



SensibleAI Agents

Admin Guide

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 6

 Finance Analyst Configuration 7

 Finance Analyst Agent Overview 7

 Finance Analyst Setup 7

 Scope 7

 Accessing Finance Analyst Configuration 7

Table of Contents

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
Accounting Period	18

Table of Contents

Overview	18
Cube Context	20
Overview	20
Cube Configuration	20
Dimension Defaults	21
Overview	21
Dimension Defaults	22
Dimension Members	24
Overview	24
Member Details	25
Bulk Actions	25
Business Context	27
Overview	27
Guidelines: What to Include	28
Guidelines: What Not to Include	28

Table of Contents

Report Row Templates	30
Overview	30
Creating a Template	30
Recommendations	31
Finance Analyst Analysis Plans	32
Accessing Analysis Plans	32
Analysis Plans Page	33
Plan Detail	34
Creating a Plan	35
Editing a Plan	37
Scheduled Analyses	37
Activities	38
Example Questions (Finance Analyst)	40
Rules	40
Narrative Summarizer	41
End User Behavior (Admin-Relevant Summary)	41

Table of Contents

Chat Settings	41
Query Results	42
SensibleAI Agents Documents Overview	43
Home Page and Global Options	43
Global Settings	44
Uninstall	45
Library	45
Clusters	46
Files	48
Requests	50
Data Connectors	52
Key Concepts	52
Source Configurations (Admin Only)	53
Connectors	54
Events	56
Validation	59

Table of Contents

Sync	60
SharePoint Configuration	61
Confluence Configuration	61
Search Agent Overview	63
Cluster Configuration	63
Access Configuration	64
Analytics	64
Deep Analysis Agent Overview	64
AI Control Tower Monitoring	64
AI Control Tower Overview	65
Chat Explorer	66
Credits	66
Credit Limits	67
Credit History	68
Analytics	69
AI Platform Connections	73

Table of Contents

Claude	73
ChatGPT	74
Microsoft Copilot	74
Gemini	75
Evaluations	76
Overview	76
Purpose	77
Configurations	78
Runs	79
System	80
Application Logs	80
Audit Logs	81
Settings	83
Roles	84
Agents Settings	85
ROI & Analytics	85

Table of Contents

Support86

 Document Storage87

 Audit Logs87

 Upgrade Portal87

 Feature Flags88

Previous Releases89

 Platform Version 989

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, 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 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)

Overview

- 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](#).

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, Access, Agents, ROI and Analytics).
 - b. SAI Support: All Administrator permissions; elevated access to Data Explorer (Explore); full access to the System section, including Document Storage, Audit Logs, Upgrade Portal, and Support Task Execution.

Overview

- 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:

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, 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:

Finance Analyst Configuration

- 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)
- 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

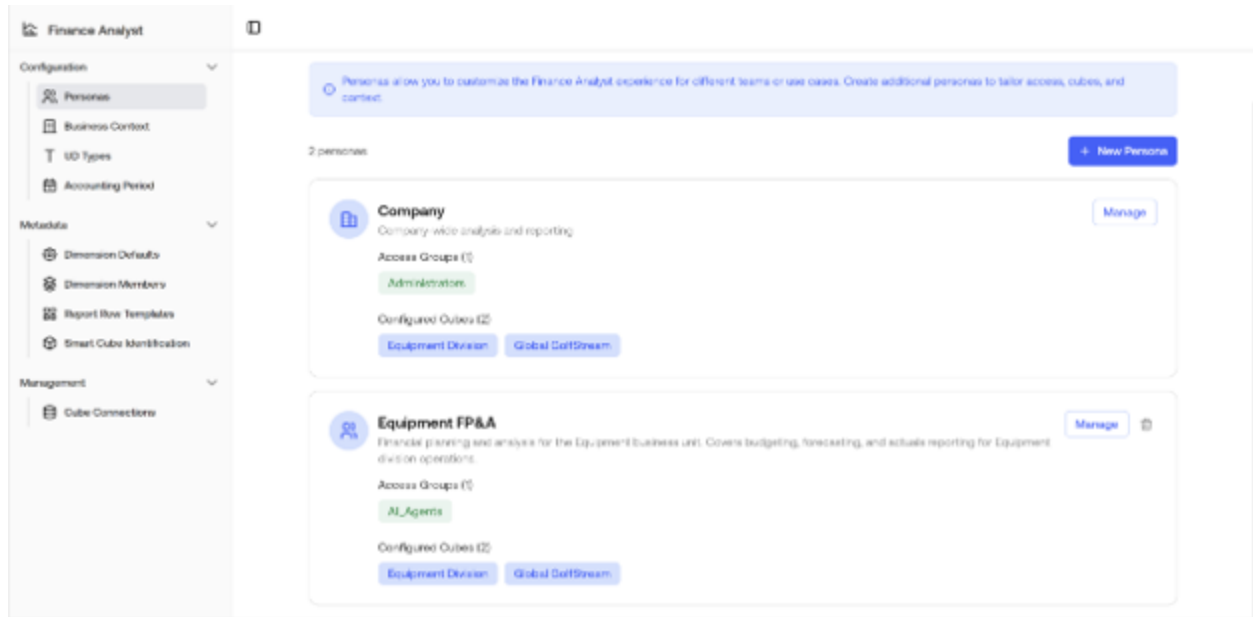
Actions are staged and executed together using Preview Sync.

