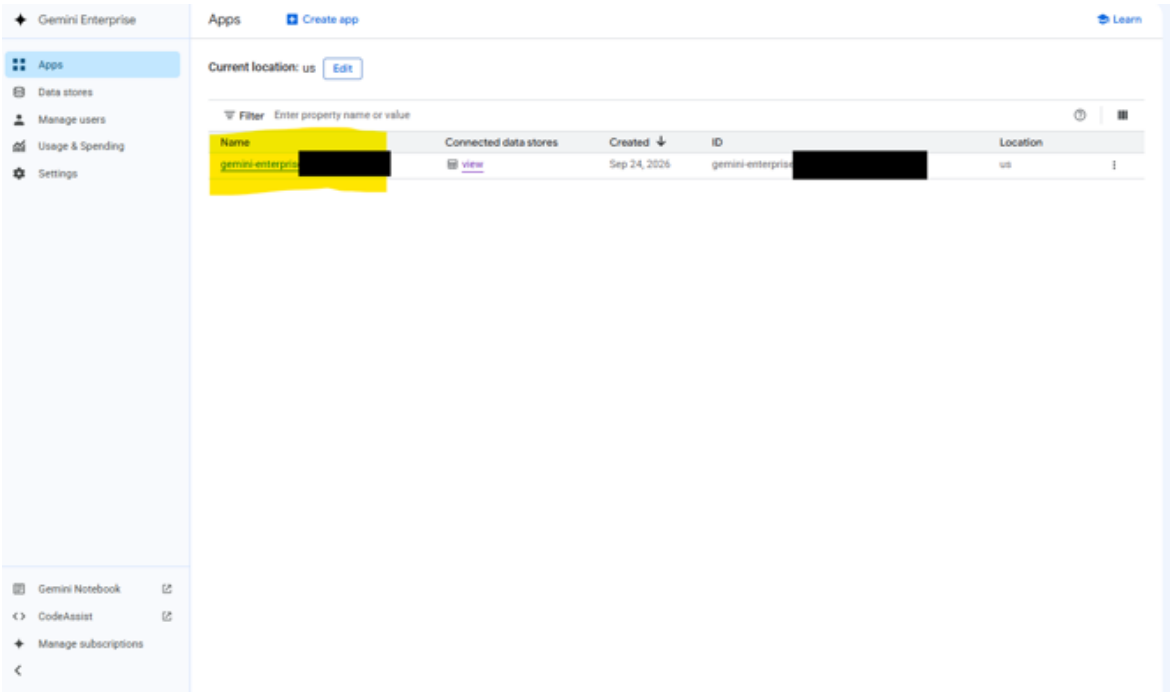
8. Select **All Actions**, and then select **Enable actions**.



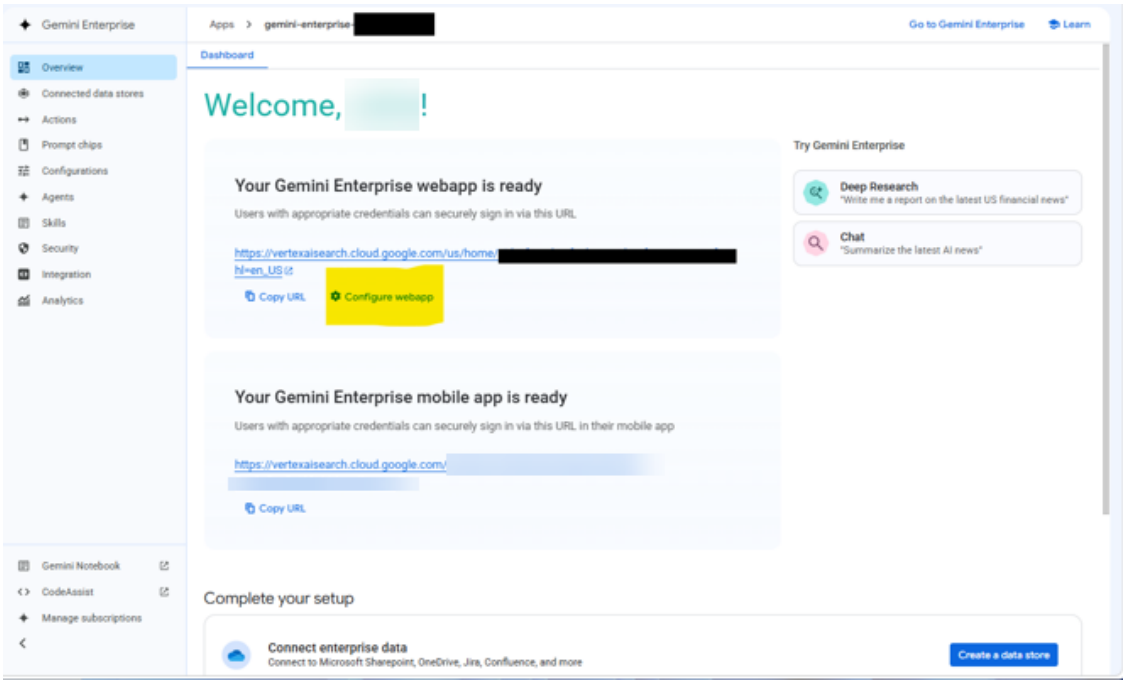
9. Confirm the MCP connector is listed on the **Data Stores** page, and then select **Gemini Enterprise**.



