

# Skylar Advisor

Version 2.8.0

# Table of Contents


**Getting Started with Skylar Advisor** ..... **1**

What is Skylar Advisor? ..... 2

Logging In to the Skylar AI User Interface ..... 2

The Skylar Advisor User Interface ..... 3

Defining Your Persona and Role Information ..... 5

**Scheduling Reports for Skylar Advisor** ..... **8**

Prerequisites ..... 9

How Report Scheduling Works ..... 9

Enabling and Running the Corpus Health Report ..... 9

Finding Your Scheduled Reports ..... 13

Scheduling the Remaining Day-One Reports ..... 14

Handoff ..... 14

Business Services ..... 15

Trends ..... 15

Spend Value ..... 15

Quick-Reference Cadence Table ..... 15

Expected Behavior During Initial Setup ..... 15

**Managing Documents and Data on the Knowledge Center Page** ..... **17**

The Knowledge Center Page ..... 18

The Active View on the Corpus Tab ..... 19

The Corpus Metrics Panel ..... 21

Using Web Search ..... 22

Best Practices for Uploading Documents to the Corpus ..... 23

Setting Up a New Corpus (Corpus Formation) ..... 24

Finding Knowledge Gaps in the Corpus (Gap Mining) ..... 25

Planning Your Corpus Health Baseline ..... 25

Part 1: Mining an Advisory for Gaps ..... 25

Part 2: Filling Gaps with the Info Gaps View ..... 27

Part 3: Verifying the Loop Closed ..... 28

What to Expect During Your First Month ..... 29

Signs Corpus Formation Is Complete ..... 29

Expected Behavior During Corpus Formation ..... 29
Understanding Corpus and Web Search Ownership ..... 30
Building Your Corpus (Owner) ..... 31
What to Include First ..... 31
What to Leave Out ..... 32
Keeping Your Corpus Healthy ..... 32
The Three Problem Types ..... 33
Removing and Monitoring Corpus Content ..... 33
KB Health: What Is Earning Its Place ..... 33
Info Gaps: What Is Missing ..... 34
What to Remove ..... 34
Administering Remote Knowledge with Web Search ..... 34
Recommended Maintenance Cadence ..... 35
Common Failure Modes ..... 35
**Using the Ask Skylar Feature ..... 37**
The Ask Skylar Page ..... 38
Using the Ask Skylar Feature ..... 38
Additional Options on the Ask Skylar Page ..... 39
**Using the Agentic Control Plane Page ..... 42**
The Agentic Control Plane Page ..... 43
Additional Agentic Features ..... 44
MCP Services ..... 45
**Managing Activities on the Dashboards Page ..... 48**
What is an Advisory? ..... 50
What is an Investigation? ..... 51
What is a Prediction? ..... 51
The Dashboards Page ..... 52
Filters ..... 53
Trend Analysis ..... 53
Activities and Options on the Recent Activity Tab ..... 54
Archiving an Activity ..... 55
Bookmarking an Activity ..... 55

Creating a Knowledge Base Article Based on an Activity ... 56
Using the Notes Feature ... 57
The Services, Device Groups, Issues, and Organizations Tabs ... 57
Viewing Basic Event Details on a Tab ... 58
Viewing More Event Details on a Tab ... 60
Viewing User Activity on the Customer Insights Page ... 62
**Working with an Advisory ... 64**
Before You Work an Advisory ... 64
Reading the Answer Badge ... 65
The Advisory Workflow at a Glance ... 65
Orienting Yourself in the Events ... 65
The Four Summary Tiles ... 66
Reading the Event List ... 66
Viewing the Event Flow Diagram ... 66
Working an Advisory From Claim to Close ... 67
Choosing the Right Tool for the Question ... 70
Escalating an Advisory ... 70
Common Mistakes When Working an Advisory ... 71
Checking Your Own Performance ... 71
**Integrating Data from Other Products ... 73**
Sharing Information with Slack and Microsoft Teams Integrations ... 74
Creating an Integration with Slack or Teams ... 74
Sending a Conversation to Slack or Teams ... 74
Importing ServiceNow Data into Skylar Advisor ... 75
Minimum Version Requirement ... 75
Supported Data Synchronization ... 75
Knowledge Base Filtering ... 76
Synchronization Approach ... 76
Available Controls ... 76
**Managing Events on the Data Lens Page ... 77**
The Data Lens Page ... 78
The Signals Tab ... 78

The Events Tab ... 79
Filtering the Data Lens Page ... 80
**Supervising with Skylar Advisor** ... **81**
What Supervisors Watch ... 82
The Two Supervision Loops ... 82
The Shift Loop ... 82
Reading the Queue ... 82
Finding the Three Faults ... 82
Unblocking and Rebalancing ... 83
Handing Off ... 83
The Weekly Loop ... 83
Reviewing Your People ... 83
Reading the Numbers Before You Use Them ... 84
Reviewing Your Knowledge ... 84
Reducing Noise ... 84
Approving Gated Tool Runs ... 85
Reporting Upward ... 85
Supervisor Failure Modes ... 86
**Managing OpenTelemetry (OTel) Alerts on the Alert Builder Page** ... **87**
Configuring the Skylar OTel Collector ... 88
Configuration ... 88
After Configuration ... 90
Considerations for OTel Data ... 90
The Alert Configuration Page ... 90
The Alert Definitions Tab ... 91
Editing Alerts ... 91
Viewing Triggered Alerts ... 92
The Alert Builder Tab ... 93
Alert Builder Aggregations ... 94
Notes for This Release ... 97
Data and Scope ... 97
Alert Conditions ... 97

Alert Timing ..... 97
The Alert Builder Page ..... 97
OTel Collector Installation and Operation ..... 98
**Experimenting with New Features on the Preview Features Page ..... 99**
The Preview Features Page ..... 100
Preview Feature "Graduates" ..... 101
Agentic ..... 101
Agentic: Topic Specialists ..... 102
Agentic: MCP Apps ..... 103
Agentic: Trusted Certificates ..... 103
Integrations: ServiceNow Outbound ..... 103
Integrations: Alert Builder ..... 104
Work Items: Metrics Enhanced Advisories ..... 105
Work Items: Advisory Suppression ..... 106
Work Items: Entity Exclusion ..... 107
Work Items: Data Lens Events ..... 108
Work Items: SLO Docs ..... 108
Work Items: Auto-close ..... 108
Platform: KB Generated Approval Flow ..... 109
Platform: Knowledge Sentry ..... 110

# Chapter 1

# Getting Started with Skylar Advisor

## Overview

***Skylar Advisor*** is an AI-powered application within the Skylar AI ecosystem that provides proactive, tailored guidance based on real-time data sent from your Skylar One system to Skylar AI. You can interact with Skylar Advisor by reviewing and following its proactive suggestions, called "Advisories." Additionally, you can query Skylar Advisor for advice and use it to perform deeper investigations of the issues and events in your environment.

For an overview of the Skylar AI platform, see *Introduction to Skylar AI*.

To view the latest release notes, see the *Skylar AI Release Notes*.

This chapter covers the following topics:


*What is Skylar Advisor?* ..... 2

*Logging In to the Skylar AI User Interface* ..... 2

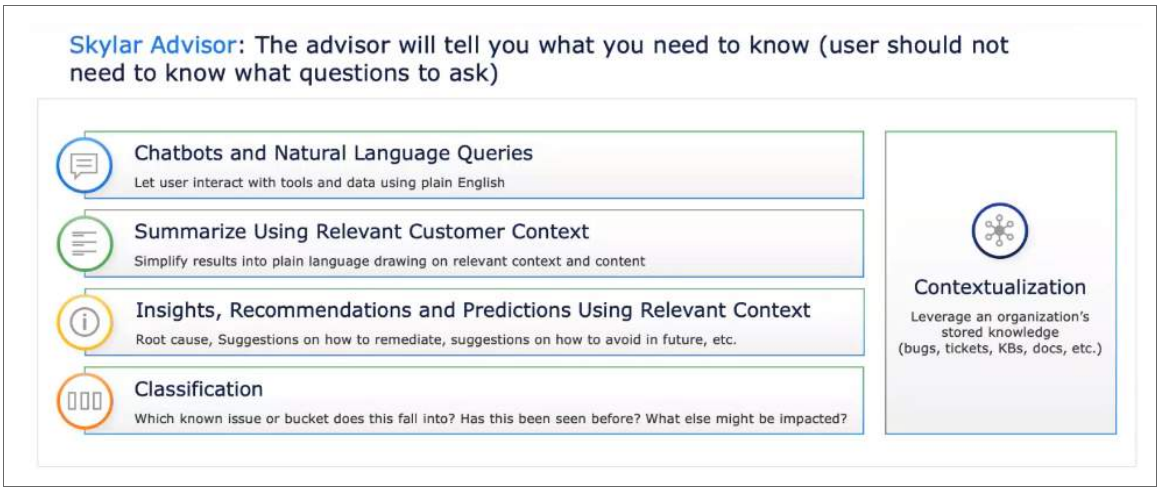
*The Skylar Advisor User Interface* ..... 3

*Defining Your Persona and Role Information* ..... 5

# What is Skylar Advisor?

Skylar Advisor is an AI-powered application that offers tailored guidance based on real-time data sent to Skylar AI from Skylar One, which lets you make informed decisions quickly and efficiently.

What Skylar Advisor is *not* is a "chatbot" or "assistant" that you have to query to get answers. Skylar Advisor uses its gathered knowledge to give you the potential answers before you get around to typing a question. Skylar Advisor can also forecast for issues that might occur, and create requests for upgrades.




Similar to large-language models (LLMs) like ChatGPT or Microsoft Copilot, you can also use the Ask Skylar feature in Skylar Advisor to start a conversion with Skylar AI to get more information.

# Logging In to the Skylar AI User Interface

If you know the URL of your Skylar AI system, you can go directly to that location. You can also access Skylar AI components from a link in Skylar One.

**NOTE:** The login process might vary slightly, depending on the versions of Skylar AI, Skylar One, and the Skylar One AP2 user interface currently running in your environment.

To log in to Skylar AI from Skylar One:

1. From Skylar One, go to the **Skylar AI** page ( ) and click the **[Visit]** button for the Skylar AI component you want to use, such as Skylar Analytics or Skylar Advisor. If you are not currently logged in to Skylar AI, the Skylar AI sign-in page appears.
2. If you need to log in to Skylar AI, type your email address and password and click **[Continue]**. The Skylar AI landing page appears.

3. If your Skylar AI site requires multi-factor authentication, but you have not set it up, you will need to scan the QR code that displays or enter the "secret key". This is a one-time-only step.
4. Click the name of the Skylar AI component you want to use, such as ***Advisor***, ***Analytics***, or ***Skylar Settings***.

# The Skylar Advisor User Interface

The Skylar Advisor user interface includes the following main pages and features:

- The **Ask Skylar** feature lets you query Skylar Advisor for customized guidance based on real-time data sent from Skylar One, which empowers you to make informed decisions quickly and efficiently. You interact with the **Ask Skylar** feature in much the same way as you would with ChatGPT, Claude, and other LLMs. For more information, see *Using the Ask Skylar Feature*.
- The **Dashboards** page is the default page that displays when you open Skylar Advisor. This page contains a list of your work and all of your interactions with Skylar Advisor. This page also lets you work with events imported from Skylar One, organized by business service, device group, issue type, or organization. For more information, see *Managing Activities on the Dashboards Page*.
    - The **[Reports]** tab on the **Dashboards** page lets you set up and schedule reports in Skylar Advisor. Available reports include "Corpus Health", "Handoff", "Business Services", "Trends", and "Spend Value". For more information, see *Scheduling Reports for Skylar Advisor*.
    - For more information about the process of moving from an unclaimed Advisory to a closed one, see *Working with an Advisory*.
    - If you are a NOC supervisor or shift lead holding Owner-level permissions, you can use Skylar Advisor to keep up with your team's work on Advisories. For more information, see *Supervising with Skylar Advisor*.
- The **Knowledge Center** page lets you upload your own documents to the Corpus to train Skylar Advisor on your company's specific workflows, policies, and processes. For more information, see *Managing Documents and Data on the Knowledge Center Page*.
- The **Data Lens** page lets you view details about the Skylar One events and event groups ("signals") that have been received by Skylar Advisor so that you can perform further investigations on them. For more information, see *Managing Signals on the Data Lens Page*.
- The **Toolbox** page ( > lets you view and test all tools from the Model Context Protocol (MCP) servers that have been enabled in Skylar Advisor. For more information, see *MCP Services*.
- The **Alert Configuration** page ( > Advisor > Alert Configuration) displays the data gathered by your OpenTelemetry (OTel) Collector and exported to Skylar Advisor. For more information, see *Managing OpenTelemetry (OTel) Alerts on the Alert Builder Page*.
- The **Customer Insights** page ( > Advisor > Customer Insights) displays data about the activities of all users on the Skylar Advisor instance. For more information, see *Viewing User Activity on the Customer Insights Page*.

- The **Preview Features** page ( > Account Settings > Preview Features) lets users toggle on and experiment with new features. For more information, see *Experimenting with New Features on the Preview Features Page*.
- The **Agentic Control Plane** page ( > Advisor > Agentic Control Plane) lets you monitor any AI agents and tools that you are using with Skylar Advisor. For more information, see *Using the Agentic Control Plane Page*.

> **TIP:** For information about the other pages available from the Skylar AI menu ( ), see *Configuring Skylar AI System Settings*.

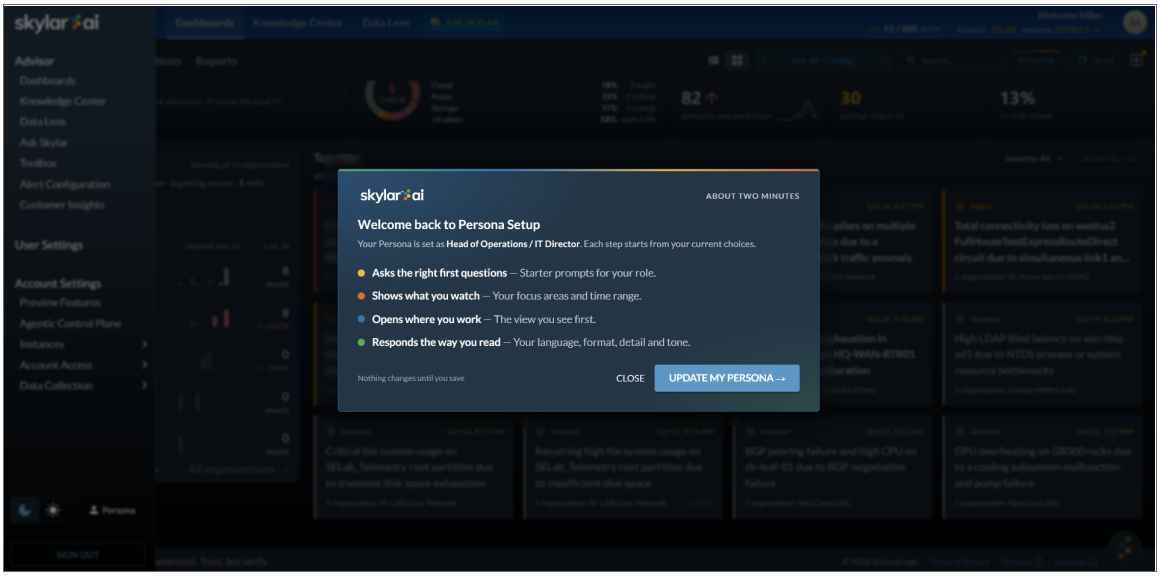
Use the following buttons and icons to help you navigate the Skylar Advisor user interface:

- The Skylar AI menu ( ) at top left gives you access to all of the Skylar Advisor pages, along with **User Settings** (your user information), **Account Settings** (information specific to this account), and the **Persona** page. The **User Settings** link takes you to the same **Profile** page that displays when you select ***Skylar Settings*** from the login page.
- If you have multiple instances of Skylar Advisor running, you can switch between those instances by clicking the drop-down next to ***Instance*** at the top right.
- The user icon ( ) at top right displays the email address and role for the current user in the user interface. You can click the **[Sign Out]** button to sign out of this session.
- The **[List]** ( ) and **[Dashboard]** ( ) buttons in the top navigation on the **[Recent Activity]** tab of the **Dashboards** page let you change how Skylar Advisor displays your activities.
- The **[Persona]** button on the **[Recent Activity]** tab takes you to the **Persona** page, where you can edit your personal settings for Skylar Advisor.
- Most Skylar Advisor pages and tabs have a **[Help Ctr]** button at top right or a Help Center icon ( ) in the small navigation bar on the left. The **Help Center** pane contains links to short videos that explain the latest features for this version of Skylar Advisor, and it also has a link to the Skylar AI product documentation site. The **Help Center** pane displays a yellow **[Updated]** button for new content, and the smaller **Help Center** icon on the left also displays a small yellow circle for new content.
- If a page contains a list, the enter selection mode icon ( ) adds check boxes to items in the list so that you can perform bulk actions to multiple items at once, such as bookmarking or archiving an item. Click the icon again to hide the check boxes.
- The notifications icon ( ) at the bottom left of the window lets you view all your recent notifications related to activities and warnings for that page. You can click **[Clear All]** to delete all past reminders. When you move to another page or tab, the list of notifications is updated for the new page.
- The **[Queue]** button on the bottom right of the **[Recent Activity]** tab displays a brief summary of your activities, including events, signals, and questions when you hover over it.
- The **Versions** link in the footer of any page displays the version numbers for Skylar Advisor and other applications. From the footer, you can also click links to view the Terms of Service and to view information about licenses and open-source packages.

# Defining Your Persona and Role Information

One of the first tasks you should work on when you get started with Skylar Advisor is to set up your Persona information with the **Persona Setup** wizard ( > Persona).

Your persona information lets Skylar Advisor know what your role is in your company, the language you will use for interacting with Skylar Advisor, and how you want Skylar Advisor to structure your feedback. You can also select your preferred tabs and layout for the **Dashboards** page for a more personalized experience.




To define your Skylar Persona using the **Persona Setup** wizard:

1. In the Skylar Advisor user interface, click the Skylar AI menu ( ) and click the **[Persona]** button at bottom right. The **Persona Setup** dialog appears.
2. On the **Persona Setup** dialog, click **[Personalize Skylar]** or **[Update my Persona]** if this is not your first visit to this page. The **Persona Setup** wizard appears.
3. The wizard presents you with a series of questions that you need to answer:

- ***Which sounds most like you?*** Select from one of the three roles, and then select a job title. You can edit the job description to better fit your role as needed.

    - *I Operate*. You are an engineer or an analyst, solving problems and closely using Skylar Advisor. If you select this option, you will have a **Your Work** pane and a **Team Work** pane at the top of the **Dashboards** page to track your activities. You can click **Coach Me** to get advice about your current workload.

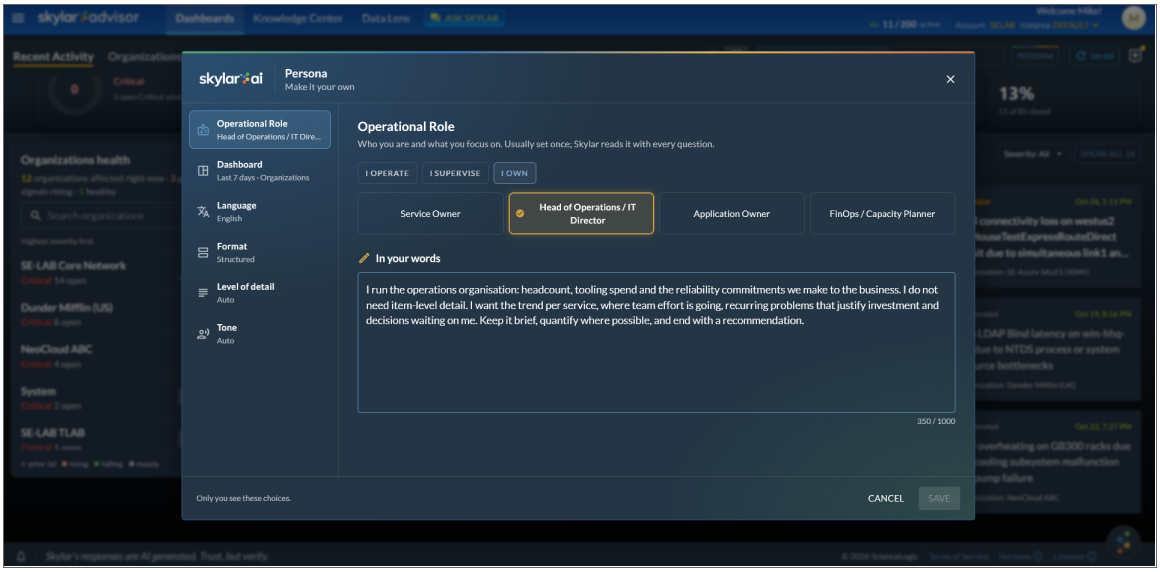
    - *I Supervise*. You are a manager with a as group of operators working for you. If you select this option, you will have a **Team Load** pane, a **Team Work** pane, and a **Health** pane at the top of the **Dashboards**page to track your team's activities. You can also edit team thresholds and have Skylar Advisor help coach your team from these panes.

    - *I Own*. You are an administrator or service owner responsible for maintaining Skylar Advisor. If you select this option, you will have a **Business Risk** pane and a **Health** pane at the top of the **Dashboards** page to track your activities. You can also create a report about the activities.

- ***What should your dashboard open on?*** Select the default time range, the tabs you want to display on the **Dashboards** page, and the number of activities to show up on each page.
- ***What language should responses use?*** Select your preferred language for your responses from Skylar Advisor
- ***What format should responses use?*** You can let Skylar Advisor choose the format for you (*Auto*), or have Skylar Advisor format responses in paragraph format or in a structured format like bulleted or numbered lists. Click each option to view as sample response.
- ***How much detail, and in what tone?*** Select the depth of detail you want for your responses. You can let Skylar Advisor choose the depth or granularity, or have Skylar Advisor give either detailed or brief responses. You can also let Skylar Advisor choose the tone, or have Skylar Advisor give either more formal or casual responses.
- ***What view should Skylar show you first?*** Select which type of data you want to display by default on the **Dashboards** page. The list of options is based on your role.

4. After the final step, click **[Save and Show My View]**. The **Dashboards** page appears with your customizations.

If you are on the **Dashboards** page, you can click the **[Persona]** button at the top right to return to the **[Persona]** page. The **Skylar Persona** page appears, but without the setup wizard that appears if you click the Skylar AI menu ( ) and click the **[Persona]** button:




**TIP:** You can return to the **Persona Setup** wizard or the **[Persona]** page at any time to update your persona information if you are not satisfied with the responses you are currently getting from Skylar Advisor.

# Chapter

# 2

# Scheduling Reports for Skylar Advisor

## Overview

This chapter covers how to set up and schedule reports in Skylar Advisor. These reports include:

- Corpus Health
- Handoff
- Business Services
- Trends
- Spend Value

**IMPORTANT:** If you are setting up Skylar Advisor for the first time, perform the activities in this chapter *before* you build out your Corpus. For more information about the Corpus, see *Prerequisites for Corpus Formation*.

This chapter covers the following topics:


*Prerequisites* ..... 9

*How Report Scheduling Works* ..... 9

*Enabling and Running the Corpus Health Report* ..... 9

*Finding Your Scheduled Reports* ..... 13

*Scheduling the Remaining Day-One Reports* ..... 14

*Quick-Reference Cadence Table* ..... 15

*Expected Behavior During Initial Setup* ..... *15*

# Prerequisites

The following items are required to set up and schedule reports:

- Skylar AI version 2.6.0 or later
- On the **Preview Features** page ( > Account Settings > Preview Features), enable the **Agentic** preview feature. For more information, see *The Preview Features Page*.

# How Report Scheduling Works

While Skylar Advisor answers the questions you ask, scheduled reports answer the questions you might not think to ask. That includes questions about the health of your Corpus, how a shift ran, or the value your organization is getting from Skylar Advisor. You can define when the reports run, so this information reaches you automatically instead of depending on somebody remembering to check.

Skylar Advisor generates reports using MCP tools that ship with the platform. You can run any tool on-demand from the **Toolbox** page ( > Advisor > Toolbox) to see what that tool might return, and you can schedule a tool so its report runs on a recurring basis.

Scheduled reports appear on the **[Reports]** tab of the **Dashboards** page, and the schedules appear on the **Scheduler** view within that tab.

Skylar Advisor includes five report tools that you can schedule during the initial setup:

1. **Corpus Health**. Is our Corpus and **Knowledge Center** page complete?
2. **Handoff**. How did this shift run?
3. **Business Services**. Are our business services healthy?
4. **Trends**. What events are occurring across our Advisories?
5. **Spend Value**. How is Skylar Advisor benefiting your organization?

Each report tool accepts an adjustable "lookback" time frame. The windows described in this chapter are recommended defaults, not limits.

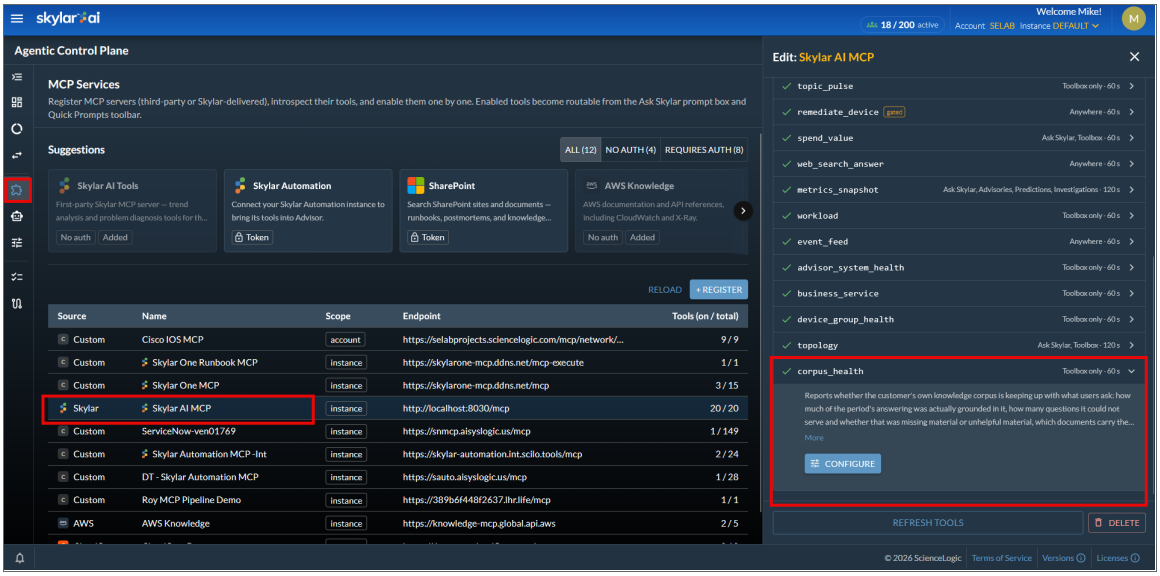
# Enabling and Running the Corpus Health Report

The "Corpus Health" report tells you whether your Corpus is complete and representative of the problems your team has been working on. If your system is in a steady state, this report replaces the need to manually mine each Advisory for knowledge gaps in your Corpus.

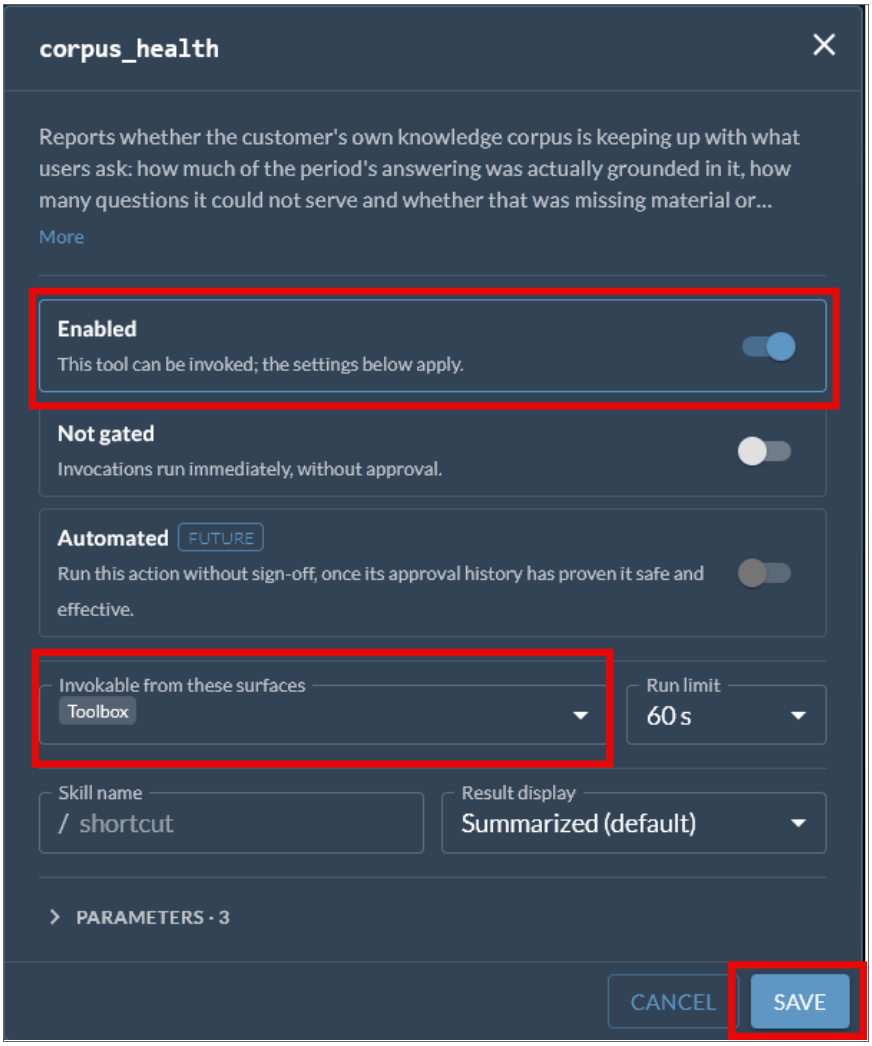
To enable, run, and schedule the "Corpus Health" report:

1. For Skylar AI 2.6.0 and older: On the **Preview Features** page ( > Account Settings > Preview Features), enable **MCP Apps** (turn the toggle blue) on the **[Agentic]** view.

2. On the **[MCP Services]** tab of the **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane > MCP Services), select the "Skylar AI MCP" service from the list at the bottom. The **Edit: Skylar AI MCP** pane appears on the right.

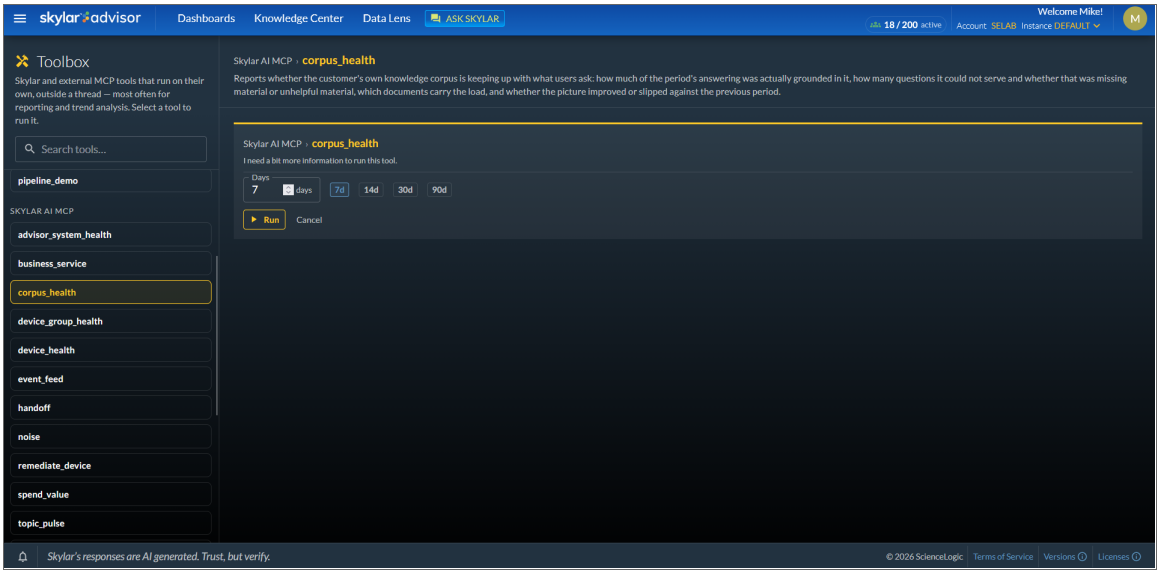


3. From the **Edit: Skylar AI MCP** pane, select the **corpus_health** tool and click the **[Configure]** button. A **corpus_health** modal appears.



