



Skylar AI Product Documentation

Version 2.1.0: Skylar Advisor and Skylar Analytics

Table of Contents

Introduction to Skylar AI	9
What is Skylar AI?	10
Features of Skylar AI	10
Components of Skylar AI	10
Data Analyzed by Skylar AI	11
Configuring Skylar One for Skylar AI	12
Creating a Service Connection	13
Enabling Skylar AI for One or More Organizations	14
For Older Versions: Running the Skylar Management Tool	14
Enabling Skylar AI Event Policies	16
Importing ServiceNow Data into Skylar AI	17
Configuring Skylar AI System Settings	18
Logging In to the Skylar AI User Interface	19
Navigating the Skylar AI User Interface	19
Overview of Authentication in Skylar AI	20
Role-Based Access Control in Skylar AI	20
Elements of Role-Based User Accounts in Skylar AI	21
Configuring Multi-Factor Authentication	23
Configuring SSO Authentication with SAML	24
Creating Access Tokens for Users	25
Managing Dashboards	26
Other Dashboard Options	26
Update All Datasets	27
Installing and Configuring Skylar Advisor and Skylar Analytics	28
Cluster Requirements for Skylar Analytics and Skylar Advisor	29
Cluster Requirements for Skylar Analytics	29
Recommended Environment Settings for Skylar Analytics	29
Additional Cluster Requirements for Skylar Advisor	30
Third-Party Chart Dependencies for Skylar Analytics and Skylar Advisor	30
GPU and Kubernetes Configuration for Skylar Advisor	30
GPU Cluster Requirements and Configuration for Skylar Advisor	31

GPU Sizing Guide	31
NVIDIA GPU Operator	31
GPU Operator Deployment Summary	31
GPU Sharing Model: MIG + Time-slicing	31
Documentation	32
MIG configuration	32
H100 MIG ConfigMap	32
H200 MIG ConfigMap	33
Time-slicing Configuration	33
H100 Time-slicing ConfigMap	33
H200 Time-slicing ConfigMap	34
GPU Allocation Summary (8-GPU node)	34
Troubleshooting the Skylar Advisor Configuration	35
Verifying NVIDIA Driver Version	35
Seeding Model Weights	35
Verify seeded data	36
Required Infrastructure Dependencies	37
Ingress Controller	37
OpenTelemetry Operator	37
Recommended Infrastructure Components	38
Load Balancing	38
DNS and TLS Requirements	38
Optional Monitoring Integration	38
Registry Access Setup	38
Step 1: Obtain Registry Credentials	39
Step 2: Sign into Helm Registry (Required)	39
Step 3: Configure Kubernetes Registry Access	39
Step 4: Configure TLS Certificate (Optional)	39
Step 5: Download the Skylar AI Chart	40
Installation Process	40
Step 1: Environment Configuration	40
Step 2: Skylar Advisor Configuration	41

Ingress Controller-Specific Annotations	42
Step 3: Deploy the Skylar AI Platform	43
Scaling Profiles	43
Cloud-Specific Considerations	43
AWS Deployments	43
GCP Deployments	44
Azure Deployments	44
Validation and Access	45
Verify Deployment Status	45
Access the Platform	45
Monitoring Integration (Optional)	46
Prometheus Integration	46
Troubleshooting	46
Common Issues	46
Image Pull Failures	46
TLS Certificate Issues	46
DNS Resolution Issues	46
Ingress Controller Issues	47
Pod Startup Failures	47
Database Connection Issues	47
Large Upload/Download Issues	47
Cloud-Specific Issues	47
AWS	47
GCP	47
Azure	47
Useful Commands	47
Deployment Status Commands	47
Secret Verification Commands	48
Infrastructure Verification Commands	48
Registry Authentication Issues	48
TLS Certificate Issues	49
Ingress Controller Issues	49

Security Considerations	49
Security Best Practices	49
Third-Party Security Responsibilities	50
Backup and Recovery	50
Critical Data Components	50
ClickHouse Data Volumes	50
PostgreSQL Data Volumes	50
Backup Strategy Options	50
Volume-Level Backups	50
Application-Level Backups	51
Hybrid Approach	51
Backup Verification	51
Regular Testing Requirements	51
Monitoring and Alerting	51
Configuration Backup	51
Helm and Kubernetes Configuration	51
Support	52
Service Provider Administration for Skylar AI	53
First Login as a Service Provider User	54
Provisioning an Account	54
Adding an ODBC User	56
Introduction to Skylar Advisor	58
What is Skylar Advisor?	59
Logging In to the Skylar AI User Interface	59
Interacting with Skylar Advisor	61
The Skylar Advisor User Interface	62
Defining Your Persona Information	63
Managing Activities and Events on the Dashboards Page	65
The Recent Activity Tab	66
Viewing Content by Activity Type	68
Archiving an Activity	68
Bookmarking an Activity	69

Creating a Knowledge Base Article Based on an Activity	69
Using the Scratch Pad Feature	70
The Device Groups, Issues, and Organizations Tabs	70
Viewing Basic Event Details on a Tab	71
Viewing More Event Details on a Tab	72
Loading and Reviewing Your Documents on the Corpus Page	73
The Main Corpus Page	73
The Analytics Tab	76
Managing Events on the Event Lens Page	77
Investigating Advisories and Events	77
Filtering the Event Lens Page	78
Using the Ask Skylar Feature	79
Features of Ask Skylar	79
Additional Options	80
Viewing Team Analytics	82
Introduction to Skylar Analytics	85
What is Skylar Analytics?	85
Mapping Skylar One Dynamic Application Object Names to Skylar AI Columns	86
Skylar Analytics: Data Visualization and Data Exploration	89
What is Data Visualization?	90
Working with Datasets in Data Visualization	91
Components of a Dataset	91
Viewing the List of Datasets	92
Viewing the Contents of a Dataset	93
Viewing Dashboards and Charts in Data Visualization	93
Logging In to the Data Visualization Component	94
Default Skylar Analytics Dashboards	95
Recommended Datasets	97
Viewing Skylar Analytics Dashboards	98
Creating and Customizing Dashboards and Charts	100
Creating a Dashboard	100
Adding Contextual Cross-filtering to a Dashboard	105

Customizing a Dashboard	107
Icons for Chart Metrics	108
Customizing the Default Column Names for Charts	108
Adding and Upgrading Dashboards and Datasets	111
Other Dashboard Options	111
Update All Datasets	112
Data Exploration: Exporting Data to Skylar AI from Third-party Tools	112
Configuring Data Exploration with Power BI	112
Additional Resources for Skylar Analytics (Apache Superset Training)	115
Skylar Analytics: Anomaly Detection	117
What is Anomaly Detection?	118
How Anomaly Detection Works	118
Viewing Graphs and Data for Anomaly Detection	119
Enabling Thresholds and Alerts for the Anomaly Chart	120
Enabling Anomaly Detection Events for Specific Metrics	122
Viewing the Devices Being Monitored for Anomaly Detection Alerts	122
Enabling Anomaly Detection Events on the Anomaly Detection Page	123
Enabling Anomaly Detection Events on the Device Investigator Page	124
Creating an Event Policy for Anomalies	124
Using Anomaly-related Events to Trigger Automated Run Book Actions	126
Skylar Analytics: Predictive Alerting	128
What is Predictive Alerting?	129
How Predictive Alerting Works	129
Viewing Predictive Alerts in Skylar One	129
Using Predictive Alerts to Trigger Automated Run Book Actions	132



Getting Started with Skylar AI

Version 2.1.0

Chapter

1

Introduction to Skylar AI

Overview

Skylar AI is a software services suite powered by artificial intelligence (AI) that is designed to automatically manage and anticipate IT incidents. Skylar AI reasons over telemetry and the stored knowledge of an organization to deliver accurate insights, recommendations, and predictions.