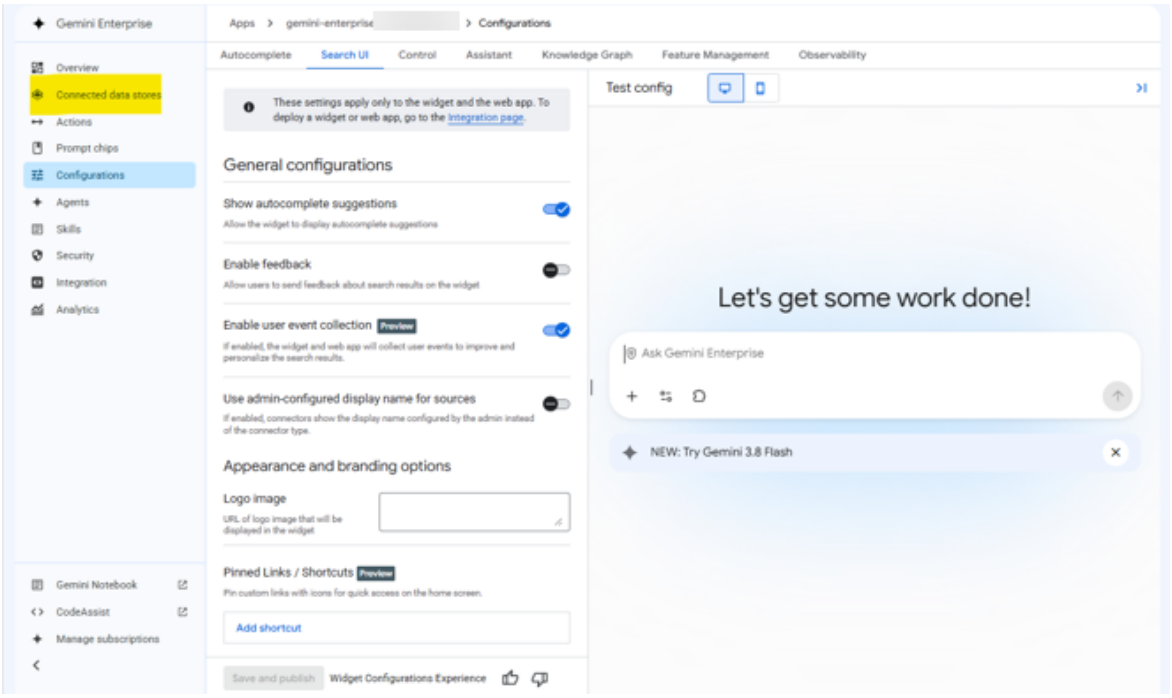
10. Select the app name, for example gemini-enterprise-...



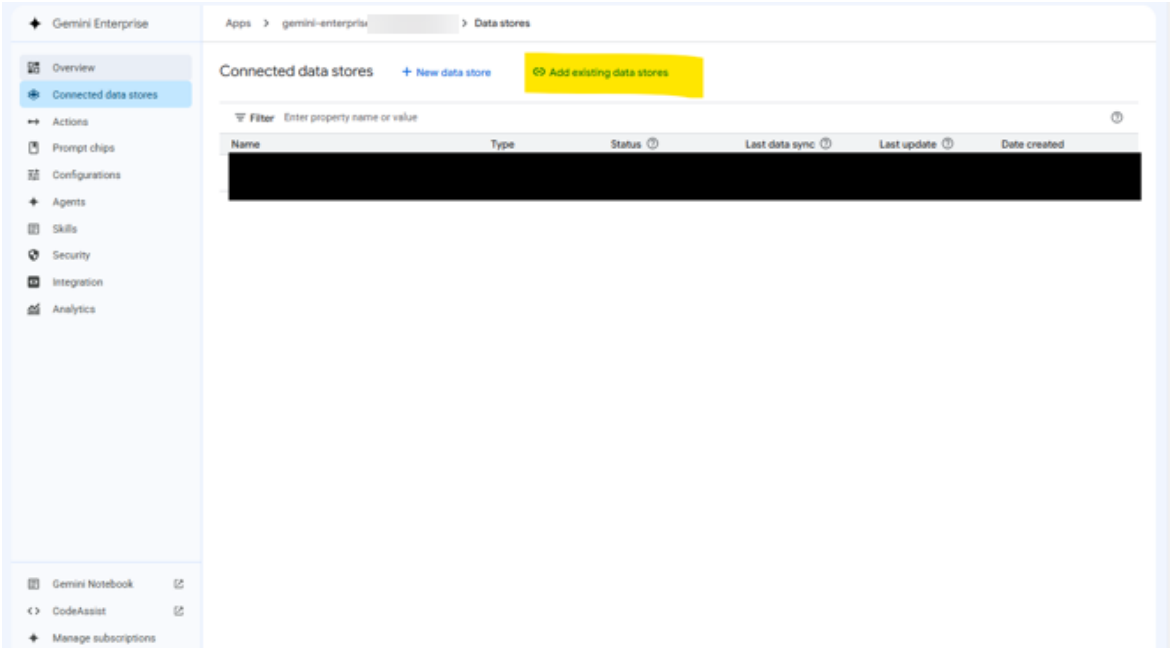
11. Select **Configure webapp**.