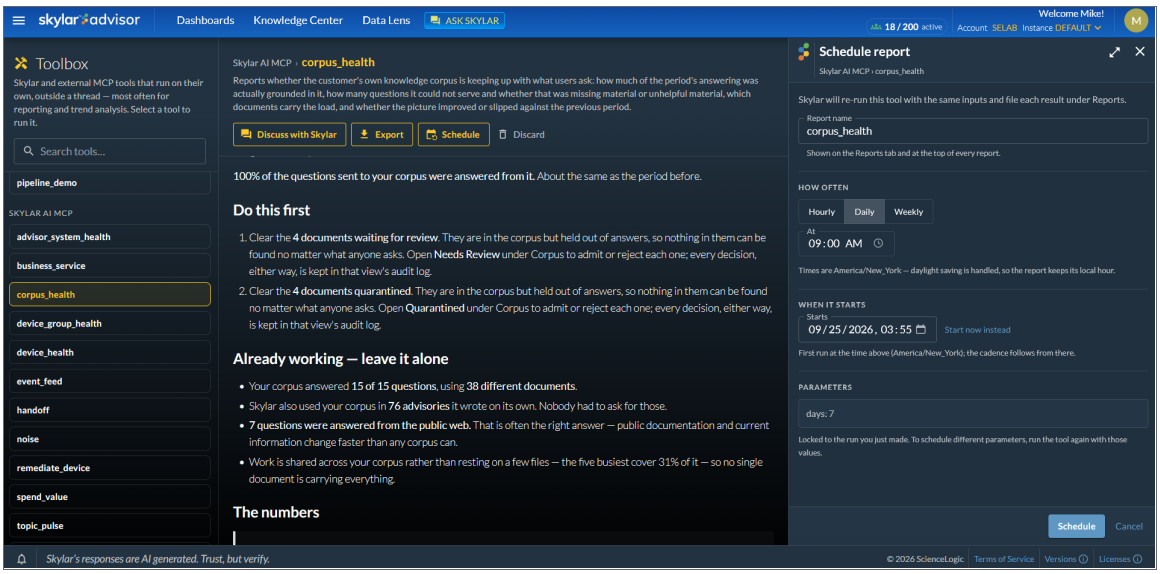

4. On the **corpus_health** modal, make sure that ***Enabled*** is selected, and then set ***Invokable from these surfaces*** to *Toolbox*.
5. Click **[Save]**.
6. As needed, repeat steps 2-5 for each of the remaining four report tools. These report tools are listed as ***handoff***, ***business_services***, ***trends***, and ***spend_value*** in the **Edit: Skylar AISkylar AI MCP** pane.

7. Go to the **Toolbox** page ( > Advisor > Toolbox) and select the **corpus_health** tool from the list on the left. A **Skylar AI MCP** pane appears to the right.




8. Select "7" or click `7d` in the ***Days*** field and click **[Run]**. Skylar Advisor generates a report describing the Corpus completeness for that time frame. If your Corpus is still empty, this run still gives you a baseline to compare against once you complete corpus formation.

9. At the top of the report, click the **[Schedule]** button. The schedule settings appear.




10. Select a cadence for the schedule, using the following suggestions from ScienceLogic:

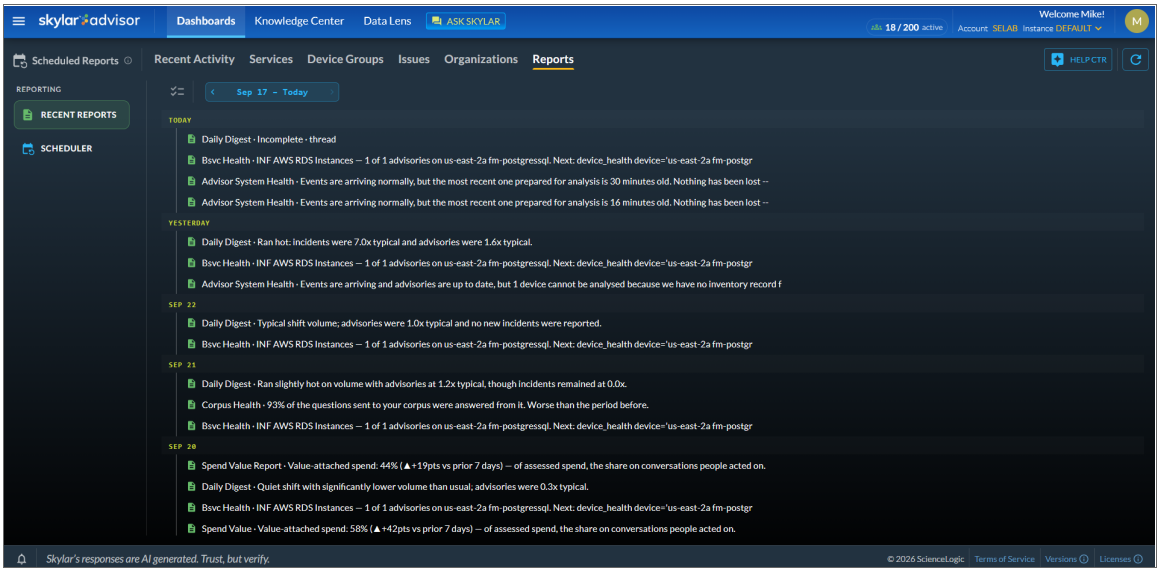
    - ***First 30 days***. Select a daily cadence. Corpus content changes quickly while you complete corpus formation, so a daily report keeps the feedback loop tight.

- ***After 30 days***. Relax the cadence to weekly as corpus health and completeness improve. ScienceLogic recommends a Sunday evening time slot so the report is ready for you Monday morning. You can edit a schedule at any time from the **Scheduler** view of the **[Reports]** tab.

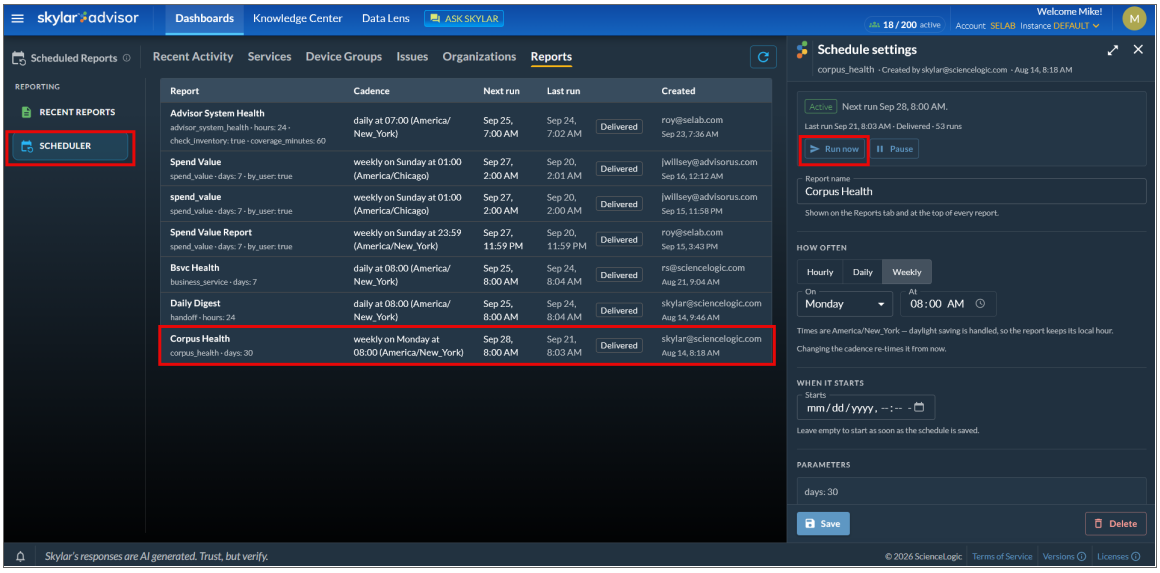
# Finding Your Scheduled Reports

To locate your scheduled reports:

1. Go to the **Dashboards** page. A **[Reports]** tab appears at the top, displaying a list of your scheduled reports.



2. Click the **Scheduler** view, the second item in the left section of the **[Reports]** tab. Your "Corpus Health" schedule appears in the list.




3. Select the report from the **Scheduler** view to open the **Schedule settings** pane, and then click **[Run now]** so you have a report to read today. After this first run, reports generate automatically on the cadence you defined.
4. Review the report on the **Recent Reports** view.

> **NOTE:** The **[Reports]** tab is empty the first time you schedule a report. Reports appear only after the schedule runs, so use the one-time run in the **Scheduler** view to generate a report immediately.

# Scheduling the Remaining Day-One Reports

Enable and schedule the remaining four day-one reports using the same pattern: enable the related MCP tool, run it once from the **Toolbox** page ( > Advisor > Toolbox) to see what it returns, and then click the **[Schedule]** button on the report to set its cadence.

## Handoff

- **What it tells you**. How the shift ran: whether it was quiet or high-activity, how many new Advisories opened, how many remain open at the end of the shift, and what the incoming shift needs to know.
- **Suggested Cadence**. Daily, or every 8 hours if that matches your shift schedule.
- **Naming**. Name the schedule *Handoff* if it runs once per shift, or *Daily Digest* if it runs once a day. You can rename a report schedule later, but the new name does not apply to reports the schedule already generated.

### Business Services

- **What it tells you**. The health of your business services. With a 7-day lookback, the report examines all Advisories and Investigations formed during the period, the rate of change on them, and which business services are in an unhealthy state.
- **Suggested Cadence**. Daily. You can also run this report on demand at any time.

### Trends

- **What it tells you**. Trends across all Advisories, predictions, and Investigations.
- **Lookback**. 30 days by default. As with all report tools, you can change the lookback window.
- **Suggested Cadence**. Weekly.

### Spend Value

- **What it tells you**. How your team is using Skylar Advisor and the value it is generating for them.
- **Suggested Cadence**. Weekly for the first 30 days, then monthly.

## Quick-Reference Cadence Table

The following table summarizes the recommended lookback and cadence for each day-one report:

| Report Name | Lookback (Default) | Cadence – First 30 Days | Cadence – Steady State |
|---|---|---|---|
| Corpus Health | 7 days | Daily | Weekly (Sunday, 8:00 p.m.) |
| Handoff / Daily Digest | Shift or day | Daily or every 8 hours | Same |
| Business Services | 7 days | Daily | Daily (plus on demand) |
| Trends | 30 days | Weekly | Weekly |
| Spend Value | 30 days | Weekly | Monthly |

## Expected Behavior During Initial Setup

Early "Corpus Health" reports might show an incomplete Corpus. This is expected during the first 30 days of setup as you complete Corpus formation. Completeness improves week over week. For more information about setting up your Corpus, see *Managing Documents and Data on the Knowledge Center Page*.

Also, the "Trends" and "Spend Value" reports might look thin during the first week. Both reports analyze accumulated history, so they provide more insight as Advisories, Investigations, and usage accumulate. For more information about these activities, see *Managing Activities on the Dashboards Page*.

# Chapter 3

# Managing Documents and Data on the Knowledge Center Page

## Overview

The **Knowledge Center** page in Skylar Advisor includes the **[Corpus]** tab, where you can upload and manage your own documents to further train Skylar Advisor on your company-specific workflows, policies, and processes.

You can also use the ***Research: Web*** option in the Ask Skylar prompt box to allow Skylar Advisor to perform an on-demand web search of specific websites and web sources instead of using the documents in the Corpus. Answers include source links and retrieval times, and they are never stored nor indexed.

This chapter covers the following topics:


*The Knowledge Center Page* ..... *18*

*Using Web Search* ..... *22*

*Best Practices for Uploading Documents to the Corpus* ..... *23*

*Setting Up a New Corpus (Corpus Formation)* ..... *24*

*Planning Your Corpus Health Baseline* ..... *25*

*Understanding Corpus and Web Search Ownership* ..... *30*

*Building Your Corpus (Owner)* ..... *31*

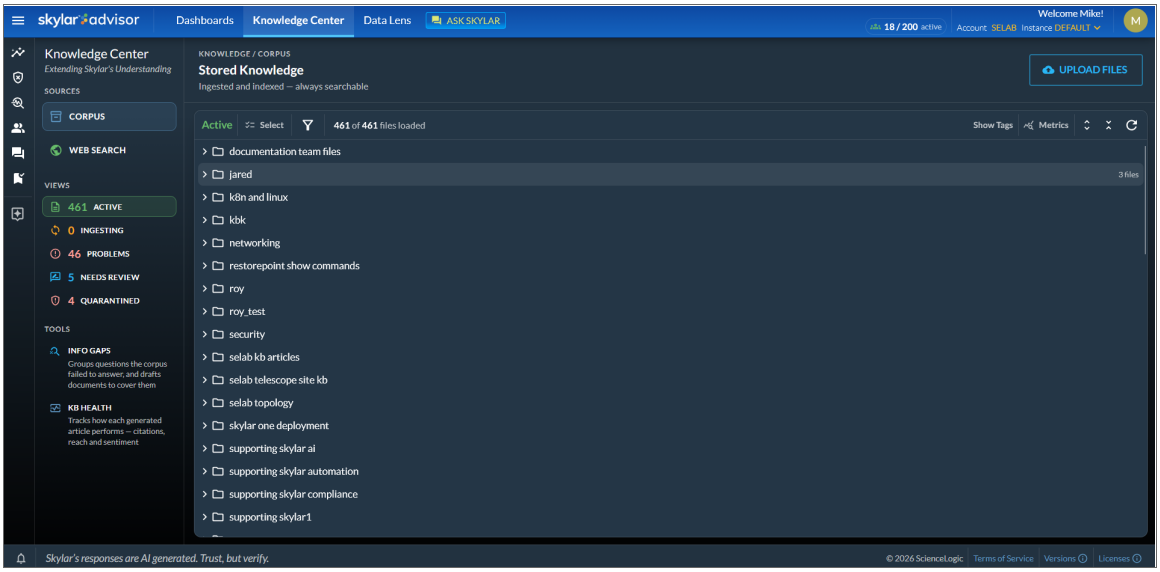
*Keeping Your Corpus Healthy* ..... *32*

*Removing and Monitoring Corpus Content* ..... *33*

*Administering Remote Knowledge with Web Search* ..... 34
*Recommended Maintenance Cadence* ..... 35
*Common Failure Modes* ..... 35

# The Knowledge Center Page

Skylar Advisor uses the documents that you and your team add to the **[Corpus]** tab of the **Knowledge Center** page to answer Ask Skylar questions and generate Advisories.



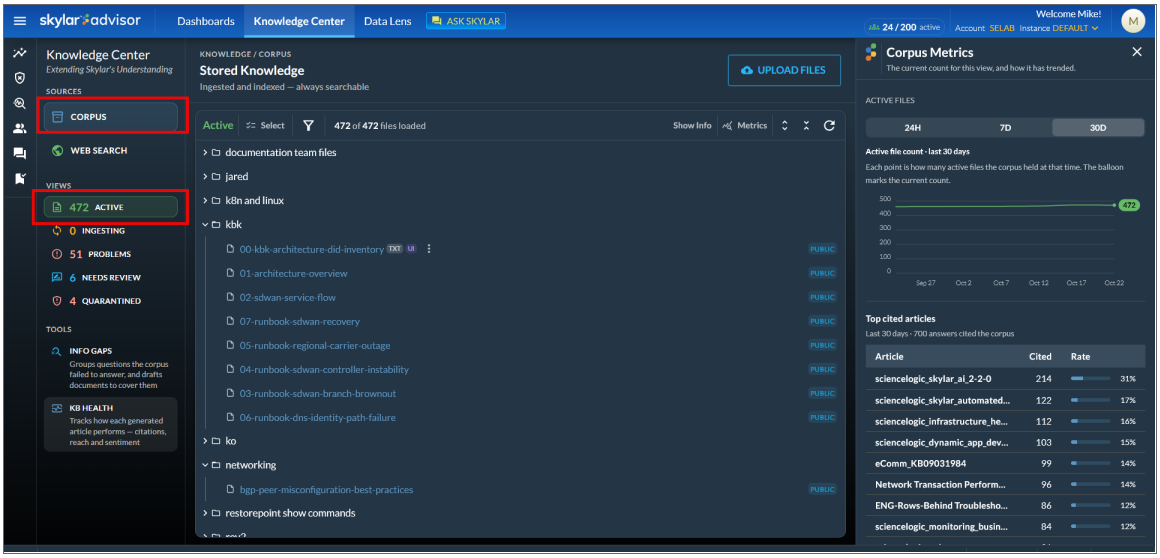

To get the most out of Skylar Advisor, ScienceLogic recommends adding Knowledge Base articles, product documentation, support tickets, release notes, and other similar documents. The more documents you upload to the **[Corpus]** tab of the **Knowledge Center** page, the more that Skylar Advisor learns about what is important in your environment.

**IMPORTANT:** Until you upload documents to the **[Corpus]** tab of the **Knowledge Center** page, Skylar Advisor uses only the general knowledge on which it was trained for its answers. ScienceLogic strongly suggests you start uploading documents specific to your environment, workflows, and best practices as soon as possible to make the feedback from Skylar Advisor more useful.

For more information, see *Best Practices for Uploading Documents to the Corpus*.

## The Active View on the Corpus Tab

The **[Active]** view on the **[Corpus]** tab of the **Knowledge Center** page displays a list of the documents that you have uploaded for your knowledge sources. This list is organized by most recently uploaded first, and you can further organize your documents by creating folders. You can add as many folders as needed.




You can monitor the files in each state of being uploaded on the following views on the **[Corpus]** tab of the **Knowledge Center** page:

- **[Active]**. These documents are completely uploaded and can be used by Skylar Advisor.
- **[Ingesting]**. These documents are in the process of being uploaded and ingested by Skylar Advisor. The countdown timer on the right side of the tab displays how long until the list is updated again, and it starts over every 15 seconds, or if the timer is clicked.
- **[Problems]**. These documents could not be added to the Corpus, along with an explanation for why it was not added. In some cases, the file might have been the wrong file type, empty, or corrupt, or the same content was already uploaded.
- **[Needs Review]**. The *Knowledge Sentry feature* flagged these documents as potentially unsafe. An administrator should review these documents to make sure they are safe to download.
- **[Quarantined]**. The Knowledge Sentry featured determined that these documents might be dangerous, and it did not upload those documents to the Corpus. These documents should be reviewed carefully before downloading them.

You can upload .pdf, .txt, .doc, .ppt, .ppx, .xls, and .xlsx files to the Corpus. You can also upload Confluence documents and generic JSON documents. The maximum file size for a document is 31 MB.

**NOTE:** Ingesting documents can take time, as Skylar Advisor creates an "embedding" that lets it search the document. It also adds tokens to the document. A rough estimate for ingesting documents while onboarding is about 2,500 documents per hour.

On the **[Corpus]** tab, you can:

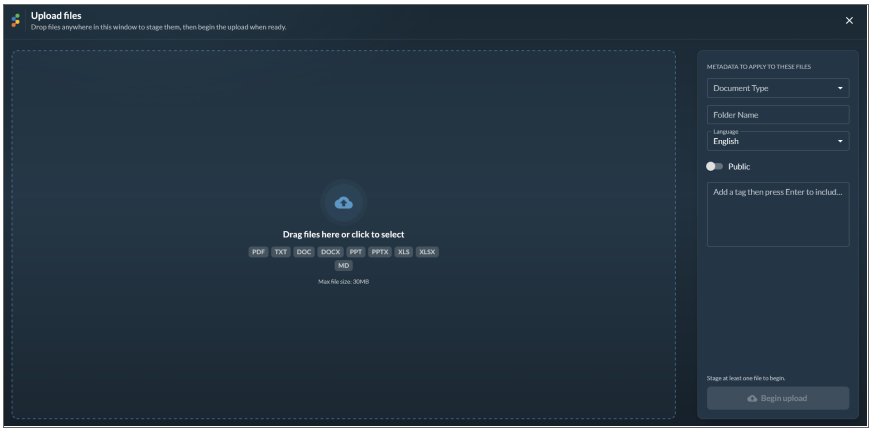
- Click the Select icon ( ) to add checkboxes next to each document. You can click a checkbox for one or more documents, and then click the Remove icon ( ) to delete those documents. You can select a folder to automatically remove all nested files and subfolders in a single delete action. You can also use the Select icon to edit tags for multiple files and folders. Enabled bulk deletion of by folder in the Corpus.
- Click the Show/Hide Filters icon ( ) to filter the list of documents by titles, tag names, documentation type, or visibility (public, private, or all). Click the icon again to hide the list of filters.
- If you have your documents organized in folders, you can click the Expand All icon ( ) to open all the folders. Click the Collapse All icon ( ) to close all the folders.
- Click the refresh icon ( ) to update the content on the **[Corpus]** tab.
- Click the ellipsis icon ( ) to view more info about a document, including metadata. You can also delete a file from this menu.
- Click the name of a document to open that document in a new modal.
- Click the **[Upload Files]** button to upload a document. You can add metadata and make folders, language, and tags for additional search context.

**IMPORTANT:** Before you start adding documents to the Corpus for Skylar Advisor, review the following topics to make sure you only upload the most useful documents to your Corpus:

- *Best Practices for Uploading Documents to the Corpus*
- *Setting Up a New Corpus (Corpus Formation)*

To add one or more files to the **[Corpus]** tab of the **Knowledge Center** page:

1. On the **[Corpus]** tab, click the **[Upload Files]** button. The **Upload Files** page appears.



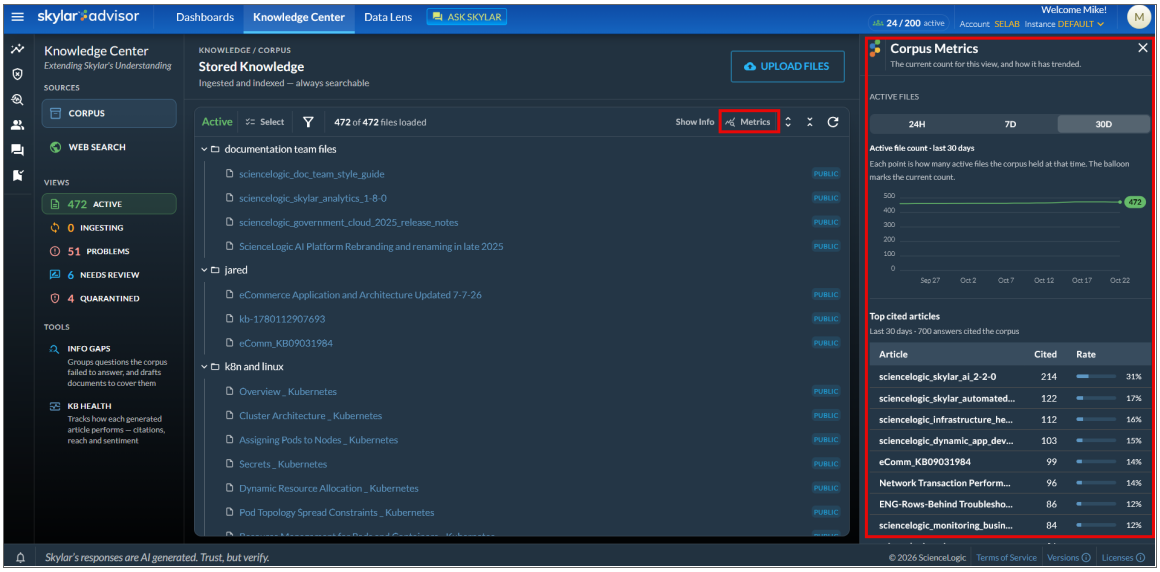

2. Drag and drop one or more files onto the window, or select the files to add them.
3. To help Skylar Advisor locate and learn each document, select the document type, type the folder name, and select the language for that document.

4. Click the ***Public*** toggle to turn it blue to make the document available to all organizations using this instance, or turn it white to make it available only to the organizations to which you have access.
5. To add specific tags, type in a tag below the ***Public*** toggle and press **[Enter]** for each tag. To see a tag for a document in a list, click the **[Show Tags]** / **[Hide Tags]** toggle in the header of the list.
6. Click **[Upload Files]**. Skylar Advisor takes a few moments to process the document you just uploaded. When the document is processed, it appears on the **[Corpus]** tab .
7. Click the refresh icon ( ) to update the content on the **[Corpus]** tab as needed. You can monitor the status of the document upload in the **[Views]** section of the **[Corpus]** tab.

**NOTE:** If the same document is added to the **[Corpus]** tab , Skylar Advisor highlights that file in yellow. You can click the ellipsis icon ( ) to delete the duplicate document.

## The Corpus Metrics Panel

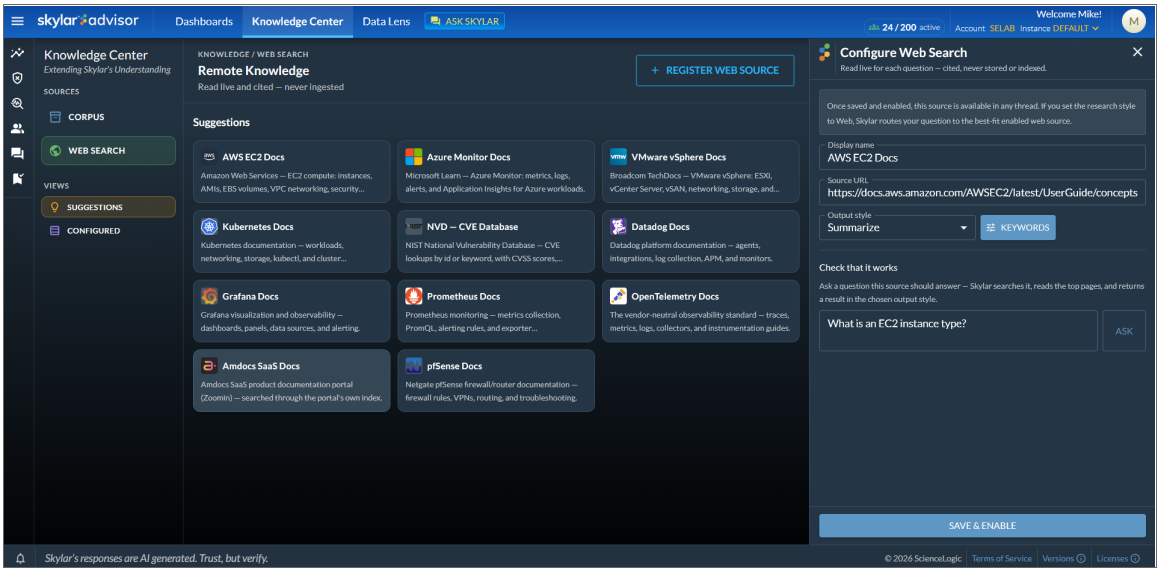
Clicking the **[Metrics]** button at the top right of the **[Corpus]** tab opens the Corpus Metrics panel. This panel contains a graph of active documents that have recently been uploaded. It also displays a list of the uploaded documents that have been used (or "cited") the most by Skylar Advisor to answer questions.



# Using Web Search

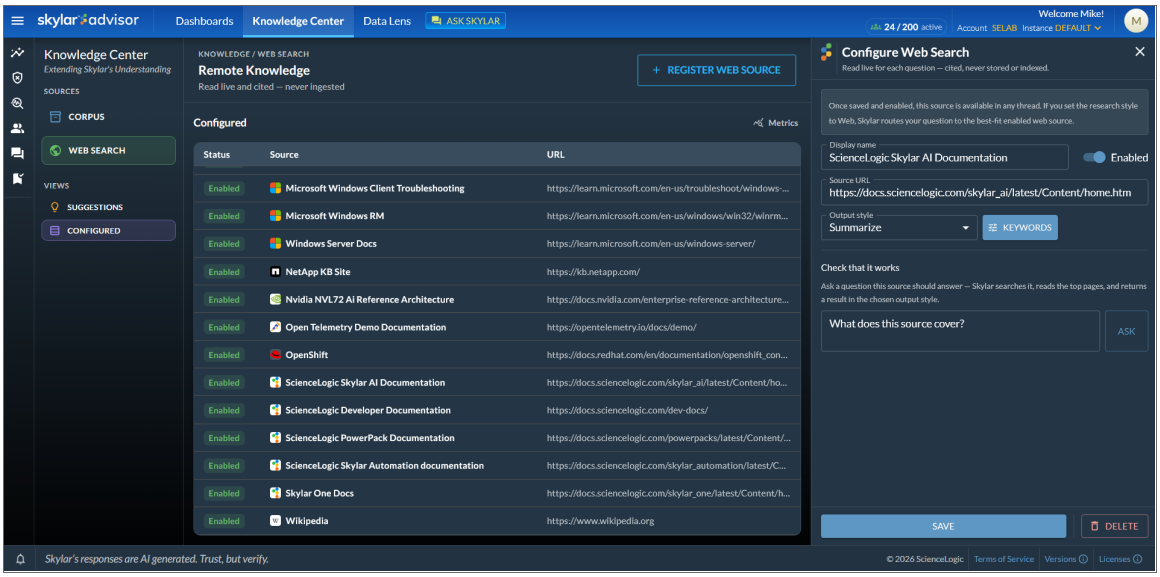
The **[Web Search]** view on the **Knowledge Center** page lets you perform an on-demand web search of specific websites and web sources instead of using the documents in the Corpus. Answers include source links and retrieval times, and they are never stored nor indexed.

You can set up the web sources on the **[Web Search]** view on the **Knowledge Center** page (Knowledge Center > Web Search > Suggestions):

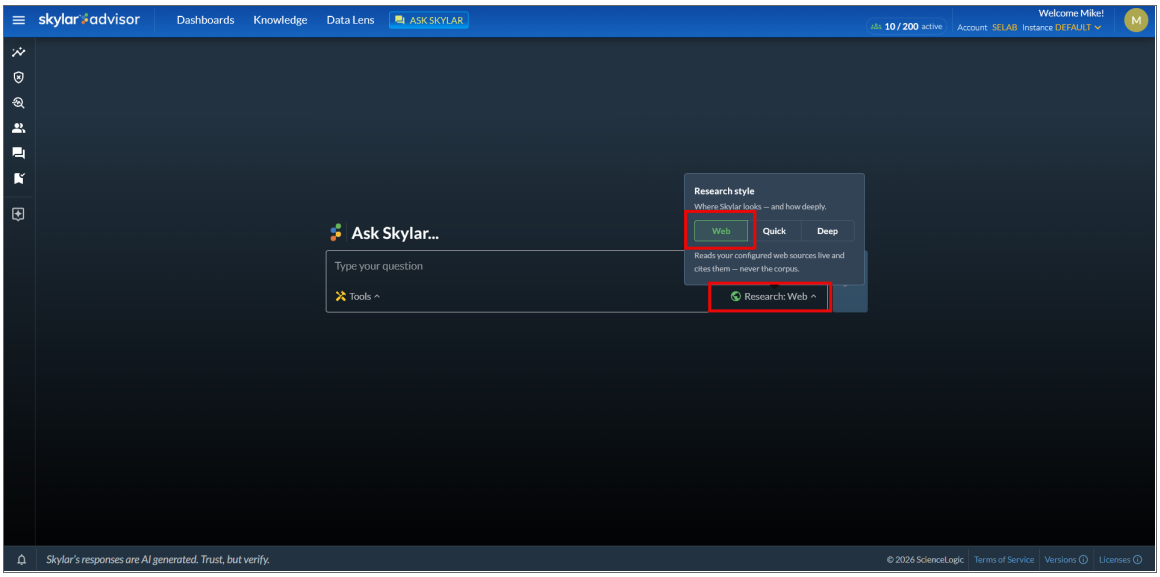



On the **[Web Search]** > **[Suggestions]** view, you can configure a suggested web source, such as ***Windows Server Docs***, to point Skylar Advisor at that source for answers. You can configure five different ScienceLogic Product Documentation sites as well as third-party sites.

The **[Web Search]** > **[Configured]** view shows the web sources that have been configured for web searches. You can edit each web source to customize it for your searching needs.




You will need to click ***Research: Web*** in the Ask Skylar prompt box to use this feature:




# Best Practices for Uploading Documents to the Corpus

Skylar Advisor learns from the operational knowledge provided by your organization. The quality of the documents in your knowledge corpus directly influences the accuracy and usefulness of the guidance from Skylar Advisor. As a result, when loading documents to the **Knowledge Center** page in Skylar Advisor, you should not upload every document you have. The best results come from a set of focused documents instead of a large amount of slightly related documents.