Skylar AI includes Skylar Advisor or Skylar Analytics, or both if you have licenses for both applications.

The following table lists the minimum and recommended versions of ScienceLogic software required for Skylar AI and its components:

Product	Minimum Version	Recommended Version
Skylar One	12.3.2	12.5.1 or later
AP2	8.18.43-81 (Jelly Bean)	8.29.41 (Nougat) or later

NOTE: Skylar Advisor currently requires Skylar Analytics. For more information, see the [Skylar AI Release Notes](#).

This chapter covers the following topics:

What is Skylar AI?	10
Configuring Skylar One for Skylar AI	12
Importing ServiceNow Data into Skylar AI	17

What is Skylar AI?

Autonomic IT leverages artificial intelligence (AI), automation, and data to intelligently self-manage an entire IT stack. Autonomic IT drives autonomous businesses with rapid decision-making, cost-optimized scalability, and innovative experiences that empower organizations to focus on core innovation. The Skylar AI platform, which includes Skylar Analytics and Skylar Advisor, helps customers with their journey towards Autonomic IT.

Skylar AI is a software services suite powered by artificial intelligence (AI) that is designed to automatically manage and anticipate IT incidents. Skylar AI reasons over telemetry and the stored knowledge of an organization to deliver accurate insights, recommendations, and predictions.

Skylar One collects data and leverages Skylar AI to learn the patterns for a particular device metric over a period of time. Skylar uses the resulting data to build a device metric-specific model that is used to define a scope of expected behavior as well as anomalous data points.

Features of Skylar AI

Skylar AI is the engine that powers several different software components. The components in the Skylar family of services share the following characteristics:

- **Reactive.** When something fails, Skylar AI tells you in plain language what happened and how to fix it with relevant context.
- **Predictive.** Skylar AI alerts you in advance to an expected out-of-capacity condition.
- **Proactive.** Skylar AI accurately answers any question asked of it with context drawn from company knowledge sources, such as bugs, support tickets, Knowledge Base articles, and Product Documentation, and recommends next steps.

Skylar AI integrates seamlessly with the Skylar One platform and other IT management tools. You can interact with Skylar AI through these familiar environments, where it enhances existing workflows with AI-driven insights and automation capabilities. Skylar AI can send you alerts and notifications, which can be customized to suit individual preferences or organizational needs. These alerts help you stay informed about potential issues, ongoing incidents, or opportunities for optimization.

Components of Skylar AI

The Skylar AI family of services includes the following components:

- **Skylar Advisor**, an AI-powered feature within the Skylar AI ecosystem that provides proactive, tailored guidance based on real-time data.
- **Skylar Analytics**, an advanced reporting and custom analytics service that combines AI-powered analytics with deep data exploration and visualization.

Data Analyzed by Skylar AI

The following image shows the flow of data into and out of Skylar One and Skylar AI:



The following list contains some of the types of data that Skylar One can send to the Skylar AI engine, where the data is analyzed and used by Skylar Advisor and Skylar Analytics:

- Alert and event logs
- Asset data
- Availability data collected by Skylar One
- Business Service health, availability, and risk metrics from Skylar One
- Class-Based Quality-of-Service (CBQoS) metadata and CBQoS time series data
- Custom attributes for devices and for interfaces (interfaces require Skylar One version 12.3.14 or later, or 12.5.4 or later).
- Device category and device group data
- DCM(+R) relationships and topology
- Dynamic Application mapping and performance data, and journals
- Email round trip data
- Event data, including event policy name
- Interface billing metrics, metadata, and tags
- Internal Collection processes, collection services, and CBQoS
- Journal Dynamic Application data
- Metadata for web content, SOAP/XML transaction, and domain name monitors
- Presentation objects, labels, and metrics collected by Dynamic Applications that include SNMP index collection objects
- Process and service data
- Skylar One Agent data, including Gen 1 (Skylar One Distributed Environment) and Gen 3 (Skylar One Extended Architecture) agents
- Thresholds
- Topology data for L2, L3, CDP, LLDP, and ad-hoc relationships between devices

Configuring Skylar One for Skylar AI

NOTE: If you or your site administrator needs to install Skylar AI in your environment, see [Appendix A: Installing and Configuring Skylar AI](#).

Before you can start using Skylar AI components, you will need to perform the following configurations in Skylar One to enable the export of data from Skylar One to Skylar:

- [Create a Service Connection](#)
- [Enable Skylar Analytics for One or More Organizations](#)
- [Enabling Skylar AI Event Policies](#)

After you perform these configurations, you can access Skylar Analytics, Skylar Advisor, and other key Skylar AI components from the **Skylar AI** page () in Skylar One.

For information about setting up users, user groups, and user roles, see [Configuring Skylar AI System Settings](#).

IMPORTANT: If you are connecting Skylar AI with a Skylar One system, ScienceLogic strongly recommends that you always use the most recent Skylar One and AP2 releases. Using the most recent releases will ensure that your Skylar AI system has access to the latest datasets and features. For more information, see the [Skylar One Platform and AP2 Release Notes](#).

Creating a Service Connection

If you are using AP2 Mochi or later with your Skylar One system, you can create a service connection for the Skylar AI engine on the **Service Connections** page (Manage > Service Connections) in Skylar One. ScienceLogic strongly recommends upgrading to Mochi or later. For more information, see the [AP2 Mochi release notes](#).

The service connection enables communication between your Skylar One system and Skylar AI. This process replaces the [Running the Skylar Management Tool](#) process in previous releases of Skylar Analytics and Skylar One.

To create a Skylar AI Engine service connection:

1. In Skylar One, go to the **Service Connections** page (Manage > Service Connections).
2. Click **Add Service Connection** and select *Skylar AI Engine*. The **Create Skylar AI Engine Credential** window appears.
3. Complete the following fields:
 - **Name.** Type a name for the new service connection.
 - **API Key.** Add the access token for Skylar AI, which you can generate on the **Access Tokens** page ( > Instances > Access Tokens) on the Skylar AI User Settings. For more information, see [Creating Access Tokens for Users](#).
 - **Skylar AI Engine URL.** Add the URL for your Skylar AI system.
4. Click **[Save]**. The service connection is added to the **Service Connections** page, and a modal displays a link to the **Organizations** page, where you can enable Skylar Analytics for one or more organizations. See the following procedure for more information.
5. Refresh or reload the browser to add all updates to Skylar One.

NOTE: Newer releases of Skylar One include a **Status** and **Status Updated** column, along with a **Service Check** column that displays a **[Run Test]** button for "Skylar AI Engine" service connection types. Click **[Run Test]** to run a script to check the status of the Skylar AI connection and display the results in a modal.

Enabling Skylar AI for One or More Organizations

You will need to select one or more organizations in Skylar One that will share data with Skylar AI. This data will come from all of the devices in a selected organization. By default, the Skylar AI features are disabled.

You can see which organizations are currently sending data to Skylar AI by going to the **Organizations** page (Registry > Accounts > Organizations) and looking at the **Skylar AI Status** column for the organizations.

To enable Skylar AI with Skylar One organizations:

1. In Skylar One, go to the **Organizations** page (Registry > Accounts > Organizations) and click the check box for one or more organizations.
2. In the **Select Action** drop-down, select *Send Data from Selected Orgs to Skylar AI* and click **[Go]** to start sending data about the selected organizations to Skylar AI. The **Skylar AI Status** column for the selected organizations changes to *Enabled*.
3. If you want to override the hard-coded default options for exporting metadata to Skylar AI, you can add the JSON values in the **Skylar Options (JSON)** text field on the **Behavior Settings** page (System > Settings > Behavior). The values entered must be in the same JSON structure as those that appear in the **config.py** file:

```
DEFAULT_OPTIONS = { "metadata": { "intervals": {"snapshot": 60,
"cleared_events": 5}, "snapshot": { "batch_sizes": { "asset_basic":
500, "asset_ip_config": 10000, "perf_index_label": 5000, "device_
config": 20000, }, }, "cleared_events": { "query_range": 60, }, } }
```

NOTE: You can use the "Skylar: Failed" event policy in the "Skylar One Default Internal Events" PowerPack to raise an event in Skylar One if Skylar AI fails. If you installed Skylar One version 12.5.1 from an ISO file, this PowerPack will be included with the installation. If you are upgrading to Skylar One version 12.5.1, you will need to download the most recent "Skylar One Default Internal Events" PowerPack from the **PowerPacks** page on the [ScienceLogic Support Center](#) (Skylar One > PowerPacks) and then install the PowerPack.