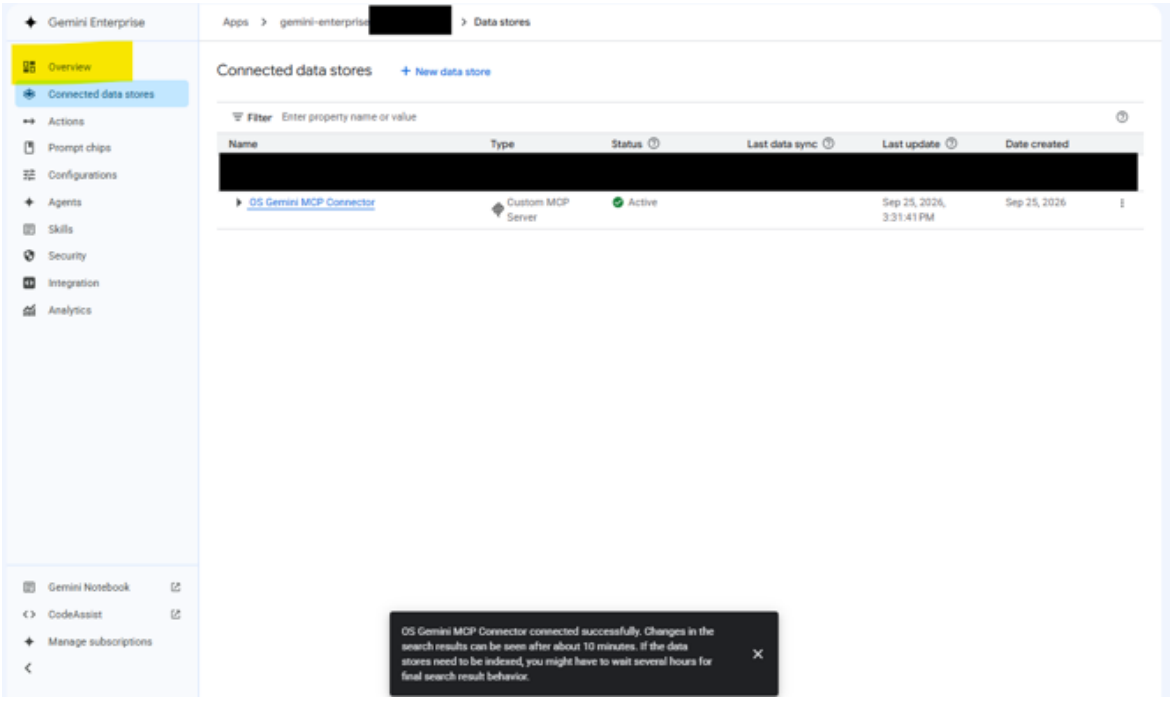
12. Select **Connected data stores**.



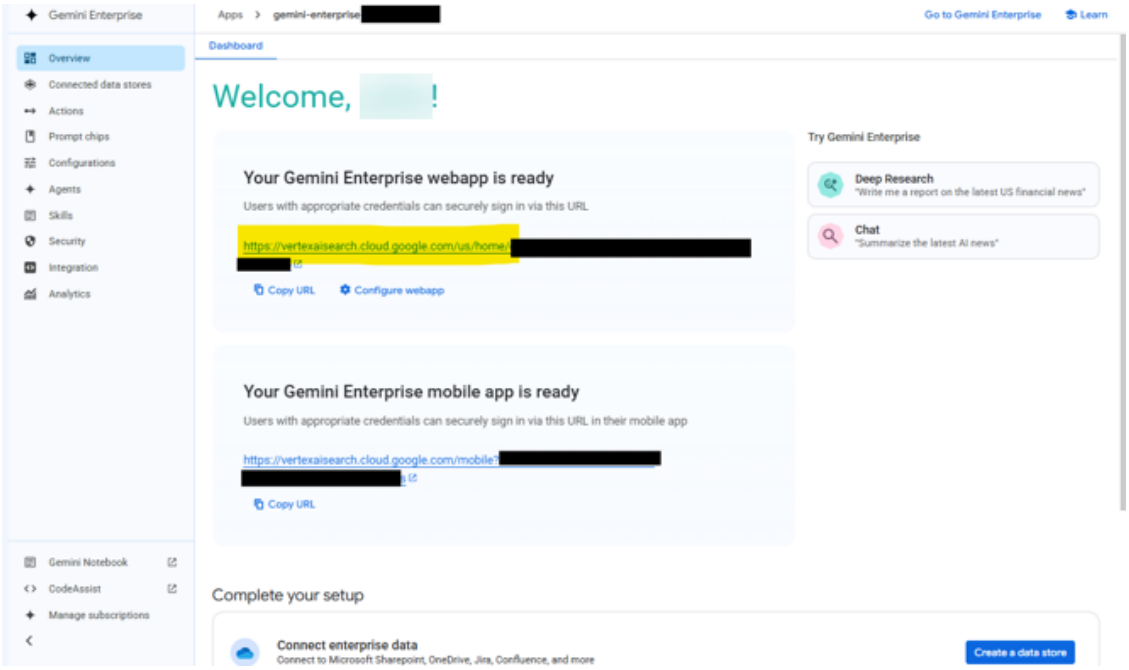
13. Select **Add existing data stores**.