Not all documentation contributes equally to the effectiveness of Skylar Advisor. Start by uploading documents that answer the following operational questions:

1. What is this system supposed to do?
2. What does failure look like?
3. How do you fix it?

Upload the following types of documents first:

- **Run books or operational playbooks**. Examples of this type of document include incident response run books, service restart procedures, recovery workflows, failover procedures, and troubleshooting guides. These documents contain explicit remediation steps, which Skylar Advisor can surface as actionable guidance.
- **Incident postmortems**. Upload post-incident reports from recent or high-impact outages. Make sure these reports include the root cause analysis, symptoms observed, and resolution steps. These documents let Skylar Advisor identify patterns and known failure scenarios.
- **System architecture overview**. Include a document describing the architecture of your core systems, including major components, service dependencies, and data flows. This document helps Skylar Advisor understand how systems interact when diagnosing issues.
- **Service dependency documentation**. Upload documentation that maps how services depend on one another, such as upstream/downstream services, database dependencies, and external API dependencies. These documents enable Skylar Advisor to identify cascading failures.
- **Troubleshooting guides**. Add step-by-step guides for diagnosing common issues. Examples include database connection errors, authentication failures, and network latency issues. These guides provide diagnostic workflows for Skylar Advisor.
- **Known Issues list**. Upload a document that records recurring or unresolved problems. This list should include symptoms, temporary workarounds, and conditions that trigger the issue to help Skylar Advisor repeat incidents faster.

For best results, focus first on uploading the documents for your three to five most critical services. These systems typically generate the majority of alerts and incidents.

The following sections provide a more in-depth process for setting up a new Corpus in Skylar Advisor.

## Setting Up a New Corpus (Corpus Formation)

By default, Skylar Advisor knows the observability space in general. However, Skylar Advisor does not know your environment, such as your run books, your procedures, or the way your team resolves issues. A few hours of deliberate work in your first week can turn generic answers into environment-specific answers that your whole team can use.

The Corpus is your private Knowledge Base for Skylar Advisor. When Skylar Advisor answers a question, it tries to ground that answer in Corpus documents. An answer grounded in the Corpus displays a **[Verifiable]** tag, with links back to the documents that informed it. An answer with no Corpus backing displays an **[Unverifiable]** tag, meaning Skylar Advisor is drawing on its general knowledge instead of your environment. For more information, see *Additional Options on the Ask Skylar Page*.

## Finding Knowledge Gaps in the Corpus (Gap Mining)

Day one deliberately starts with an empty Corpus. Skylar One is already emitting events into the data lake for Skylar Analytics, so signals form and Advisories generate even with nothing in the Corpus. Every **[Unverifiable]** tag on those early advisories points to a ***gap***. You can address that knowledge gap in your Corpus by asking questions. This process is often called "gap mining."

The **[Info Gaps]** view on the **[Corpus]** tab of the **Knowledge Center** page analyzes only the questions you type. Advisories, Investigations, and Predictions do not feed it directly. If you only review Advisories but do not use Ask Skylar to ask questions, *Part 2* of this task returns an empty list. Mining questions in *Part 1* is the intake for Part 2.

This task follows a repeatable pattern:

1. Work an Advisory the way you normally would.
2. Let Skylar Advisor log every question it could not answer from the Corpus.
3. Fill in those gaps with real sources of truth.
4. Repeat this process.

# Planning Your Corpus Health Baseline

> **NOTE:** This section covers Corpus formation for an operator on the **Knowledge Center** page. If you are solely responsible for the **Knowledge Center** page, see also *Building Your Corpus*.

Use this task to build out your Corpus during your first week with Skylar Advisor. No prior Corpus work is required. Budget roughly one hour per Advisory, and plan to work through 5 to 10 Advisories using this process.

The "Corpus Health" report measures how well-grounded answers were over the questions asked during its lookback window. Capture your baseline after you complete Part 1 and before you backfill gaps in Part 2, below, when the number of questions asked is high and the gap share is near 100%. This before baseline is what makes a later comparison meaningful, and it is the chart a team lead can show to leadership.

Report lookback windows make exact timing non-critical, but reading a near-total-gap report makes the later comparison clear. For more information about setting up the report, see *Enabling and Running the Corpus Health Report*.

## Part 1: Mining an Advisory for Gaps

Budget about 15 minutes per Advisory for this part.

**NOTE:** If signals are forming but no Advisories exist yet, go to a heatmap on the **Dashboards** page and generate an ad-hoc Investigation from an event on a single device in a major or critical severity state. Discuss that Investigation with Ask Skylar exactly as you would an Advisory; every step below applies the same way. For more information, see *Viewing Basic Event Details on a Tab*.

To mine an Advisory for Corpus gaps:

1. Go to the **Dashboards** page and open an Advisory ( ) in a problem space your team knows well. Subject-matter familiarity matters, because you judge the quality of the answers.
2. Read the Advisory and note the **[Unverifiable]** tag. The Advisory reads plausibly but carries no **[Verifiable]** link, because nothing in the Corpus yet informs it. This is expected.
3. Use Ask Skylar to ask questions with the Quick Prompts below the Advisory, or write your own questions if you know the subject matter. Each response also returns **[Unverifiable]**. This is still expected.

   **NOTE:** Not every prompt feeds the gap analysis in Part 2. Some suggestions run agents or external tools instead of consulting the Corpus. Also, free-text questions that route to the agentic or data-lake path, such as "What servers were provisioned in 2022," also bypass the Corpus. The questions that count are knowledge questions: "How, why, and what is the procedure?" Spend your time on knowledge questions, or Part 2 returns less than the work you put into it.

4. Work the problem to conclusion, asking as many questions and follow-up questions as you would during a real shift. Every unanswerable knowledge question you ask becomes raw material for the gap analysis in Part 2.

   **TIP:** Click the **[Notes]** button to add comments or context to the conversation without spending a question on them. Notes do not trigger a response from Skylar Advisor. For more information, see *Using the Notes Feature*.

   **TIP:** Click the thumbs-down icon ( ) on any answer that is not well-grounded. The **[Info Gaps]** view treats a down-voted answer as gap signal and prioritizes it, giving you a second channel alongside the questions themselves.

5. Click the **[KB Gen]** button to create a KB article from the conversation if it reached a genuine, environment-specific conclusion, while the reasoning is fresh. If the conversation mostly repeated ground that a prior conversation already covered, skip KB generation and let the **[Info Gaps]** view cluster the related questions in Part 2 into one article instead of several overlapping ones. For more information, see *Creating a Knowledge Base Article Based on an Activity*.

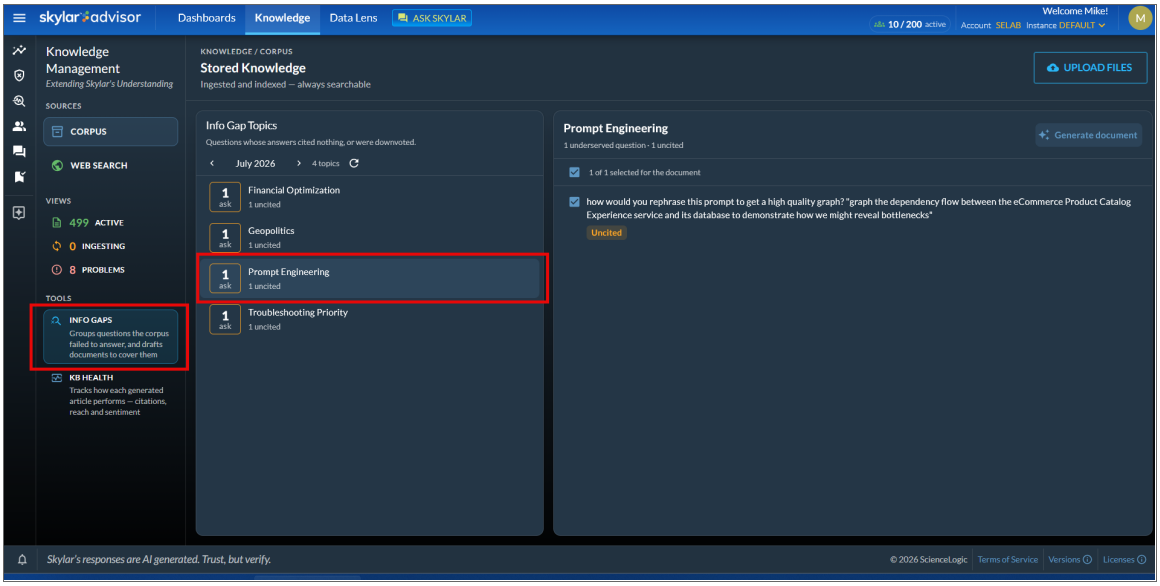
6. Repeat these steps across five to ten advisories. This process builds the question history that the **[Info Gaps]** view needs.

> **NOTE:** If you are a subject-matter expert, go to Ask Skylar and ask the general knowledge questions you expect your operators to ask, independent of any Advisory or Investigation, such as "What is the procedure for X?" These questions land in the same question history, so any unanswerable questions surface in the **[Info Gaps]** view alongside the Advisory-driven gaps.

## Part 2: Filling Gaps with the Info Gaps View

Budget about 20 minutes per Advisory for this part.

The **[Info Gaps]** view analyzes every question asked of Skylar Advisor that the Corpus could not answer. This view creates the ledger of what your Corpus is missing, based on situations where the documentation in the Corpus is insufficient or earned a down vote ( ). Skylar Advisor presents tools to cover the gaps, or you can generate a new document that addresses the issue.




To fill gaps using the Info Gaps view:

1. Go to the **[Corpus]** tab on the **Knowledge Center** page and select the **[Info Gaps]** view. The unanswerable questions from your Part 1 sessions appear in the list.
2. Select one or more questions. Skylar Advisor groups the questions you select into topic areas.
3. Fill each gap using one of two sources of truth:
   - Ask Skylar Advisor to generate a KB article representing the topic content.
   - Upload an existing KB article that your team already owns directly to the Corpus. For more information, see *Best Practices for Uploading Documents to the Corpus*.

   Both approaches are valid. An existing article is usually faster and more authoritative where one exists, while a generated article covers ground where nothing is written down.

> **NOTE:** A filled gap does not leave the **[Info Gaps]** list immediately. The list reflects a windowed history of when questions were asked, and a gap you backfilled yesterday keeps showing as long as the window still covers the original question, because a better, cited answer is never re-scanned against it. Judge your progress by whether new unanswerable questions are shrinking week over week and by the cold test in Part 3, not by the length of this list.

## Part 3: Verifying the Loop Closed

Budget about 10 minutes per Advisory for this part.

This is a cold-response test. If you re-ask the same questions inside the Advisory where you first asked them, Skylar Advisor can draw on three sources:

- the conversation it just had with you
- the KB article written from that conversation
- Corpus retrieval

Asking the same questions cold in Ask Skylar removes the first two sources, so a **[Verifiable]** answer there could only come from your Corpus. This test also better represents the real reader: an operator weeks from now, in a different context, asking a similar question about a similar problem.

To verify that the loop closed:

1. Open two browser windows side by side: one on Ask Skylar, and one on the Advisory. In the Advisory, open the options menu on a question and select *Copy*, paste the question into Ask Skylar, and run it. Work down the thread question by question. Expect a response with a **[Verifiable]** link, because it traces back to a document in the Corpus.

   Expect a few misses. A cold question is harder to serve than one asked with Advisory context, so some answers that would pass inside the original thread might fail here. A miss is signal, not a setback: it tells you the article content or phrasing does not retrieve well on its own. Revise the article and ask again. The **[KB Health]** view can help identify articles that need revision. For more information, see *Corpus Tools*.

2. Click each **[Verifiable]** tag and inspect the documents in the **Facts from Corpus** pane. Confirm a clear match between the document you added and the response you received. Retrieval might answer a question using a different Corpus document than the one you expected, so this check confirms that your article is actually carrying the answer. A **[Verifiable]** tag on a weak match is worse than an honest **[Unverifiable]** tag; if the match is weak, revisit the source material. For more information, see *Additional Options on the Ask Skylar Page*.

   > **NOTE:** The **Facts from Corpus** pane lists documents in relevance order, with the document that Skylar Advisor scored as most relevant listed first. If you disagree with that ranking, it usually means you have overlapping documents, or the question was too broad. Neither case is necessarily an error, but both are worth flagging and investigating.

## What to Expect During Your First Month

For your first month of Corpus formation, review the following checklist:

- During weeks 1 through 4, keep running this loop from the previous procedure as new types of problems arrive and Advisories form around them. This is normal.
- As volumes increase, the loop changes shape. Early on, you personally work every Advisory and see every miss. That does not scale: once Advisory volume grows past what you can work yourself, misses happen in questions you never see. The **[Info Gaps]** view becomes essential at this point, because it watches every Corpus-routed knowledge question across the team and surfaces the unanswerable ones for you.
- The **[KB Health]** view also starts contributing at this stage. It analyzes KB reach: where an article is working, where it is not, and what to improve. The **[KB Health]** view is enabled by the same **Corpus Tools** preview feature that enables the **[Info Gaps]** view, so it is already available. Where the **[Info Gaps]** view tells you what the Corpus is missing, **[KB Health]** tells you if what you already wrote is earning its keep.
- Over time, the need for this hands-on activity diminishes. The Corpus becomes whole and representative of the problems your team works, and gap mining becomes occasional instead of daily.
- In steady state, mine gaps by watching the scheduled "Corpus Health" report instead of working Advisory by Advisory. For more information, see *Enabling and Running the Corpus Health Report*.

## Signs Corpus Formation Is Complete

These tasks do not need to run indefinitely. Your Corpus is effectively complete when all three of the following are true:

- You have worked roughly ten Advisories through the full loop: mined, filled, and cold-verified.
- New unanswerable questions are visibly tapering. Judge this by the week-over-week volume of new questions entering the **[Info Gaps]** view, not by the length of the list, which keeps showing filled gaps until its window rolls past the original questions.
- The "Corpus Health" report is trending up against your baseline. Compare the latest report to the near-total-gap baseline you captured after your first mining pass; grounding should climb week over week.

When all three are true, move from the hands-on loop to the scheduled "Corpus Health" report and the **[KB Health]** view. If any one of them stalls, such as flat new-question volume or a plateaued "Corpus Health" report, work a few more Advisories in the problem spaces the reports point to.

## Expected Behavior During Corpus Formation

Every early answer returns **[Unverifiable]**. This is expected with an empty Corpus. It is the signal you are mining, not a defect.

Early Advisories read like general observability knowledge, because that is exactly what they are until the Corpus informs them. They become environment-specific as the Corpus fills. For more information, see *What is an Advisory?*

The **[Info Gaps]** view keeps accumulating new questions after you backfill gaps. New problem types generate new gaps, so this loop is designed to be re-run, and volume tapers over the first 30 days.

A gap you already filled might still show in the **[Info Gaps]** list. This is expected: the list is a windowed history of when questions were asked, not a live to-do list, and it clears as the window rolls past the original questions. Trust the cold test in Part 3 and the week-over-week new-question volume instead of list length.

## Understanding Corpus and Web Search Ownership

If you are the owner of the **Knowledge Center** page, you are responsible for both of its parts: the Corpus and Web Search:

- The Corpus is the documents you ingest and index, or stored knowledge.
- Web Search is the documentation sites that Skylar Advisor reads live and never stores, or remote knowledge.

Both parts are located on the **Knowledge Center** page, and both are yours to maintain.

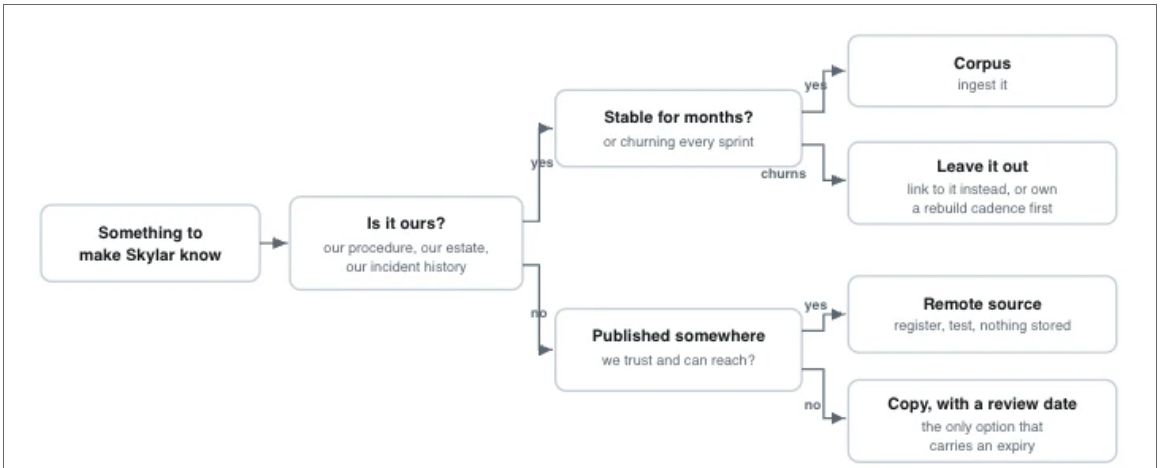
> **NOTE:** The **[Info Gaps]** and **[KB Health]** views require the **Corpus Tools** preview feature. The **Web Search** feature and its source administration require the **Web Search** preview feature and Owner-level permissions or higher, so some of the views described in this topic might not appear in your environment. For more information, see *Corpus Tools* and *Web Search*.

The Corpus and Web Search are not competing stores. They cover different kinds of truth, and choosing correctly between them is most of this task.

| | Corpus (Stored Knowledge) | Web Search (Remote Knowledge) |
|---|---|---|
| **What it holds** | Your documents, ingested and indexed. | Pointers to documentation sites, read live at question time. |
| **Ownership** | Yours. Private, and it works with no network access. | Someone else, always current. |
| **Maintenance cost** | You keep it true. Nothing else does. | The vendor maintains it. You maintain the pointer. |
| **Goes stale?** | Yes, and that is the main risk you manage. | No. Nothing is stored, so nothing goes stale. |
| **Answer tag** | **[Verifiable]** | **[Web Cited]** |

**IMPORTANT:** If the knowledge is yours and would still be true in three months, add it to the Corpus. If someone else publishes and maintains it, register it as a remote source in **Web Search** instead of keeping a copy you have to update as the source changes.

Refer to the following flowchart for more information about where to store your knowledge:




The two middle outcomes are the ones usually missed. Knowledge that is yours but churns every sprint does not belong in a store you re-ingest by hand — and a vendor document you cannot read live is a copy with an expiry date, which is fine only if someone owns the review.

# Building Your Corpus (Owner)

**NOTE:** This section covers Corpus formation for an owner of the **Knowledge Center** page.

Early in Corpus formation, the temptation is to add volume: load every document and let retrieval sort it out. Retrieval cannot sort it out. A Corpus with two answers to the same question is worse than a Corpus with none, because Skylar Advisor finds one of the answers and an operator acts on it.

## What to Include First

Start by uploading documents that answer three questions:

1. What is the system supposed to do?
2. What does failure look like?
3. How do you fix it?

| Include | Why It Earns Its Place |
|---|---|
| Run books and standard procedures | The highest-value content you own. It directly answers "What do I do now?" No vendor publishes your approach to problem solving. |
| Postmortems and resolution notes | Feeds "Has this happened before, and how was it fixed?" which is a question the web can never answer. |
| Your architecture and topology | What depends on what in your estate. Turns a device alarm into a business-impact statement. |
| Escalation paths and ownership | Small, stable, and constantly needed at 3:00 a.m. |

| Include | Why It Earns Its Place |
|---|---|
| Service-level agreements | Per-service response and resolution clocks, and what a breach costs. These documents are contractual, specific to you, and stable between renewals, and they turn "this is critical" into "you have forty minutes." |
| Your deployment's configuration decisions | Not the vendor's manual. The parts where you deviated from it, and why. |

### What to Leave Out

- **Public vendor documentation**. This is the most common mistake. It is large, maintained by someone else, and goes stale the day after you ingest it. Register it as a remote source instead. For more information, see *Administering Remote Knowledge with Web Search*.
- **Anything that changes faster than you re-ingest it**. If nobody owns re-ingesting a document monthly, it will be wrong within a quarter.
- **The same procedure in multiple formats**. Pick the authoritative version. Duplicate documents do not add coverage; they split it.
- **Unresolved tickets, chat exports, and raw logs**. These documents carry high volume and low signal, and they dilute retrieval for everything else.
- **Drafts and unconfirmed information**. The Corpus is the source of truth. Anything in it carries your authority whether you intended it to or not.

> **TIP:** Aim for the 20 questions your operators ask most often, answered once each, authoritatively. That approach outperforms a thousand documents nobody can navigate, and it gives you a baseline to measure against once the **[Info Gaps]** view starts reporting what is missing.

## Keeping Your Corpus Healthy

Under the **[Corpus]** tab of the **Knowledge Center** page, the **[Active]**, **[Ingesting]**, and **[Problems]** views show the operational state of your Corpus. Check them on a cadence, not only when someone reports an issue.

| View | What It Tells You |
|---|---|
| **[Active]** | What is live and searchable right now. This is the real Corpus: not what you uploaded, but what survived ingestion. |
| **[Ingesting]** | Documents in flight. A document that sits here far longer than its peers is worth checking. |
| **[Problems]** | What did not ingest, grouped by cause and severity. Anything in this view is content you believe you have and do not. |

### The Three Problem Types

| Type | Severity | What to Do |
|---|---|---|
| Duplicates | Minor | The same content is already present. Usually harmless, but a rising count means people are loading overlapping material, which is worth a conversation among the people who curate Corpus material. |
| Failed to ingest | Major | The document did not make it in. Read the status message on the row, fix the source, and re-upload. This is the pile that silently shrinks your coverage. |
| Corrupted | Critical | The file itself is broken. Re-source it from the original instead of retrying; a retry of a corrupted file is still a corrupted file. |

**NOTE:** An empty **[Problems]** view is not the same as a healthy Corpus. The **[Problems]** view only reports documents that failed on the way to being ingested. A document that ingested successfully but is now outdated appears as **[Active]** and looks identical to your best content. For more information, see *Removing and Monitoring Corpus Content*.

## Removing and Monitoring Corpus Content

After your Corpus is established, your task shifts from addition to editing. The **Corpus Tools** preview feature provides two views, **[KB Health]** and **[Info Gaps]**, that tell you what to act on from opposite ends.

### KB Health: What Is Earning Its Place

Every generated KB article carries citation and sentiment metrics, plus two scores: ***KB Impact***, which measures how often an article is cited weighted by how strongly, on a zero to 100 scale; and ***Influence***, which measures that strength per citation, so a rarely used article that is decisive when it is used still scores high.

| What You See | What It Usually Means |
|---|---|
| Never cited | One of three things: the article is wrong, unfindable, or about something that stopped happening. All three are worth ten minutes of review. |
| Cited often, low influence | The article keeps coming up but rarely decides the answer. It is usually too broad; consider splitting it. |
| Cited rarely, high influence | Your best content. Leave it alone, and consider whether its subject deserves more coverage. |
| Negative sentiment | Operators are marking answers built on this article as unhelpful. Review it against current reality before anything else in this table. |

## Info Gaps: What Is Missing

The **[Info Gaps]** view collects the questions in its window whose answers cited nothing from the Corpus or were down-voted, and clusters them into topics. For each topic, you can generate a covering document, edit it, save it, ingest it, and then replay the original questions against the updated Corpus to confirm that the gap closed.

Always replay the questions after you ingest a covering document. Replaying is what makes this a learning loop instead of a backlog; a covering document that does not change the answer is a document you now maintain for no benefit.

## What to Remove

- **Superseded procedures**. When you replace a run book, delete the old one the same day. Two live answers to the same question is the failure mode you are guarding against.
- **Never-cited articles about problems you no longer have**. If the underlying system changed and the fix no longer applies, the article is now historical only.
- **Content whose subject moved to a vendor site you now read live**. Replace the copy with a remote source and delete the copy from the Corpus.
- **Duplicates accepted during a bulk upload**. These duplicates do not surface as errors; they surface as inconsistent answers.

Deleting content is part of this task, not a failure of it. A Corpus that only grows becomes an average of everything you have ever believed, and retrieval has no way to know which of two contradictory documents is current. Use the **[KB Health]** and **[Info Gaps]** views to decide what to remove.

# Administering Remote Knowledge with Web Search

Under the **Web Search** feature on the **Knowledge Center** page, you register documentation sites that Skylar Advisor reads live at question time. Nothing is ingested and nothing is stored; answers cite the exact pages read and when they were read. Registering and editing sources requires Owner-level permissions or higher.

> IMPORTANT: You need to click ***Research: Web*** in the Ask Skylar prompt box before you ask your questions.

To register and maintain a remote source:

1. On the **[Suggestions]** view, select a pre-verified source, or register your own **https** documentation site by hand. Selecting a suggested source pre-fills its name, URL, a sample question, and a tuned self-description.

2. Run the live test before you save the source. The test classifies the site and answers a real question from it, including search, page reads, and citations, before the source is saved. Judge the quality of the answer, not just whether the test returned a result. A source that indexes content but answers off-topic is worse than no source, because it wins questions it should not answer.
3. Review and edit the keywords for the source. This field decides which questions reach which source. Skylar Advisor derives keywords automatically, but your edits are usually an improvement. Use vocabulary that is distinctive to that source and absent from your other sources; overlapping descriptions are the most common cause of a question landing on the wrong site.
4. Choose the output style for the source: *Summarize* composes a written answer from the retrieved pages, and *Raw* returns the retrieved excerpts verbatim, still cited. Use *Raw* where the record itself is the answer, such as vulnerability data.
5. Review the citation rate for each source monthly on the **[Configured]** view. This view reports how many of the month's web answers cited each source. If a source is cited in nothing, its description is likely inaccurate, or the source is unnecessary. Correct the description once; if the citation rate stays at zero, remove the source.

> **NOTE:** Grow your list of remote sources slowly. Source selection compares each question against every source's self-description and keywords, so ten overlapping descriptions blur the distinctions that routing depends on. Two well-described sources outperform six mediocre ones.

## Recommended Maintenance Cadence

The following table summarizes a recommended cadence for maintaining your Corpus and Web Search sources:

| Cadence | Do This |
|---|---|
| Weekly | Clear the **[Problems]** view: corrupted documents first, then documents that failed to ingest. Skim the **[Ingesting]** view for anything stuck. |
| Monthly | Work the **[Info Gaps]** view top topic down, and replay each. Review **[KB Health]** for never-cited and negative-sentiment articles. Check remote-source citation rates and health. |
| Quarterly | Prune. Walk the oldest documents in the **[Active]** view and ask of each: is this still true, is it still ours, and is it still needed? Delete rather than defer. |
| On every incident of note | Review the KB article the operator generated. It is drafted from their thread, so it inherits their gaps, and it will be trusted next time. |

## Common Failure Modes

The following table lists failure modes to avoid when maintaining your Corpus and Web Search sources:

| Pattern | Why It Costs You |
|---|---|
| Ingesting vendor documentation | You take on maintenance of something maintained better elsewhere, and it starts going stale immediately. |
| Measuring the Corpus by document count | Coverage is what matters. The **[Info Gaps]** view measures coverage; document count measures effort. |
| Adding without removing | Contradictory documents accumulate, with no way for retrieval to tell which one is current. |
| Generating covering documents without replaying | You cannot tell whether the gap closed, and you have added maintenance either way. |
| Leaving ingestion problems to accumulate | Every row is content you believe you have and do not. It shows up as a bad answer weeks later. |
| Adding many remote sources at once | Overlapping self-descriptions send questions to the wrong site, and the whole feature looks unreliable. |
| Treating operator KB articles as finished | These articles are drafted from one thread by someone mid-incident. Review is the difference between an asset and a liability. |

Keep the following points in mind:

- Content that is yours and stable belongs in the Corpus; content that is someone else's and current belongs in a remote source.
- Clear the **[Problems]** view weekly.
- Close gaps in the **[Info Gaps]** view, and replay the questions to confirm each gap closed.
- Delete content as deliberately as you add it. The Corpus remains the source of truth only while someone is willing to remove what is no longer true.

# Chapter 4

# Using the Ask Skylar Feature

## Overview

You can query Skylar Advisor for advice using the Ask Skylar feature to perform deeper investigations of your issues and events from Skylar One.

Skylar Advisor offers tailored guidance based on real-time data from Skylar One, helping you make informed decisions quickly and efficiently.

The Ask Skylar prompt box saves you from staring at dashboards that other people made to solve your problems. Instead, it lets you ask your own questions over a data lake that holds your metrics and a knowledge base that contains data on how your company solves problems.

This chapter covers the following topics:

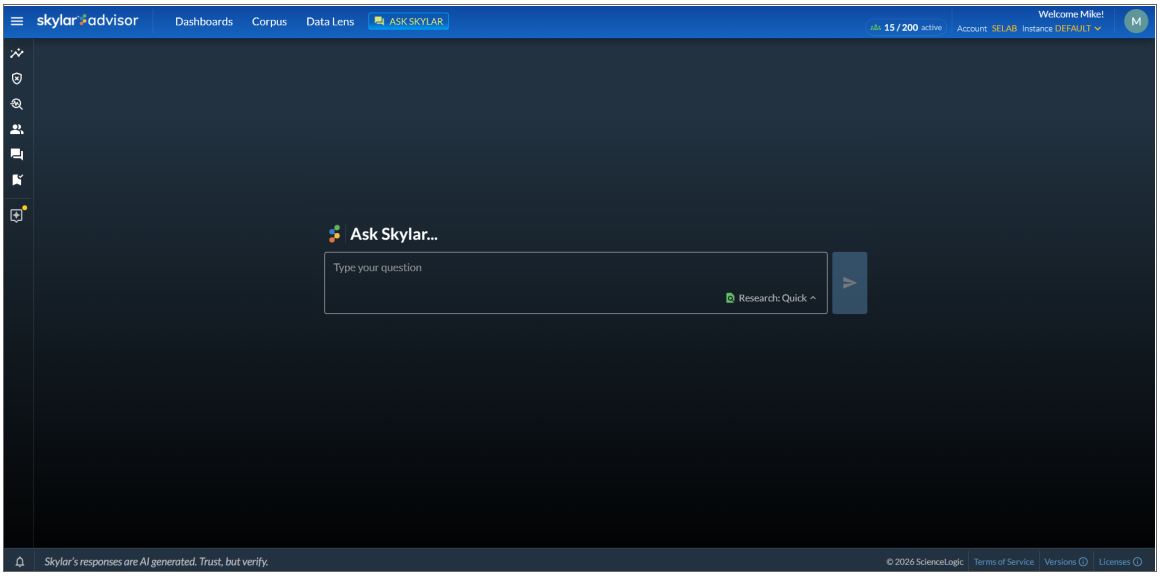

*The Ask Skylar Page* ..... *38*

*Using the Ask Skylar Feature* ..... *38*

*Additional Options on the Ask Skylar Page* ..... *39*

# The Ask Skylar Page

The **[Ask Skylar]** button lets you query Skylar Advisor for tailored guidance based on real-time data, which empowers you to make informed decisions quickly and efficiently.




You can ask Skylar Advisor any question, and Skylar Advisor will respond by drawing on everything it has learned from the documents uploaded to the **Corpus** page to provide you with the most accurate answers to your questions.

# Using the Ask Skylar Feature

To use the **Ask Skylar** feature:

1. Click the **[Ask Skylar]** button at the top of the Skylar Advisor user interface or at the bottom of an existing activity from the **[Recent Activity]** tab of the **Dashboards** page.
2. In the **Ask Skylar** prompt box, type your question for Skylar Advisor.

> **TIP:** You can drag and drop one or more text (.**txt**) files onto the Ask Skylar prompt box so you can ask questions using those files as context. For example, you could upload a text file containing log information, and then type your questions about logs in the Ask Skylar prompt box.

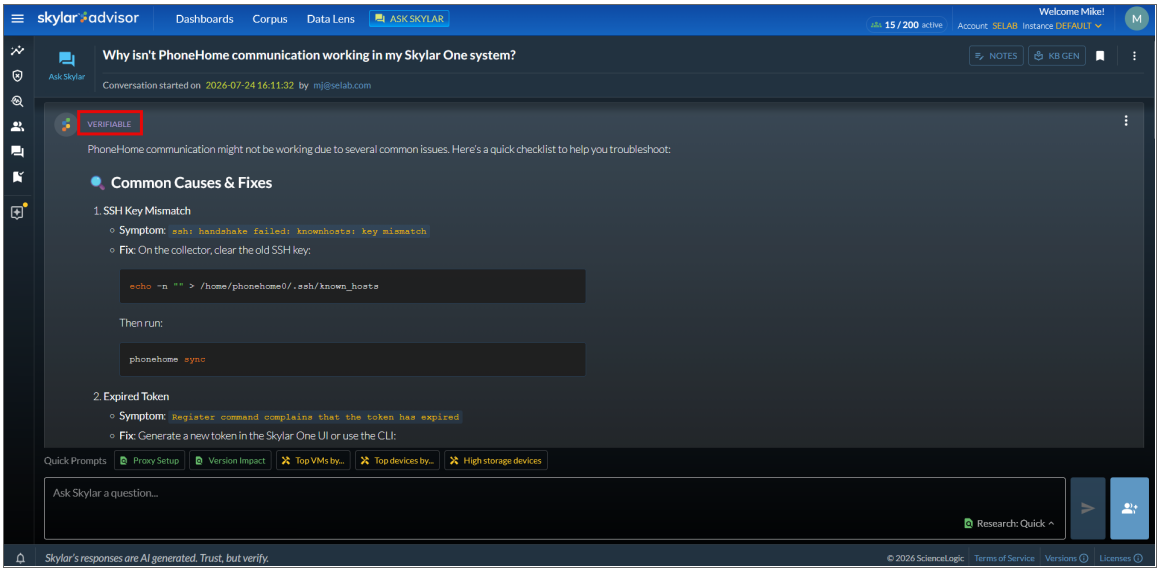
3. Before you click the send icon ( ), you can use the **Tools** options to select an agentic AI tool to handle your question.