NOTE: A generic, internal alert definition is generated in Skylar One to indicate whenever a Skylar AI connection is successful or encounters a problem.

For Older Versions: Running the Skylar Management Tool

If you are using a version of AP2 before Mochi, you will need to set up Skylar AI with the steps below for the Skylar SL1 Management Tool instead of the **Service Connections** page in Skylar One. ScienceLogic strongly recommends that you upgrade to Mochi. For more information, see the [AP2 Mochi release notes](#).

The Management Tool configures Skylar One data and Skylar One processes, and it starts monitoring the Skylar connection and configuration. The script is named `sl-otelcol-mgmt.py`, and it is included in the `sl-otelcol` RPM package.

To run the Skylar SL1 Management Tool:

1. Use the following command to run the Management script on the Database Server (a Skylar One Central Database or a Skylar One Data Engine):

```
sudo sl-otelcol-mgmt.py -vv skylar --skylar-all --skylar-endpoint
"<URL_for_skylar_system>" --skylar-api-key "<skylar-access-token>" --
ap2-feature-flags
```

where:

- `<URL_for_skylar_system>` is the URL for your Skylar AI system
- `<skylar-access-token>` is the access token for Skylar AI, which you can generate on the **[Access Tokens]** tab of the **Skylar Settings** page. For more information, see [Creating Access Tokens for Users](#).

This command configures the OpenTelemetry Collector, restarts services that export data, and checks that connectivity to the supplied endpoints is healthy.

You can also use the following configuration options if needed:

- `--verify-cert false`. Allows users in on-premises environments to connect to Skylar AI using self-signed certificates.
- `--ca-bundle /<path>/bundle.pem`. Allows users to specify a path to a `.pem` file and assign it to the `REQUESTS_CA_BUNDLE` environment variable.
- `--skylar-disable`. Stops all Skylar AI exports and services. This flag performs the same operations as the pause command (see step 3, below) and also removes any Skylar AI pages from the Skylar One user interface.

NOTE: If you have already run setup before and are not changing the connection details, you do not need to include `--skylar-endpoint "<URL_for_skylar_system>" --skylar-api-key "<skylar-access-token>"`.

In addition, `--ap2-feature-flags` is only needed the first time you install Skylar AI.

After successfully running the script, on the **System Logs** page (System > Monitor > System Logs), you will see "Info" messages for each configuration change (filter on `sl-otelcol-mgmt`). You will also see "Major" system log messages whenever connectivity fails for the Skylar endpoint or the OpenTelemetry Collector.

After data streams into the Data Visualization dashboards, and other Skylar AI components, they will populate with data. Please note that this process might take several minutes.

2. If you have run the setup script before, run the following command to enable Skylar AI and make sure that everything is working as expected:

```
sudo sl-otelcol-mgmt.py -vv skylar --skylar-all
```

3. If you need to pause Skylar AI, run the following command:

```
sudo sl-otelcol-mgmt.py -vv skylar
```

Pausing sets all Skylar AI toggle fields to disabled, restarts the event engine and data pull services to reflect the changed configuration, stops Skylar One managed services such as the Metadata Exporter, Alerts Poller, and **sl-otel-mgmt.timer**, and stops and disables the **sl-otelcol** systemd service.

4. To check the status of the installation, run the following command:

```
sudo sl-otelcol-mgmt.py -vv status
```

You should look for the following messages in the output:

```
----- checking feature toggles
SL_EXPORT_EVENTS = False
SL_EXPORT_METRICS = True
SL_EXPORT_CONFIG = True
----- checking services
sl-otelcol is enabled and running
----- checking connectivity
checking: Skylar endpoint is healthy
checking: local OTELCOLO endpoint is healthy
```

NOTE: If you need to turn off the Skylar AI connection, run the following command:

```
sudo sl-otelcol-mgmt.py -vv skylar --skip-status-service
```

5. Go to the previous procedure to specify the organizations you want to use for exporting data to Skylar.

Enabling Skylar AI Event Policies

In addition, the Predictive Alerting and Anomaly Detection components of Skylar Analytics require the "Skylar Analytics Event Policies" PowerPack. This PowerPack includes the Skylar One event policies from the "Skylar - Predictive events", "SL1: Skylar Anomaly Score Event Monitoring", and "SL1: Anomaly Index Event Monitoring" PowerPacks.

Older versions of this PowerPack were named "Skylar Predictive Analysis".

To install the "Skylar Analytics Event Policies" PowerPack:

1. Search for and download the PowerPack from the **PowerPacks** page at the [ScienceLogic Support Center](#) (Skylar One > PowerPacks, login required). Alternatively, you can use the link provided by ScienceLogic, if applicable.
2. In Skylar One, go to the **PowerPacks** page (System > Manage > PowerPacks), click **[Actions]**, and then click **[Import PowerPack]**.
3. Browse and select the downloaded PowerPack and click **[Import]**.
4. On the next screen, click **[Install]** and, when prompted for confirmation, click **[OK]**.
5. To confirm that the PowerPack was installed properly by go to the **Event Policies** page (Events > Event Policies) and type the word "predictive" into the **Name** search field. You should see a number of "Predictive Alerting" event policies.

The "Skylar Analytics Event Policies" PowerPack replaces the event policies in the following PowerPacks:

- Skylar Analytics
- Skylar - Predictive events
- SL1: Anomaly Score Event Monitoring
- SL1: Anomaly Index Event Monitoring

If you are still using event policies from those PowerPacks, ScienceLogic recommends that you complete the following steps to avoid disrupting existing events and Skylar AI links for predictive alerts:

1. In Skylar One, open each of the older PowerPacks on the **PowerPacks** page, go to the **Event Policies** tab, and click the remove or bomb icon (💣) for each event policy to remove the older policies.
2. Import and install the latest "Skylar Analytics Event Policies" PowerPack.
3. Edit the "Skylar Analytics Event Policies" PowerPack and click the lightning icon (⚡) next to all of the event policies that start with "anom" or "pred" are aligned to the new PowerPack.
4. Delete the older PowerPacks.

NOTE: Going forward, upgrades and installs will not require you to repeat these steps.

Importing ServiceNow Data into Skylar AI

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

The Skylar Advisor SyncPack retrieves KB article content from ServiceNow and sends those articles to Skylar AI for ingestion.

For more information, see the [Skylar Advisor SyncPack](#) manual.

Configuring Skylar AI System Settings

Overview

This chapter covers how to configure the various settings for Skylar AI products by using the **Skylar Settings** page to update your user profile, edit the user interface theme, and specify your password and multi-factor authentication settings for Skylar AI.

You can also use the Skylar AI menu () at the top left of every page to access additional pages for the current Skylar AI application, along with User Settings and Account Settings for your Skylar AI applications.

The content in this chapter corresponds with the user interface for Skylar AI version 2.0.0 or later.

NOTE: If you or your site administrator needs to install Skylar AI in your environment, see [Appendix A: Installing and Configuring Skylar AI](#).