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.

Configure UD Types

Manage UD Types

Select Dimension Type

UD1

UD2

UD3

UD4

UD5

UD6

UD7

UD8

Configure UD Types for UD1

Dimension	UD Type
Motors Cost Centers	Cost Center
MainUD1	Not Used
RootUD1Dim	Not Used
Equipment Cost Centers	Cost Center
Equipment Cost Centers Budget	Cost Center

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.”

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:

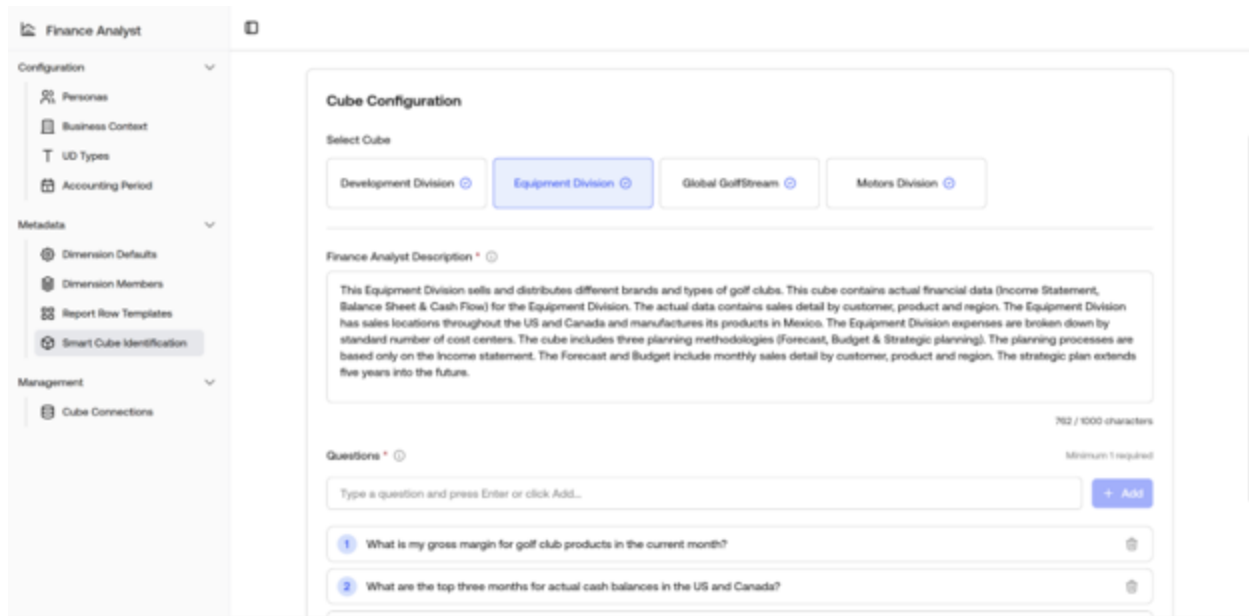
1. Cube Configuration (this page)
2. Persona-level mode selection (configured in Persona settings)

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 personas)
- Persona level (overrides company defaults)

Dimension Defaults

Configure default dimension members for Finance Analyst queries

① Dimension defaults are used to populate dimension values when Finance Analyst cannot infer a value from the user's question or available business context. This ensures a complete data intersection.

Select Persona

 Company

Dimension Defaults

Select Dimension Type

Entity

Flow

Scenario

IC

Origin

UD1

UD2

UD3

UD4

UD5

UD6

UD7

UD8

Dimension Name	Default Member	
Development	GSD	
Equipment	NAE	
Software	GSW	



Dimension Defaults

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

Recommendations

- Set company defaults first, and override only where needed.
- Defaults typically align with “Top” or “None.”
- Revalidate defaults after cube ingestion.

Configurable Dimension Defaults

Select Dimension Type

Entity

Flow

Scenario

UD1

UD2

UD3

UD4

UD5

UD6

UD7

UD8

Dimension Name	Default Member	
GlobalCorporate	All Orgs	
Development	GSD	
Equipment	GSE	
Motors	GSM	

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.