   The **Tools** options only appear on the Ask Skylar prompt box if you enable one or more MCP servers on the **[MCP Services]** tab of the **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane > MCP Services). For more information, see *MCP Services*.

4. Before you click the send icon ( ), you can use the **Research** options to specify the depth of research that Skylar Advisor should use when generating its answer. Select *Quick* for faster, less-technical answers. Select *Deep* to tell Skylar Advisor to perform a thorough analysis of the corpus. This option is more time-consuming, but you might find it useful for more complex questions.

   If you have the **Web Search** Preview Feature enabled, you can select *Web* from the **Research** options, and Skylar Advisor will only search the websites specified on the **Web Search** > **Configured** view of the **Knowledge Center** page.

5. Click the send icon ( ). Skylar Advisor displays its progress and a timer, along with a **[Stop]** button if you would like to cancel the query:




6. If you would like more information about any text in an answer, highlight and right-click the text and select *Quick Take*. A **Quick Take** pane appears, displaying additional information about that text.

7. To view additional features you can use with Ask Skylar, see the *following topic*.

# Additional Options on the Ask Skylar Page

Along with the answer to your question, Skylar Advisor also displays the following buttons and options:

- **Quick Prompts**. Based on the original question, Skylar Advisor displays additional prompts in green boxes that you can click for additional information. You can hover over a Quick Prompt to see more information about it.

  If you have the *Agentic* preview feature enabled, you can click one of the yellow agentic tools buttons in the **Quick Prompts** section to select an AI agent to help with your Ask Skylar questions and conversations.

  If you are using the **Ask Skylar** feature in an Advisory or Investigation detail record, this section displays an ***Events*** link that you can click to view the Skylar One events that were used to generate this Advisory.

- **Collaborate with others** icon (). Opens the **Collaborate** pane, where you can share your Skylar AI conversation with another user. Click the **[Copy Link]** button to generate a URL that you can share with another user. You can also click the **[Slack]** or **[Teams]** buttons to send this conversation to Slack or Microsoft Teams; for more information, see *Sharing Information with Slack and Microsoft Teams Integrations*.

  After you share the conversation, it is also added to the **Shared Conversations** section on the **[Recent Activity]** tab of the **Dashboards** page. You might need to click the refresh icon () to see the most recent questions and answers in this conversation.

- **[Verifiable]** tag. This tag displays at the top of the answer when the information for that answer comes from a document in the Corpus. Click the **[Verifiable]** tag to open the **Facts from Corpus** pane, which lists the documents from the Corpus used for the answer, along with a link to the document.

  You can also click the ***Show Raw Facts*** toggle to display the content used for the answer. If an answer did not come from the Corpus but from the general knowledge of Skylar Advisor, this tag displays as **[Unverifiable]**.

- **[Verified (via Agents)]**. This tag displays at the top of the answer when the information for that answer comes from an AI agent instead of documents from the **Corpus** page. You can click this tag to open the **Response Audit** pane, which lists the number and type of tokens used for that response.

- Thumbs up and thumbs down icons (). At the bottom of the answer, these buttons let you vote for or against on the quality of the response from Skylar Advisor. You can use the prompt box that appears after you click thumbs up or down to explain your answer.

> **TIP:** Try to provide as much feedback as possible, both positive or negative, as that feedback will continue making Skylar Advisor smarter.

- **[Notes]** button. Opens the **Scratch Pad** pane, where you can add notes in plain text or Markdown format. For more information, see *Using the Scratch Pad Feature*.

- **[KB Gen]** button. Tells Skylar Advisor to create a draft Knowledge Base article based on that conversation. For more information, see *Creating a Knowledge Base Article Based on an Activity*.

- Add bookmark icon (). Adds the conversation to the **Bookmarks** section on the **[Recent Activity]** tab of the **Dashboards** page.

- The "More options" menu ():
    - *Copy Message*. Copies the question you asked Skylar Advisor into the clipboard.
    - *Thread Audit*. Displays the **Thread Audit** pane with details about the conversation, including the claimed status, the duration of the conversation, and the full token spend for the conversation thread for cost and performance review. The bottom half of the **Thread Audit** pane contains links to any part of the conversation.
    - *Refresh Messages*. Updates the content on the current page.
    - *Select messages to archive*. Lets you choose which messages you want to archive.

# Chapter 5

# Using the Agentic Control Plane Page

## Overview

The **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane) lets you monitor any AI agents and tools that you are using with Skylar Advisor.

On this page, you can view a list of agents, jobs, tools, and health, along with latency information, audit trails, agent limits and bounds, and token spend over time. You can also design and manage multi-step agent flows.

To use the features on the **Agentic Control Plane** page, enable the **Agentic** option on the **Preview Features** page ( > Account Settings > Preview Features > Integrations > Web Search).

This chapter covers the following topics:

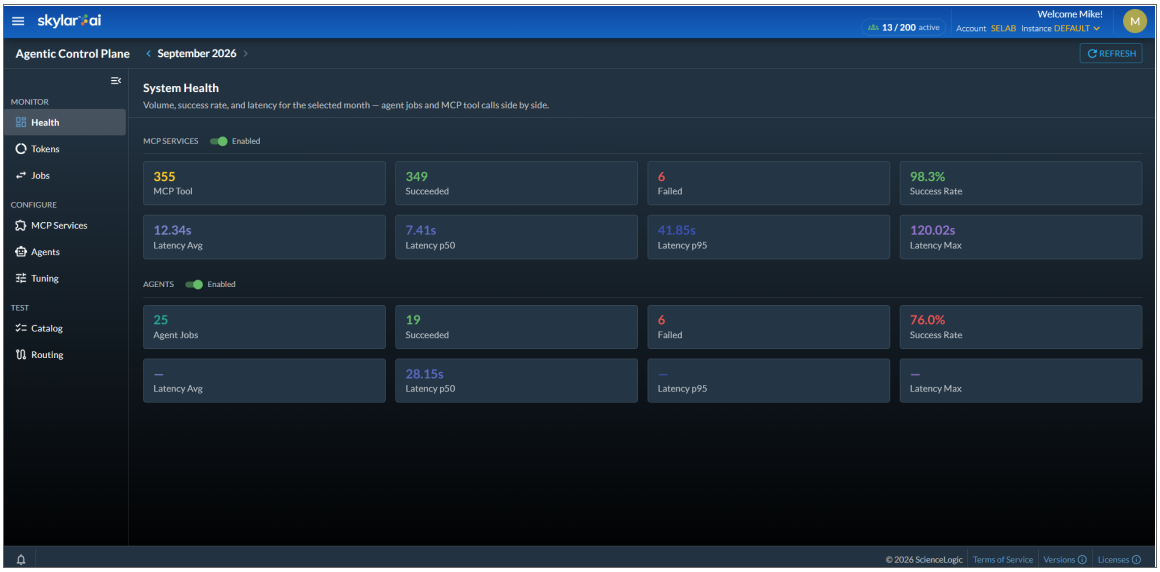
*The Agentic Control Plane Page* ..... *43*

*Additional Agentic Features* ..... *44*

*MCP Services* ..... *45*

# The Agentic Control Plane Page

The **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane) lets you monitor any AI agents that you are using with Skylar Advisor. On this page, you can view a list of agents, jobs, tools, and health, along with latency information, audit trails, agent limits and bounds, and token spend over time.




When you enable an agentic tool, Skylar Advisor routes the relevant questions to that tool instead of searching the documents on the **Corpus** page. Each agentic tool comes pre-loaded with a list of sample questions.

The **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane) contains the following tabs:

- **[Health]**. Displays a system-wide snapshot of the AI agent activity in Skylar Advisor for the selected month: volume, success rate, and latency, compared across agent jobs and direct tool calls. For example, in the **Agents** section, if the **Success Rate** is below 90%, an agent might be failing.
- **[Tokens]**. Displays the token consumption for the selected month: totals, per-call averages, range and efficiency, and spend by hour of day.
- **[Jobs]**. Every action that uses an AI agent gets a Job ID, which is listed on this page. You can search by job ID, description, or agent to inspect status, duration, and token usage. Select a job to view more information about that job in an **Agent Audit** pane; on this pane you can click **[View Anonymize Request]** to hide personal data, questions, and device information while still showing resource usage and the complete audit trail. You can also click the **[Analyze]** button for traces, logs, and SQL queries to troubleshoot agent behavior.
- **[MCP Services]**. Lets you register and preview external MCP servers that contain a variety of AI tools you can use to analyze your data. When you enable an MCP server, it is added as a resource on the **Tools** menu of the **[Ask Skylar]** prompt box and the **Quick Prompts** toolbar.

- **[Agents]**. Lists the different AI agents that Skylar Advisor runs to answer a question. You can enable and disable existing agents, and in future releases you can register new agents and add them to the list of agents. Select an agent in the list to view more information for that agent, including a description of what the agent does, and its privileges, token limits, and time limits.

> **NOTE:** If you have ***Topic Specialists*** enabled, you can enable, edit, and create specialized AI agents that begins work on an issue as soon as it arrives. For more information, see *Agentic: Topic Specialists*.

- **[Tuning]**. Choose how a question reaches the Skylar Advisor agents. *Eager* sends more questions to the agents for higher spend, while *Cautious* sends more to the standard answer for lower spend. *Balanced* is the default, and it is recommended, as it sends data and chart requests, plus everyday device and metric lookups, to the Skylar Advisor agents.
- **[Agent Tools]**. Lists the different types of agentic tools you can use with Skylar Advisor. The agent tools are grouped into **Device Telemetry** (for a single device) and **Fleet Ranking** (for all devices, VMs, or interfaces). Select a tool to view more information about that tool, including sample questions for Ask Skylar.
- **[Catalog]**. Displays an inventory of the tools and agents enabled in this Skylar Advisor instances and a list of collision reports.
- **[Routing]**. Shows the path that Skylar Advisor took to generate its answers to your question, displaying the tools used and why it used those tools.

## Additional Agentic Features

Selecting the **Agentic** option on the **Preview Features** page also enables the following features:

- Additional device performance data in an Advisory, Investigation, or Prediction, or in an Ask Skylar answer. In these activities, the name of a device will be highlighted in blue text. When you click the device name, a pop-up menu displays options for *Memory*, *Disk*, and *CPU*. Click one of the options for more information, which appears as an answer in an Ask Skylar prompt box.

  You can click the **[View Graph]** button to see a visualization of the data flow that includes events, devices, device groups, and business services, where relevant.

- **AI Agent Quick Prompts**. In addition to the standard green Quick Prompts, you can also click one of the yellow Quick Prompts to run an AI Agent tool related to the Ask Skylar conversation. Based on the original question, Skylar Advisor displays additional prompts in green boxes that you can click for additional information.

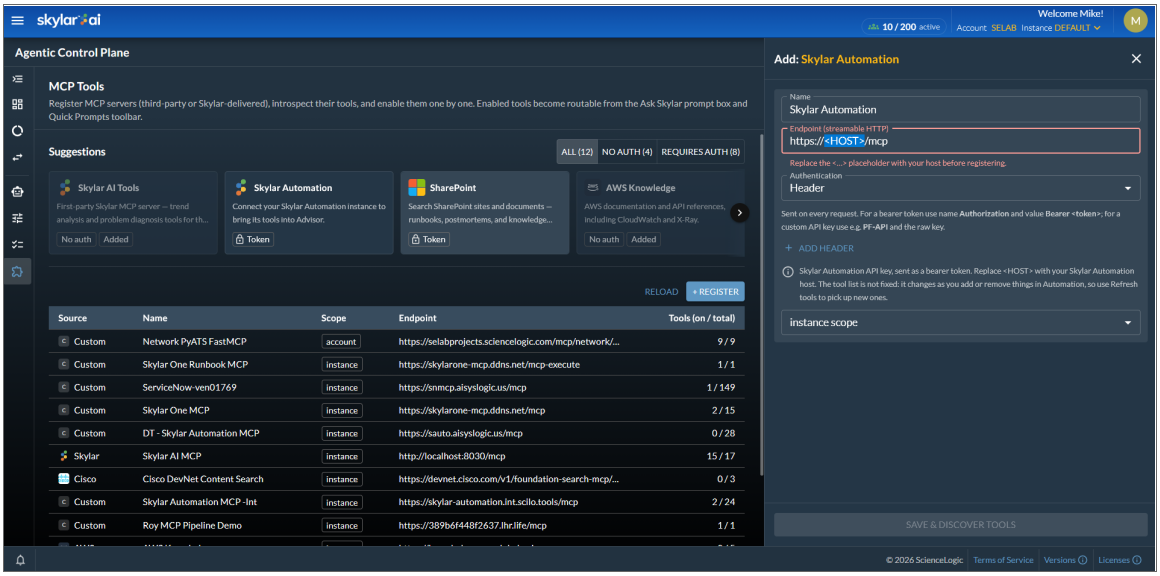
  Questions that were answered using agents have a **[Verifiable (Via Agents)]** tag. Clicking this tag opens a **Response Audit** pane that lists the token usage, the length of the Ask Skylar conversation, and an **Audit Logs** section with links to the key sections of the conversation. Conversations also have an **[Audit]** button, which you can click to view more details on the **Thread Audit** pane.

**NOTE:** To control and enforce the behavior of the Skylar Advisor AI agents, ScienceLogic has placed two gates in front of every agentic capability. The first gate is based on the release, where ScienceLogic only releases capabilities that have been vetted as production-ready, and anything unpublished is invisible to the agents. The second gate is based on customer opt-in, where a Skylar Advisor administrator must deliberately enable an agent before users can access it. As a result, a Skylar Advisor AI agent does not run unless it has cleared both the ScienceLogic release gate and the explicit opt-in by the site administrator.

# MCP Services

The **[MCP Services]** tab on the **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane > MCP Services) provides access to a set of ScienceLogic and external MCP servers that contain a variety of AI tools you can use to analyze your data.

You can register (add) a new MCP server and enable (configure) the AI tools in an MCP server so that you can preview those new features.

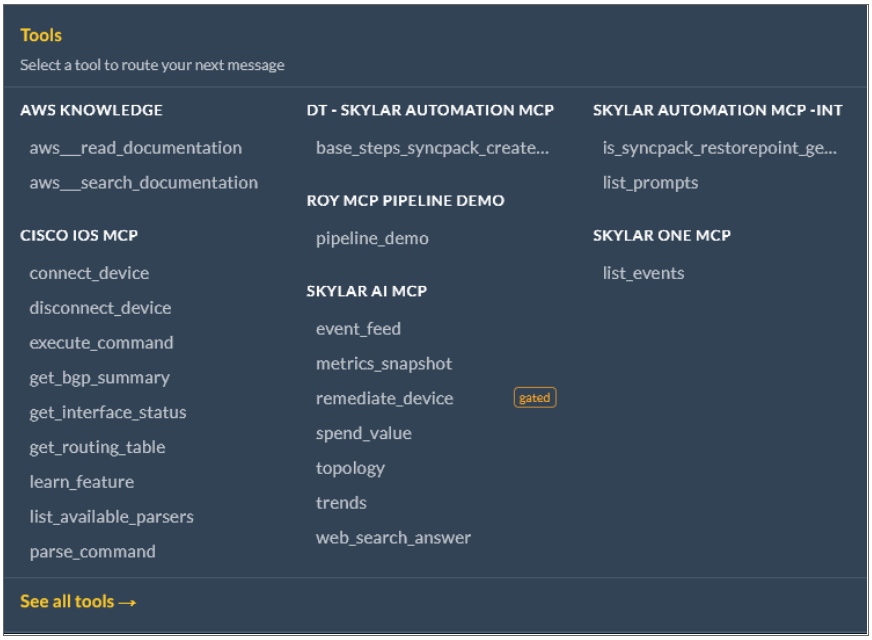



You can add the "Skylar Automation" MCP service from the list of suggested integrations. This tool lets you connect a Skylar Automation instance and bring the Skylar Automation MCP tools directly into the Skylar Advisor interface. For more information, see *Configuring the Skylar Automation MCP Server*.

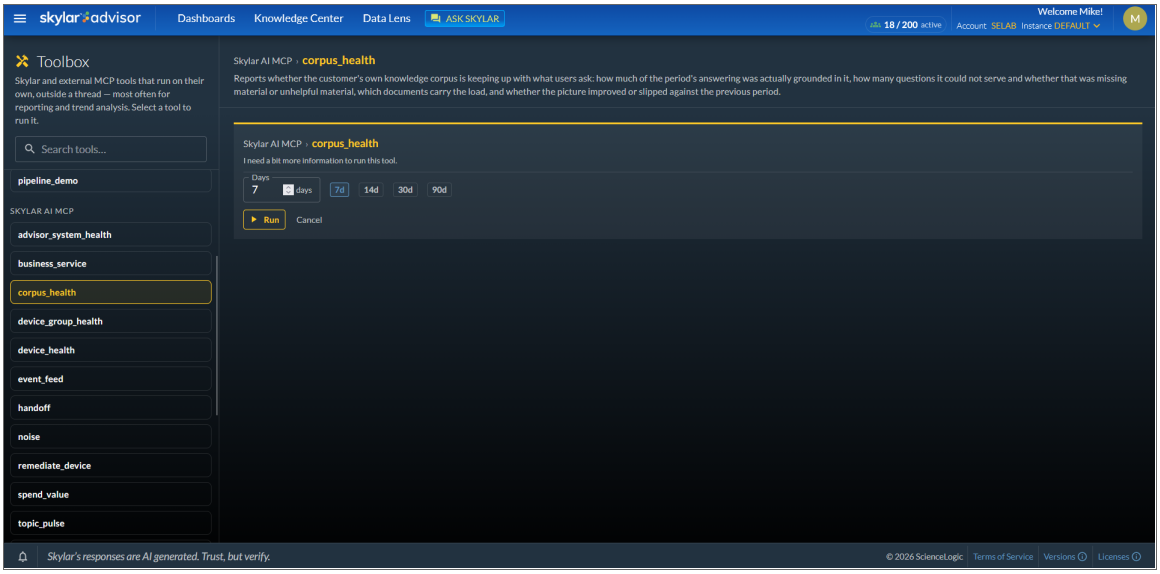
**NOTE:** If a tool on this page has a **[gated]** tag, then that tool must be part of an approval workflow before you can use it.

The following pages and features are available when you have MCP services enabled.

- The **Tools** menu of the **[Ask Skylar]** prompt box and the **Quick Prompts** toolbar contain AI tools from the MCP servers that you have enabled. These options expand on what Skylar Advisor can query and act on beyond the uploaded documents on the **Corpus** page.



- The **Toolbox** page ( > Advisor > Toolbox). This page lets you view and test all tools from the MCP servers that have been enabled in Skylar Advisor. Select a tool from the list on the left to view more information about the tool, including a short description and the time frame for the data used by that tool.




Click the **[Run]** button to run the tool. If you find an issue or other data that you would like to explore further, click the **[Discuss with Skylar]** button to send that information to an Ask Skylar conversation.

**NOTE:** The "Skylar AI MCP" server is a ScienceLogic MCP server that includes a number of diagnosis and analytics tools. The other MCP server tools listed on the **MCP Tools** page, such as Microsoft Learn, AWS Knowledge, and Cisco DevNet Content Search, are public servers that have been curated by ScienceLogic.

Chapter

# 6

# Managing Activities on the Dashboards Page

## Overview

The **Dashboards** page is the default page that displays when you open Skylar Advisor. This page contains a list of all your interactions and notices related to Skylar Advisor,

These activities include "Advisories", "Investigations", and "Predictions", which will be explained in this chapter.

- **Advisories.** A detailed report where the operator did nothing but observe: Skylar AI found something that did not look right and raised a report you can drill into, read, act on, and discuss in a conversation with Ask Skylar.
- **Investigations**. You and Skylar AI working a known problem in coordination: you found the device, and now you ask questions using Ask Skylar to form a picture, a root cause, and a remediation step.
- **Predictions.** A variation of Advisories for problems that have not yet materialized: spend the resource now (extend that disk) and prevent save the outage later.

This page also includes links to your conversations with Skylar Advisor and any artifacts that you or Skylar Advisor generated.

This chapter covers the following topics:


*What is an Advisory?* ..... 50

*What is an Investigation?* ..... 51

*What is a Prediction?* ..... 51

*The Dashboards Page* ..... 52

*Activities and Options on the Recent Activity Tab* ..... 54

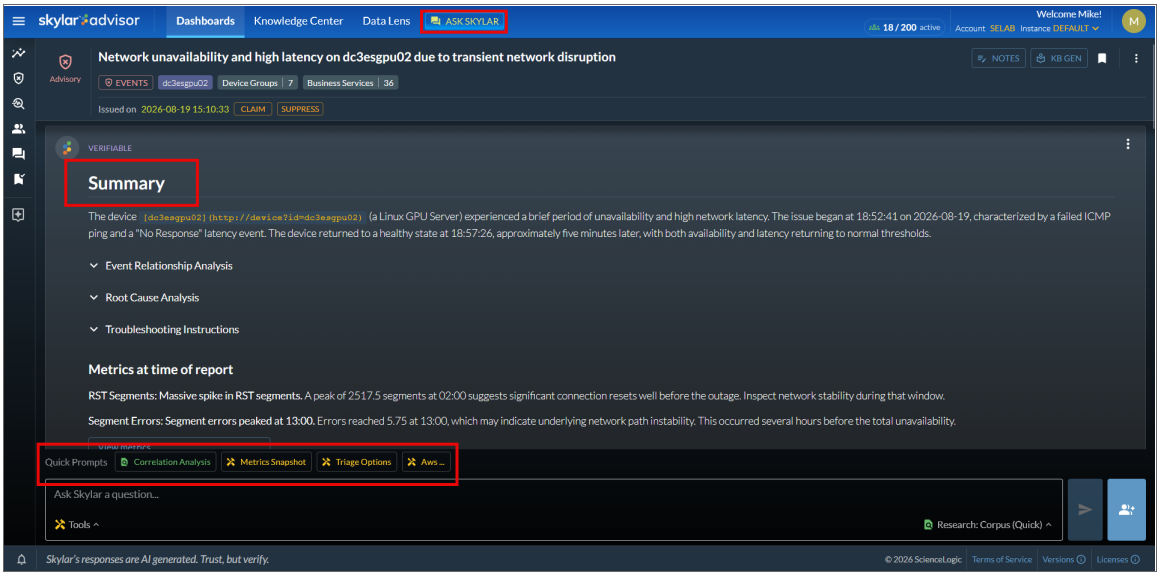
*Viewing User Activity on the Customer Insights Page* ..... *62*

# What is an Advisory?

***Advisories*** are events that Skylar Advisor considers to be the most important issues for you. They help uncover and provide resolution guidance for critical problems. Advisories are automatically generated in Skylar Advisor by looking for meaningful patterns across events instead of reacting to every alert in isolation.

Advisories are designed to help you focus your attention so you can take informed action instead of chasing event noise. An Advisory is a synthesized insight, not a raw alert. An Advisory represents a pattern of related issues (an "event group") that is impacting multiple devices or components and is persisting long enough to indicate real operational impact. Correlation is based on event severity, event message patterns, and metadata, such as entity, status, or service.

When you click an Advisory on the **Dashboards** page, you can view additional details about it, including a summary of the issue, root cause analysis, troubleshooting information, and more. You can also use the **Ask Skylar** feature when you have an Advisory open to do more research on the Advisory. To help you with your questions, Skylar Advisor will generate a set of ***Quick Prompts***, which are additional prompts related to the event.




At the top of an Advisory, you can click the **[Events]** tag to open the **Related Events** panel. This panel lists additional information about all of the devices and events related to this Advisory, including root cause. You can click ***View graph*** on the panel to view the flow of the events leading to the Advisory.

Under the **[Events]** tag, you can click the **[Claim]** button to assign this Advisory to yourself. For more information, see *Working an Advisory From Claim to Close*.

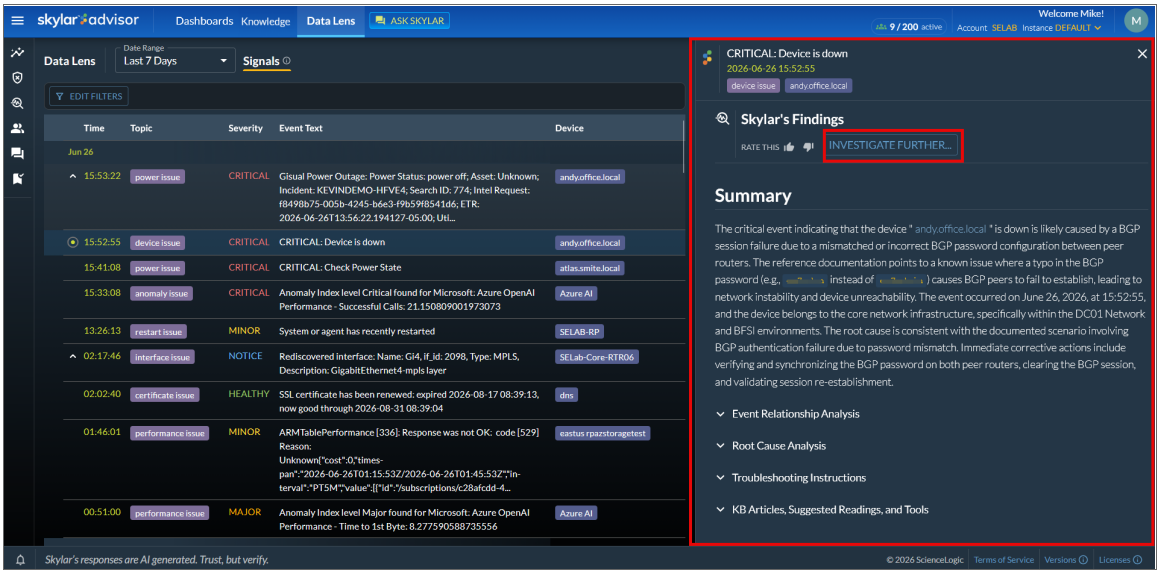
When Skylar Advisor creates an Advisory, it also sends data about that Advisory to Skylar One as an alert. If you have an event policy set up to handle these kinds of alerts, Skylar Advisor will generate an event that includes the Advisory title, summary, and a direct link to the Advisory in Skylar Advisor for streamlined investigation.

When you close an Advisory, Skylar Advisor sends a "Healthy" alert to Skylar One to enable automated event clearing.

# What is an Investigation?

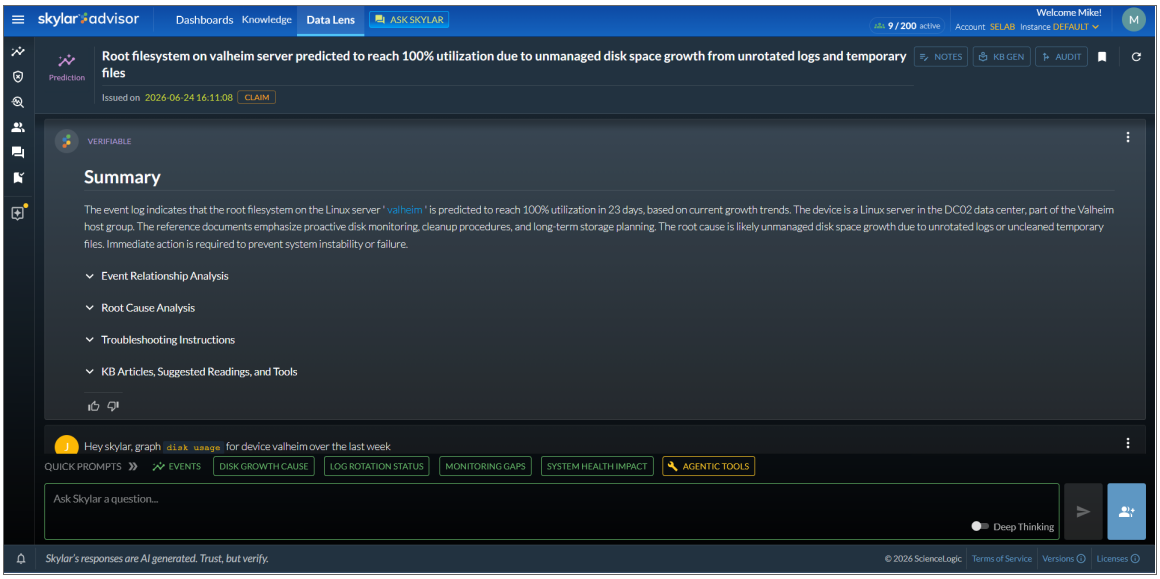
An ***Investigation*** is an on-demand event analysis that you can choose to generate for an event or event group.

You can create an Investigation by clicking the **[Start Investigation]** button or the **[Investigate Further]** button for an Advisory or an event from the **Data Lens** page. Skylar Advisor will then run an analysis on that event. The Investigation detail page lists the devices involved with the issue, along with any impacted device groups and business services.




# What is a Prediction?

In Skylar Advisor, a ***Prediction*** is a potential event that might occur in the near future if no action is taken. Predictions in Skylar Advisor are based on data from Skylar One and Skylar Analytics.

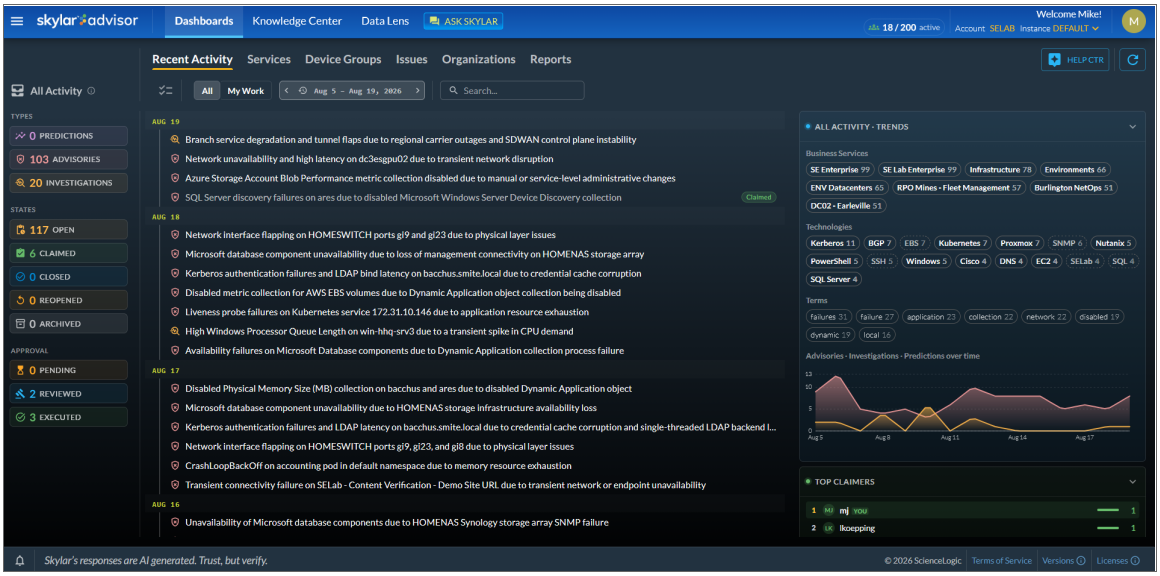



Skylar Advisor will label an Advisory as a "Prediction" if that Advisory was based on Predictive Alerting. Currently, Skylar Advisor generates predictions only from file system data.

# The Dashboards Page

The **Dashboards** page contains a list of all your interactions related to Skylar Advisor, including Predictions, Advisories, Investigations, and conversations with Skylar Advisor.

This page provides access to filtered lists of alerts and events, organized by business service, device group, issue type, or organization. You can also access reports on this page.



**IMPORTANT:** The content and layout of this page is based on the options you chose on the **Persona** page. For more information, see *Defining Your Persona and Role Information*.