This chapter covers the following topics:

Logging In to the Skylar AI User Interface	19
Navigating the Skylar AI User Interface	19
Overview of Authentication in Skylar AI	20
Creating Access Tokens for Users	25
Managing Dashboards	26

Logging In to the Skylar AI User Interface

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

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

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.

Clicking the **[Visit]** button for Skylar Predictive Alerting and Skylar Anomaly Detection opens new pages for those components in Skylar One.
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**.

Navigating the Skylar AI User Interface

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

- Click the Skylar AI menu icon () at top left to access additional pages for the current Skylar AI application, along with the **User Settings** page, the **Account Settings** pages, and the **Persona** page (Skylar Advisor only). The **User Settings** link takes you to the same **Profile** page that displays when you select **Skylar Settings** from the login page.
- Click the "Skylar AI" icon at top left to return to the Skylar AI login page.
- If you have multiple instances of Skylar AI running, you can switch between those instances by clicking the drop-down next to "Viewing Instance" at the top right.
- Click the user icon () at top right to view the email address and role for the current user in the Skylar AI user interface. On this drop-down menu, you can click the **[Sign Out]** button to sign out of this session.
- Click the **Versions** link in the footer of any page to view version information for the Skylar Settings user interface. From the footer, you can also click links to view the Terms of Service and information about licenses and open-source packages.

Overview of Authentication in Skylar AI

Authentication for Skylar AI has the following features:

- Multi-tenant support, including a Super User login for host management .
- Multiple instances that represent separate domains of data access within an account.
- Pre-defined roles for access control.
- Email and password (local accounts) authorization by default, and Security Assertion Markup Language (SAML) single sign-on (SSO) authorization configured as needed.
- Access tokens for integration with external tools.

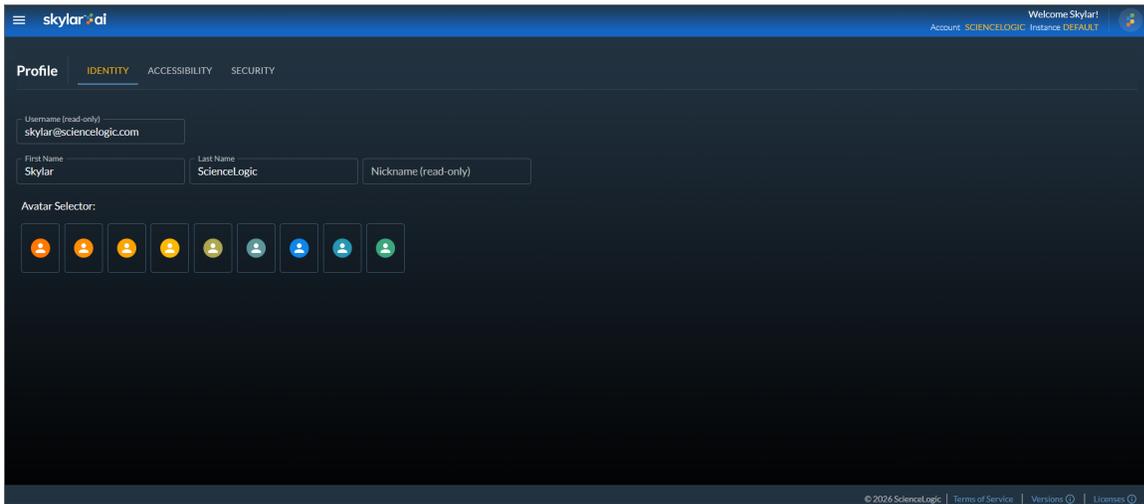
When a user is logged in to a Skylar AI component, that session uses the following rules:

- Email domains and SAML are configured per account (tenant).
- The first login for any new user starts with a prompt to create a new password.
- Logging into a user session requires either an email and password combination or a successful SAML2 redirect workflow.
- User passwords must be at least 15 characters long.
- New user passwords must be different from the last five passwords for this user.
- Users will be prompted to reset their passwords every 60 days.
- Users cannot change their passwords twice in 24 hours, and an error message displays if they try to do so.
- User sessions that have been idle for 15 minutes are automatically terminated. Administrator user sessions that have been idle for ten minutes are automatically terminated. An Admin user can adjust the timeout value on the **Authentication** page ( > Account Access > Authentication) in User Settings.
- If a user has three failed login attempts within a 15-minute interval, the user's account is locked for 15 minutes. An administrator user can unlock that user account from the **Edit User** dialog on the **Users** page ( > Account Access > Users) in User Settings.
- User accounts that have not been active for 35 days are automatically locked. Administrators can configure the inactivity timeout by using the **Inactive User Timeout** option on the **Authentication** page ( > Account Settings > Account Access > Authentication).

Role-Based Access Control in Skylar AI

To access the role-based access control settings, log into the Skylar AI user interface and click **Skylar Settings**. You can also access this page when you are logged in by clicking the Skylar AI menu icon () at top left and selecting *User Settings*.

The following image displays the **Profile** page in Skylar Settings:



Elements of Role-Based User Accounts in Skylar AI

An **account** in a Skylar AI system represents a complete Skylar AI configuration for a company. You can have multiple **instances** of Skylar AI in a single account.

An account contains a combination of the following:

- **Instances.** An instance is a logical store for account data. In other words, an instance is a complete Skylar AI system with its own set of login credentials and user settings. Examples of instances include a production instance, a QA instance, and a testing instance. An account can contain multiple instances. A **user** can view only the instances that are specified on the **groups** to which that user is a member. If only one instance is available, you will use the instance labeled "default".

On the **Available Instances** page ( > Instances > Available Instances) in Skylar Settings, you can view a list of instances for the current user. An admin user can also access the ODBC connection information for an instance, which contains the Microsoft Open Database Connectivity (ODBC) host, password, port, and user information for Data Exploration using ODBC on the ODBC Users page (Analytics Admin > ODBC Users > ODBC Connection Info).

If your system is using more than one instance, you will be able to select an instance after you log into Skylar Analytics.

- **Access Tokens.** You can add access tokens to connect Skylar AI with Skylar One or a third-party application. 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 to with an access token. You can set an expiration date for an access token, and you can also regenerate a token if needed.

On the **Access Tokens** page ( > Instances > Access Tokens) in Account Settings, you can view and add access tokens. For more information, see [Using Access Tokens for Users](#).

- **Users.** Everyone using Skylar AI should have their own user accounts. A user must belong to at least one **group**.

On the **Users** page ( > Account Access > Users) in Account Settings, you can view, edit, and add users for an account, and you can also reset the password for a user.

- **Groups.** A group controls which areas of Skylar AI a user can access. User groups are configured with a **role** and either a list of specific instances or *All* instances. If you select *All* instances, any instances that are created later are aligned with this group. Users can belong to more than one group. The active role for a user is based on the highest privilege from the groups aligned with that user.

On the **Groups** page ( > Account Access > Groups) in Account Settings, you can view, edit, and add user groups for an account.

- **Roles.** A role controls what features a user can access. You assign a role by creating or editing a user, and then aligning a group to that user. The active role for a user is based on the highest privilege from the groups aligned with that user.

The types of roles include the following:

- **Super User.** Assigned to the **admin** user that can manage all user accounts. The default login is **skylar@sciencelogic.com**. The Super User role can create and manage customer accounts, manage multiple instances, and set up SAML authentication for a customer.
- **Service Provider.** This role lets you provision new accounts and set up SSO for accounts. This role cannot edit the user with the Super User role.
- **Owner.** This role lets you monitor user management and user access, including the creation and assignment of instances. The **Owner** role also has the privilege to reset a user password.
- **Admin.** This role lets you perform day-to-day configuration tasks, including integrations and customization. The **Admin** role has access to the **SQL Lab** page in Skylar Analytics, which gives the Admin user access to all of the data in the instance. An Admin user can also add, edit, and delete users.
- **Editor.** For a future release, this role will let a user edit (create, update, and delete) objects, particularly incident type metadata.