The screenshot shows the 'Finance Analyst Configuration Portal' interface. At the top, there is a 'Select Persona' dropdown menu with 'Company' selected. Below this are two dropdown menus: 'Dimension Type' set to 'Account' and 'Dimension' set to 'Equipment Accounts'. The main area is divided into two panels. The left panel, titled 'Members', contains a tree view with a 'Select' button above it. The tree structure is: 'None' (expanded), 'AcctTop' (expanded), 'Income Statement' (expanded), and '99999 - Net income' (selected). Other items under 'Income Statement' include 'Balance Sheet', 'Cash Flow', 'Supplemental', and 'Alternate Rollup'. The right panel, titled 'Member Details', contains four fields: 'Name' (value: 99999), 'Description' (value: Net income), 'Finance Analyst Description' (empty text area with a '0 / 500 characters' limit), and 'Used by Finance Analyst' (a toggle switch currently set to 'Yes').

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.

The screenshot shows a configuration interface for 'Equipment FP&A'. At the top, there is a 'Select Persona' dropdown menu with 'Equipment FP&A' selected. Below this is a checkbox labeled 'Inherit Company context' which is checked. Underneath is a section titled 'Business Context & Rules (Inherited from Company)' containing two lines of text: 'When a user asks to see a category breakdown (such as by region), use the Tree expansion, unless a different granularity is specified or indicated by a follow-up.' and 'When a user asks to see financial data for strategic customers, default to showing the aggregated value with no expansion, unless otherwise specified.' At the bottom, there is an 'Additional Context for Equipment FP&A (Optional)' section with a text input field containing the placeholder 'Add additional context specific to Equipment FP&A'.

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:

Business Context

- 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.

Report Row Templates

- **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.

Finance Analyst Analysis Plans



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.

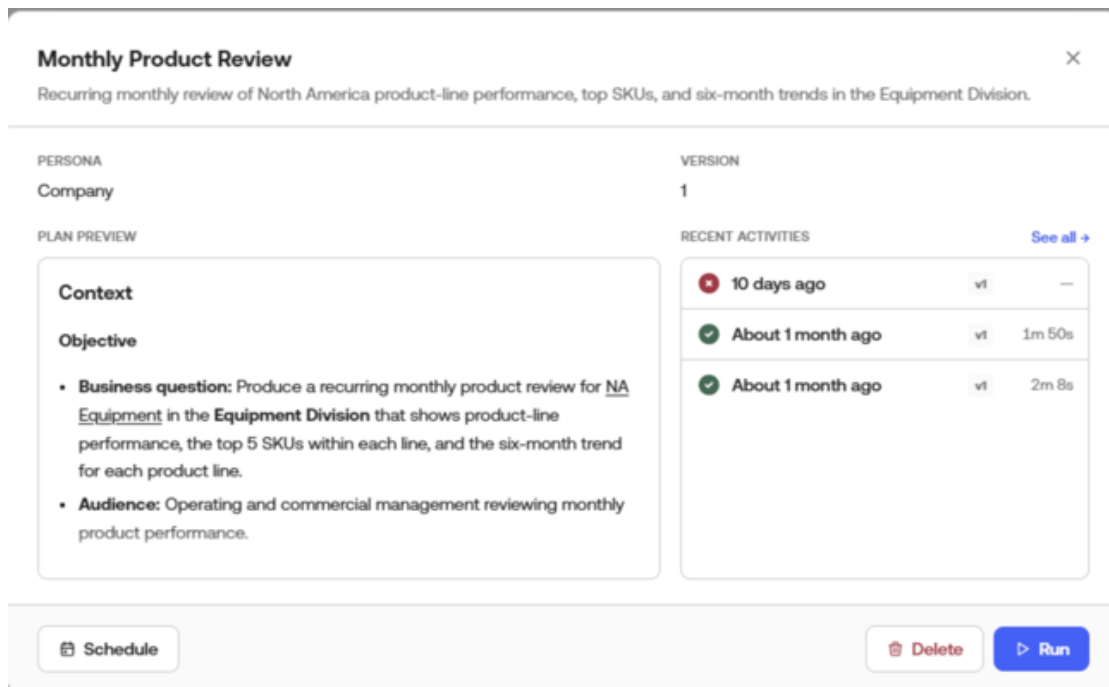
Select a plan to open its detail view. Review the following fields:

- **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:

Finance Analyst Analysis Plans

- **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:

Finance Analyst Analysis Plans

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.

Finance Analyst Analysis Plans

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.