## Filters

You can use the filters at the top of the **[Recent Activity]** tab on the **Dashboards** page to focus on the information displayed on this page:

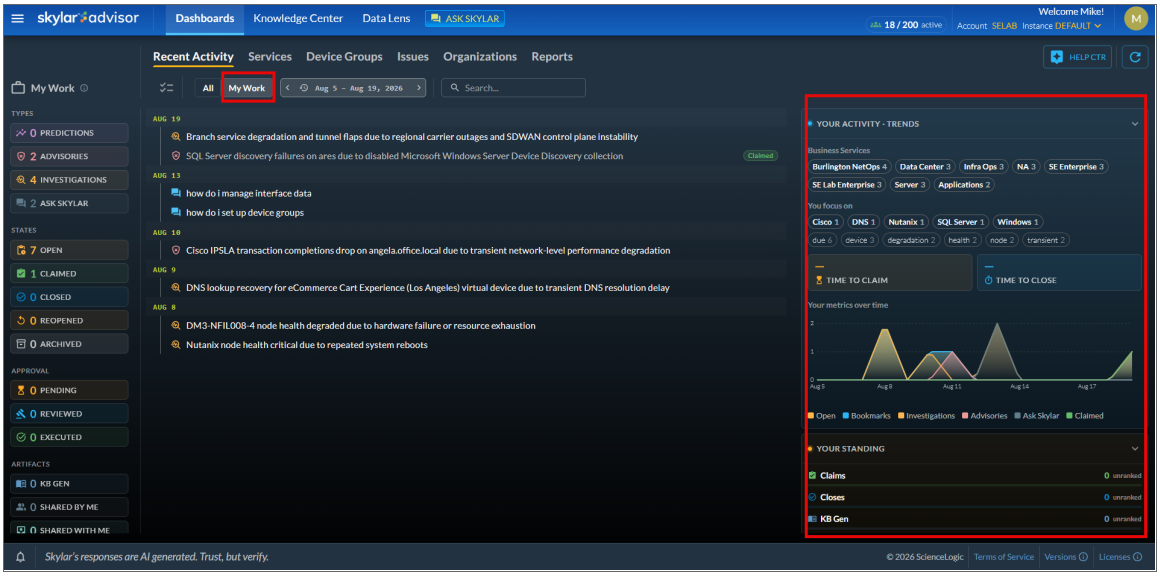
- **All**. Shows all activities that occurred during the current time frame. On the left, a set of buttons display a numbered list of your activities, organized by type.
- **My Work**. Shows all activities that you investigated, claimed, closed, asked, shared, or were shared with you. This filter always displays your Ask Skylar questions and answers. On the left, a set of clickable views display a numbered list of your activities and a numbered list of any artifacts you have generated, such as Knowledge Base (KB) articles, shared conversations, and bookmarked activities.
- **Time frame**. Click the dates on this field to update the time frame of the data display. lets you move week-by-week through your past activities. You can update the time range, from six hours to 14 days, or you can select a custom range.
- **Search**. Type search information into this text box to display activities that match your search text, or click the **[Filters]** button to select additional filters.

**TIP:** All of the tabs on the **Dashboards** page have a **[Help Ctr]** button at top right. Click this button to open a pane with short training videos that walk you through the various features of Skylar Advisor, including a short video of the latest features for this version of Skylar Advisor.

## Trend Analysis

The **[Recent Activity]** tab on the **Dashboards** page contains a trend analysis pane that charts conversation volume, outcomes, and trends over time for a quick summary of your Skylar Advisor activities. The pane contains a list of commonly used terms and most active users. It also charts conversation volume, outcomes, and trends over time, so you can see how Skylar Advisor is being used at a glance.

The **Trends** section contains quick filters that you can use to find similar, related activities. Based on the filter that you selected, the graph and other data below this section update with the filtered data.




# Activities and Options on the Recent Activity Tab

The **[Recent Activity]** tab on the **Dashboards** page displays the following activities and options in the **All Activity/My Work** pane on the left:

- ***Predictions*** (). Events that could potentially occur soon if nothing is done about them, based on data from Skylar One and Skylar Analytics. Currently, Skylar Advisor generates predictions only from file system data.
- ***Advisories*** (). Events that Skylar Advisor considers to be the most important issues for you. They help uncover and provide resolution guidance for critical problems. Advisories are automatically generated by Skylar Advisor, and they directly correlate to alerts and events from Skylar One.
- ***Investigations*** (). Ad-hoc event analyses that you can choose to generate for an event or event group. You create an Investigation by clicking the **[Start Investigation]** button for an Advisory on the **Data Lens** page.
- ***Shared Conversions*** (). A list of interactions with Skylar Advisor using the **Ask Skylar** feature that have been shared with you by other Skylar Advisor users.
- ***Ask Skylar*** () The **Ask Skylar** text box includes the questions you asked Skylar Advisor using the Ask Skylar feature. You can ask Skylar Advisor "plain English" questions just like you would for any other large-language model (LLM) like ChatGPT or Claude. To continue a conversation from the list, click a question.

The pane includes a **States** section that lists activities that are Open, Claimed, or Closed, for example, and an **Approval** section for activities that are being reviewed. The My Work pane also includes an **Artifacts** section for any items that you have created, like KB articles or bookmarks, and any shared items.

## Archiving an Activity

You can archive an activity to remove it from the list on the **[Recent Activity]** tab. Please note that Skylar Advisor never deletes activities or conversations.

To archive or remove one or more activities that you no longer want to see on this tab:

1. On the **[Recent Activity]** tab, click the enter selection mode icon ( ) to enable it (turn it blue).
2. Click the check box for any activity that you want to remove from the list. To select all activities, click the check all icon ( ).
3. Click the archive selected icon ( ). After you confirm your action, the activities are removed from the tab.
4. To view your archived activities, click the **[Filters]** button and select *Archived*. On this page you can also "unarchive" an activity that was archived.

## Bookmarking an Activity

You might want to bookmark an activity that has important data so you can find them quickly. This can be especially helpful for activities that are more than three months in the past, because you would need to use filters to find activities that are over three months old.

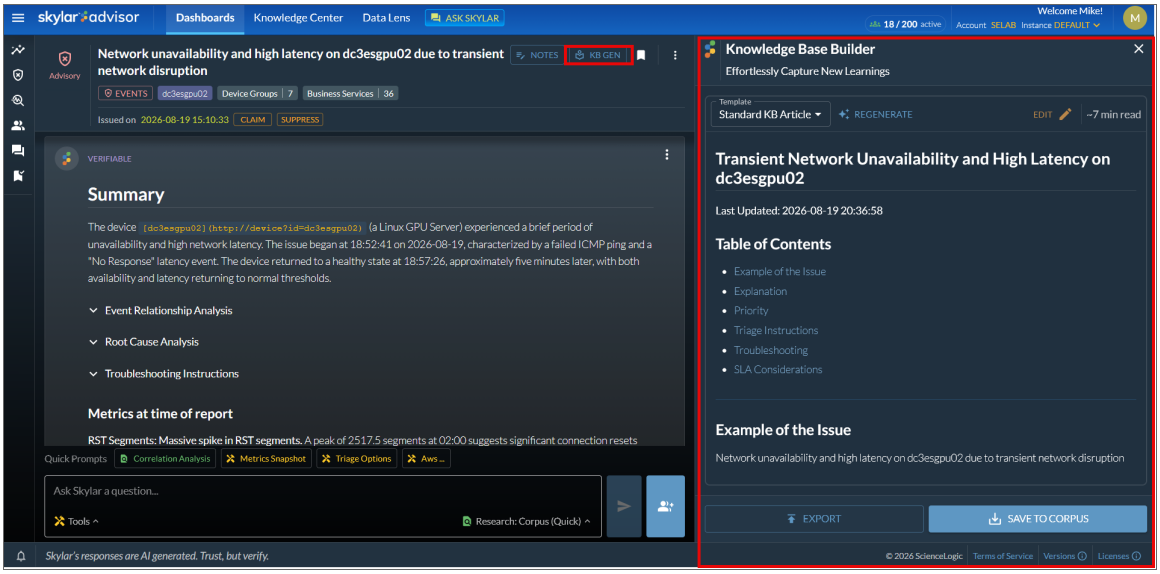
To bookmark one or more activities:

1. On the **[Recent Activity]** tab, click the enter selection mode icon ( ) to enable it (turn it blue).
2. Click the check box next to any activity that you want to bookmark. To select all activities, click the check all icon ( ).
3. Click the bookmark selected icon ( ). The activities are removed from the tab.
4. To view your archived activities, click the **[Filters]** button and select *Archived*.

**TIP:** On an activity detail page, you can also click the bookmark icon ( ) at top right to add or remove a bookmark.

# Creating a Knowledge Base Article Based on an Activity

On an activity detail page for an Advisory, Investigation, Prediction, or a Conversation, you can click the **[KB Gen]** button to have Skylar Advisor create a draft Knowledge Base (KB) article based on that activity.



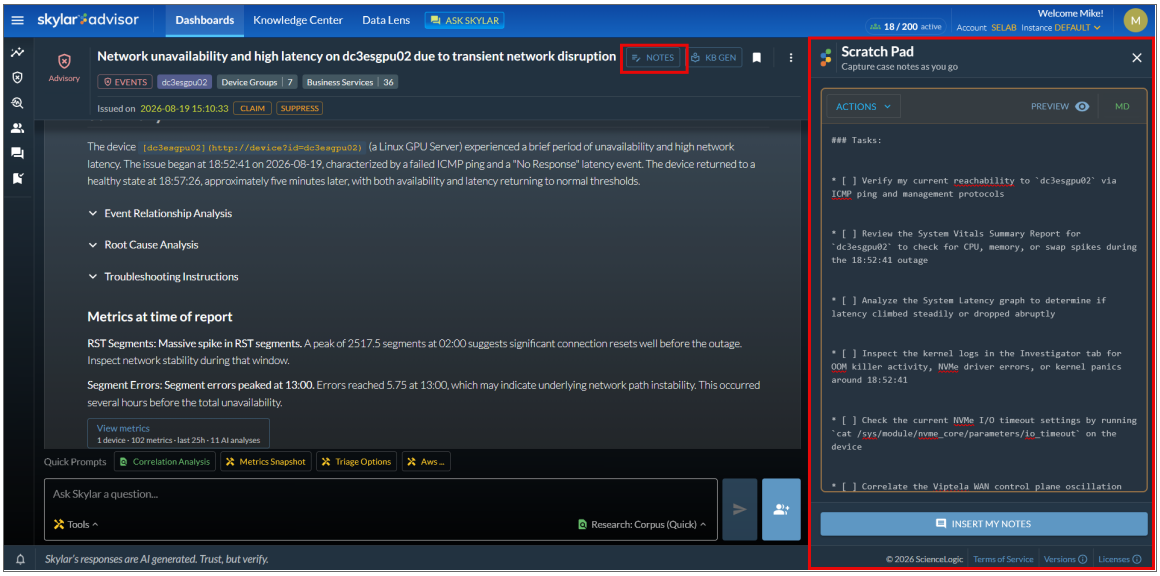

After Skylar Advisor creates the KB article, you can:

- Click **[Regenerate]** to create the article again. You might want to regenerate an article if you asked Skylar Advisor additional questions or added a comment.
- Click **[Edit]** to edit the article in Markdown format. In edit mode, you can click **[Preview]** to see how your text looks with formatting, and you can click **[MD]** for a list of basic Markdown commands.
- Click **[Export]** to download the article as a PDF.
- Click **[Save to Corpus]** to save the article to the **Corpus** page.

All KB articles created by Skylar Advisor include a "[Generated by Skylar AI]" tag at the end of the article.

## Using the Notes Feature

On an activity detail page for an Advisory, Investigation, Prediction, or Conversation, you can click the **[Notes]** button to open the **Scratch Pad** pane, where you can add notes in plain text or Markdown format.




On the ***Actions*** menu, you can:

- Click *Task List* to generate a checklist of actions based on the selected activity.
- Click *Incident Resolution* to generate an Incident Resolution Summary document, which contains Root Cause Analysis, Impact, and other notes.

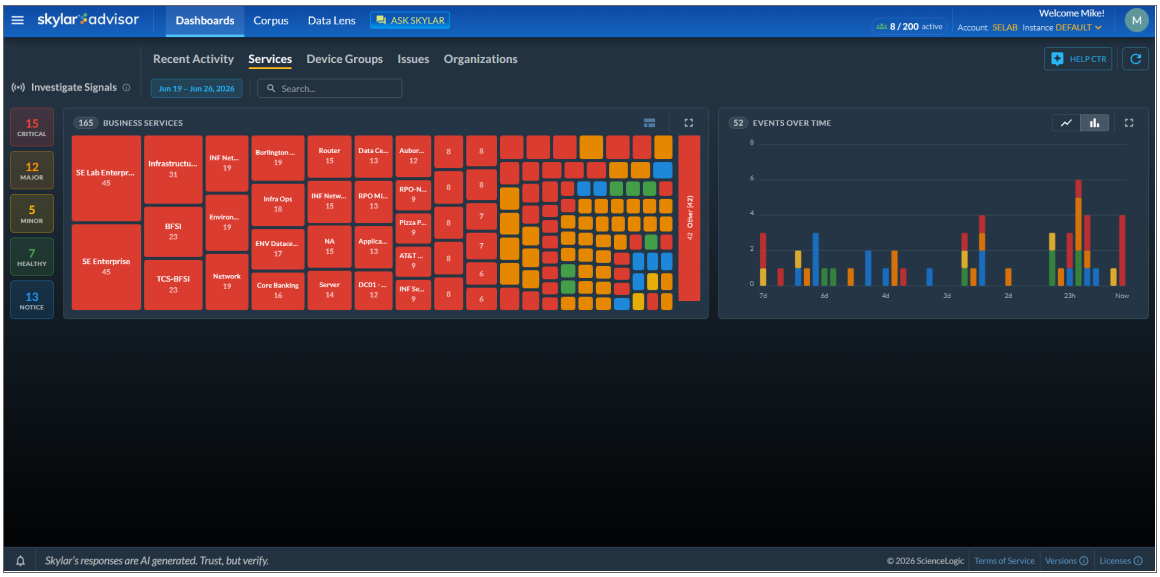
Click **[Preview]** to see how your text looks with formatting. Click **[MD]** for a list of basic Markdown commands.

When you are done, you can click **[Insert My Notes]** to add the contents of the Scratch Pad to the selected activity. Your notes will appear as a comment in the conversation.

## The Services, Device Groups, Issues, and Organizations Tabs

The remaining tabs the **Dashboards** page display event information in different ways, but the various icons and buttons on those tabs work in the same way. There are four tabs:

- The **[Services]** tab displays a list of events for Skylar One business services.
- The **[Device Groups]** tab of displays a list of events organized by Skylar One device group.
- The **[Issues]** tab displays a list of events organized by issue type.
- The **[Organizations]** tab of the **Dashboards** page displays a list of events organized by Skylar One organizations.




**TIP:** If you do not have any events for a specific tab, that tab will not display on the **Dashboards** page.

## Viewing Basic Event Details on a Tab

You can use the icons on the left of the **[Services]**, **[Device Groups]**, **[Issues]**, and **[Organizations]** tabs to filter the list of events on these tabs to only show one type of severity, such as all Critical events for your organizations or customers. You can click multiple buttons as needed. Click the button or buttons again to view all the events again.

These tabs display a "heatmap" view ( ) that shows the events organized by severity in a color-coded graph, with Critical events displaying in red, Major events in orange, and so on. Click the "show treemap" view icon ( ) to view the events without the heatmap colors.

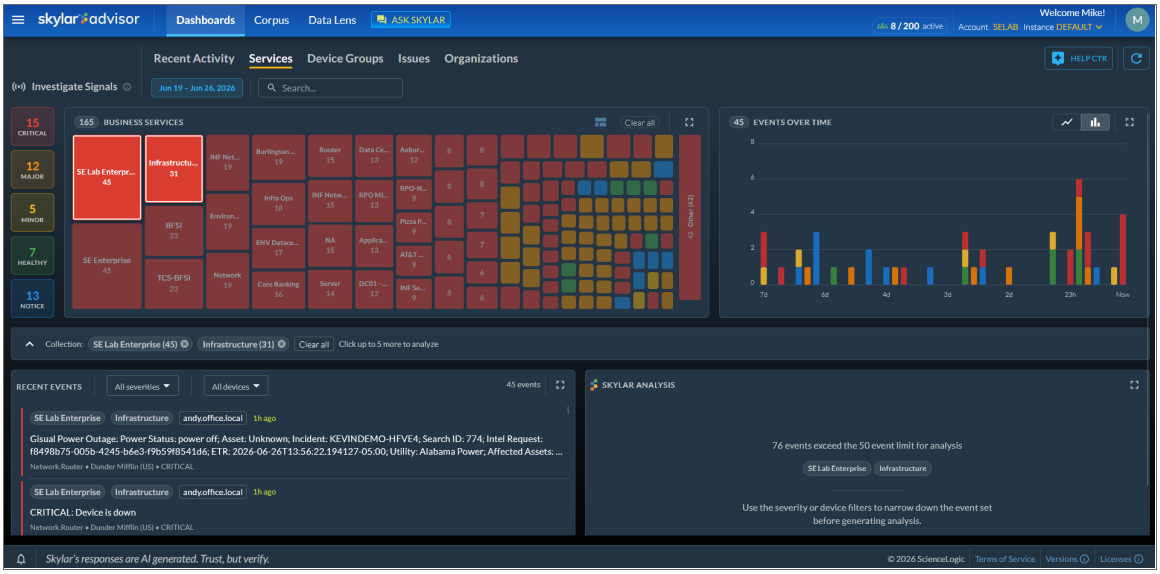
When you mouse over a set of events organized by device group, issue, or organization, a pop-up displays a summary of the events for that set of events.

The **Events Over Time** pane displays events on a timeline based on either severity ( ) or by volume ( ).

**TIP:** For any of the panes on these tabs, you can click the "expand to fullscreen" icon () to make that pane full-screen. Click the "exit fullscreen" icon () to return the pane to its default size again.

## Viewing More Event Details on a Tab

When you select one or more sets of services, device groups, issues, or organizations (based on the tab you selected), the **Recent Events** pane and the **Skylar Analysis** pane appear at the bottom of the page. Also, the **Events Over Time** pane updates to display date only for the selected items.




On the **Recent Events** pane, you can:

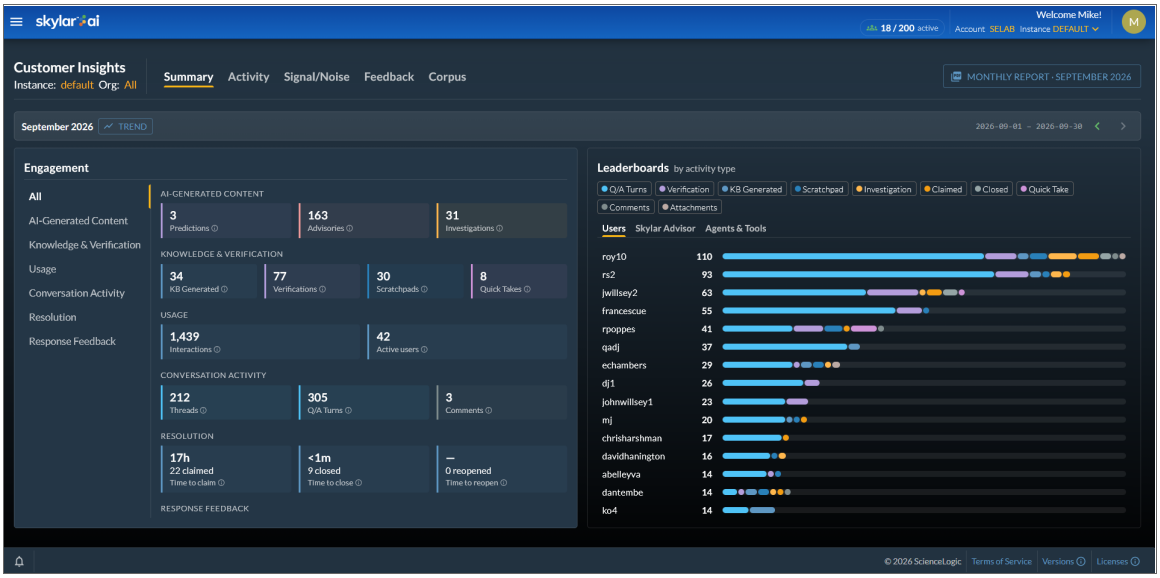
- Review the events from the item you selected above (in this example, all the "Server" events and the "Test-kgc" events).
- Click a device IP or device name to show only events for that device.

On the **Skylar Analysis** pane, you can:

- Click **[Generate Analysis]** to have Skylar Advisor create a new analysis of the items you selected above.
- When the analysis is done, if you need more information, you can click **[Investigate This Problem with Skylar]** to do a deeper investigation of the issue and create a new Investigation. The new Investigation is added to the **[Recent Activity]** tab of the **Dashboards** page.

# Viewing User Activity on the Customer Insights Page

In the Skylar Advisor user interface, the **Customer Insights** page ( > Advisor > Customer Insights) is an interactive, read-only page that displays data about the activities of all users on that system and instance.

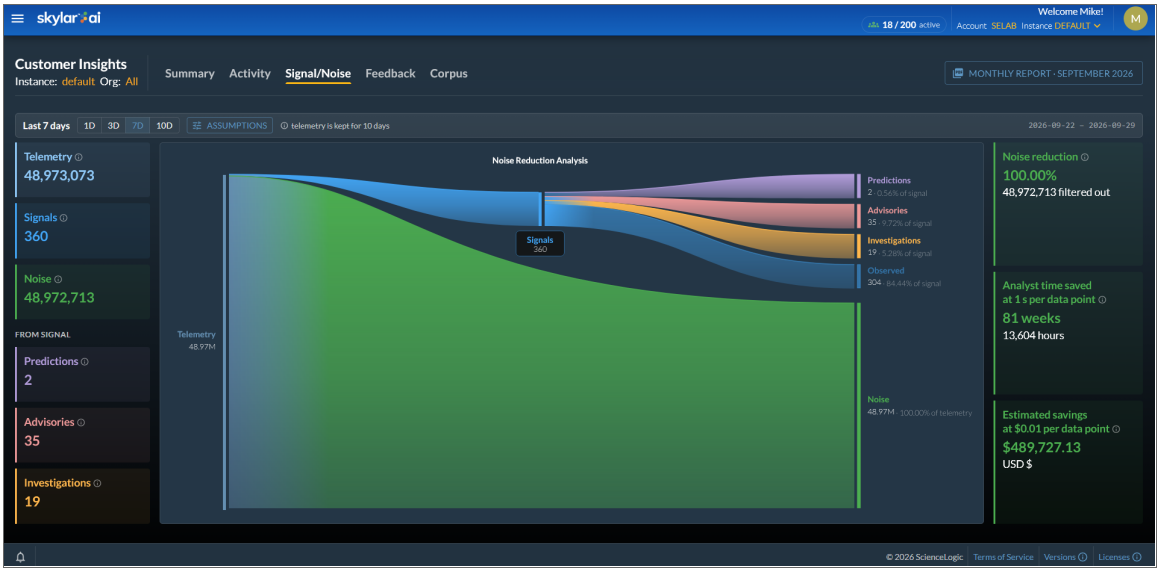



**NOTE:** Currently, this page displays up to 50 users at a time.

The following tabs on the **Team Analytics** page let you view how all users are interacting with Skylar Advisor:

- **[Summary]**. In the **Engagement** section, displays the number of conversations, question-and-answer turns, and comments by users, as well as AI-generated content like Advisories, Investigations, and Quick Takes, and the number of thumbs-up and thumbs-down votes. The right panel lists the most active users, along with the number of conversations each user had.
- **[Activity]**. Charts all of the various user activities, such as Predictions, Advisories, Investigations, Question-Answer Turns, Agentic Turns, and more. The graph displays the value delivered by each activity type. You can click a metric in the table on the left to filter the chart.

- **[Signal/Noise]**. Displays a variety of data about the "signal" (useful data) and the "noise" (unhelpful data) generated in this account by Skylar Advisor. Signals can become Advisories, Investigations, or Predictions.




The middle of the **[Signal/Noise]** tab includes a visual Sankey chart that displays the useful signal versus the noise generated. You can also track the estimated savings as a result of the noise reduction in Skylar Advisor.

- **[Feedback]**. Displays the specific feedback users added when they clicked the thumbs-up and thumbs-down votes.
- **[Corpus]**. Displays the number of active, queued, and failed documents on the **Corpus** page in various time frames, with line graphs to show trends.

# Chapter

# 7

# Working with an Advisory

## Overview

Use the steps in this chapter as a standard path from an unclaimed Advisory to a closed one, including the Skylar Advisor features that help at each step and the checks that keep the record defensible. Keep this topic open beside your first few advisories. If you supervise a team of operators, see the related chapter, *Supervising with Skylar Advisor*.

This chapter covers the following topics:


*Before You Work an Advisory* ..... 64

*The Advisory Workflow at a Glance* ..... 65

*Orienting Yourself in the Events* ..... 65

*Working an Advisory From Claim to Close* ..... 67

*Choosing the Right Tool for the Question* ..... 70

*Escalating an Advisory* ..... 70

*Common Mistakes When Working an Advisory* ..... 71

*Checking Your Own Performance* ..... 71


## Before You Work an Advisory

An Advisory or Prediction is an account of something that Skylar Advisor detected, written at the moment it fired. Keep in mind the following:

- An Advisory describes the past. It is a snapshot of the moment it was raised, not a live view. The older it is, the less you should assume it still reflects current reality.
- The thread is the record. Everything you do in the Advisory, including claiming, asking questions, running a tool, adding a comment, and closing, is written into the thread with your name and a timestamp. You do not keep a separate work log; the thread is the work log, and that is what the next person reads.

## Reading the Answer Badge

Skylar Advisor labels where each answer came from. Check the badge before you act on anything.

| Badge | What It Means for You |
|---|---|
| **[Verifiable]** | Built from facts found in your Corpus. Click the badge to see the source documents. |
| **[Verified (via Agents)]** | The Skylar Advisor AI agents queried your monitoring data directly to produce it. Trustworthy for numbers. |
| **[Web Cited]** | Read live from documentation sites your administrator approved. Click it to see exactly which pages were read, and when. |
| **[Tool Call]** | Produced by a tool rather than written prose. Treat the output as the tool's result. |
| **[Unverifiable]** | Do not act on this alone. It came from general model knowledge, not from your data or documents. Confirm it before you change anything. |

IMPORTANT: If an answer is marked **[Unverifiable]** and you are about to act on it, stop and check it first, either with a tool, with the metrics, or by asking a more senior team member. Plausible and correct are not the same thing, and the badge tells you which one you have.

# The Advisory Workflow at a Glance

Working an Advisory follows three phases:

1. Orienting yourself in the report and its events.
2. Doing the investigative work.
3. Closing the loop so the record is complete. For more information, see *Working an Advisory Step by Step*.

# Orienting Yourself in the Events

An Advisory is exactly what the name says: advice. Skylar Advisor noticed something and is telling you about it. The events are what it noticed. Open them early, before you form a theory of your own. The events show you the scale, the duration, and the sequence in about thirty seconds, and everything you do afterward is better aimed because of it.

In the thread header of an Advisory, Prediction, or Investigation, an **[Events]** button opens the Events panel beside the conversation instead of over it, so you can read the report and the evidence together. This is the same **[Events]** tag described in *What is an Advisory?* What you get is the exact feed Skylar Advisor analyzed: not a fresh query, and not everything that happened, but the rows the Advisory was drawn from.

## The Four Summary Tiles

| Tile | What It Tells You, and What It Does |
|---|---|
| Time Span | How long the run of events covers, with the actual date range beneath it. A burst over four minutes and a grind over nine hours are different problems. |
| Devices | How many distinct devices are involved, your first read on blast radius. Click the tile to copy the device names as CSV. |
| Event Ids | How many distinct event types fired. One ID repeating is a different story from nine IDs interleaving. Click the tile to copy them as CSV. |
| Event Flow | Not a statistic, but a control. Click **View graph** to open the event flow diagram described below. |

## Reading the Event List

Below the tiles, you have the option to switch between the whole feed and the part that matters, when Skylar Advisor identified one:

- **All events**. Every event Skylar Advisor analyzed. Rows that belong to the causal chain are numbered and lightly tinted, so you can find them while scrolling a long list.
- **Causal chain**. Only the events the Advisory rests on, in order. On a noisy Advisory, this is the difference between eighteen rows and three.

The first step in the chain carries a **ROOT CAUSE** tag that marks that event is the key event for the Advisory. Skylar Advisor watches live traffic for that same event to raise the Advisory again, so it is often the event to fix if you want the Advisory to stop recurring.

Each row carries a severity band, the device, the event text, and its timestamp. The panel expands to full screen from its header when the side panel gets tight, and the header copies every device ID in the feed to your clipboard in one click, which is useful for pasting into a ticket or another tool.

## Viewing the Event Flow Diagram

Click **View graph** to open the event flow diagram in full screen. The diagram draws the Advisory as three to four columns, left to right:

1. ***Events***. What fired, ranked by severity, with repeat occurrences counted rather than repeated, and the causal chain's key events marked.
2. ***Devices***. Which machines those events landed on, each tagged with its class and category and carrying its own event count.
3. ***Device Groups***. Where relevant, lists the device groups for the impacted devices.

4. ***Business services***. What those devices serve. Where service membership is not populated, the column falls back to device groups instead of sitting empty. Business services appear only when available.

A timeline runs across the diagram, positioning each event in time, so you see order as well as connection.

Connectors are colored by device rather than by severity: on the common case where every event shares one severity, coloring by severity would give you a single-hue picture with nothing to follow. Coloring by device lets you trace one machine's path from its events through to the services it affects.

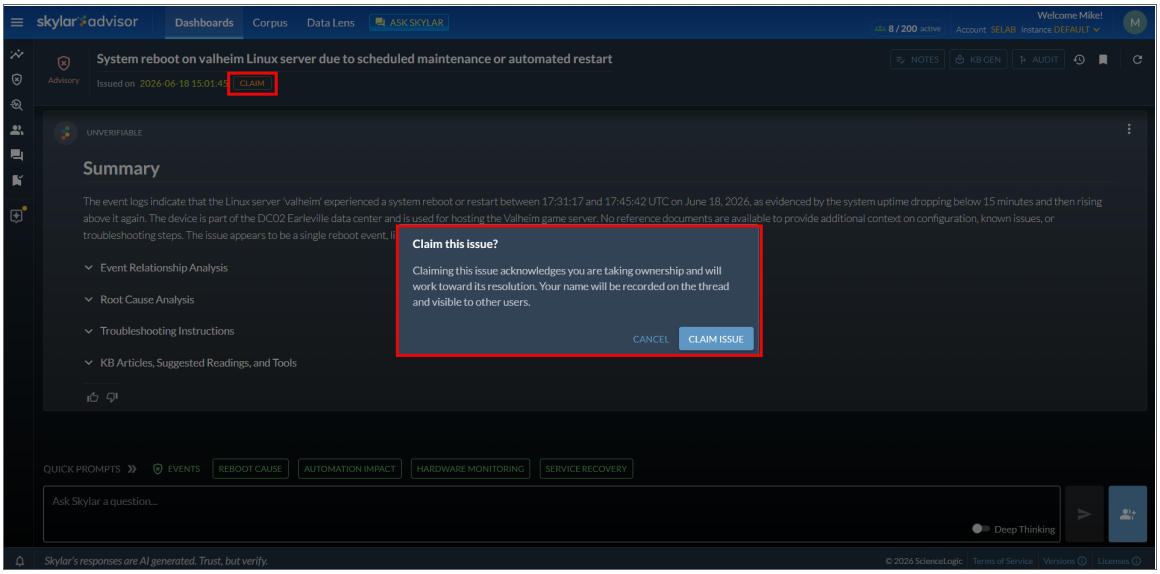
The diagram includes a toggle at the top that contains the same all-events and causal-chain switch described above, so you can strip it back to the spine.

When you read the diagram from left to right, it tells you what happened, to what, and who is affected, in one view. That is the fastest way to know whether you are looking at one unhealthy device or the leading edge of something wider.

## Working an Advisory From Claim to Close

To work an Advisory from claim to close:

1. Open the Advisory ( ) and click **[Claim]** in the thread header. This stamps your name on it so nobody duplicates your work, and it starts the clock on your ownership. If the Advisory is already claimed by someone else, the header states this; check with that person before you start working it.



2. Read the contents of the Advisory end-to-end before doing anything else. Note the affected devices, the time window, and what Skylar Advisor says it observed. Then open the events behind the report, as described in *Orienting Yourself in the Events*. The report is what Skylar Advisor makes of the events; the events are what it saw.

   If the Advisory includes charts or metric attachments, read them for shape rather than single values: is the line climbing, spiking, flat but high, or already back down? You are checking whether the shape matches the problem described.

   Where a chart contains significant movement or a notable moment, Skylar Advisor attaches an AI analysis. Read it; it names what moved and when, which is usually faster than finding it yourself. You can filter the panel to show only charts with an AI analysis, or show every analysis at once. These charts are interactive, not static images, and the **Metrics Snapshot** pane also provides:

   - ***Full screen***. Displays a common cursor across every chart, so one hover marks the same instant everywhere, which is how you see what moved first.
   - ***Metrics by category***. Groups CPU, memory, disk, and network so they read as groups instead of one chart at a time.
   - ***Brush to zoom***. Drag across a chart to set a time range. Use the arrow keys to move the window, and click **[Reset]** to return to the full range.
   - ***Snapshot walker***. Steps between older and newer snapshots once the thread holds more than one.
   - ***Stitched timeline***. Merges snapshots into one continuous view: a problem starting, peaking, and clearing in a single chart. Pairs with the closing snapshot you take before you close the Advisory.
   - ***Table view and CSV download***. Provides the same data as a table you can export, for copying values into the Scratch Pad or sending them elsewhere.