NOTE: For this release of Skylar AI, the **Editor** and **Viewer** roles are the same. In future releases, these roles will be further defined.

- **Viewer.** For a future release, this role will give a user read-only access to Skylar AI. **Viewer** users can edit their own profiles.

- **Authentication.** Each Skylar AI system is configured by the **Owner** user by default for email authentication, which uses an email address and password combination. An **Owner** user can also set up authentication with a shared Identity Provider through the SAML2 protocol. If you enable single sign-on (SSO) with SAML, users that log in with the specified domain will be redirected to the SAML provider for this account.

On the **Authentication** page in Account Settings (☰ > Account Access > Authentication), you can configure SAML for this account. For more information, see [Configuring SSO Authentication with SAML](#).

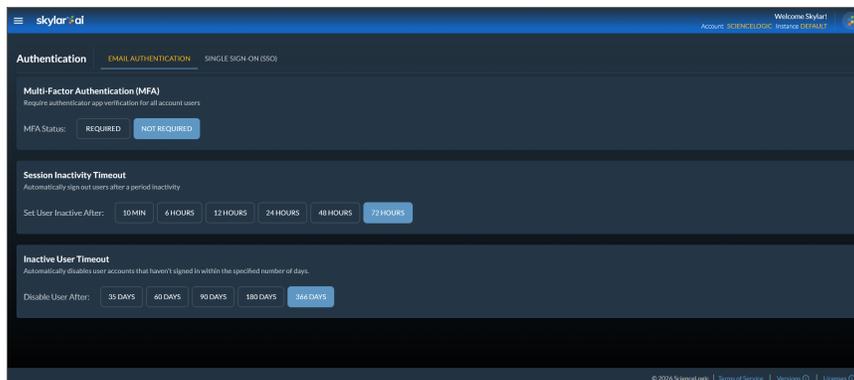
Configuring Multi-Factor Authentication

Skylar AI owner users can enable multi-factor authentication (MFA) for all users on a Skylar AI instance. On a site where MFA is required, users will need to scan a QR code or type in a secret key to set up MFA the next time they log in (if MFA is not already set up for their user accounts). Users will need to download an authenticator application like Google Authenticator or Authy to finish setting up MFA.

NOTE: MFA is set to Required by default for all new Skylar AI systems.

To require MFA for all users:

1. As an owner user, go to the **Authentication** page and go to the **[Email Authentication]** tab (☰ > Account Access > Authentication > Email Authentication):



2. Click *Required* next to the **MFA Status** field. The **Require Multi-Factor Authentication** modal appears.
3. Click **[Require MFA]**. All account users will need to set up an authenticator app on their next sign-in.

Users on a Skylar AI instance that does not have MFA required can set MFA for their account:

1. If needed, click the "Skylar AI" icon at top left to return to the Skylar AI login page and switch to **Skylar Settings**.

2. In Skylar Settings, go to the **[Security]** tab of the **Profile** page (User Preferences > Profile > Security).
3. Click the **[Configure MFA]** button.
4. Follow the instructions on the **Configure Multi-Factor Authentication** page to set up MFA.

TIP: Current MFA users can click the **[Reconfigure MFA]** button on the **[Security]** tab of the **Profile** page to change their MFA settings.

5. An owner user can click the *Reset MFA* option for a user on the **Users** page ( > Account Access > Users); for example, if a user loses his or her phone.

NOTE: A Super User can turn off MFA for all other users on a Skylar AI instance.

Configuring SSO Authentication with SAML

Users with the **Owner** role can configure single sign-on (SSO) authentication with SAML for their accounts. When SSO authentication with SAML is enabled, all logins for that customer will be authenticated by the SAML identity provider, such as Auth0, Okta, or JumpCloud.

If you have an issue with authenticating, you can contact ScienceLogic to disable SAML for the account and potentially reset the owner's local (non-SAML) password if needed.

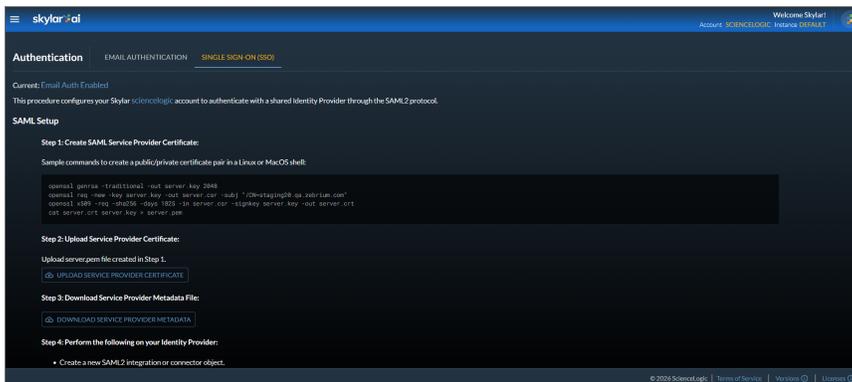
Before you can set up SSO authentication with SAML in Skylar AI, you will first need to create your user groups with your SAML identity provider if you do not already have them set up. Be sure to use the same names for your user groups with your SAML provider and with Skylar AI.

Do not switch the account to SAML until you have confirmed that the owner of the account has properly configured their SSO provider to recognize the Skylar platform.

To set up SSO Authentication with SAML in the Skylar AI user interface:

1. Go to the **Groups** page ( > Account Access > Groups) in Account Settings and click **[Add Group]**. The **Add Group** dialog appears.
2. Type a name for the group, select a role of *Admin*, and select one or more instances.
3. Click **[Add]**. The group is added to the **Groups** page.

4. Go to the **Authentication** page (☰ > Account Access > Authentication), click the **[Single Sign-On (SSO)]** tab, and review the instructions for SAML setup:



5. Follow steps 1-7 from the **[Single Sign-On (SSO)]** tab.

TIP: For step 7 on the **[Single Sign-On (SSO)]** tab, after you click the **[Set Authentication Style]** button, you can select *Enable SAML Test Mode for 10 minutes* to test the new authentication configuration. If the authentication works as expected, you can come back to step 7 and select *SAML* to make the configuration permanent.

Creating Access Tokens for Users

You can use the **Access Tokens** page in Skylar Settings (Instances > Access Tokens) to add access tokens to connect Skylar AI with Skylar One or a third-party application. A Skylar access token is used for authentication in place of an API key.

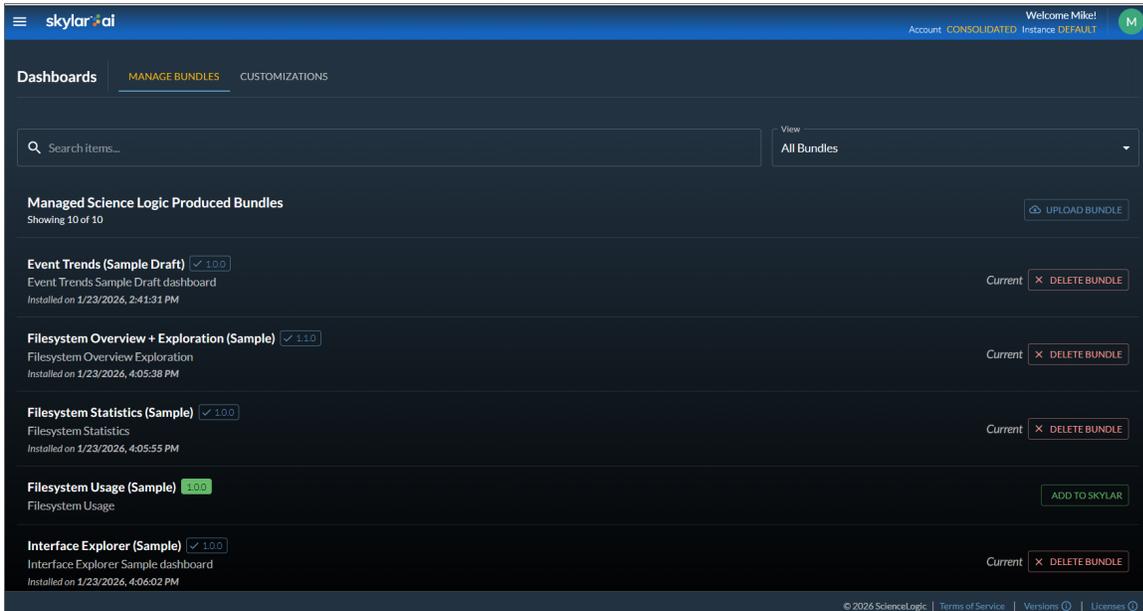
You can set an expiration date for an access token, and you can also regenerate a token if needed.