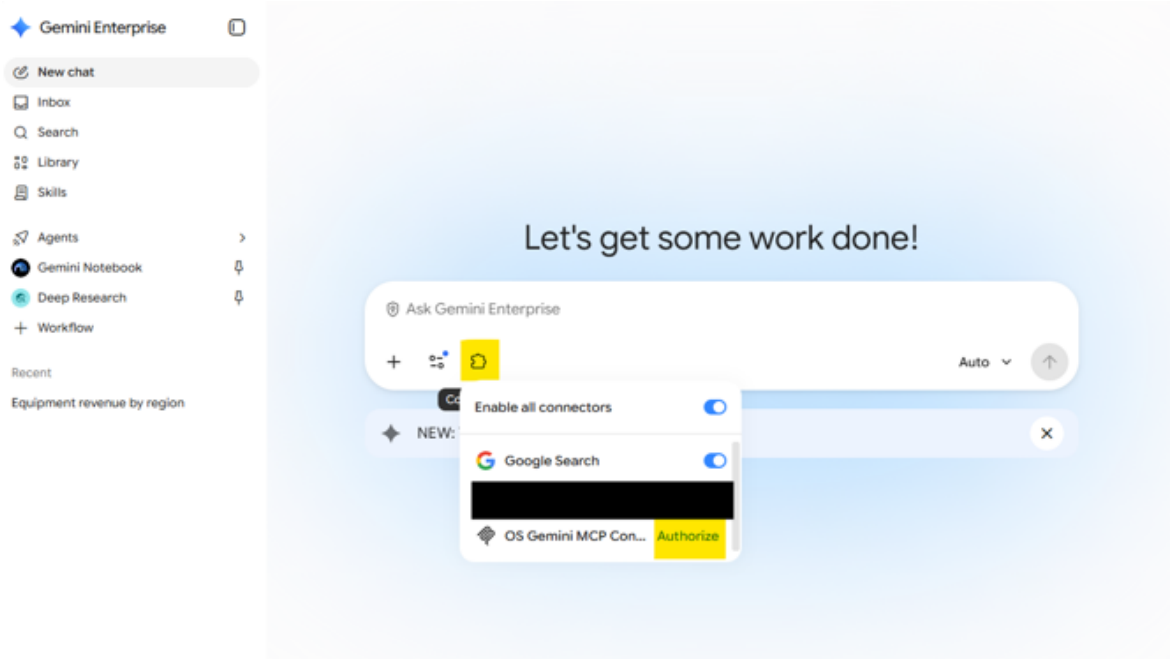
14. Select your MCP connector, and then select **Overview**.



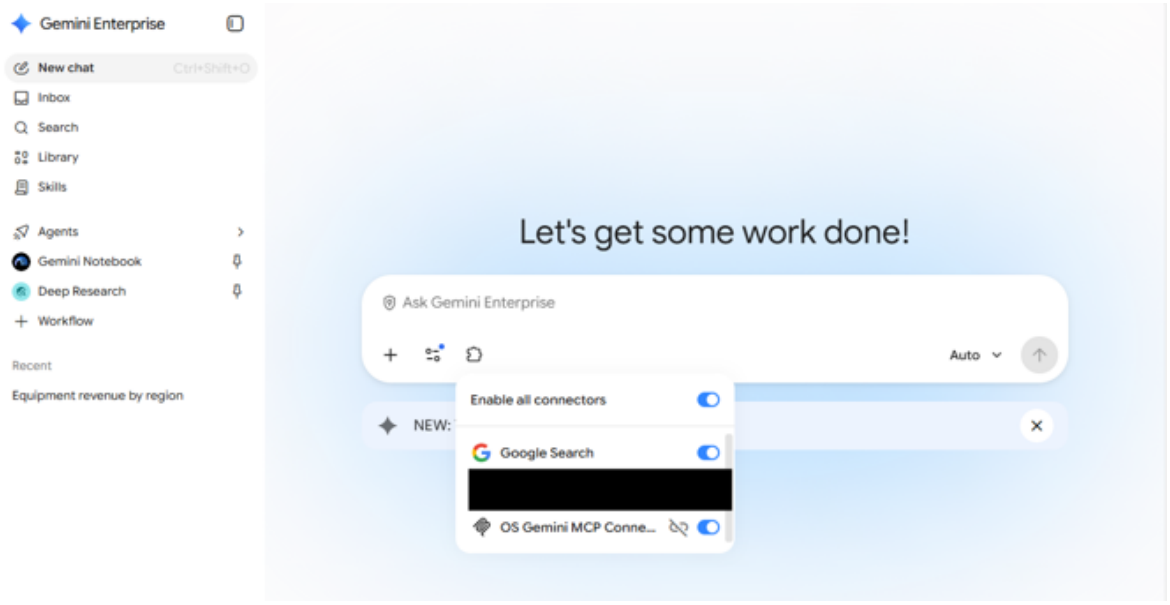
15. Select the link to your application.