     For more information, see *Metrics Enhanced Advisories*.

3. Run the ***metrics_snapshot*** tool from the **Tools** menu on the Ask Skylar prompt box. This tool pulls current telemetry for the affected devices. This step is not optional if more than two hours have passed since the Advisory formed, because it is what stops wasted work. You are answering one question: is this still happening right now?

   - ***Still happening***. Continue to the next step.
   - ***Recovered on its own***. Do not close the Advisory yet. Note what you found, note that the problem self-recovered, and record what the metrics show now. A problem that resolved itself often returns.
   - ***Worse than the Advisory described***. Escalate before you continue, because the problem looks more severe than it did when the report was generated. Raise team awareness by adding colleagues to the conversation, or through the Slack or Teams integrations where they are enabled.

4. Click the **[Notes]** button to open the Scratch Pad and write out what you intend to check, in order. Even three lines is worth writing, because it stops you circling and becomes your work note later. Skylar Advisor also provides an action that drafts this task list for you from the live thread. For more information, see *Using the Notes Feature*.

5. Work your tasks using your normal tools plus whatever Skylar Advisor suggests. Write the outcome of each check in the Scratch Pad as you finish it, including the checks you ruled out. "Checked X, not the cause" is one of the most valuable lines you can leave behind.

   As you work, you can use the following features to help your investigation:

   - ***Quick Take***. Select any text in the thread and ask for a quick take on it. This is the fastest way to resolve an unfamiliar term, error code, or acronym without losing your place.
   - ***Quick Prompts***. The suggested follow-up questions under an answer. They are generated from the current conversation, and they are usually the cheapest next step.
   - ***Tools***. See *Choosing the Right Tool for the Question*.
   - ***Web search***. If enabled, ask what vendor documentation says about a setting or error code. Answers cite the pages read and when they were read.
   - ***Colleagues***. Add someone to the thread, or share it to Slack or Teams if those integrations are configured. They arrive with the full history instead of a summary you have to retype. For more information, see *Sharing Information with Slack and Microsoft Teams Integrations*.

6. When you find the root cause, insert your task list and work notes into the thread as a comment. Until you insert them, your Scratch Pad notes exist only on your machine. For more information, see the warning at the end of this section.

7. Use the *Incident Resolution* action in the Scratch Pad to draft a summary of what happened, what caused it, and what fixed it. Read the summary before you insert it, because it is drafted from the thread and can only be as accurate as what you wrote down. For more information, see *Using the Notes Feature*.

8. If metrics were involved, run ***metrics_snapshot*** tool once more so the thread ends with evidence that the system returned to normal. This is what turns "I fixed it" into something the next shift and any review can verify.

9. Generate a KB article from the thread, and read it before saving it. It is drafted from the conversation, so confirm that the root cause and the fix are stated correctly. This article is what makes the next occurrence faster, for you or for someone else. Which type of article you generate depends on your team's best practice. For more information, see *Creating a Knowledge Base Article Based on an Activity*.

10. Click **[Close]** in the thread header. If you close an Advisory too early, click **[Reopen]**; reopening is normal, and it is better than leaving an incorrect record.

    When you close an Advisory, Skylar Advisor sends a "Healthy" alert to Skylar One, which enables automated event clearing in Skylar One.

> **IMPORTANT:** The Scratch Pad auto-saves in your browser, not in the thread. If you close the tab, switch machines, or hand off mid-shift without inserting your notes as a comment, that work is not visible to anyone else. Insert your notes as you reach milestones, not only at the end.

## Choosing the Right Tool for the Question

Tool availability depends on what your administrator has enabled. If a tool described here is missing from your **Toolbox** ( > Advisor > Toolbox) or the Ask Skylar prompt box, it has not been enabled for your instance.

| When You Are Asking | Reach For | It Gives You |
|---|---|---|
| "Is this still happening now?" | ***metrics_ snapshot*** | Current telemetry for the affected devices, as a chart you can attach. |
| "Has this happened before, and how was it fixed?" | ***precedent*** | Prior occurrences and how they were resolved, from history and the Corpus. |
| "What should I try first?" | ***triage*** | Several approaches compared, with a recommended next action. |
| "Is it just this device, or wider?" | ***correlate*** | The blast radius: what else is affected, and the likely common cause. |
| "Is this device generally unhealthy?" | ***device_health*** | The device's recurring issues, whether they are improving or worsening, and suggested fixes. |
| "What does this error code or setting mean?" | Web search, or Quick Take | Vendor documentation read live and cited, or a fast plain-language explanation. |

**NOTE:** Some tools are configured to require sign-off before they run. If a tool requires approval, an approval request appears in the thread instead of a result, and it waits for someone in the approver group. This is expected behavior; continue with other tasks and return to it later. You cannot approve your own request. For more information, see *MCP Tools*.

## Escalating an Advisory

Escalate an Advisory when you have exhausted your task list, when the impact is growing, or when the fix needs a change you are not authorized to make. Escalating early is not a failure. Escalating without any findings wastes the next person's time.

Before you hand off an Advisory, confirm that the thread contains:

- Your task list, with what you checked and what you ruled out. Other users cannot see your Scratch Pad, so insert your notes as a comment to leave a record of your work before you escalate.
- A current metrics snapshot, so the next person sees the state now rather than the state at Advisory time.
- The blast radius, if you established it, and who is affected.
- Anything you changed, so nobody has to guess whether the system was touched.

Add the next person to the thread, or share it to Slack or Teams. Leave the Advisory claimed by you until that person picks it up, so it does not appear unowned.

## Common Mistakes When Working an Advisory

| Mistake | Why It Hurts |
|---|---|
| Working an Advisory without claiming it | Two people work the same problem, and neither knows. |
| Trusting an hours-old Advisory as current | You investigate a problem that already cleared, or miss that it got worse. |
| Acting on an **[Unverifiable]** answer without confirming it | The answer might be plausible and wrong. Confirm it against your data first. |
| Keeping notes only in the Scratch Pad | Nobody else can see them, and they do not survive a handover. |
| Recording only what worked | The next person repeats every check you already ruled out. |
| Closing without a recovery snapshot | Nothing in the record shows that the system actually returned to normal. |
| Saving a KB article without reading it | A wrong article is worse than none, because it will be trusted next time. |

Keep the following four points in mind:

- Claim an Advisory before you touch it.
- Prove that a problem is still happening before you work it.
- Record what you ruled out, not just what you found.
- Document your notes in the thread, not just the Scratch Pad.

## Checking Your Own Performance

Skylar Advisor keeps a profile of your own activity, one calendar month at a time. Reviewing it is worth five minutes at the end of a shift or a week.

To review your activity profile:

1. Click your avatar anywhere in a thread to open your profile page.
2. Use the stepper in the header to move one calendar month at a time. The stepper does not step into the future, so the current month is always the newest view, and it is still filling up.
    - ***Time to Claim***. The average gap between an Advisory being raised and you claiming it.
    - ***Time to Close***. The average time from your claim to your close.
    - ***Ask Skylar***. New questions you put to the Corpus, a measure of upskilling.
    - ***Turns***. Total questions asked while solving problems across threads.
    - ***KB Generated*** and ***Shared***. What you left behind for other people.

3. Read the tiles across the top of the page, which summarize the month: Below the tiles, an activity-over-time chart shows the shape of the month, and a **Standing** panel ranks you against the rest of your organization on claims, closes, KB generation, and shares, showing your position (for example, #N of M) and a bar comparing you to the top contributor.
    - ***Activity summary***. A factual account of what you worked on and your pace, and anything that stood out. This section deliberately contains no advice.
    - ***Areas for improvement***. Two or three focus areas, each naming the metric behind it and suggesting a next step.
4. Click the **[Coach]** button in the profile header to open two sections: Both sections are grounded in your actual numbers for the month, not general advice. If the month has very little activity, coaching is skipped rather than invented.

> **NOTE:** Coaching is currently an Owner-level control, so most operators see the trends and the standing but not the **[Coach]** button. The advice is still available to you; ask your team lead or platform owner to open it with you on your profile.

Keep the following in mind when you read your numbers:

- **Time to Claim is mostly about visibility.** If it is climbing, the usual cause is that advisories are not surfacing where you are looking, not that you are working slowly. Raise this rather than absorbing it.
- **A long Time to Close is not automatically bad.** Hard problems take time. It is worth attention when it is long and you generated no KB article, because that combination usually means what you learned did not get captured.
- **Standing is relative to your team, not an absolute.** A middling rank in a strong team is not a failing. The panel shows you where to aim, not a score.

> **IMPORTANT:** Do not work the metrics. Closing quickly, skipping the recovery snapshot, and skipping the KB article make these numbers look better and make the record worse. The numbers exist to help you spot your own patterns; the thread is what people actually rely on.

# Chapter 8

# Integrating Data from Other Products

## Overview

This chapter covers how to integrate Skylar Advisor with Slack and Microsoft Teams, and to share conversations and import external data.

This chapter covers the following topics:

*Sharing Information with Slack and Microsoft Teams Integrations* ..... 74

*Importing ServiceNow Data into Skylar Advisor* ..... 75

# Sharing Information with Slack and Microsoft Teams Integrations

On the **Integrations** page ( > Account Settings > Account Access > Integrations), you can create integrations with Slack and Microsoft Teams using webhooks. These integrations let users share an Ask Skylar conversation with other users by using the **[Share a Conversation]** icon ( ) and the **Collaborate** pane.

## Creating an Integration with Slack or Teams

Before you can create the integration in the Skylar Advisor user interface, you will first need to locate a webhook in either Slack or Teams. To locate the webhook, see the documentation for Slack or Teams, or ask Skylar AI using the **Ask Skylar** feature.

To create an integration with Slack or Teams:

1. In Skylar Advisor, go to the **Integrations** page ( > Account Settings > Account Access > Integrations) and click **[Add Integration]**. A drop-down appears.
2. Select either *Slack* or *Microsoft Teams*. An integration pane appears to the right.
3. Complete the following fields:
   - ***Integration Name***. Type a name for this integration.
   - ***Enabled***. To use this integration immediately, leave this option at the default (blue). To disable this integration, click the toggle to turn it white.
   - ***Advisor Instances***. Select one or more instances in your Skylar Advisor system where you want to be able to share conversations.
   - ***Webhook URL***. Paste the webhook URL you found for Slack or Teams.
4. Click **Save**.
5. To test the connection, click the **[Test Connection]** button at the bottom of the integration pane.

## Sending a Conversation to Slack or Teams

To send a conversation to a Slack or Teams channel:

1. From the Ask Skylar prompt box for an existing conversation, click the **Collaborate with others** icon ( ) at the bottom right. A **Collaborate** pane appears.
2. Click the **[Slack]** or **[Teams]** button. A **Send to** window appears.
3. Complete the following fields:
   - **Select Channels**. For Teams only, select a channel from the list.
   - **Title**. Type a title for the conversation.
   - **Summary**. Update the existing summary as needed.

- **[Compact Summary]**. Click this button to have Skylar Advisor optimize the summary before sending it to the channel.

4. Click **[Send to Slack]** or **[Send to Teams]**. The conversation is sent to the specified channel, with a link to return back to the Ask Skylar conversation. The conversation is also added to the **Shared** section on the **[Recent Activity]** tab of the **Dashboards** page. You might need to click the refresh icon () to see the most recent questions and answers in this conversation.

# Importing ServiceNow Data into Skylar Advisor

If you use ServiceNow to track incidents, alerts, or cases, or if you want to sync your ServiceNow Knowledge Base (KB) articles, you can use the "Skylar Advisor" SyncPack to import that ServiceNow data into Skylar Advisor using Skylar Automation.

The "Skylar Advisor" SyncPack lets you retrieve the content of KB articles from ServiceNow and send those articles as PDF attachments for Skylar Advisor to ingest. You can also synchronize the following ServiceNow data types:

- Incidents
- Cases
- Alerts (Event Management)
- Change Requests

To enable these features, you need to create an access token in Skylar Settings. For more information, see *Creating Access Tokens for Users*.

When you have the Skylar Advisor access token, you need to add it to a configuration object in Skylar Automation along with the SyncPack. For more information, see the *Skylar Advisor SyncPack* manual.

## Minimum Version Requirement

The "Skylar Advisor" SyncPack requires "ServiceNow Base" SyncPack 4.0.0 or later.

The SyncPack does not require a specific version of the ServiceNow release.

## Supported Data Synchronization

If you use ServiceNow to track incidents, alerts, or cases, or if you want to sync your ServiceNow Knowledge Base (KB) articles, you can use the "Skylar Advisor" SyncPack to import that ServiceNow data into Skylar Advisor using Skylar Automation.

To enable these features, you need to create an access token in Skylar Settings. For more information, see *Creating Access Tokens for Users*. When you have the Skylar Advisor access token, you need to add it to a configuration object in Skylar Automation along with the SyncPack. For more information, see the *Skylar Advisor SyncPack* manual.

The following table lists the ServiceNow data types that you can synchronize with Skylar Advisor:

| Data Type | Supported | Direction | Notes |
|---|---|---|---|
| Knowledge Base | Yes | Unidirectional | Includes PDF attachments under 4 MB |

| Data Type | Supported | Direction | Notes |
|---|---|---|---|
| articles | | | each. |
| Incidents | Yes | Bidirectional | Advisories can create incidents in ServiceNow. |
| Cases | Yes | Unidirectional | |
| Alerts (Event Management) | Yes | Unidirectional | |
| Change Requests | Yes | Unidirectional | |
| Problems | No | | Not yet available. |

## Knowledge Base Filtering

Skylar Advisor supports ServiceNow encoded query filters for knowledge base synchronization.

### Synchronization Approach

- **Initial historical load**. Pulls all existing Knowledge Base articles for onboarding and backfill scenarios.
- **Maintenance sync**. Picks up records created or updated within the configured lookback period. This lightweight sync can run at a frequency you agree on with ScienceLogic.

### Available Controls

Current controls are structural rather than content-based. You can scope which parts of the ServiceNow estate are ingested, but you cannot apply every filter within that scope.

| Filter or Control | Supported | Details |
|---|---|---|
| Organization or business unit | Yes | Ingestion is partitioned by ServiceNow domain and company. Unmapped records default to a catch-all organization, so confirm organization mapping before onboarding. |
| Domains or folders | Yes | Can be used to scope ingestion. |
| Age | Yes | Can be used to limit content by age. |
| Tags | No | Not currently available. |
| File types | PDF only | PDF attachments must be under 4 MB each. |

# Chapter 9

# Managing Events on the Data Lens Page

## Overview

The **Data Lens** page in Skylar Advisor lets you view details about the Skylar One events that have been received by Skylar Advisor. These inbound events are grouped together by relevance and listed on the **[Signals]** tab.

When grouping events for the **[Signals]** tab, Skylar Advisor performs semantic event correlation to group together related events based on how it understood the events. Skylar Advisor also considers related attributes like device groups.

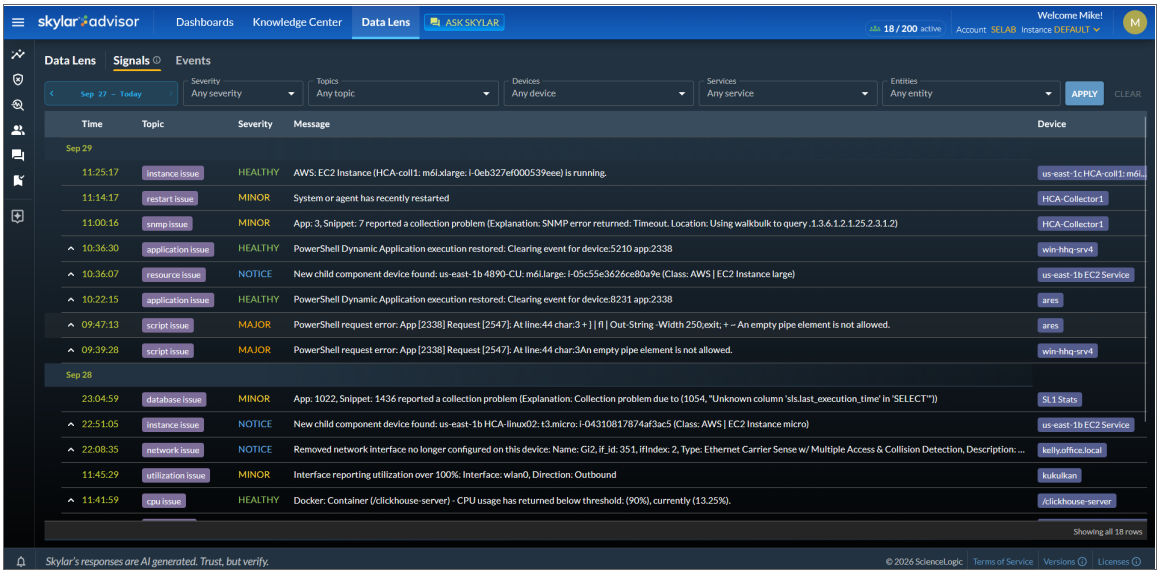
This chapter covers the following topics:

*The Data Lens Page* ..... *78*

*Filtering the Data Lens Page* ..... *80*

# The Data Lens Page

The **Data Lens** page lets you view details about groups of events that have been received by Skylar Advisor. Skylar Advisor performs semantic event correlation to group together related events based on how it understood the events.




Not all events are listed on this page, but just related events organized and grouped by Skylar Advisor.

## The Signals Tab

By default, the **Data Lens** page displays the **[Signals]** tab, which contains a list of *signals*: a group of correlated events from your monitored environment that need you to review or triage. These signals can be rolled up into an Investigation.

**NOTE:** You can also look at the **[Signal/Noise]** tab on the **Customer Insights** page ( > Advisor > Customer Insights > Signal/Noise) to view the ratio of "signal to noise," or the amount of useful data versus the amount of unhelpful data generated in this account by Skylar Advisor. Signals can become Advisories, Investigations, or Predictions. The middle of the **[Signal/Noise]** tab includes a visual Sankey chart that displays the different types of useful signals versus the noise generated. You can also track the estimated savings as a result of the noise reduction in Skylar Advisor.

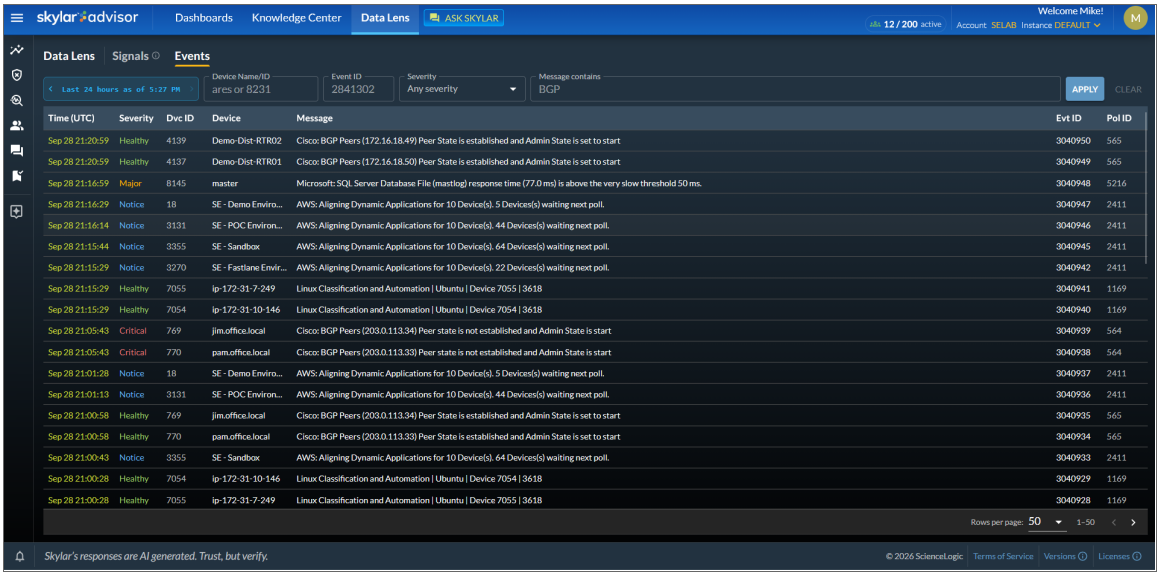
Events are listed on the **Data Lens** page with the most recent events at the top. You can adjust the ***Date Filter*** drop down to change the time frame of the events. For more information, see *Filtering the Data Lens Page* .

In the list of events, the up arrow icon ( ) to the left of a event indicates the event is part of a group. Click the up arrow icon ( ) to view the related events that make up the group.

Click an event on the **Data Lens** page to open a **Summary** pane on the right with more information. If you want to dig deeper into that event, click the **[Start Investigation]** button on the pane to generate a new Investigation activity. You will then be able to interact with and share the Investigation.

## The Events Tab

If you enable the **Data Lens Events** option from the **Preview Features** page ( > Account Settings > Preview Features), the **Data Lens** page will include an [Events] tab that displays a table of all events.

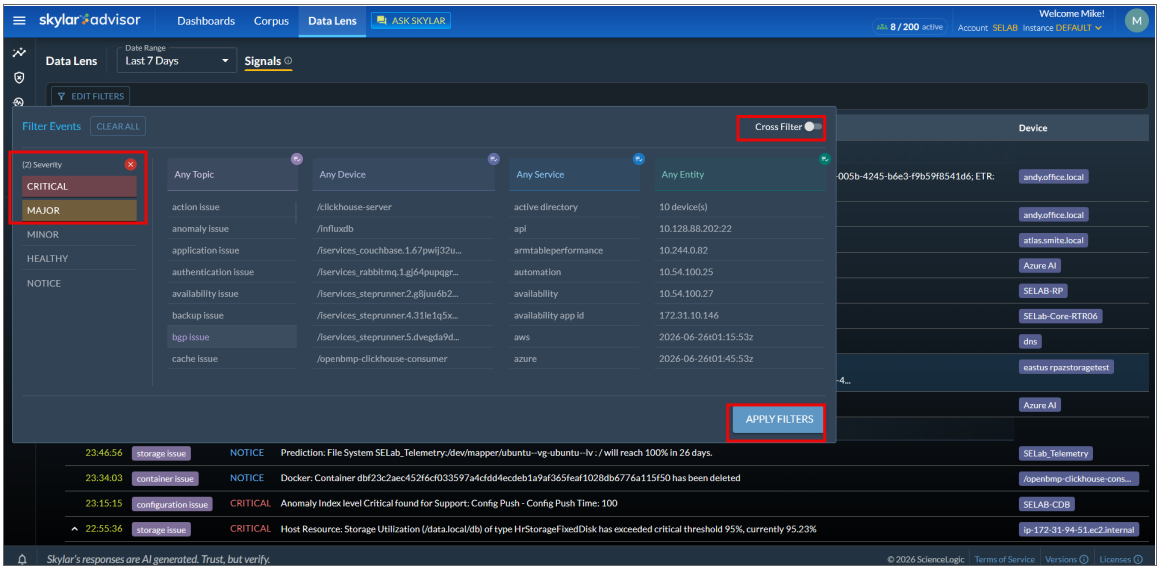



Click an event to view more information about that event.

You can use the filter boxes at the top of the table to sort the list of events. For more information, see *Filtering the Data Lens Page*.

# Filtering the Data Lens Page

On the **Data Lens** page, events are shown with the most recent event at the top. You can use the filter boxes at the top of the table to view more or less information on this page.




To filter this page, click in one or more of the filter boxes at the top:

- Change the time frame by clicking the date range at the top left.
- Under ***Severity***, select one or more severities to show on the **Data Lens** page.
- Narrow down the selections by typing in the "Any" text boxes at the top of each column.
- Enable cross filtering with the ***Cross Filter*** toggle at the top right. When enabled (blue), filter options are constrained based on your current selections.
- Click **[Apply]** to save your filters.

**TIP:** After you save your filter setup, the filters that are currently in use display at the top of the page.

# Chapter 10

# Supervising with Skylar Advisor

## Overview

Use this chapter if you are a NOC supervisor or shift lead holding Owner-level permissions. An operator's unit of work is one Advisory. As a supervisor, you unit of work is the whole set of Advisories: whether work is moving, whether your people are growing, and whether the queue itself is getting smaller.

This chapter is organized around those three goals, and it is a companion to *Working an Advisory*, the operator task.

This chapter covers the following topics:


*What Supervisors Watch* ..... 82

*The Two Supervision Loops* ..... 82

*The Shift Loop* ..... 82

*The Weekly Loop* ..... 83

*Approving Gated Tool Runs* ..... 85

*Reporting Upward* ..... 85

*Supervisor Failure Modes* ..... 86

# What Supervisors Watch

Everything your operators do in Skylar Advisor is written into their threads with a name and a timestamp. You almost never need to ask what happened; read the thread, then ask about the judgment behind it, not the facts. For more information, see *Before You Work an Advisory*.

Holding the Owner role means you can see things your operators cannot: the **[Coach]** panel on their profiles, the Agentic Control Plane page, and tool and source configuration. Where this task sends you somewhere your operators cannot reach, the text states this clearly.

# The Two Supervision Loops

Supervising with Skylar Advisor runs on two loops with different goals:

- The ***shift loop*** keeps today's queue moving. For more information, see *The Shift Loop*.
- The ***weekly loop*** reduces next month's volume: growing your people, improving your knowledge, and cutting noise at the source. For more information, see *The Weekly Loop*.

If you only ever run the shift loop, nothing structurally improves.

# The Shift Loop

The following sub-topics cover the typical steps in a shift.

## Reading the Queue

Go to the **[Recent Activity]** tab on the **Dashboards** page and set the scope to the whole organization instead of **My Work**. **My Work** is your own participation; for supervising, you want everyone's. For more information, see *The Dashboards Page*.

Set the time frame to the shift. The presets cover the usual spans, and a custom range covers the rest. Then use the state filters in the left rail: *Open*, *Claimed*, *Closed*, *Reopened*, and *Archived*. Those five states are your whole queue view. For more information, see *Claim and Close*.

## Finding the Three Faults

Almost everything that needs your attention shows up as one of the following:

- ***Unclaimed and aging***. Filter to *Open*. The oldest advisories appear at the bottom. An Advisory nobody has claimed is the one failure mode where nothing at all is happening. If this list is consistently long, the problem is capacity or visibility, not effort.

- ***Claimed but stalled***. Filter to *Claimed* and look for threads with an old claim and no recent messages. The thread tells you instantly whether the operator is stuck, waiting on an approval, or has moved on and forgotten it.
- ***Reopened***. Filter to *Reopened*. Every one is a closure that did not hold. Read the closing notes: was the fix wrong, or was the Advisory closed without a recovery snapshot? These are your highest-value coaching material, because the evidence is already in the thread.

### Unblocking and Rebalancing

Approve anything waiting on you. For more information, see *Approving Gated Tool Runs*. Add yourself or a second operator to threads that are stuck; adding a colleague brings them the full history instead of a retyped summary. Where the impact is growing, escalate rather than letting an operator keep working it alone.

Use the other **Dashboards** page tabs to answer who a problem is actually affecting: the **[Services]** and **[Service Tiers]** tabs for business impact, the **[Device Groups]** and **[Organizations]** tabs for where the problem is concentrated, and the **[Issues]** tab for what keeps recurring. For more information, see *The Services, Device Groups, Issues, and Organizations Tabs*.

### Handing Off

If your administrator has enabled MCP Tools, run the `handoff` tool from the **Toolbox** page at the end of the shift. It reports what is new, what is still open and who owns it, what was resolved, and what to watch, fleet-wide, over the window you choose. Read it before you send it: it summarizes what your operators wrote down, so a thin thread produces a thin handoff, which is itself worth knowing. For more information, see *Handoff*.

> **TIP:** Ask the thread before you ask the person. Asking an operator "what is happening with this one" costs them their focus and gets you a summary. The thread gets you the facts in less time. Save the conversation for the judgment call: why they ruled something out, or why they escalated when they did.

## The Weekly Loop

The following sub-topics cover the typical steps in a weekly loop.

### Reviewing Your People

Click any operator's avatar to open their profile, and step it to the month you want. You see their ***Time to Claim*** and ***Time to Close*** averages, how much they used Ask Skylar, how many KB articles they generated, how much they shared, an activity-over-time chart, and a **Standing** panel that ranks them against the rest of the organization. For more information, see *Checking Your Own Performance*.

Click the **[Coach]** button for two grounded sections: a factual ***Activity summary*** with no advice in it, and ***Areas for improvement***, which lists two or three focus areas, each naming the metric behind it and a suggested next step. Both sections are drawn from that month's real numbers, and coaching is skipped when there is too little activity to say anything meaningful.

> **NOTE:** Your operators cannot see the **[Coach]** button. It is an Owner-level control, so an operator opening their own profile sees the trends and their standing but not the advice. If you want an operator to act on the coaching, open it with them; a conversation grounded in the advice lands better than the numbers alone.

### Reading the Numbers Before You Use Them

| What You See | What It Usually Means |
|---|---|
| Time to Claim climbing | A visibility or capacity problem before it is an effort problem. Ask where the operator is watching from before you ask why they are slow. |
| Time to Close long, KB Generated zero | The most actionable pattern on the page. The operator is solving hard problems and capturing none of it, so the team relearns it each time. |
| Ask Skylar high, Turns low | Someone using Skylar Advisor to learn rather than to close. Early in a ramp, this is exactly right. |
| Shared near zero | Working alone. Sometimes discipline, sometimes reluctance to ask; worth a conversation either way. |
| A low standing in a strong team | Often nothing. Standing is relative to your organization, not an absolute bar. |

## Reviewing Your Knowledge

Two surfaces under the **Knowledge Center** page answer the same question from opposite ends: is what your team knows actually serving it?

- The **[KB Health]** view shows how each generated article is performing. An article that nothing ever cites is either wrong, unfindable, or about something that stopped happening. For more information, see *KB Health: What Is Earning Its Place*.
- The **[Info Gaps]** view shows what the Corpus is failing to serve. For more information, see *Info Gaps: What Is Missing*.

Protect the replay step when your team closes an Info Gaps topic. Replaying the original questions is what turns "we should document that" into a loop that actually finishes, and it is one of the most effective things you can do to make next month's shifts shorter.

## Reducing Noise

Your operators work what arrives. You are positioned to reduce what arrives at all. If your administrator has enabled MCP Tools, run the following tools from the **Toolbox** page, where a tool runs on its own outside any conversation. The **Toolbox** page is also where you export a report; for more information, see *Reporting Upward*.

| Tool | Use It to Answer |
|---|---|
| `noise` | Which sources are noisiest and which alerts keep re-firing, with specific tuning suggested, such as deduplication, thresholds, or suppression. |
| `trends` | Where the team's effort actually went, with period-over-period movement. The honest answer to "what are we spending our time on." |
| `topic_pulse` | Whether a technology or symptom is rising or cooling across the fleet. |
| `correlate` | Whether a run of separate advisories was really one thing. For more information, see *Choosing the Right Tool for the Question*. |
| `spend_value` | Token spend for the period, optionally broken out per user, and how much of it produced value. For more information, see *Spend Value*. |

## Approving Gated Tool Runs

You can configure some tools to require sign-off before they run. For more information, see *MCP Tools*. When an operator reaches one of these tools, a request appears in their thread carrying the tool, its parameters, who asked, and their written justification. You can approve the request, or deny it with a structured reason: wrong target, insufficient justification, timing, or other.

The following rules are enforced by the server, not by the interface:

- You cannot approve your own request. If you raised it, someone else signs it off.
- Approver membership is checked server-side. Being able to see the thread does not make you an approver; you must belong to that tool's approver group.
- Approval expires. A request that waits too long lapses instead of being decided against a situation that has already moved on.
- Approval is not permanent. Once granted, the run stays available to the operator for a bounded window; after that, it needs approval again.

> **NOTE:** There is no separate approval inbox. A pending request is visible inside the thread where it was raised, and also on the **Dashboards** page, under the *Pending* filter in the left navigation.

## Reporting Upward

You can export anything you run from the **Toolbox** page as an HTML page, a PowerPoint deck, a PDF, a Word document, or Markdown. Headings become slides in the deck, and tables stay intact. As a result, your weekly report is usually a tool run plus a review pass on the report Skylar Advisor generated for you, not a document you assemble by hand.