To create an access token:

1. Log in to Skylar AI and go to the **Access Tokens** page (☰ > 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*.
 - **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.

Managing Dashboards

A user with an owner role can manage the Skylar Analytics dashboards on the **[Manage Bundles]** tab of the **Dashboards** page ( > Analytics > Settings > Dashboard Mgmt):



You can search for dashboard bundles and sort the list of bundles by *All Bundles*, *Installed*, *Not Installed*, and *Updates Available*.

Owner users can install dashboards with **(Sample)** in their names by clicking the **[Add to Skylar]** button. These dashboards display a variety of visualization or chart configurations to show users different ways to display data, and users can reference these dashboards as examples. Many of the **(Sample)** dashboard layouts display multiple visualizations of the same raw data that would not typically be used at the same time on a production dashboard. These charts can also be copied and modified as needed, saving development time when building new dashboards.

ScienceLogic recommends that you keep the original versions of these **(Sample)** dashboard as unpublished draft dashboards and use them only for reference. For more information, see [Adding and Upgrading Dashboards](#).

Other Dashboard Options

You have the following options on the **[Manage Bundles]** tab of the **Dashboards** page:

- **[Upload Bundle]**. Upload a .zip file containing a Skylar Analytics dashboard that has been exported from a Skylar Analytics system. This feature lets you share dashboards that were created and then exported by other users, for example.
- **[Add to Skylar]**. Click this button to install a new dashboard for Skylar Analytics

- **[Upgrade Now]**. Click this button to upgrade an existing dashboard. *Current* displays next to a dashboard to show that you are running the most recent version of that dashboard.
- **[Delete Bundle]**. Remove this bundle from the Skylar Analytics system.

Update All Datasets

In addition, you can use the **[Sync Skylar Datasets]** button on the **[Customizations]** tab on the **Dashboards** page to update all of your datasets based on Skylar One PowerPacks, including PowerPacks that have been updated in Skylar One.

If all datasets have been updated, the button does not appear, and the text "Datasets are current" appears instead. This button is only available to owner users in Skylar AI.

Appendix

A

Installing and Configuring Skylar Advisor and Skylar Analytics

Overview

Deploying Skylar Advisor and Skylar Analytics requires Harbor for the registry and Helm for deployment. Skylar Advisor and Skylar Analytics are deployed using the same Helm umbrella chart, which contains over 20 microservices, databases, and supporting infrastructure components.

This chapter covers how to deploy Skylar Advisor and Skylar Analytics on a Kubernetes infrastructure using Harbor, Helm, and the user interfaces for Skylar Advisor and Skylar Analytics.

IMPORTANT: If you are installing only Skylar Analytics, look for "Important" notes like this at the beginning of a topic to determine if that topic is relevant to installing Skylar Analytics. Currently, you cannot use Skylar Advisor without Skylar Analytics.

This chapter covers the following topics:

<i>Cluster Requirements for Skylar Analytics and Skylar Advisor</i>	29
<i>GPU and Kubernetes Configuration for Skylar Advisor</i>	30
<i>Required Infrastructure Dependencies</i>	37
<i>Recommended Infrastructure Components</i>	38
<i>Registry Access Setup</i>	38
<i>Installation Process</i>	40

<i>Validation and Access</i>	45
<i>Troubleshooting</i>	46
<i>Security Considerations</i>	49
<i>Backup and Recovery</i>	50
<i>Support</i>	52

Cluster Requirements for Skylar Analytics and Skylar Advisor

Cluster Requirements for Skylar Analytics

- **Kubernetes:** Version 1.32 or later
- **Storage:** Configured EBS StorageClass capable of dynamic PV/PVC provisioning. Storage cannot be NFS-backed, as this is not supported by ScienceLogic databases.
- **Storage Capacity:** Recommended 1 TB or more total storage capacity. This requirement varies by tenant datapoints per minute (DPM) requirements.
- **Networking:** Currently, only IPv4 is supported. IPv6 is not supported at this time.
- **Deployment Options:**
 - Self-hosted Kubernetes clusters, such as bare metal or VMware
 - Cloud-managed Kubernetes services (EKS, GKE, AKS)

Recommended Environment Settings for Skylar Analytics

The following table lists the recommended configuration for your Skylar Analytics environment based on your expected volume of data, measured in datapoints per minute (DPM):

Size	CPU Cores	Memory (GB)	Disk (GB)	DPM Range
Small	8-10	32	500	0 - 30,000
Medium	25-30	176	650	30,000 - 215,000
Large	45-50	336	700	215,000 - 300,00

Additional Cluster Requirements for Skylar Advisor

IMPORTANT: If you are installing only Skylar Analytics, you can skip this topic.

- **Storage:** For Skylar Advisor, ScienceLogic recommends that you use RWX Storage configurations only for Skylar Advisor services. This is in addition to the EBS storage class required for Skylar Analytics.
- **Kubernetes:** Version 1.35 with ContainerD 2.1 or cri-o 1.32+ runtime.
- **GPU resources provisioning:** You need to dynamically allocate GPU resources that can be correctly provisioned to Kubernetes nodes under the **[Resources]** tab. Examples would be NVIDIA time-slicing or VMware vGPU.
- **GPU Driver.** Version 570.195.03 with CUDA 12.8.

For additional configuration information, see [GPU and Kubernetes Configuration for Skylar Advisor](#).

Third-Party Chart Dependencies for Skylar Analytics and Skylar Advisor

Skylar AI includes several third-party Helm charts from Bitnami, Chainguard, and other providers:

- **Bitnami Charts:** ClickHouse, PostgreSQL HA, and associated components
- **Maintenance Notice:** ScienceLogic validates and updates third-party chart versions with each Skylar AI release. However, customers are responsible for:
 - Security patching of third-party components between Skylar AI releases
 - Vulnerability management for non-Skylar AI components
 - Understanding the security posture of included third-party dependencies

For a complete list of external dependencies and their versions, review the **Chart.yaml** file in the Skylar AI chart. You will download this file in [Step 5: Download the Skylar AI Chart](#), below.

GPU and Kubernetes Configuration for Skylar Advisor

IMPORTANT: If you are installing only Skylar Analytics, you can skip this topic and go to [Required Infrastructure Dependencies](#).

This section describes the required GPU layout and Kubernetes configuration for deploying Skylar Advisor on an NVIDIA GPU-enabled Kubernetes cluster. It assumes basic familiarity with Kubernetes, Helm, and ConfigMaps.

GPU Cluster Requirements and Configuration for Skylar Advisor

GPU Sizing Guide

The following GPU requirements are based on the size of your Skylar AI deployment:

- Small deployments: 4 NVIDIA RTX 6000 GPUs
- Medium deployments: 4 NVIDIA H200 GPUs
- Large deployments: 8 NVIDIA H200 GPUs

NVIDIA GPU Operator

The NVIDIA GPU Operator manages the complete life cycle of GPU enablement on Kubernetes nodes, including NVIDIA drivers, container runtime integration, device plug-ins, GPU feature discovery, and monitoring components. Using the Operator eliminates the need to manually configure GPU nodes and ensures that GPUs are consistently exposed to Kubernetes workloads.