The screenshot displays the 'Dashboard - SensibleAI Agents Finance' interface. The top navigation bar includes 'Finance Analyst', 'Analysis Plans', 'Scheduled Analyses', and 'Activities'. The 'Activities' tab is active, showing a list of tasks and a detailed view of a report.

Back to Activities

Created one plan to one month-over-month variance, so one only needs to manually support ranking and trend context rather than explain a large monthly relocation. I'm now pulling the SKU rankings for each line and will tie them back to the six-month trajectory.

Built 4 cube views, used 4 tools.

I've ruled out a large June step-change: growth was broad-based across every line and the SKU moves are small relative to line size, which points to steady underlying demand rather than a single-product disruption. I'm calculating top-SKU concentration and six-month change rates now so the final report can distinguish sustained momentum from simple scale effects.

Used 9 tools.

I've prepared the full monthly product review for NA Equipment -- open the canvas for the report.

NAE Product Review
Recurring monthly product review for NA Equipment covering June 2025 product-line.

Ask a follow-up question...

Company

Report

Generated from: Monthly Product Review

NAE Product Review

Monthly product review for NA Equipment in the Equipment Division, anchored on June 2025 actual product revenue.

Headline Takeaway

June product revenue increased across all four tracked product lines, led by **Mach Speed** at \$31,397,891, followed by **GolfStream Epic** at \$28,188,464, **Acme Nexus Series** at \$12,862,690, and **Club Parts** at \$6,203,264. Month over month, each line was positive, but none met the report's material commentary threshold of 5% and \$100K, with growth ranging from 1.36% to 3.75%.

Product Line Summary

Product line	Jun 2025 revenue	May 2025 revenue	MoM variance	MoM variance %
Mach Speed	\$31,397,891	\$30,863,617	\$434,274	1.40%
GolfStream Epic	\$28,188,464	\$27,739,870	\$378,596	1.36%

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:

- Library
- Connectors
- Deep Analysis
- Admin (administrators only)



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

Global Settings

Configure the following settings:

- (Security) Power User Role
- (Security) Base User Role
- (Timezone) Display Time Format
- (Timezone) Daylight Savings
- Backend Endpoint Parameter
- Chatbot Endpoint Parameter

Global Settings

Global Options

☒ Global Settings ☐ Uninstall ☐ Solution Info

Global Settings

(Security) Power User Role:
Administrators

(Security) Base User Role:
Administrators

(Timezone) Display Time Format:
(UTC) Coordinated Univers...

(Timezone) Daylight Savings:
No

Backend Endpoint Parameter:
https://v8aisg-agents-dev3.onestreamtest.com/we

Chatbot Endpoint Parameter:
https://v8aisg-agents-dev3.onestreamtest.com

Close Save

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.

Library

The Library page has the following tabs: Clusters, Files, and Requests.

Clusters

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:** can access cluster information
- **Owner:** member permissions, approve/reject, and manage requests

Create a Cluster

To create a cluster, complete the following steps:

SensibleAI Agents Documents Overview

1. Go to **Library**> **Cluster**.
2. Select the **Create** button.



3. Enter a name (required) and description (optional). See [Cluster Naming Rules](#).
4. Select a Cluster Type.
5. Select Owners.
6. Set Required Approvals: Users who have the ability to approve or reject requests on a given cluster.
7. Select Members.
8. Review and submit.

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

From Cluster Detail, 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 Owners can delete a cluster. Deletion requires approval.

Files

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

Upload Requests

Upload to OneStream File System.

Upload files using File Explorer (Windows or Modern Browser Experience). The recommended destination is the user folder or Favorites within it.

Review the following limits:

- 85MB per file (including files in .zip)
- 1000 files per upload request

To create an upload request, complete the following steps:

1. Go to **Library > Files**.
2. Select **Upload**.
3. Enter a name (required) and description (optional)
4. Select clusters for access.
5. Review and submit.

After submission, files are staged/scanned. Every file submitted to the system is security scanned, then cluster Owners approve or reject before indexing.

Duplicate File Handling

Duplicate names degrade response quality and retrieval diversity.

System behavior during upload:

- Overwrite Allowed: if an existing file belongs only to clusters included in the request.
- Overwrite Blocked: if an existing file belongs to clusters outside the request (rename or use Edit Request)

Edit Requests

Edit Requests modify cluster access for existing files:

- 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 Requests

Deletes a file from all clusters and removes it from the library. At least one Owner from each associated cluster must approve the deletion.

Requests

Requests support approving/rejecting pending requests and reviewing request details and history. Requests can have the following statuses and additional transitional statuses:

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

Requests to Review


SensibleAI administrators and cluster Owners can approve/reject requests. Approved requests proceed to Uploading and then Completed/Failed.