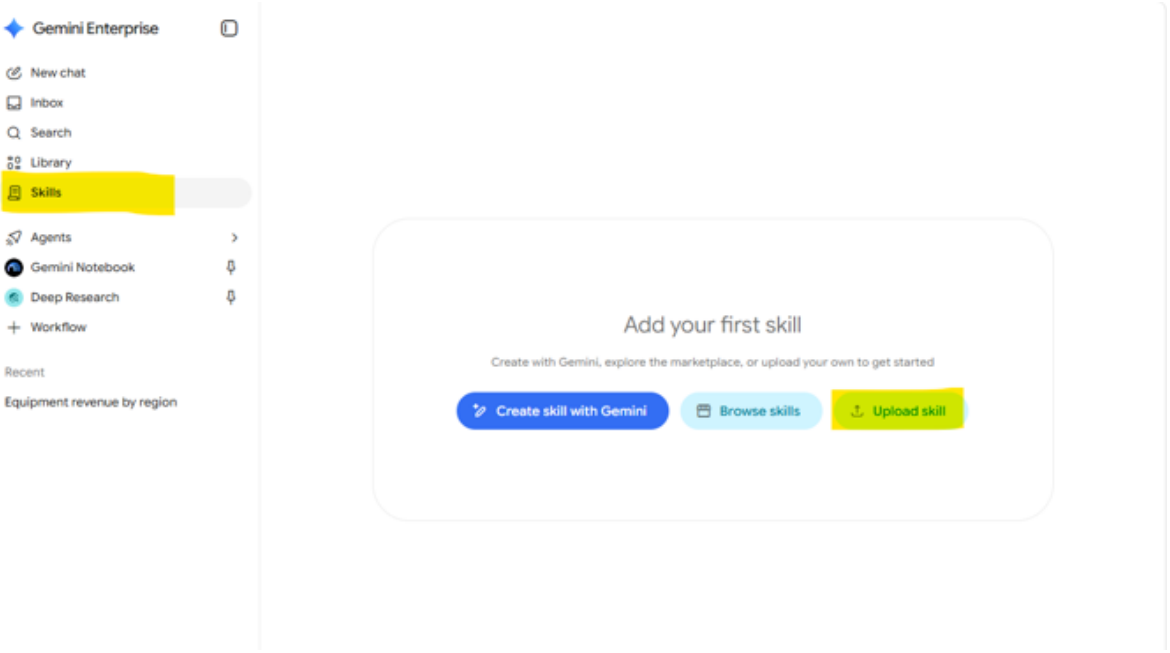
16. Select **Connectors**, and then select **Authorize** next to your Custom MCP.



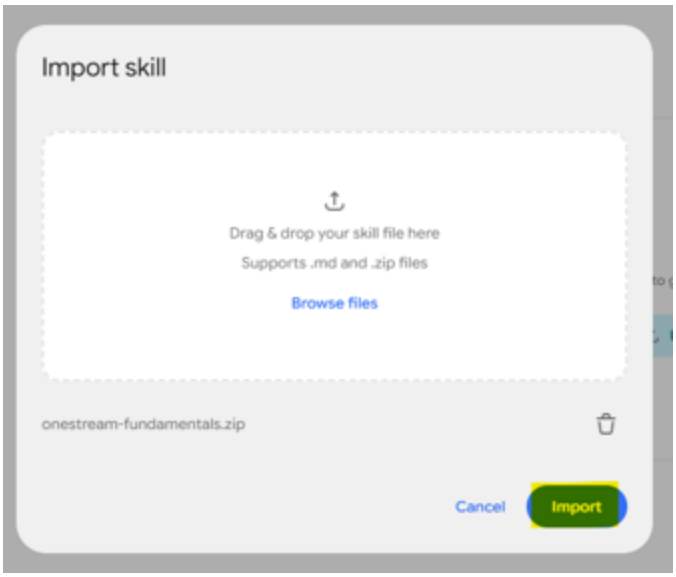
17. Verify that MCP authorization completed successfully.



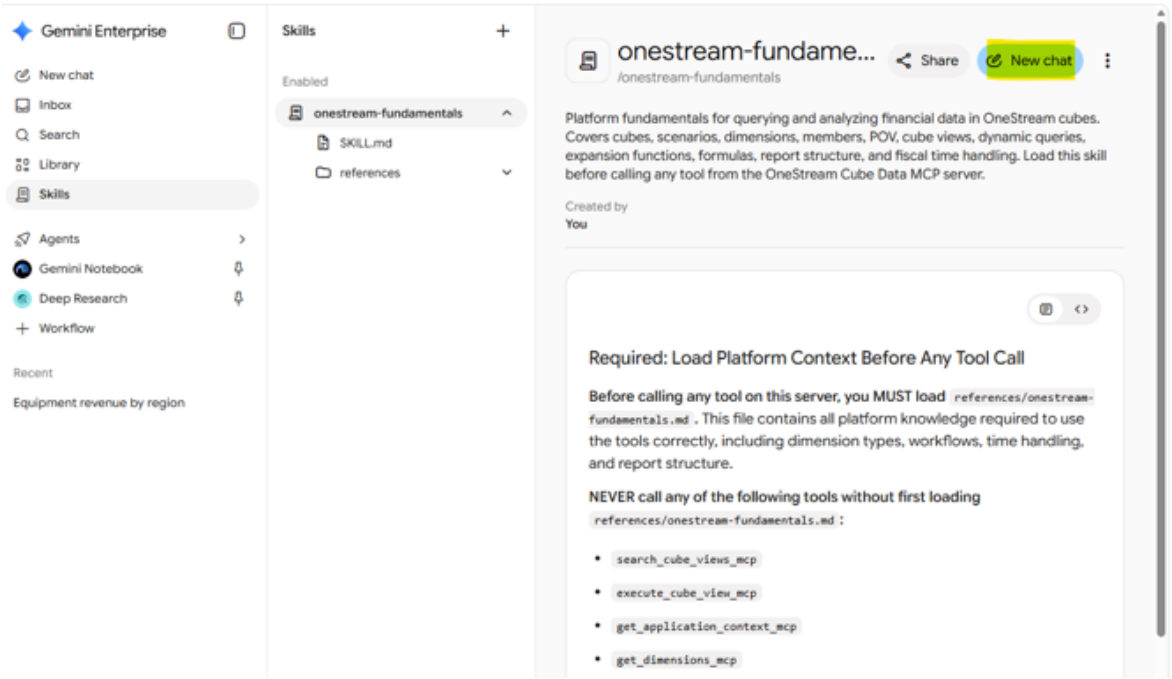
18. Select **Skills**, and then upload the Gemini skill.