For more information, see the NVIDIA GPU Operator documentation:

<https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/latest/>.

GPU Operator Deployment Summary

The following table lists the deployment details for the GPU Operator:

Item	Value
Helm chart	gpu-operator
Helm repository	https://nvidia.github.io/gpu-operator
Name space	gpu-operator
Container runtime	containerd
NVIDIA driver version	570.195.03
MIG strategy	mixed
MIG configuration	External ConfigMap
Time-slicing configuration	External ConfigMap
DCGM exporter	Enabled

GPU Sharing Model: MIG + Time-slicing

Advisor uses Multi-Instance GPU (MIG) together with time-slicing to balance isolation, utilization, and performance:

- **MIG (Multi-Instance GPU)** partitions a physical GPU into multiple hardware-isolated GPU instances, each with dedicated memory and compute resources.
- **Time-slicing** oversubscribes a GPU or MIG instance by creating multiple schedulable replicas that share the same underlying GPU resources via time multiplexing.

This design allows Skylar Advisor to:

- Reserve most GPUs for high-performance inference
- Provide shared GPU capacity for OCR and document processing pipelines

Documentation

MIG: <https://docs.nvidia.com/datacenter/tesla/mig-user-guide/>

Time-slicing: <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/latest/gpu-sharing.html>

MIG configuration

Apply one of the following ConfigMaps depending on whether the node uses H100 or H200 GPUs:

H100 MIG ConfigMap

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: mig-config-h100
  namespace: gpu-operator
data:
  config.yaml: |
    version: v1
    mig-configs:
      custom-mig:
        # GPU 0: one large MIG slice
        - devices: [0]
          mig-enabled: true
          mig-devices:
            "7g.80gb": 1

        # GPU 1: two medium MIG slices
        - devices: [1]
          mig-enabled: true
          mig-devices:
            "3g.40gb": 2

        # GPUs 2-7: full GPUs (no MIG)
```

```
- devices: [2,3,4,5,6,7]
  mig-enabled: false
```

H200 MIG ConfigMap

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: mig-config-h200
  namespace: gpu-operator
data:
  config.yaml: |
    version: v1
    mig-configs:
      custom-mig:
        # GPU 0: one large MIG slice
        - devices: [0]
          mig-enabled: true
          mig-devices:
            "7g.141gb": 1

        # GPU 1: two medium MIG slices
        - devices: [1]
          mig-enabled: true
          mig-devices:
            "3g.71gb": 2

        # GPUs 2-7: full GPUs (no MIG)
        - devices: [2,3,4,5,6,7]
          mig-enabled: false
```

Time-slicing Configuration

Time-slicing is applied on top of the MIG devices to create additional schedulable replicas.

H100 Time-slicing ConfigMap

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: time-slicing-config-h100
  namespace: gpu-operator
```

```

data:
  any: |
    version: v1
    flags:
      migStrategy: mixed
    sharing:
      timeSlicing:
        renameByDefault: false
        failRequestsGreaterThanOne: false
        resources:
          - name: nvidia.com/mig-7g.80gb
            replicas: 2
          - name: nvidia.com/mig-3g.40gb
            replicas: 4

```

H200 Time-slicing ConfigMap

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: time-slicing-config-h200
  namespace: gpu-operator
data:
  any: |
    version: v1
    flags:
      migStrategy: mixed
    sharing:
      timeSlicing:
        renameByDefault: false
        failRequestsGreaterThanOne: false
        resources:
          - name: nvidia.com/mig-7g.141gb
            replicas: 2
          - name: nvidia.com/mig-3g.71gb
            replicas: 4

```

GPU Allocation Summary (8-GPU node)




```

name: seed-skylar-models
namespace: default
spec:
  restartPolicy: Never
  containers:
  - name: aws-cli
    image: amazon/aws-cli:2
    command: ["/bin/sh", "-lc"]
    args:
    - |
      echo "Starting model sync..."
      aws s3 cp --recursive s3://<bucket>/<prefix>/ /data/skylar-ai-
models/
      echo "Model sync complete."
    env:
    - name: AWS_ACCESS_KEY_ID
      value: "<ACCESS_KEY_PROVIDED_BY_SCIENCELOGIC>"
    - name: AWS_SECRET_ACCESS_KEY
      value: "<SECRET_KEY_PROVIDED_BY_SCIENCELOGIC>"
      # Optional if required
      # - name: AWS_DEFAULT_REGION
      #   value: "us-east-1"
    volumeMounts:
    - name: skylar-ai-models
      mountPath: /data/skylar-ai-models
  volumes:
  - name: skylar-ai-models
    persistentVolumeClaim:
      claimName: skylar-ai-models

```

Verify seeded data

After the pod completes:

```
kubectl logs seed-skylar-models
```

From any pod mounting the same volume:

```
ls -lah /data/skylar-ai-models
```

Required Infrastructure Dependencies

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Ingress Controller

HTTP/HTTPS traffic routing and SSL termination:

- **Required:** An ingress controller must be installed and configured
- **Recommended:** ingress-nginx controller
- **Alternatives:** AWS Load Balancer Controller, GKE Ingress, Azure Application Gateway, Traefik, or HAProxy Ingress

OpenTelemetry Operator

Observability data collection and processing:

- **Required:** OpenTelemetry Operator with custom image
- **Image:** `registry.scilo.tools/skylar/sl-otelcol:0.16`

```
# Install OpenTelemetry Operator with custom image
helm repo add open-telemetry https://open-telemetry.github.io/opentelemetry-helm-charts
helm repo update
helm install opentelemetry-operator open-telemetry/opentelemetry-operator \
  --set "manager.collectorImage.repository=registry.scilo.tools/skylar/sl-otelcol" \
  --set "manager.collectorImage.tag=0.16" \
  --set admissionWebhooks.certManager.enabled=false \
  --set admissionWebhooks.autoGenerateCert.enabled=true
```

If you are using Cert-Manager for your certificate management, you can remove the following lines:

```
--set admissionWebhooks.certManager.enabled=false \
--set admissionWebhooks.autoGenerateCert.enabled=true
```

Recommended Infrastructure Components

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Load Balancing

A load balancer solution is needed to distribute traffic across worker nodes:

- **Self-hosted:** F5 BIG-IP, HAProxy with floating IP, or MetalLB
- **AWS:** Application Load Balancer (ALB) or Network Load Balancer (NLB)
- **GCP:** Google Cloud Load Balancer
- **Azure:** Azure Load Balancer

DNS and TLS Requirements

Required:

- **FQDN:** A fully qualified domain name pointing to the load balancer
- **TLS Certificate:** Valid TLS certificate for the FQDN, provided as a Kubernetes secret

Optional Monitoring Integration

Prometheus-based Monitoring

- **Supported:** Integration with existing Prometheus deployments
- **Benefits:** Custom metrics export from Skylar AI services
- **Configuration:** Skylar AI services expose Prometheus metrics endpoints that can be scraped for enhanced monitoring

Registry Access Setup

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Step 1: Obtain Registry Credentials

1. Navigate to the registry at <https://registry.scilo.tools/>. The Harbor landing page appears.
2. Click **[Login via OIDC Provider]**.
3. Click **[Customer Login]**.
4. Enter your Salesforce credentials provided by ScienceLogic.
5. Click your name in the top right corner and select *User Profile*.
6. Click the copy icon in the **CLI Secret** field for use in the next steps.