Troubleshooting Failed Requests

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 (admin-defined)
- **Connector:** Utilizes the source configuration to retrieve information and store it in clusters.

Source Configurations (Admin Only)

Data Connectors Connectors Source Configurations Events Refresh

Source Configurations

NAME	SOURCE TYPE	CREATED BY	CREATED TIME	NUMBER OF DATA CONNECTORS
OneStream SharePoint	SharePoint	Danilo Aramondy	Jul 28, 2025, 12:32 PM	5
AI & Ops Confluence Configuration	Confluence	Danilo Aramondy	Nov 6, 2025, 1:41 PM	5

- View details and associated connectors

Data Connectors Connectors Source Configurations Events Refresh

Source Configurations

NAME	SOURCE TYPE	CREATED BY	CREATED TIME	NUMBER OF DATA CONNECTORS
OneStream SharePoint	SharePoint	Danilo Aramondy	Jul 28, 2025, 12:32 PM	5
AI & Ops Confluence Configuration	Confluence	Danilo Aramondy	Nov 6, 2025, 1:41 PM	5

- 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.

SensibleAI Agents Documents Overview

Source Configuration: AI & Ops Confluence Configuration

Source Configuration Details

Name:

AI & Ops Confluence Configuration

Description:

N/A

Source Type:

Confluence

Active:

Yes

Created By:

Datha Aramreddy

Created:

Nov 6, 2025, 1:41 PM

Site URL:

Client ID:

Data Connectors

CONNECTOR NAME	TYPE	STATUS
Test AI Engineering ...	Confluence Space	Ready
AI Engineering Teams	Confluence Space	Ready
AI Engineering Confl...	Confluence Space	Deleted
Wernicke Release M...	Confluence Space	Ready
Applied AI Test Con...	Confluence Space	Ready
Shared User To Con...	Confluence Space	Deleted

- Delete (deletes all associated connectors)

IMPORTANT: Review associated connectors before deleting a source configuration.

Connectors

Data Connectors

ConnectorsSource ConfigurationsEvents

My Connectors

All Connectors

Connector Name

Source

Connector

Status

Owner

Created Time

Last Sync Time

Customer Success, Sharepoint, MRISupport	Sharepoint	Customer Success	Ready	Jack Allen	Aug 20, 2025, 1:48 PM	Mar 4, 2025, 4:29 PM
Shared User To Connector	Confluence	Internal AI/Op Confluence Knowledge Base Articles	Syncing	Shawn Vaidler	Jan 15, 2025, 9:45 AM	-
AI/Engineering Teams	Confluence	Internal AI/Op Confluence Knowledge Base Articles	Ready	Shawn Vaidler	Jan 15, 2025, 1:00 PM	-
Applied AI Test Connector	Confluence	Applied AI Test Cluster	Ready	Arjun Chandra	Dec 4, 2025, 4:22 PM	-
Wernicke Release Management Confluence Connector	Confluence	Wernicke Release Management	Ready	Karen Reiss	Dec 4, 2025, 4:28 PM	Mar 4, 2025, 3:07 PM
Test AI Engineering Connector	Confluence	Default	Ready	Datha Aramreddy	Nov 6, 2025, 1:41 PM	-
GES HC GES and All OS	Sharepoint	OneDrive GES	Syncing	Lara Storm	Sep 15, 2025, 3:49 PM	Nov 4, 2025, 5:18 PM
GES HC GES Only	Sharepoint	GES HC GES Only	Ready	Lara Storm	Sep 15, 2025, 3:20 PM	Sep 15, 2025, 7:22 PM
SMW Internal Kilo	Sharepoint	Internal Sensibleflow Knowledge Base	Ready	Harsh Harshad	Sep 12, 2025, 14:33 AM	Sep 12, 2025, 6:23 PM
SMW External Kilo	Sharepoint	External Sensibleflow Knowledge Base	Ready	Harsh Harshad	Sep 12, 2025, 10:49 AM	Sep 12, 2025, 6:23 PM

The Connectors tab allows users to see the Data Connectors that they have created. The All Connectors tab is only available to users with the SensibleAI Agents Administrator role. 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
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

Connector Statuses
Syncing
Deleting
Disconnected

Events

Events provide audit log entries for:

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

SensibleAI Agents Documents Overview

Data Connectors

ConnectorsSource ConfigurationsEvents

Refresh

My EventsAdd EventsEvents to Review

Search

EVENT NAME	EVENT TYPE	EVENT TYPE	EVENT STATUS	PERFORMER USER	EVENT TIME	EVENT DETAILS
Mercedes Release Management Conference Connector	Connector	Sync	Pending Review	Kavin Bous	Mar 10, 2025, 9:27 PM	File Count: 31 Rows
Customer_Success_Management_KBRepository	Connector	Sync	Completed	Jack Allen	Mar 10, 2025, 6:20 PM	File Count: 1 Rows
Customer_Success_Management_KBRepository	Connector	Sync	Completed	Jack Allen	Feb 24, 2025, 5:44 PM	File Count: 30 Rows
Customer_Success_Management_KBRepository	Connector	Validate	Completed	Jack Allen	Feb 24, 2025, 5:44 PM	-
Test co connector	Connector	Create	Completed	Debra Aramowicz	Feb 24, 2025, 5:29 PM	-
Test co connector	Connector	Create	Completed	Debra Aramowicz	Feb 24, 2025, 5:29 PM	-
Customer_Success_Management_KBRepository	Connector	Create	Completed	Jack Allen	Feb 23, 2025, 6:40 PM	-
Shared How-To Connector	Connector	Sync	Pending Review	Simon Veldhor	Feb 19, 2025, 10:40 PM	File Count: 75 Rows
Shared How-To Connector	Connector	Validate	Completed	Simon Veldhor	Feb 19, 2025, 10:40 PM	-
Shared How-To Connector	Connector	Sync	Failed	Simon Veldhor	Jan 15, 2025, 3:00 PM	File Count: 0 Rows
Shared How-To Connector	Connector	Validate	Completed	Simon Veldhor	Jan 15, 2025, 2:07 PM	-
Shared How-To Connector	Connector	Create	Completed	Simon Veldhor	Jan 15, 2025, 2:07 PM	-
AI Engineering Teams	Connector	Sync	Failed	Simon Veldhor	Jan 15, 2025, 1:43 PM	File Count: 0 Rows
Test AI Engineering Connector	Connector	Validate	Completed	Debra Aramowicz	Jan 14, 2025, 8:49 PM	-
AI Engineering Teams	Connector	Sync	Failed	Simon Veldhor	Jan 13, 2025, 6:37 PM	File Count: 0 Rows
AI Engineering Teams	Connector	Validate	Completed	Simon Veldhor	Jan 13, 2025, 6:37 PM	-

Event statuses include:

Event Statuses
Staging
Scanning
Pending Review
Syncing
Completed
Failed
Partially Failed

Connector Creation

1. Navigate to **Connectors > Create**.



2. Enter name and optional description.
3. Choose the Source Type and Connector Type.

Define > Configure > Access > Summary

 This connector links to third-party services solely for connectivity and integration purposes. It does not control data accessed or hosted by such third-party services. Please avoid the transmitting of special categories of personal data, protected health information, or credit card information. The use of connected third-party services is strictly governed by applicable third-party terms.

Connector Name

AI Product and Engineering Sharepoint

Connector Description

Source Type

SharePoint

Connector Type

SharePoint Folder

Buttons: Cancel, Previous, Next

4. Configure (existing or new source configuration, admin-only for new).

Define > Configure > Access > Summary

Source Configuration

☒ Use Existing Configuration ☐ Create New Configuration

Select Source Configuration

Select...

SharePoint Folder Details

SharePoint Site Name

SharePoint Drive Name

SharePoint Folder Path

SharePoint Authentication

Authorize access to your SharePoint account to set up the data connector.

Authorize

Cancel Previous Next

5. Select clusters for access.
6. Review and submit.

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

Validation

Run validation to test connectivity: Only Owners can validate

1. Select the connector.
2. Select Validate.
3. Review the validation event in Events.

Sync

Sync is manually triggered and applies changes to the document index.

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)

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.
- 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. Each SharePoint site must be configured separately for security purposes.

Confluence Configuration

Confluence can be connected as a data source using a source configuration and one or more connectors. Each Confluence site must be configured separately for security purposes.

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 (admin only):

Field	Description
Configuration Name	The display name for this configuration
SharePoint Hostname	The SharePoint hostname 
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 [Library](#).

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 quarterly. This page displays a daily usage chart and the following KPIs:

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

AI Control Tower Overview




Credit Limits

NOTE: This is SAI Administrator+.

This page manages company limits and user limits:

- Split evenly
- Configure details
- Sync Users

AI Control Tower Overview




Review the following notes:

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

Credit History

Credit History provides a breakdown by users or agents.

AI Control Tower Overview



Analytics


The Analytics page provides usage and ROI views across:

- Finance Analyst
- Search
- Deep Analysis

The dashboard provides the following KPIs:


AI Control Tower Overview

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




AI Control Tower Overview

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




- **Clusters:** Search-specific.



- **Documents:** Search-specific.