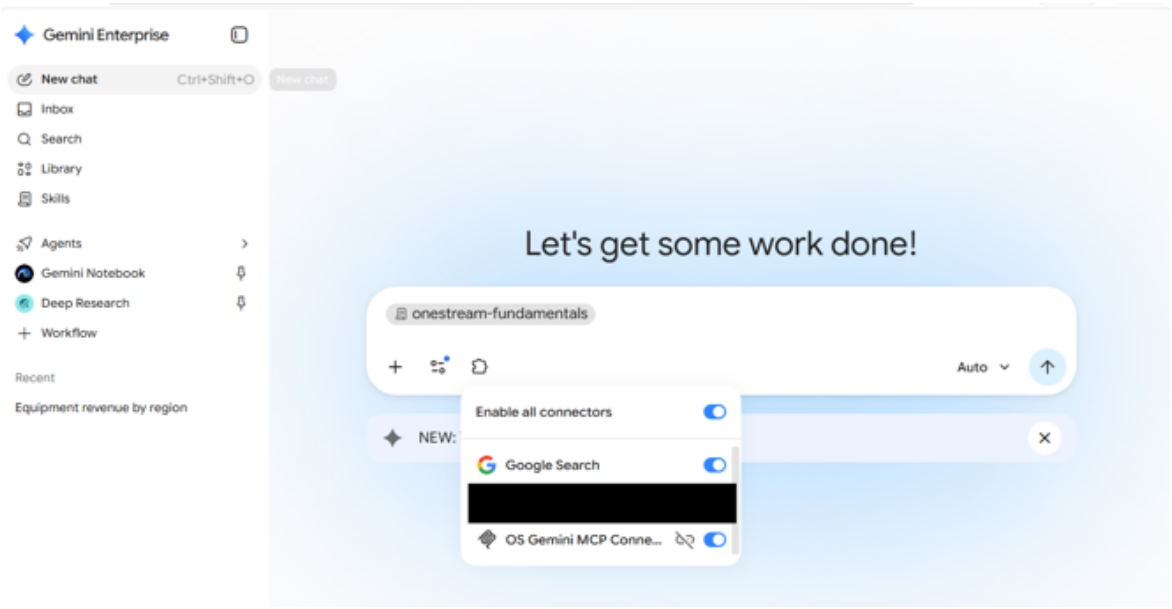
19. Select **Import**.



20. Select **New Chat**.



21. Verify that your custom MCP is enabled.



Environment Upgrade

The following steps describe the environment upgrade process:

Step 1: Review Release Notes After the Upgrade Is Released

Release notes for currently available release versions are available on the SensibleAI Agents Release Notes.

When a new upgrade is released, review the following items:

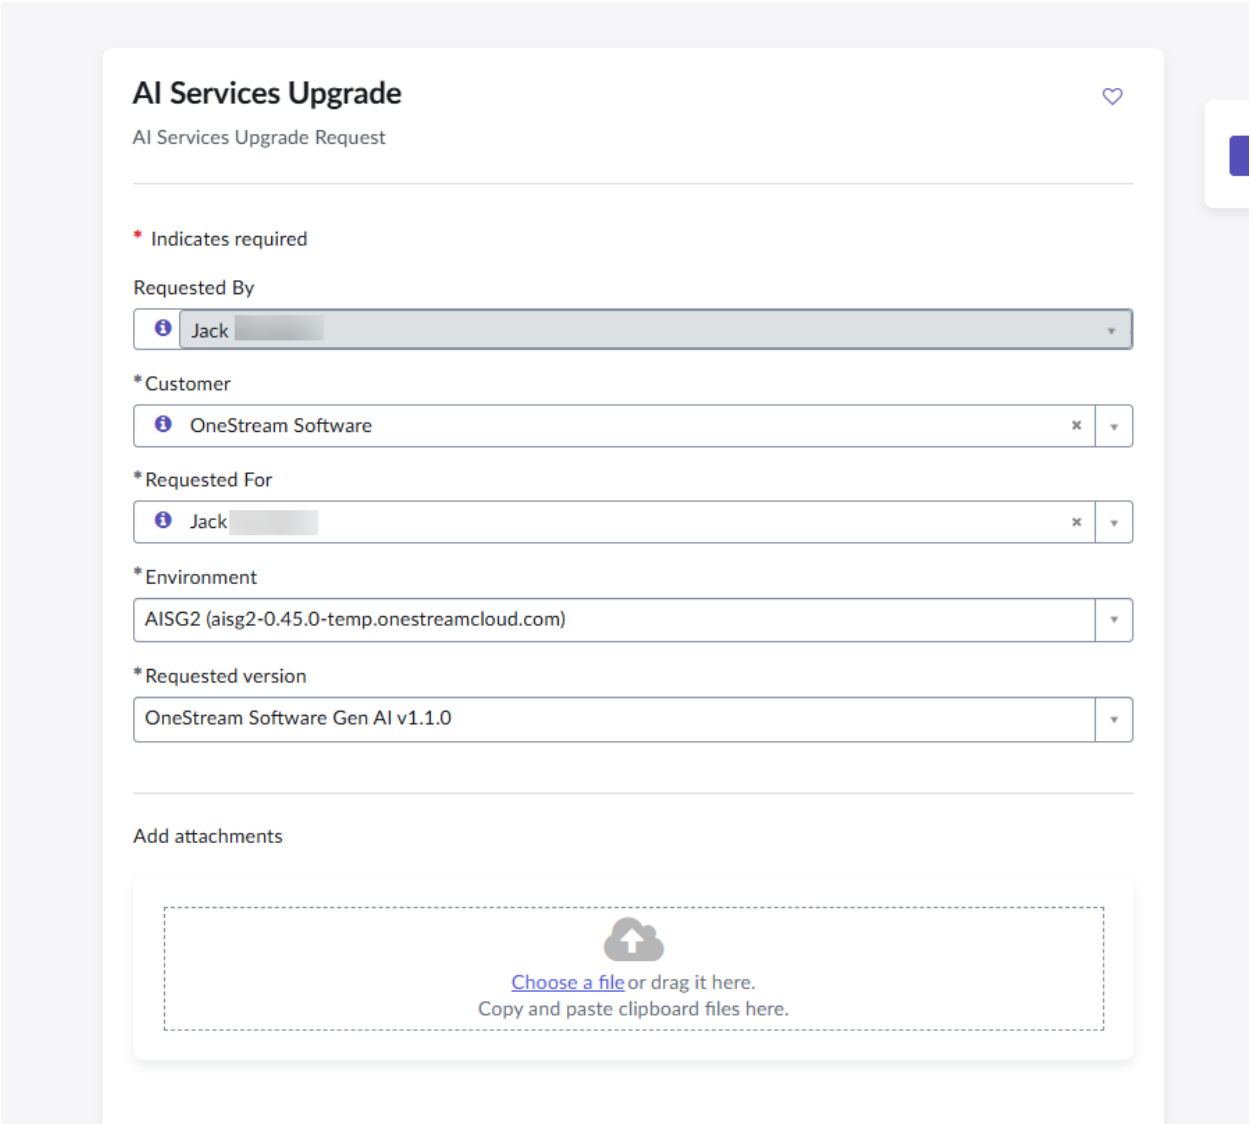
- Release Overview: Enhancements and bug fixes included in the release.
- Upgrade Guide
- Compatibility Guide

Step 2: Submit RITM

When you decide you want to upgrade, navigate to your Service Now support portal and input the following support ticket:

AI Services Upgrade | OneStream Software - ServiceNow

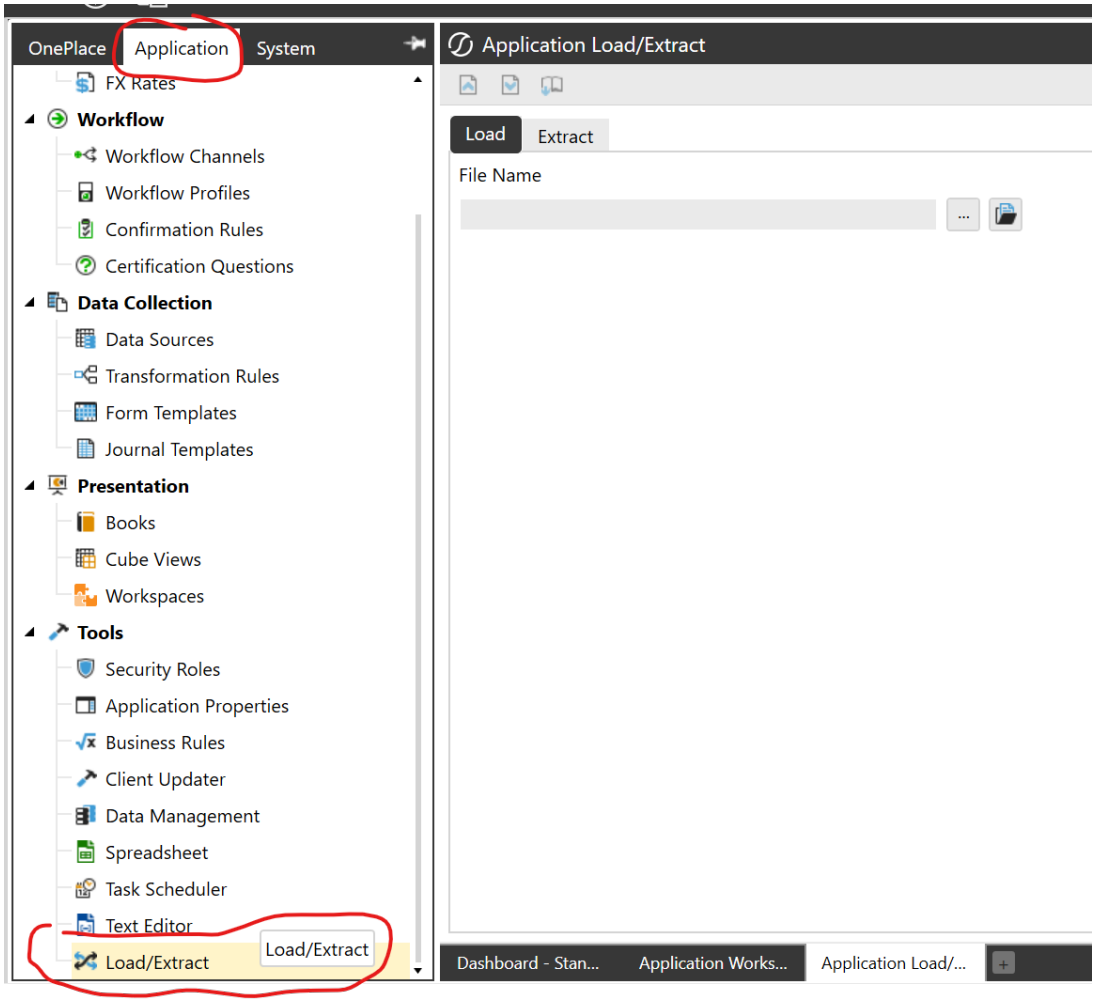
- **Requested By**: Customer organizations OneStream representative
- **Customer**: Customer organization name
- **Requested For**: Customer organizations OneStream representative
- **Environment**: Customer agents environment URL. This will have AISG in the name for dropdown and usually agents somewhere in the OneStream URL.
- **Requested Version**: The desired version for the upgrade. Advertised in notifications