The following table lists what typically travels well upward:

| The Question You Are Asked | What to Bring |
|---|---|
| "How did the week go?" | The ***handoff*** tool over a weekly window, or the ***trends*** tool for effort and recurrence with period-over-period deltas. |

| The Question You Are Asked | What to Bring |
|---|---|
| "Are we getting on top of it?" | The ***noise*** tool, and the tuning you have already actioned since last time. Movement matters more than the absolute number. |
| "What did it cost?" | The ***spend_value*** tool for the period, broken out per user if the question is about distribution. |
| "What was the business impact?" | The **[Services]** and **[Service Tiers]** tabs on the **Dashboards** page, which speak in services rather than devices. |
| "Why did this one take so long?" | The thread. It carries the timeline, what was ruled out and by whom, and the metrics on either side of the fix. |

IMPORTANT: Claim and close timings are only as honest as your team's habits. If operators close early and reopen, or close without a recovery snapshot, the timings improve while the service does not. Fix the habit before you report the number, and if you know a number is flattered, say so when you present it.

## Supervisor Failure Modes

| Pattern | Why It Costs You |
|---|---|
| Running only the shift loop | Work keeps moving and nothing gets smaller. The weekly loop is the only one that reduces next month's volume. |
| Asking operators for status the thread already holds | You break their focus to get a worse version of what you could have read yourself. |
| Coaching from the standing panel alone | Standing is relative to your team. Coach from the thread and the two timing metrics, not from a leaderboard position. |
| Treating a reopened Advisory as a failure | Reopening is the honest act. Punish it, and operators stop doing it, which costs you your best quality signal. |
| Letting KB articles accumulate unreviewed | A wrong article is worse than none, because the next operator trusts it. The **[KB Health]** view tells you which articles are actually being relied on. |
| Waiting to be notified about approvals | Monitor the *Pending* filter on the **Dashboards** page yourself. A request that expires quietly leaves the operator thinking the approval gate is broken. |

Keep the following four points in mind:

- Read the queue by state, not by scrolling.
- Chase unclaimed, then stalled, then reopened advisories, in that order.
- Protect the weekly loop when the shift gets loud.
- Watch for approval requests to keep the team moving.

# Chapter 11

# Managing OpenTelemetry (OTel) Alerts on the Alert Builder Page

## Overview

*OpenTelemetry* provides a flexible solution for transferring observability data from an application into Skylar AI. As an alternative to importing data from Skylar One, you can configure Skylar AI to ingest data from an OpenTelemetry (OTel) Collector.

You can view and create alerts based on metrics gathered by the OTel Collector.

> **NOTE:** The Skylar AI 2.8.0 preview feature only allows the collection and viewing of metrics and specific gauges and sums. More metrics will be available in future releases.

This chapter covers the following topics:


*Configuring the Skylar OTel Collector* ..... *88*

*The Alert Configuration Page* ..... *90*

*The Alert Definitions Tab* ..... *91*

*The Alert Builder Tab* ..... *93*

*Alert Builder Aggregations* ..... *94*

*Notes for This Release* ..... *97*

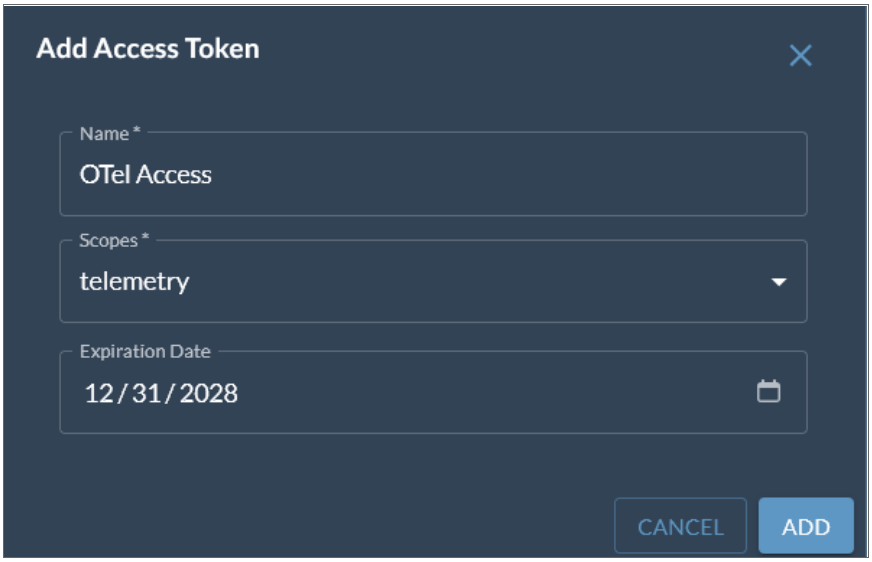
# Configuring the Skylar OTel Collector

You will need to configure and deploy an OTel Collector that collects a range of data and then converts that data into OpenTelemetry. That data gets sent to Skylar AI using ***OpenTelemetry Protocol (OTLP)***, a vendor-neutral standard that defines how telemetry data is encoded and transmitted between a source and a destination.

## Configuration

To configure an OTel Collector:

1. Log in to Skylar AI and go to the **Access Tokens** page ( > Account Settings > Instances > Access Tokens).
2. Click the **[Add Access Token]** button. The **Add Access Token** window appears.

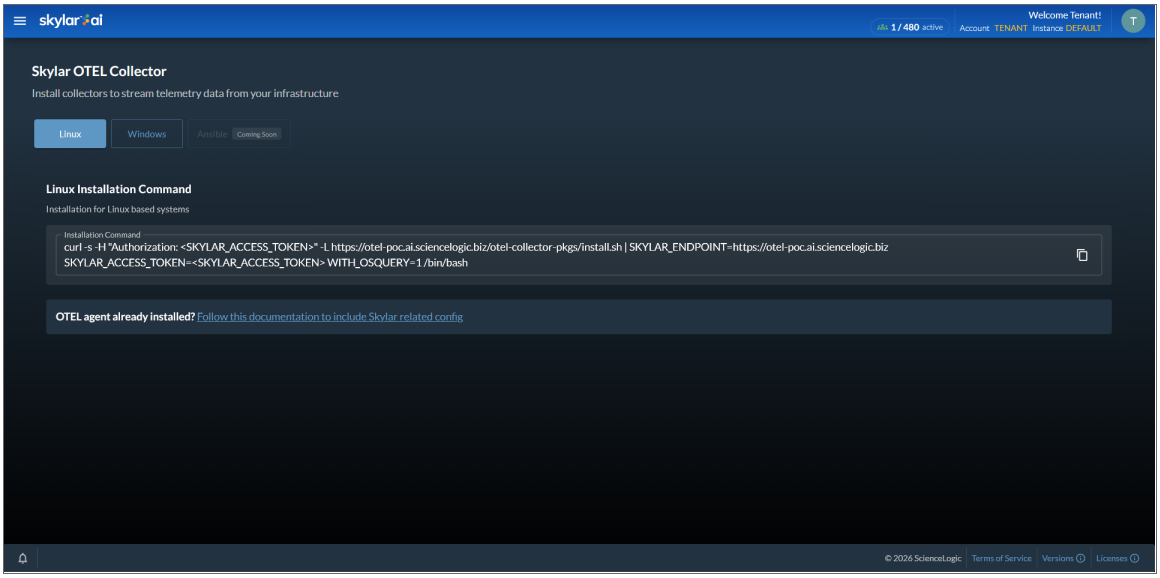



3. Complete the following fields:
   - ***Name***. Type a name for the token, such as "Skylar One Collector".
   - ***Scopes***. The scope of an access token determines which application or service you can connect to with the access token. You can select more than one scope for an access token. You will need a different access token for each Skylar AI instance you are connecting with access token.

     If you are creating this access token to *Create a Service Connection* in Skylar One , select both *sl1_connector* and *telemetry*.

     If you are creating this access token to *Configure an OTel Collector* for Skylar Advisor, select *telemetry*.
   - ***Expiration Date***. Select an expiration date.

4. Click the **[Add]** button. The access token is added to the **Access Tokens** page.
5. Click the copy icon ( ) to copy the access token to the clipboard.
6. In the Skylar Advisor user interface, go to the **Preview Features** page ( > Account Settings > Preview Features) and make sure that the **Alert Builder** option on the **[Integrations]** view is selected (blue).
7. Go to the **Skylar OTEL Collector** page ( > Account Settings > Data Collection > Skylar OTEL Collector).




8. Based on the operating system of your collector, click the **[Linux]** or **[Windows]** button and copy the installation commands.

   If your platform is not listed, you can use an OTel collector that matches your platform, based on the following configurations:

   - https://github.com/open-telemetry/opentelemetry-collector/releases
   - https://github.com/open-telemetry/opentelemetry-collector-contrib/releases

9. Run those commands on your collector. Be sure to replace `<SKYLAR_ACCESS_TOKEN>` with your new token from steps 1-5. You must be an administrator to run these commands.

**NOTE:** The `OSQUERY` command is only available for Linux, and it can be set to not install by changing the following environment variable from the copied installation command: `WITH_OSQUERY=1` should be `WITH_OSQUERY=0`.

## After Configuration

As soon as you enable the OTel Collector and log in to the Skylar Advisor user interface, you should refresh the browser. You should now see a link to the **Alert Configuration** page in the Skylar AI menu ( > Advisor > Alert Configuration).

Please note that the pages are initially inoperable as Skylar Advisor might take up to three to five minutes to become aware that the feature is enabled. After that time, the page should be functional, and the **[Alert Definitions]** and **[Alert Builder]** tabs will stop returning error messages.

> **NOTE:** If you disable the OTel Collector by turning off the preview feature, Skylar Advisor will stop processing new data. Existing data remains in the database until it meets the TTL. Alerts continue to operate, but as no new data comes in, they will never get triggered.

## Considerations for OTel Data

Keep in mind the following details about using OTel data with Skylar Advisor:

- Alerts cannot be created for OTel traces and logs.
- Alerts apply to OTel collected metrics, not data from Skylar One.
- Alert conditions cannot be combined across metrics like a Skylar One formula.
- An alert with an occurrence count greater than 1, when updated or upon creation, will not backfill historical data. This means it will need to partially fill its buffer before alerting.
- You can only have 32 alerts with an occurrence count greater than 1. Alerts with an occurrence count of 1 are unlimited.
- There are not alerts for when a device stops sending data to Skylar Advisor.
- Collecting OTel data requires an account licensed for Skylar Advisor.

# The Alert Configuration Page

The data gathered by your OTel Collector appears in Skylar Advisor on the **Alert Configuration** page ( > Insights > Alert Configuration). You can use this OTel data to generate new alerts that can then be ingested and used by Skylar Advisor.

> **NOTE:** To use the **Alert Configuration** page in version 2.8.0 of Skylar AI, you will need to enable **Alert Builder** in the **[Integrations]** view on the **Preview Features** page ( > Preview Features).

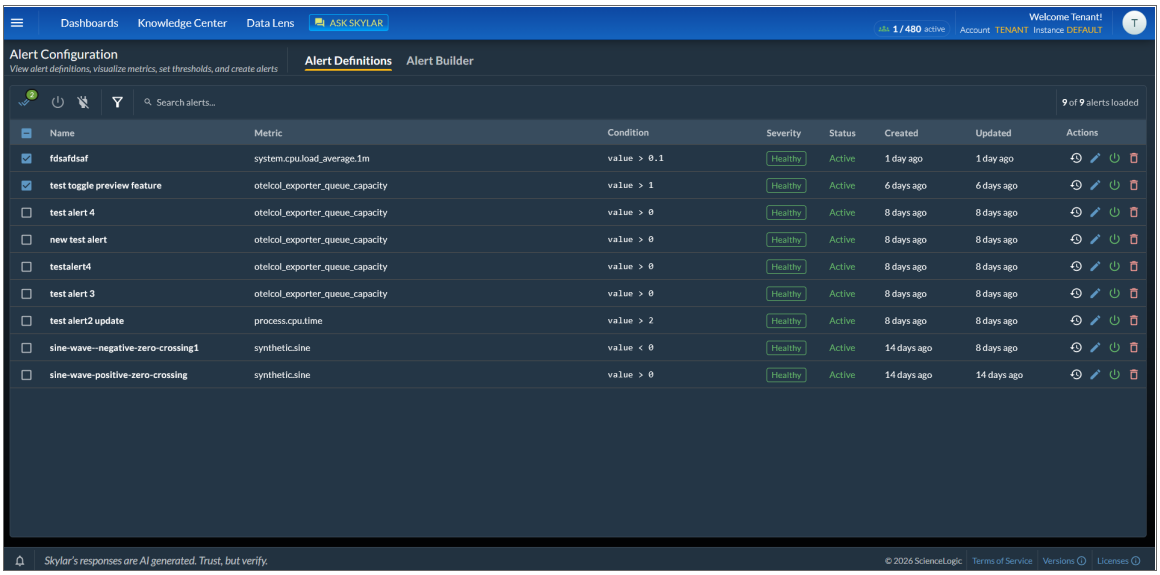
On the tabs of the **Alert Configuration** page, you can browse OTel metrics and create alerts directly from those metrics.

This page includes two tabs:

- *Alert Definitions*
- *Alert Builder*

# The Alert Definitions Tab

The **[Alert Definitions]** tab on the **Alert Configuration** page contains a list of the alerts created with OTel metrics. This tab also displays additional information for each alert, such as the conditions that will trigger the alert, and the severity and status of the alert.

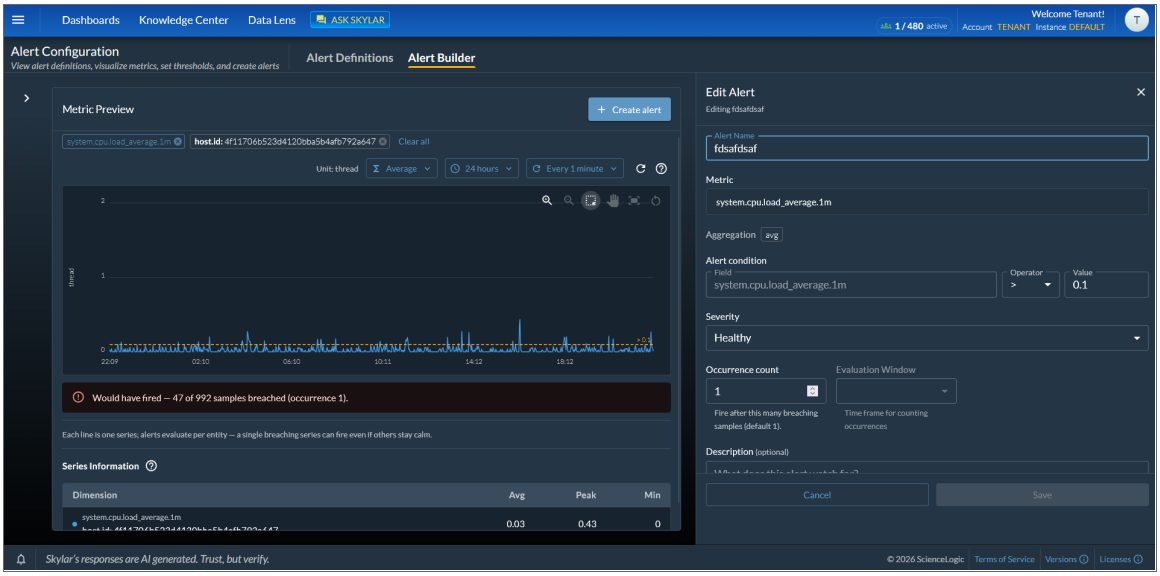



## Editing Alerts

In the **Actions** column, you click an icon to edit an alert, make an alert inactive, or delete an alert.

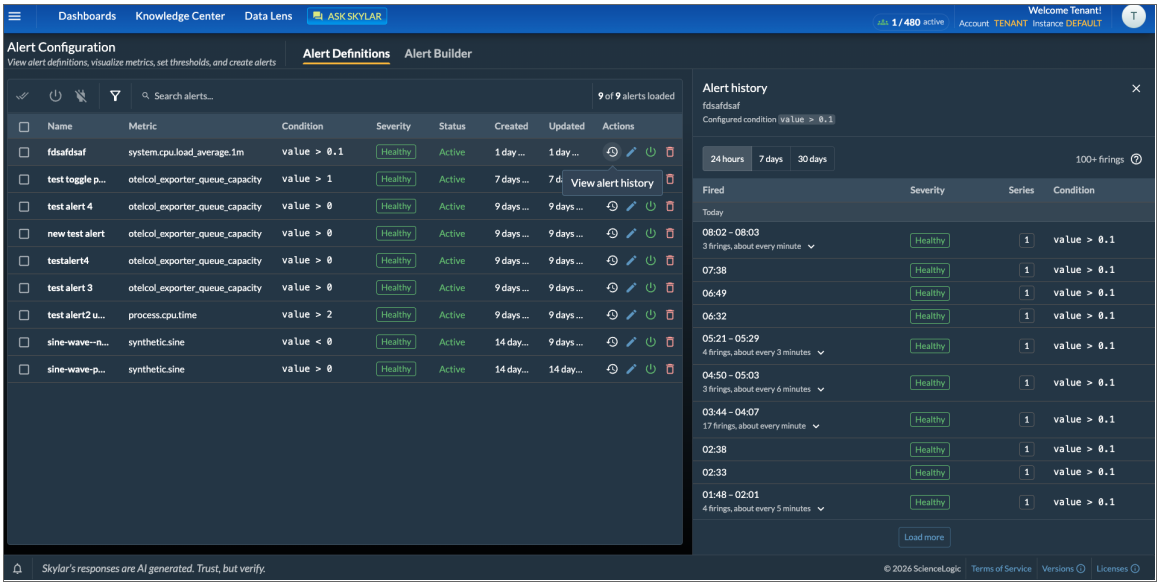
**TIP:** You can sort the columns on the **[Alert Definitions]** tab by clicking the table headings. You can also sort the list by typing search criteria into the **Search** field above the table.

When you click the edit icon ( ) for an alert, an **Edit Alert** pane appears on the right with data about that alert. For more information, see *To create an alert based on OTel metrics* .




## Viewing Triggered Alerts

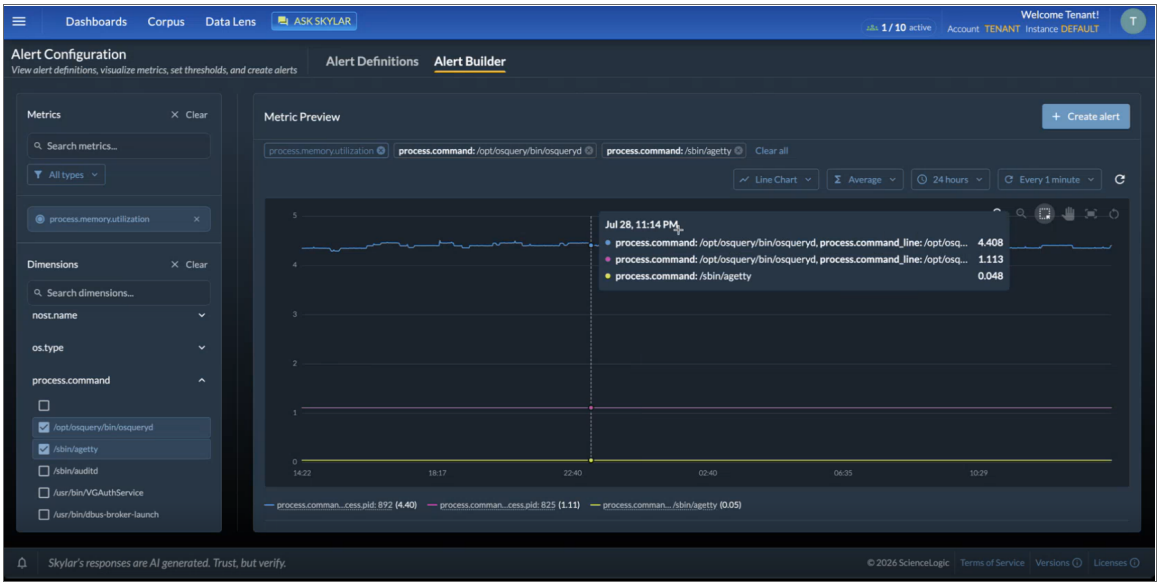
You can view more details about a triggered (or "fired") alert on the **Alert history** pane. To access this pane, go to the **Alert Definitions** page ( > Advisor > Alert Configuration) and click the "View alert history" icon ( ) for a specific alert.




This view is paginated shows up to 100 triggered alerts at a time. You can view the alert time, the severity, the condition for triggering it, and the number of series that triggered it for that time period.

# The Alert Builder Tab

The **[Alert Builder]** tab on the **Alert Configuration** page lets you search the existing OTel metrics, view a graph based on that data, and create an alert for those metrics as needed.




On the **[Alert Builder]** tab, you can create and manage threshold-based alerts for OTel metrics.

To create an alert based on OTel metrics:

1. On the **[Alert Builder]** tab, go to the **Metrics** panel on the left and search for one or more metrics to monitor for the new alert.

   You can sort the metrics by how the data is recorded:

   - *Gauge*: A value measured at a point in time, such as memory used.
   - *Sum*. A counter that accumulates, such as requests served.

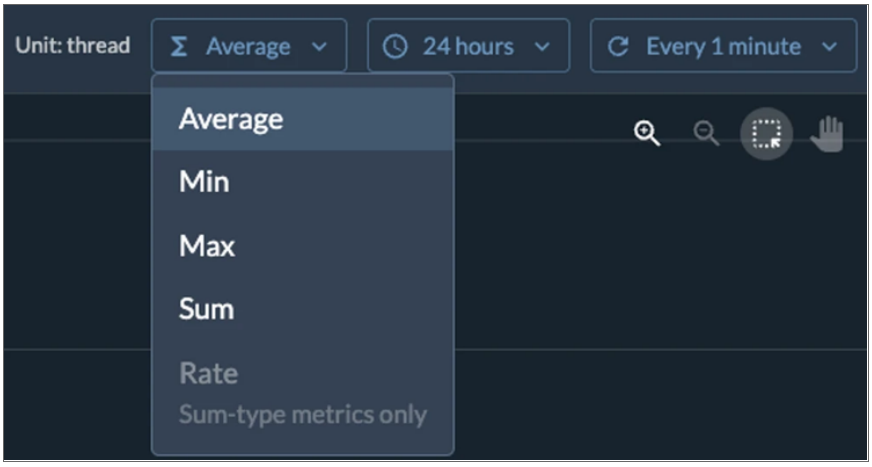
2. Select a metric and one or more dimensions for that metric. The **Metric Preview** panel updates with a graph showing the history for that metric.
3. If needed, you can adjust the settings for the graph using the filter buttons on the **Metric Preview** panel.
4. When you have located the metric and dimensions that you want to monitor, click the **[Create alert]** button. An **Alert Builder** pane appears on the right.

5. Complete the fields on the **Create Alert** pane:
    - ***Alert Name***. Type a descriptive name for this alert.
    - ***Metric***. Lists the metric you selected for this alert. This field is read-only.
    - ***Aggregation***. Displays the type of aggregation used for this alert, such as *avg* for average or *min* for minimum. This field is read-only. For more information, see Alert Builder Aggregations.
    - The following four fields let you set up a threshold for this metric:
        - ***Alert condition***. These three fields let you set up a threshold for this metric. When that threshold is crossed, the alert is generated.
        - ***Severity***. Select one of five severity levels, which are the same labels as those used by Skylar One, from low to high.
        - ***Occurrence count***. Specify how many times a threshold must be crossed ("breached") before the alert is generated. The default is 1.
        - ***Evaluation window***. The period over which the alert condition is evaluated. The alert triggers when the condition is met the number of times set in Occurrence count within this window. Available only when ***Occurrence count*** is greater than 1.
    - ***Description***. Type a short explanation of what your alert does and why. Optional.
6. Click **[Save]**. The new alert is added to the **[Alert Definitions]** tab on the **Alert Configuration**page.

## Alert Builder Aggregations

The **Alert Builder** page lets you choose an aggregation type when you are creating or editing an alert. An ***aggregation*** is an operation that compresses time series data so that a single data point represents collected data points over a defined span of time. The goal of an aggregation is for data that is being collected at a high frequency and the customer wants to reduce the amount of alerts they see over time.

The **Alert Builder** page currently offers the following aggregation options:

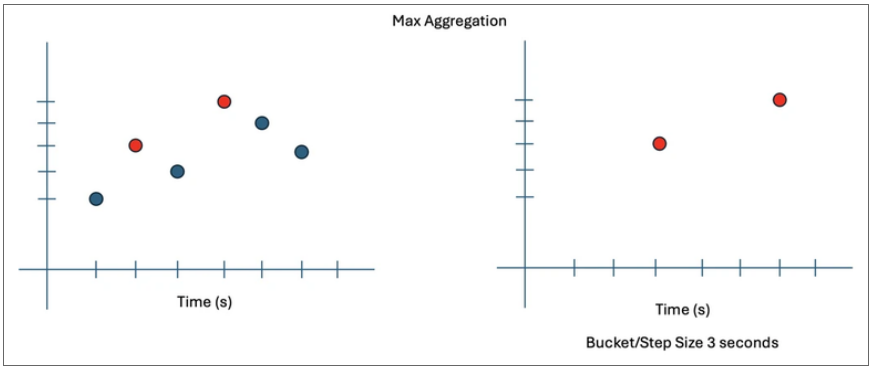



The best way to think of these spans of times is to think of them as buckets that either choose the sum, maximum, minimum or average value within the bucket.

Rate is only used for counter/sum metric types. The way that Skylar AI calculates the rate within the bucket is *(max - min) / (bucket size in seconds)*.

Typically our bucket size is fixed to 60 seconds. This true with the exception of our 24-hour time period.

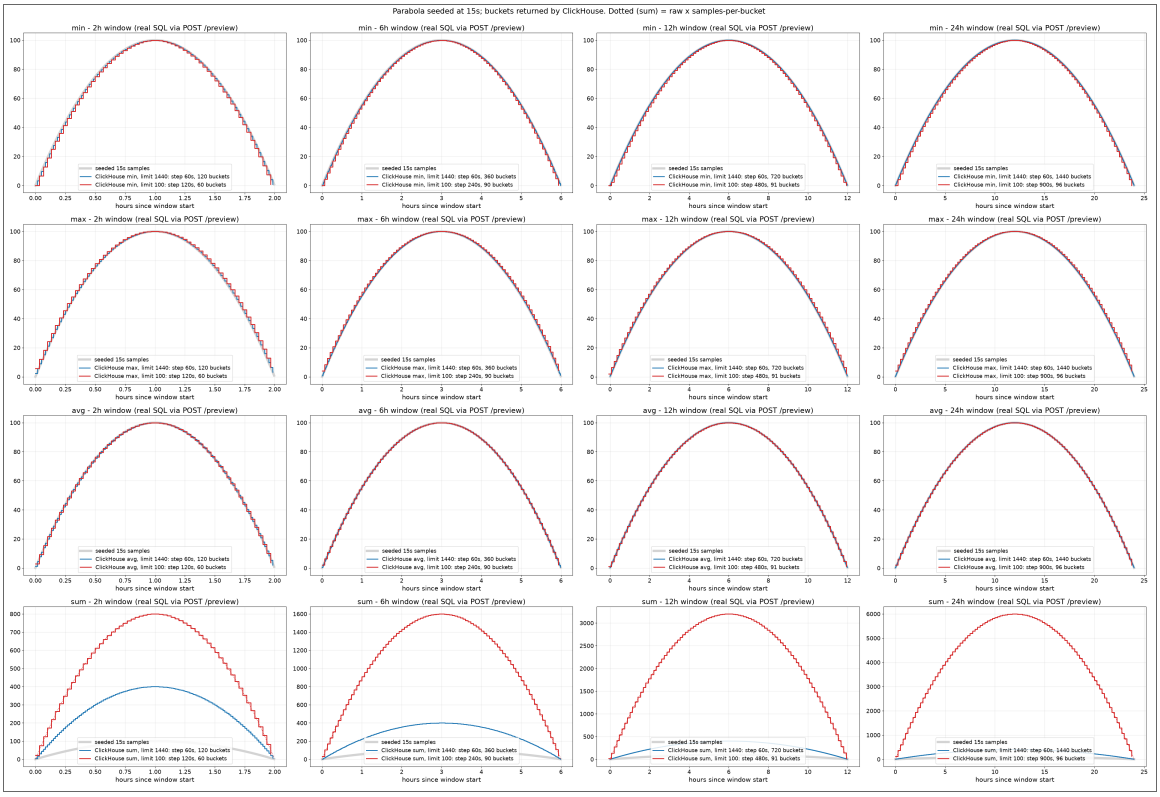
For the bucketing process, consider that we have six data points across six seconds. We have a bucket/step size of three seconds. The aggregation is set to *Max* value. For every bucket we choose the *Max* value and report that as the latest time value:

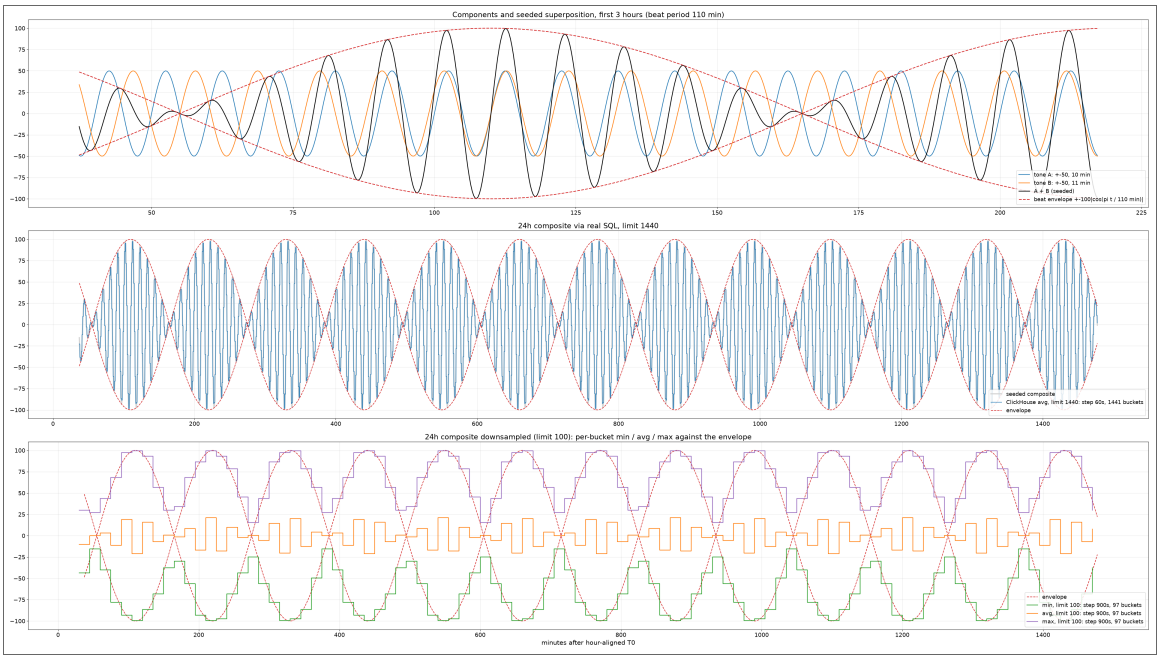



The result is fewer, but more relevant data points over time, for a clearer picture of what is going on in your environment.

**NOTE:** Alert definitions currently do not take into account aggregation. Instead, they operate solely on the instantaneous collection interval, and alerts are triggered per sample.

Below are additional examples of aggregations:



# Notes for This Release

This section lists key considerations for the 2.8.0 preview feature version of Alert Configuration and Alert Builder.

## Data and Scope

- Alerts on the **[Alert Builder]** tab are based only on OTel metrics. Traces and logs sent by the OTel Collector are not available for alert creation.
- Alerts apply to data sent by the OTel Collector. They do not apply to data imported from Skylar One.

## Alert Conditions

- Each alert has a single condition that compares the metric value to a threshold, using one of these operators: **>, >=, <, <=, ==, !=**. Multiple conditions cannot be combined in one alert.
- A limited number of alerts per instance can use an ***Occurrence count*** greater than 1. The default maximum is 32

## Alert Timing

- When the ***Occurrence count*** is 1, an alert can take up to five minutes to trigger after the condition is met.
- When ***Occurrence count*** is greater than 1, the time to trigger depends on the Occurrence window: up to ten minutes for a one-hour window, and up to two hours for a 12-hour or one-day window.
- Creating, editing, or re-enabling an alert with an ***Occurrence count*** greater than 1 resets its history (no historical backfill). The alert cannot trigger until the condition has been met the required number of times within the ***Evaluation window*** after the change.
- If Skylar AI cannot activate an alert after you save it, the alert is saved with a status of *Inactive*. Check the ***Status*** column on the **Alert Definitions** tab and re-enable the alert.

## The Alert Builder Page

- After an OTel collector starts sending data, its metrics can take several minutes to appear in the **Metrics** panel. Metrics that arrive later may not appear until the page data is refreshed.
- The "Metric Preview" graph is a visual aid. Data gaps are reported as a 0 value. Negative values are clipped.

## OTel Collector Installation and Operation

- The installation commands contain an access token. Treat the commands as confidential.
- After you upgrade a collector, verify that the collector service is running and that its endpoint and access token settings are unchanged.