AI Control Tower Overview



- **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. Each integration is initiated from the AI Control Tower Integrations page.

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.

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:

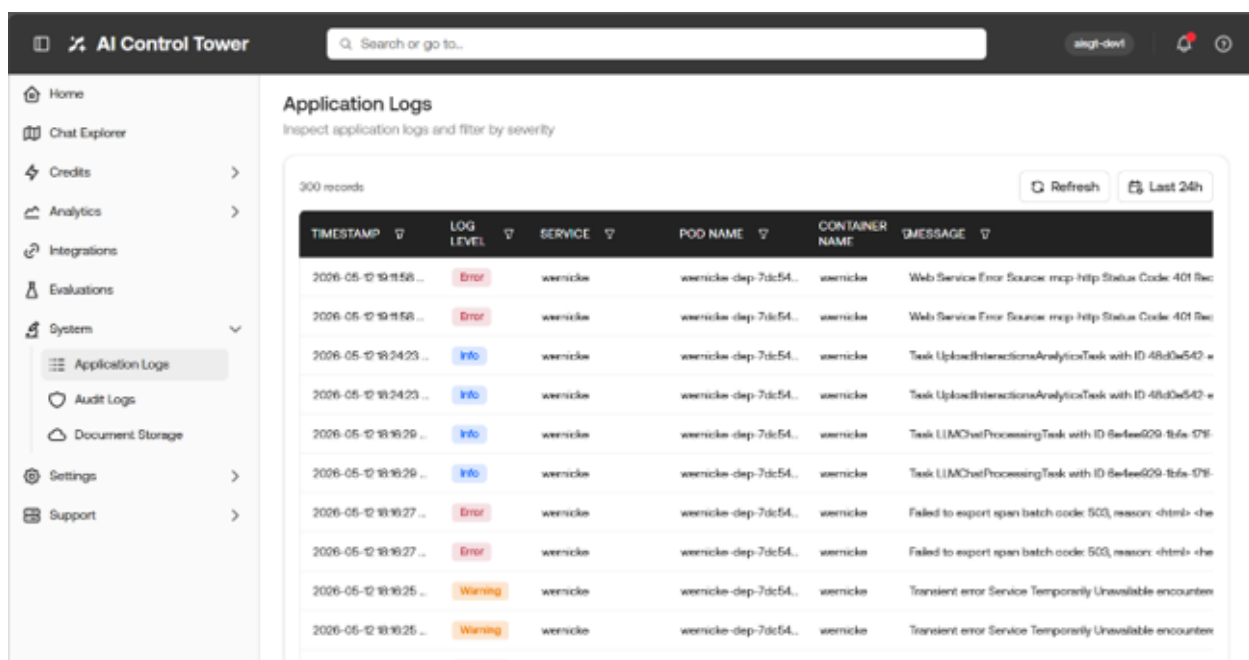
- A score
- An explanation of which elements were correct or incorrect

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: monitor application logs, review audit logs, and manage document storage.

Application Logs



TIMESTAMP	LOG LEVEL	SERVICE	POD NAME	CONTAINER NAME	MESSAGE
2026-05-12 19:11:58 ..	Error	wericke	wericke-dep-7dc54...	wericke	Web Service Error Source: map-http Status Code: 401 Sec
2026-05-12 19:11:58 ..	Error	wericke	wericke-dep-7dc54...	wericke	Web Service Error Source: map-http Status Code: 401 Sec
2026-05-12 19:24:23 ..	Info	wericke	wericke-dep-7dc54...	wericke	Task UploadInteractionsAnalyticsTask with ID 48d3e542 »
2026-05-12 19:24:23 ..	Info	wericke	wericke-dep-7dc54...	wericke	Task UploadInteractionsAnalyticsTask with ID 48d3e542 »
2026-05-12 19:16:29 ..	Info	wericke	wericke-dep-7dc54...	wericke	Task LIMChatProcessingTask with ID 8e1ee629-8fa-57f
2026-05-12 19:16:29 ..	Info	wericke	wericke-dep-7dc54...	wericke	Task LIMChatProcessingTask with ID 8e1ee629-8fa-57f
2026-05-12 19:16:27 ..	Error	wericke	wericke-dep-7dc54...	wericke	Failed to export span batch code: 503, reason: <html> <he
2026-05-12 19:16:27 ..	Error	wericke	wericke-dep-7dc54...	wericke	Failed to export span batch code: 503, reason: <html> <he
2026-05-12 19:16:25 ..	Warning	wericke	wericke-dep-7dc54...	wericke	Transient error Service Temporarily Unavailable encounter
2026-05-12 19:16:25 ..	Warning	wericke	wericke-dep-7dc54...	wericke	Transient error Service Temporarily Unavailable encounter