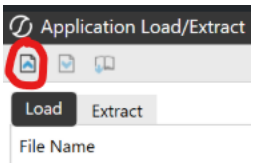
Step 3: Set up the Environment

An update will be given on the upgrade ticket once it has been completed. You can then follow these steps to finish the upgrade and set up the environment.

1. Go to Solution Exchange and download the latest Marketplace applications.
2. Load these items into the environment:
    a. In OneStream, navigate to **Application** > **Load/Extract**.



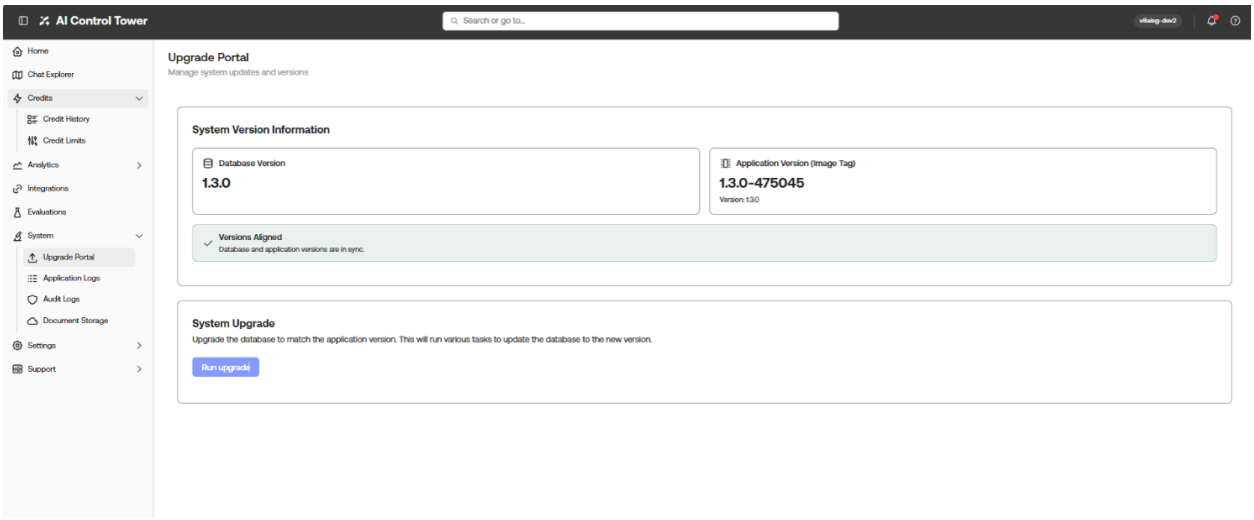
b. Load each .zip file individually and select **Load**.



3. Verify that the applications appear in OnePlace and load successfully.

Step 4: Run the Port

Navigate to **SensibleAI Control Tower** > **System** > **Upgrade Portal**.



Starting a port

Ports are kicked off from the **System Upgrade** screen.

> **CAUTION:** While a port is running, the system locks. New tasks are blocked, and a 503 response is returned.

If a step fails

Open a customer support ticket, and provide the following details:

1. **Version Task ID**
2. **Version range** (version_start → version_end)
3. **Which step/sub-step shows** FAILURE (name and step index)
4. **Task comments**: The error message attached to the failed task (shown in the activity feed/status response)
5. **Progress at the time of failure**: completed_tasks / total_tasks and completed_versions / total_versions

> **NOTE:** A screenshot is usually sufficient to capture all of these details.

Someone should then be in contact with the next steps on how to proceed.

Step 5: Confirm the Environment is Working

Verify that all functionality is working as expected.