# Chapter 12

# Experimenting with New Features on the Preview Features Page

## Overview

The **Preview Features** page ( > Account Settings > Preview Features) in Skylar Advisor lets users toggle on and experiment with new features. These features are stable, but ScienceLogic is actively developing and improving them in dialogue with users.

This chapter covers the following topics:


*The Preview Features Page* ..... 100

*Agentic* ..... 101

*Agentic: Topic Specialists* ..... 102

*Agentic: MCP Apps* ..... 103

*Agentic: Trusted Certificates* ..... 103

*Integrations: ServiceNow Outbound* ..... 103

*Integrations: Alert Builder* ..... 104

*Work Items: Metrics Enhanced Advisories* ..... 105

*Work Items: Advisory Suppression* ..... 106

*Work Items: Entity Exclusion* ..... 107

*Work Items: Data Lens Events* ..... 108

*Work Items: SLO Docs* ..... *108*

*Work Items: Auto-close* ..... *108*

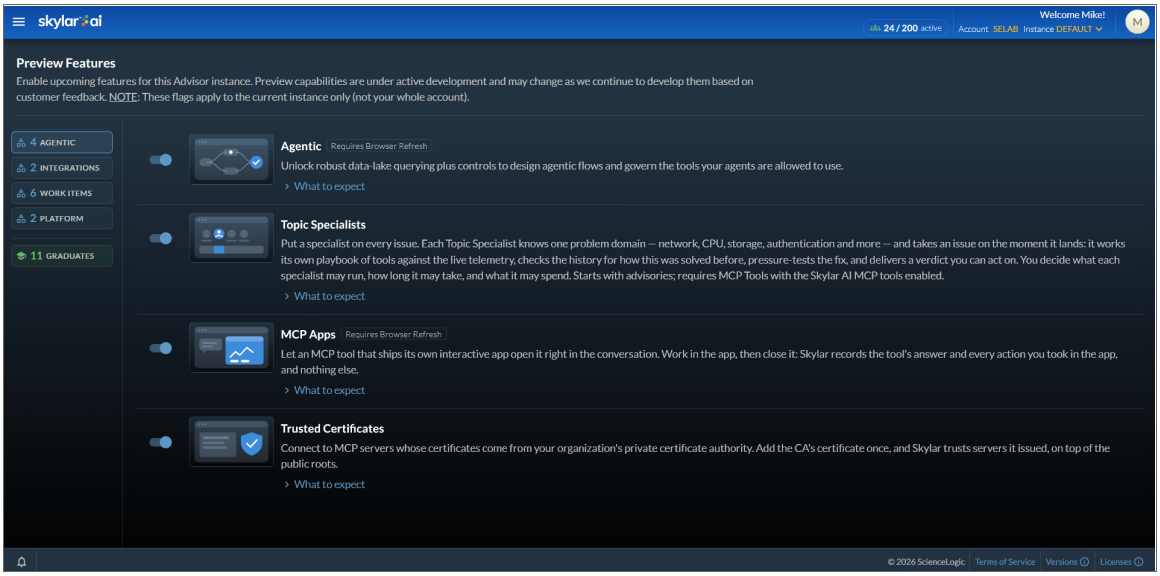
*Platform: KB Generated Approval Flow* ..... *109*

*Platform: Knowledge Sentry* ..... *110*

# The Preview Features Page

The **Preview Features** page ( > Account Settings > Preview Features) gives users the option to toggle on one or more new features that will be coming to a future release of Skylar Advisor. Please note that these features are stable, and ScienceLogic is actively developing and improving them in dialogue with users.

These toggles apply only to the current instance, not the entire Skylar Advisor account.




Each preview feature contains a What to expect section that you can click to get more information about that feature and where to use it in the Skylar Advisor user interface.

The left pane of the **Preview Features** page includes clickable views on the left that organizes the preview features by category, such as **Agentic**, **Integrations**, **Work Items**, and **Platform**. The views also list the number of features for each set of features.

**TIP:** Click the "Play" icon next to a preview feature to watch a short demo of that feature. Also, some features might require you to refresh your browser before they appear.

## Preview Feature "Graduates"

The **Preview Features** page also includes a view for ***Graduates***, features that used to be listed on the **Preview Features** page, but are now officially part of the Skylar Advisor user interface. The "graduates" are listed by release.

For Skylar AI 2.8.0, the following preview features were made permanent:

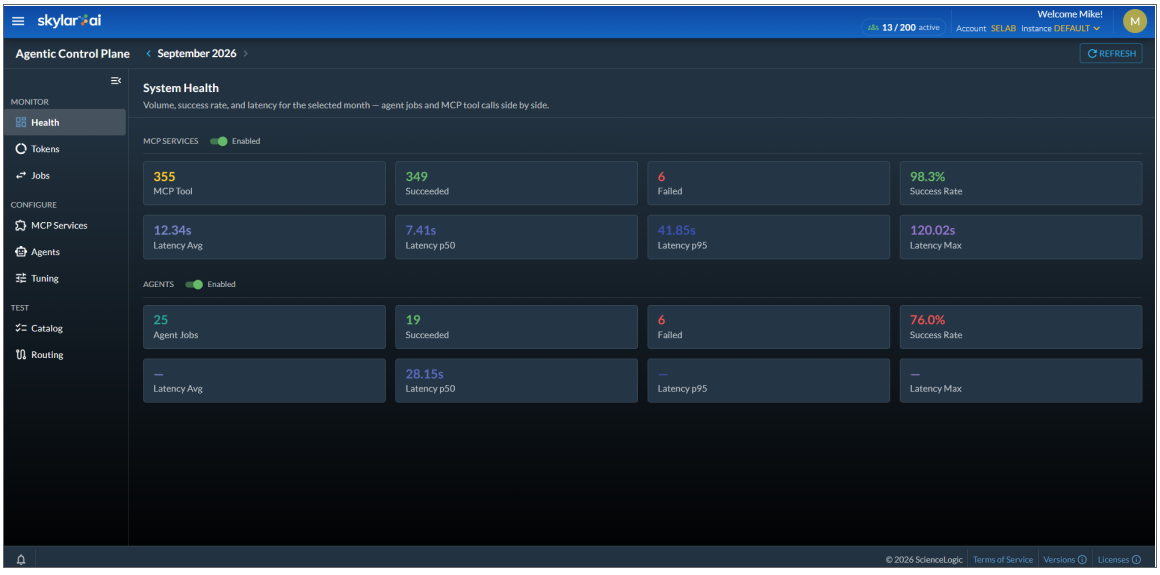
- Web Search, which is now covered in *Using Web Search*.
- Claim & Close, which is now covered in *Working an Advisory From Claim to Close*.
- Corpus Tools, which are now covered in *Filling Gaps with the Info Gaps View*.

For Skylar AI 2.7.0, the following preview features were made permanent:

- MCP Tools, which are now covered in *MCP Services*.
- Scheduled Reports, which are now covered in *Scheduling Reports for Skylar Advisor*.
- Ad-hoc File Upload, which is now covered in step 2 of *Using the Ask Skylar Feature*.
- Trend Analysis, which is now covered in *Trend Analysis*.
- Conversation Audit, which is now covered in *Additional Options on the Ask Skylar Page*

# Agentic

Selecting the **Agentic** option on the **Preview Features** page enables the **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane). This page lets you monitor any AI agents that you are using with Skylar Advisor. You can view a list of agents, jobs, tools, and health, along with latency information, audit trails, agent limits and bounds, and token spend over time.



When you enable an agentic tool, Skylar Advisor routes the relevant questions to that tool instead of searching the documents on the **Corpus** page. Each agentic tool comes pre-loaded with a list of sample questions.

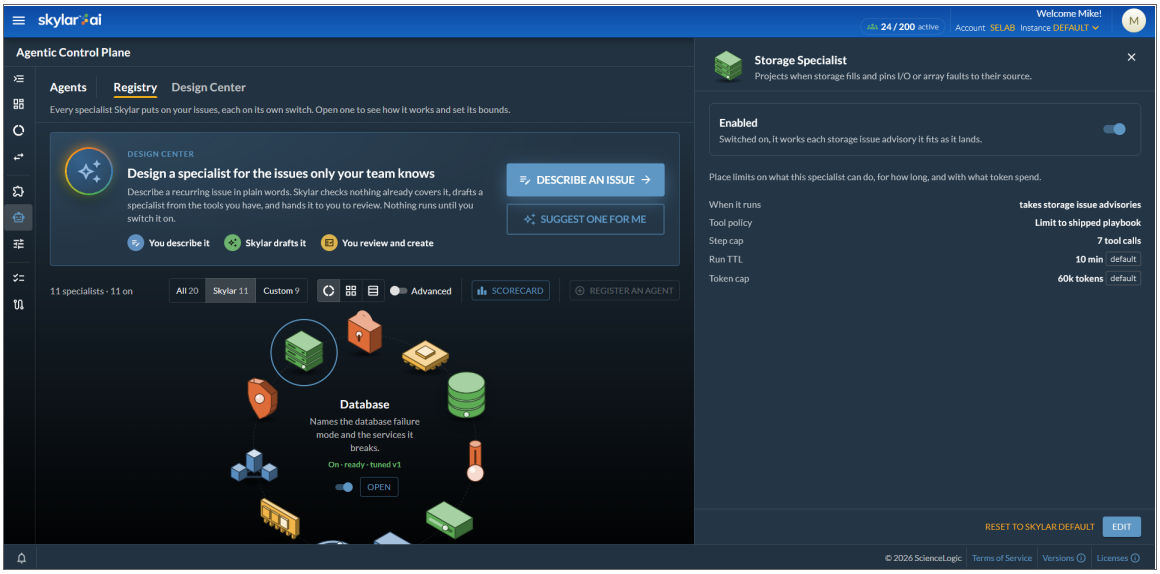
For more information, see *Using the Agentic Control Plane Page*.

# Agentic: Topic Specialists

Selecting the **Topic Specialists** option on the **Agentic** view of the **Preview Features** page enables AI agents that are trained on a specific problem domain to review every Advisory, producing a cited verdict, recommendation, or remediation plan.

With the topic specialists, each Advisory is tagged by topic, such as network, CPU, storage, database, latency, memory, authentication, security, temperature, or processor. A matching AI agent or "specialist" investigates against live telemetry and generates an Advisory report. The report combines the monitoring Advisory with the specialist's evidence: what's wrong, why, what argues against it, and what to do next, with KB and web citations.

You can monitor your topic specialists on the **[Agents]** tab of the **Agentic Control Plane** page:




You can specify which specialist to run, how long it might take, and what it may spend. By default, all topic specialists are switched off. You can enable and edit the schedule for when a specialist runs: when a matching Advisory lands, or a scheduled "sweep" of what is new in Skylar Advisor. Admin users can also design their own specialists by describing a recurring problem in plain language.

Topic specialists require that you enable the "MCP Apps" preview feature, with the "Skylar AI MCP" tools enabled on the **[MCP Services]** tab of the **Agentic Control Plane** page.

For more information, see *Agentic Control Plane*.

# Agentic: MCP Apps

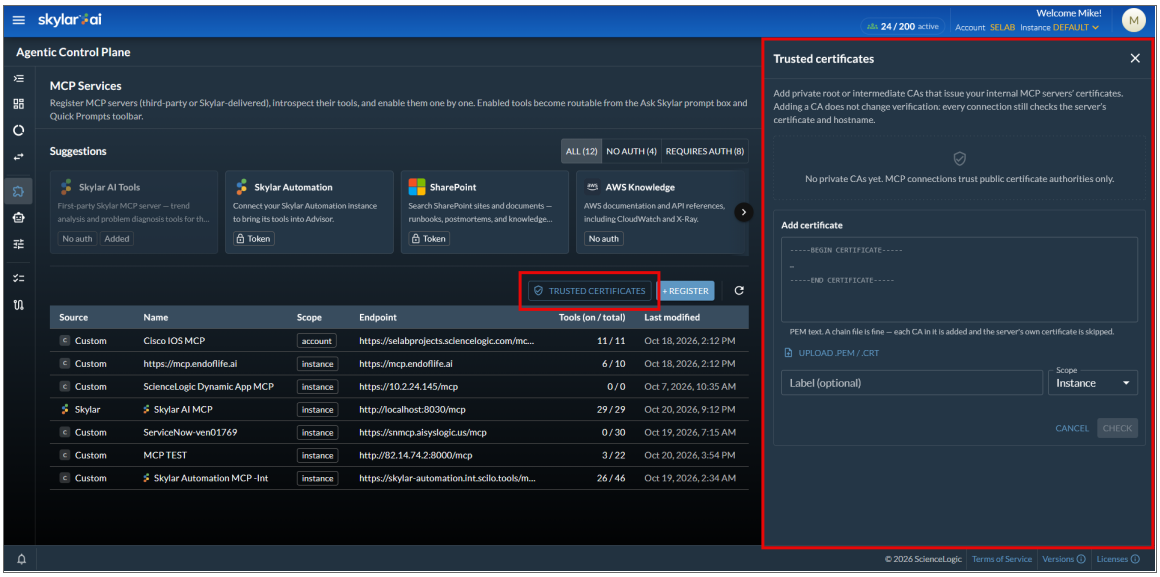
Selecting the **MCP Apps** option on the **Agentic** view of the **Preview Features** page enables an MCP tool that ships its own interactive application open it right in the conversation.

You can work in the application and then close it. Skylar Advisor records the tool's answer and every action you took in the app, and nothing else.

# Agentic: Trusted Certificates

Selecting the **Trusted Certificates** option on the **Agentic** view of the **Preview Features** page lets you connect to MCP servers whose certificates come from your organization's private certificate authority (CA). After you add the CA's certificate, Skylar Advisor trusts the MCP servers whose certificates that CA issued.

To add, review, or remove a private CA, go to the **[MCP Services]** tab on the **Agentic Control Plane** page ( > Account Settings > Agentic Control Plane > MCP Services) and click the **[Trusted Certificates]** button:




For more information, see *MCP Services*.

# Integrations: ServiceNow Outbound

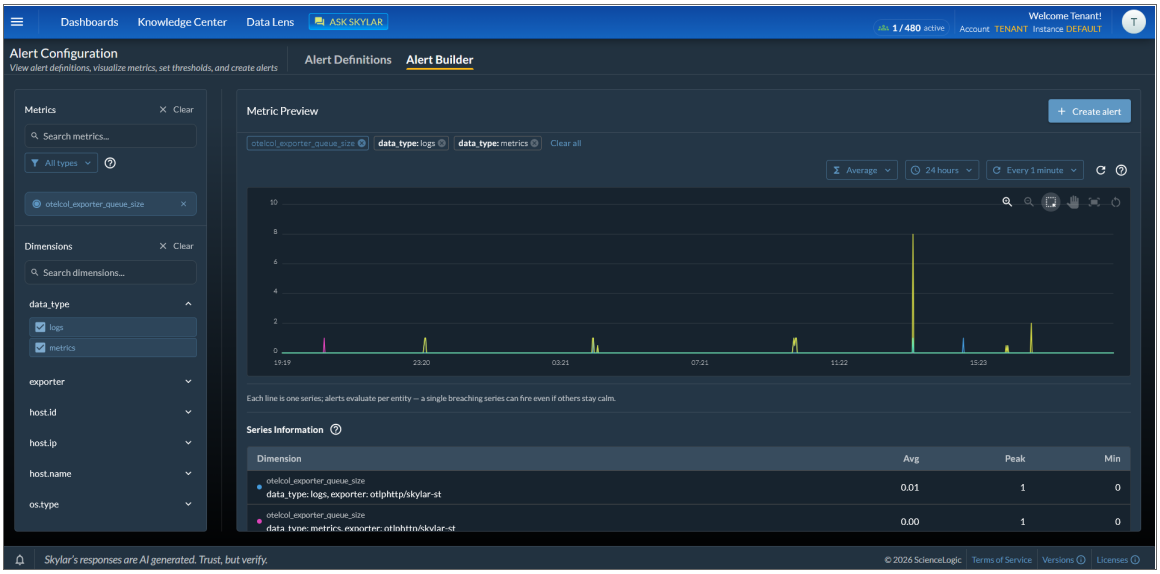
Selecting the **ServiceNow Outbound** option on the **Integrations** view of the **Preview Features** page enables Skylar AI to automatically create a ServiceNow Incident or Problem from a Skylar AI result, accelerating incident response and reducing manual operational handoffs. This feature is still under

development.

# Integrations: Alert Builder

Selecting the **Alert Builder** option on the **Integrations** view of the **Preview Features** page lets you stream your own OpenTelemetry metrics from an OTel Collector into Skylar AI. ***OpenTelemetry*** provides a flexible alternative to using Skylar One for importing observability data into Skylar AI.

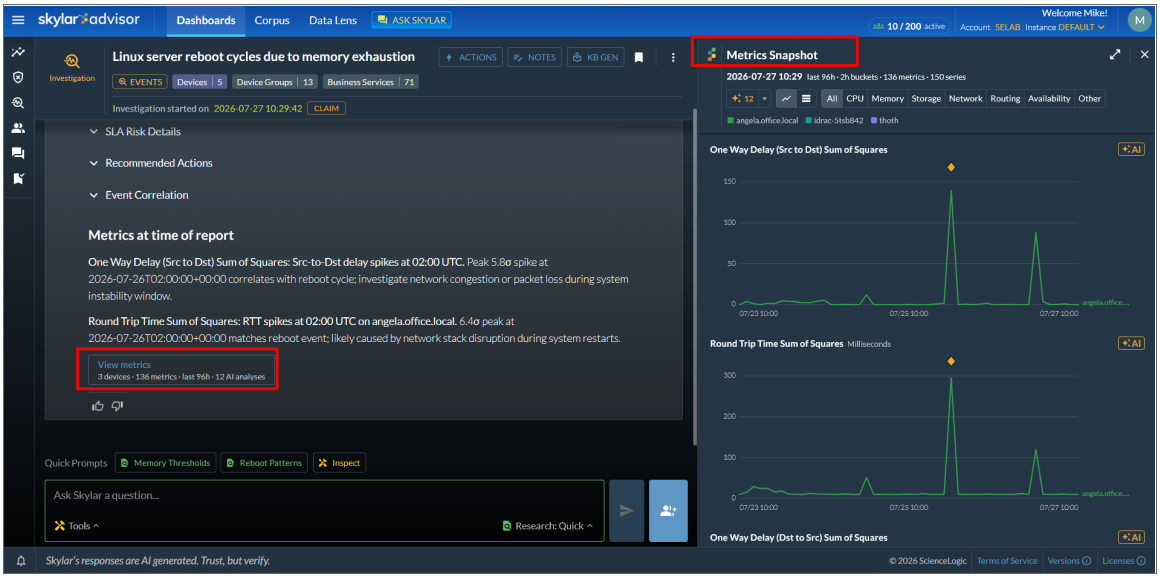
You can use the **Alert Builder** page (Alert Configuration > Alert Builder) to create alerts based on those metrics.




For more information, see *Managing OpenTelemetry (OTel) Alerts on the Alert Builder Page*.

# Work Items: Metrics Enhanced Advisories

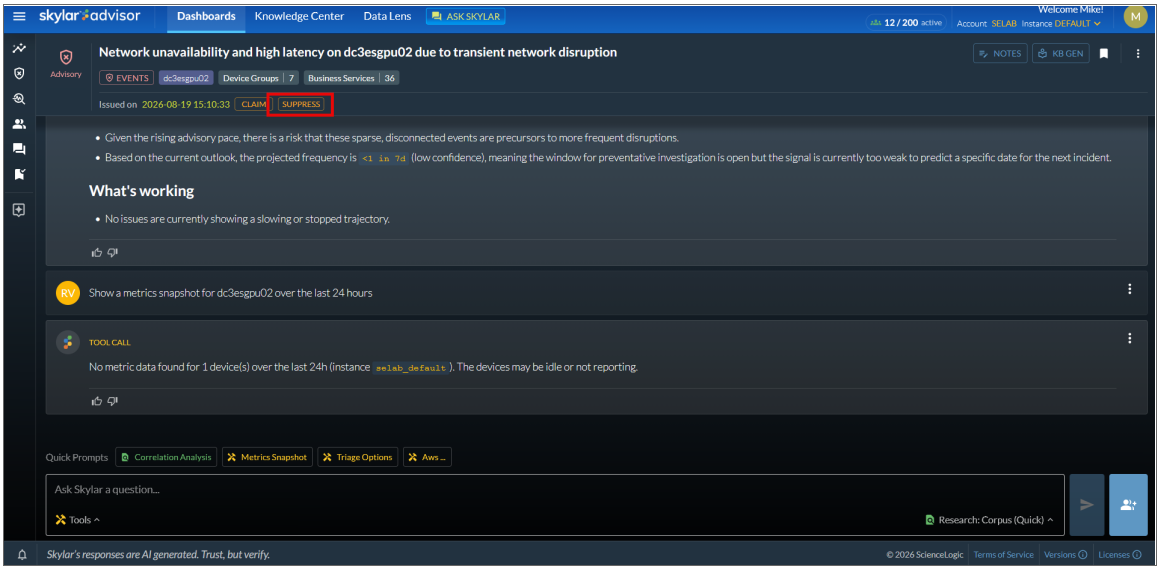
In the **Work Items** view, you can enable **Metrics Enhanced Advisories** to add more data to an Advisory or Investigation. When an Advisory or Investigation is generated, Skylar Advisor snapshots the data-lake metrics for the affected devices and adds an AI-analyzed summary to the report. When you click the **[View metrics]** button, the data displays as a chart on the **Metrics Snapshot** pane.




For more information, see *Working an Advisory From Claim to Close*.

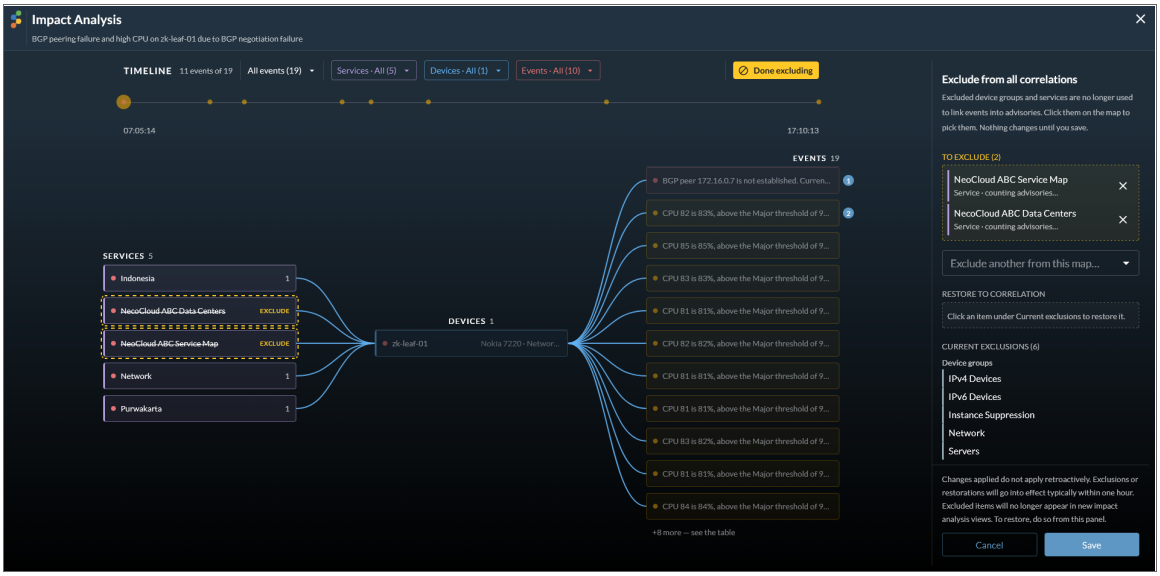
# Work Items: Advisory Suppression

In the **Work Items** view, you can enable **Advisory Suppression** to add a **[Suppress]** button to an Advisory that prevents Skylar Advisor from reporting on that kind of Advisory again.



# Work Items: Entity Exclusion

In the **Work Items** view, you can enable **Entity Exclusion** to exclude specific device groups and business services from Skylar Analytics as needed.

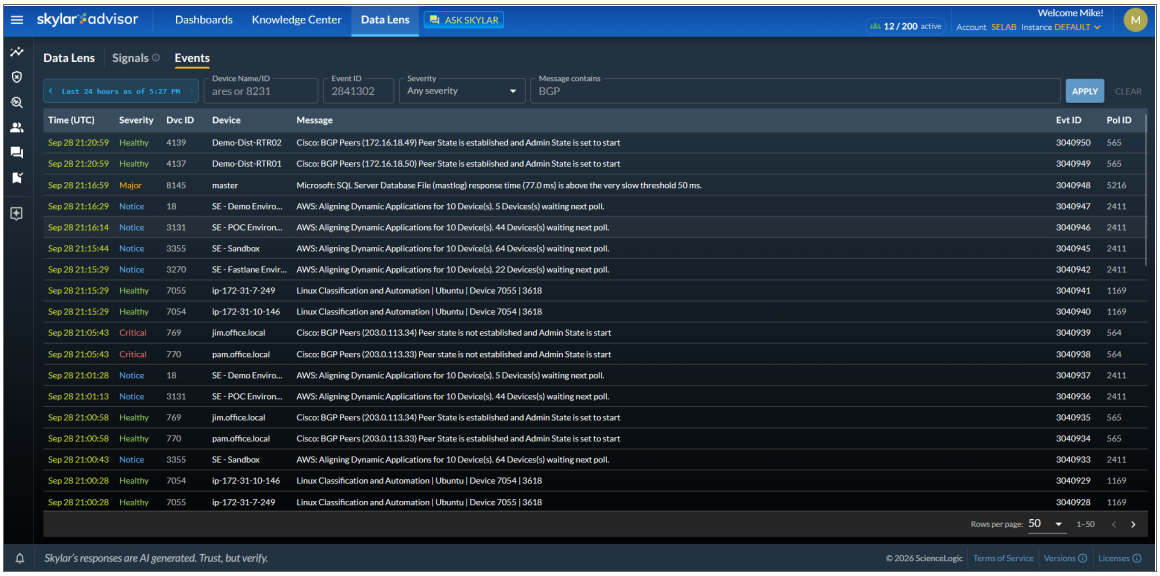



Use the **[Impact Analysis]** option in the menu at the top of an Advisory to include or exclude "catch-all" device groups and business services from event correlations. This option prevents unrelated events from being merged into a single Advisory and ensures that the entity map displays only meaningful connections.

Skylar Advisor automatically detects and excludes these catch-all groups by default, though you can manually override these exclusions or reset them to automatic via the user interface.

# Work Items: Data Lens Events

In the **Work Items** view, you can enable **Data Lens Events** to add an **[Events]** tab to the **Data Lens** page that displays a table of events. This feature provides you with better visibility into the specific events being received and considered for signal generation.




For more information, see *Managing Events on the Data Lens Page*.

# Work Items: SLO Docs

In the **Work Items** view, you can enable **SLO Docs** to track Service-Level Objectives (SLOs) based on documents that you upload to the **slo** folder in the Corpus, written in a set format.

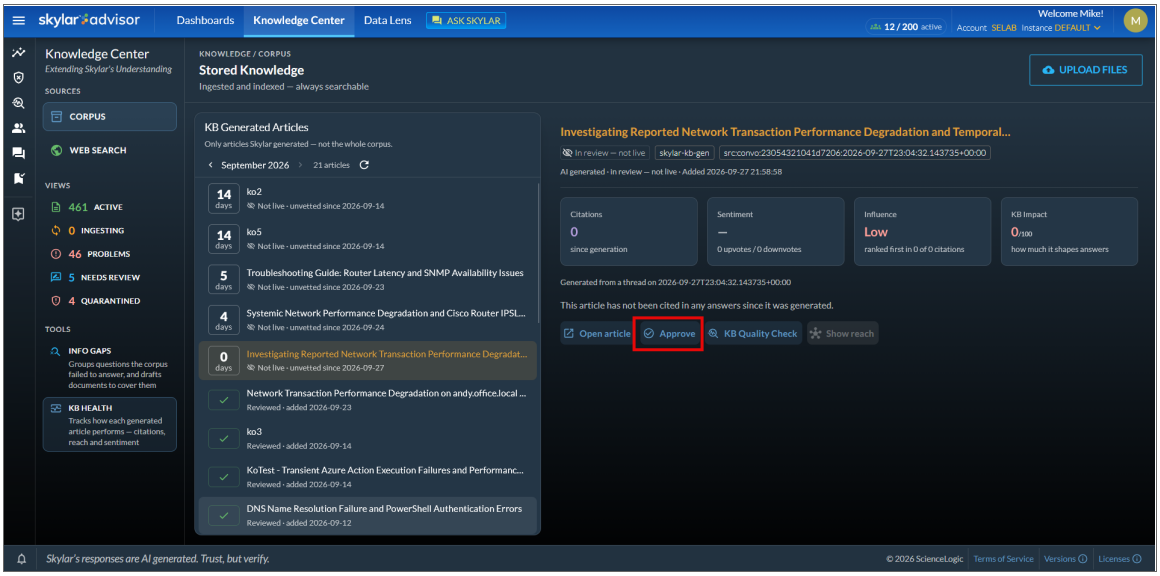
# Work Items: Auto-close

In the **Work Items** view, you can enable **Auto-Close** to enable a periodic background process that automatically closes idle Advisory, Prediction, and Investigation threads based on instance-specific thresholds. This feature helps maintain accurate workload and health metrics in the **Recent Activity** page by tagging and closing untouched threads.

# Platform: KB Generated Approval Flow

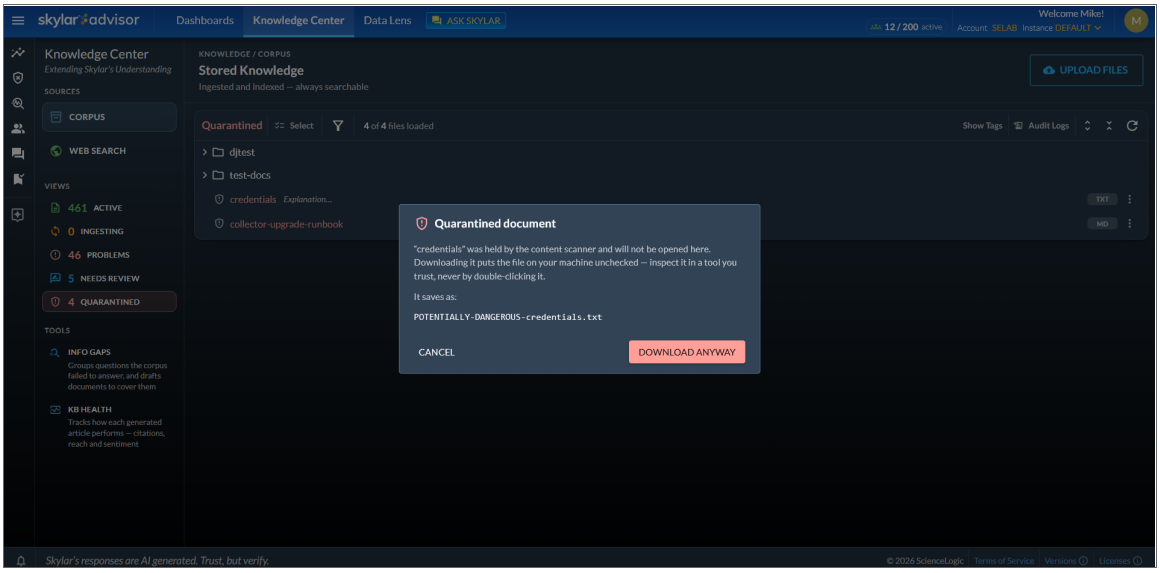
Enabling the **KB Generated Approval Flow** option on the **[Platform]** view of the **Preview Features** page enables a review workflow for Knowledge Base articles generated within the Corpus. This feature lets Subject Matter Experts (SMEs) approve or reject articles instead of immediately publishing the article. This gives reviewers a chance to read and approve the article first.

You can use the **[KB Health]** view to the **[Corpus]** tab on the **Knowledge Center** page to track citations, sentiment, influence, and impact for each generated article, with an AI quality check and a proposed rewrite. You can also click **[Approve]** to add a KB article to the Corpus.



# Platform: Knowledge Sentry

Enabling the **Knowledge Sentry** option on the **[Platform]** view of the **Preview Features** page implements a security gatekeeper layer for document uploads to the Corpus. Files are now automatically scanned for secrets, SQL injection patterns, and adversarial hidden text. Any suspicious files are quarantined until a subject-matter expert reviews them.




For more information, see *The Active View on the Corpus Tab*.

© 2003 - 2026, ScienceLogic, Inc.

All rights reserved.

ScienceLogic™, the ScienceLogic logo, and ScienceLogic's product and service names are trademarks or service marks of ScienceLogic, Inc. and its affiliates. Use of ScienceLogic's trademarks or service marks without permission is prohibited.

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information herein, the information provided in this document may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described herein at any time without notice.



800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010