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


TIMESTAMP	TYPE	ACTOR	SUBJECT	SUMMARY	DETAIL
2026-05-12 19:00:01	Sync Users AI Credit	System (Scheduled)	Sync Users AI Credit	AuditAction.UPDATE (sync_users_ai_credit)	Info
2026-05-12 19:00:01	Sync Users AI Credit	System (Scheduled)	Sync Users AI Credit	AuditAction.UPDATE (sync_users_ai_credit)	Info

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:

- Roles
- Agents
- ROI & Analytics

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 (Char Reviews and Interaction Explorer), and the Control Tower dashboard.
- **SAI Administrator:** Viewer access plus the ability to manage Credit Limits and full Settings (Application, Access, Agents, ROI and Analytics).
- **SAI Support:** All Administrator access, plus elevated access to the Data Explorer (Explore) and the entire System section (Document Storage, Audit Logs, Upgrade Portal, Support Task Execution).
- **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

AI Control Tower Overview



Support

Support includes the following pages:

- Data Manipulator
- Upgrade Portal
- Feature Flags
- Support Task Execution

Document Storage

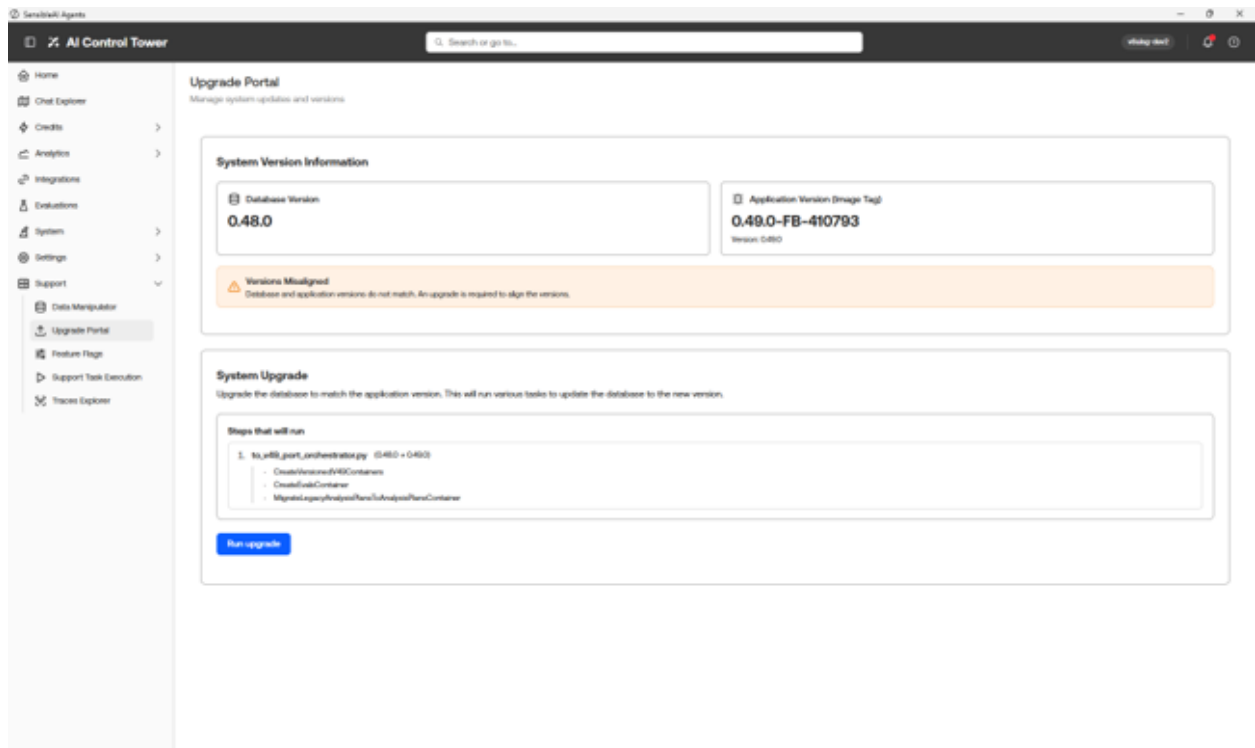
This page displays storage consumption for Total Document Storage (all uploads) and Indexed Document Storage (knowledge clusters).

Audit Logs


This page captures sensitive activity with before/after status.

Upgrade Portal

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



Feature Flags



This page includes the following feature flag:

- **web_search**: Enables web search capability for supported agents.

Previous Releases

You can find documentation for previous SensibleAI Agents releases below.

Platform Version 9

- PV9.0.0 SV1.0.0 [Guide](#) | [Release Notes](#)