Step 2: Sign into Helm Registry (Required)

You must authenticate with the Helm registry before proceeding with any chart operations. You authenticate with the registry using your Harbor CLI credentials (from the previous step):

```
helm registry login registry.scilo.tools \  
  --username <your-harbor-username> \  
  --password <your-cli-secret-from-step-1>
```

Step 3: Configure Kubernetes Registry Access

Create a Docker registry secret in your target namespace to enable image pulling:

```
# Create the deployment namespace  
kubectl create namespace skylar-production  
  
# Create Docker registry secret with Harbor CLI credentials  
kubectl create secret docker-registry harbor-creds \  
  --docker-server=registry.scilo.tools \  
  --docker-username=<your-harbor-username> \  
  --docker-password=<your-cli-secret-from-step-1> \  
  --docker-email=<your-email> \  
  --namespace=skylar-production
```

Step 4: Configure TLS Certificate (Optional)

NOTE: You can skip this step if you have automated certificate management (such as [cert-manager](#)) or deployment processes that handle TLS configuration.

Create a TLS secret with your provided certificate and private key:

```
kubectl create secret tls skylar-tls-secret \  
  --cert=path/to/your/certificate.crt \  
  --key=path/to/your/private.key \  
  --namespace=skylar-production
```

Step 5: Download the Skylar AI Chart

Use Helm to pull the Skylar AI **Chart.yaml** file from the registry (registry login required from Step 2), and then decompress the charts:

```
helm pull oci://registry.scilo.tools/skylar/skylar-charts
```

```
tar -xvf skylar-charts-x.x.x.tgz
```

Installation Process

IMPORTANT: Unless noted otherwise, the information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Step 1: Environment Configuration

Create an environment-specific override file to customize the deployment. This file should be developed in collaboration with ScienceLogic Engineering to ensure proper configuration for your environment.

Example **override.yaml** structure:

```
# File: /path/to/override.yaml  
global:  
  # Registry configuration  
  registry:  
    hostname: "registry.scilo.tools"  
    username: "<your-harbor-username>"  
    password: "<your-cli-secret-from-step-1>"  
  
  # Image pull secrets reference  
  imagePullSecrets:  
    - name: harbor-creds  
  
  # Kubernetes cluster IP range. Used for network controls on database  
  # access in dataviz clickhouse  
  kubernetes_cidr_ranges: 100.64.0.0/16
```

```

# Ingress configuration
ingress:
  hostname: "skylar.yourdomain.com" # Your provided FQDN
  protocol: "https"
  className: "nginx" # Adjust based on your ingress controller
  tls:
    - secretName: skylar-tls-secret # Your TLS certificate secret
      hosts:
        - "skylar.yourdomain.com"

# Enable/disable platform components
enablePlatform: true
enableAnalytics: true
enableMonitoring: false # Set to true if integrating with existing Prometheus

# Component-specific configurations
clickhouse:
  persistence:
    size: "500Gi"

postgresql-ha:
  postgresql:
    persistence:
      size: "100Gi"

```

Step 2: Skylar Advisor Configuration

IMPORTANT: If you are installing only Skylar Analytics, you can skip this topic and go to [Step 3: Deploy the Skylar AI Platform](#).

If you are enabling Skylar Advisor, you will want to be aware of the following requirements:

- A GPU sizing file is necessary. Work with ScienceLogic to determine the best configuration.
- Skylar Advisor requires a storageClass that allows RWX, because Skylar Advisor requires a shared filesystem across services.

For Skylar Advisor, you will need to add the following settings to the **override.yaml** file:

```
global:
  enableAdvisor: true
  enableAdvisorOCR: true

#If enabling Skylar Advisor then set StorageClass which allows RWX.
skylar-advisor-api:
  persistence:
    storageClass:
```

Ingress Controller-Specific Annotations

Skylar AI includes default ingress annotations optimized for NGINX Ingress Controller. These annotations handle proxy buffering, timeouts, and request size limits that are essential for operating Skylar AI:

Skylar AI recommends the following ingress annotations optimized for [NGINX Ingress Controller](#). These annotations handle proxy buffering, timeouts, and request size limits that are essential for Skylar AI operation:

```
nginx.org/client-max-body-size: 256m
```

```
nginx.org/proxy-buffer-size: 512k
```

```
nginx.org/proxy-buffering: "on"
```

```
nginx.org/proxy-buffers: "4 512k"
```

```
nginx.org/proxy-max-temp-file-size: 1024m
```

```
nginx.org/proxy-read-timeout: "300s"
```

Skylar Advisor Only:

```
nginx.org/proxy-hide-headers: "Content-Security-Policy"
```

```
nginx.org/location-snippets: |
```

```
  add_header Content-Security-Policy "frame-ancestors *;" always;
```

You will need to enable snippets to allow the Skylar Advisor config: <https://docs.nginx.com/nginx-ingress-controller/configuration/ingress-resources/advanced-configuration-with-snippets/>.

For Non-NGINX Ingress Controllers, you will need to override these annotations with equivalent configurations for your ingress controller.

IMPORTANT: Make sure that your ingress controller configuration supports:

- Large request body sizes (256 MB minimum)
- Extended read timeouts (300 seconds minimum)
- Proper proxy buffering for large responses

ScienceLogic is actively working on reducing these requirements with upcoming releases.

Step 3: Deploy the Skylar AI Platform

```
# Deploy skylar with scaling profile and custom overrides
helm upgrade --install skylar-production \
  oci://registry.scilo.tools/skylar/skylar-charts \
  --namespace skylar-production \
  --values envs/scaling/small.yaml \
  --values /path/to/override.yaml

#If enabling advisor add
--value /path/to/GPU_Scaling_File.yaml
```

Scaling Profiles

Choose the appropriate scaling profile based on your environment and datapoints per minute (DPM). These scaling profiles can be located within the downloaded Helm chart:

- `envs/scaling/small.yaml`: Development or testing environments (0-30,000 DPM)
- `envs/scaling/medium.yaml`: Small production deployments (30,000-215,000 DPM)
- `envs/scaling/large.yaml`: Medium production deployments (215,000-300,000 DPM)
- `envs/scaling/xlarge.yaml`: Large production deployments (300,000 - 900,000 DPM)

For additional information, see [Recommended Environment Settings for Skylar Analytics](#).

Cloud-Specific Considerations

AWS Deployments

```
# Additional AWS-specific configuration
global:
  ingress:
    annotations:
```

```

    service.beta.kubernetes.io/aws-load-balancer-type: "nlb"
    service.beta.kubernetes.io/aws-load-balancer-backend-protocol: "tcp"

# Storage configuration
clickhouse:
  persistence:
    storageClass: "gp3"

postgresql-ha:
  postgresql:
    persistence:
      storageClass: "gp3"

```

GCP Deployments

```

# Additional GCP-specific configuration
global:
  ingress:
    annotations:
      kubernetes.io/ingress.class: "gce"
      kubernetes.io/ingress.global-static-ip-name: "skylar-ip"

# Storage configuration
clickhouse:
  persistence:
    storageClass: "pd-ssd"

postgresql-ha:
  postgresql:
    persistence:
      storageClass: "pd-ssd"

```

Azure Deployments

```

# Additional Azure-specific configuration
global:
  ingress:
    annotations:
      kubernetes.io/ingress.class: "azure/application-gateway"

# Storage configuration

```

```
clickhouse:
  persistence:
    storageClass: "managed-premium"

postgresql-ha:
  postgresql:
    persistence:
      storageClass: "managed-premium"
```

Validation and Access

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Verify Deployment Status

```
# Check all pods are running
kubectl get pods -n skylar-production

# Verify ingress configuration
kubectl get ingress -n skylar-production

# Check TLS certificate is properly configured
kubectl describe ingress -n skylar-production

# Check image pull secrets are working
kubectl describe pod <any-skylar-pod> -n skylar-production | grep -A5
"Events:"
```

Access the Platform

1. Verify DNS Resolution:

```
nslookup skylar.yourdomain.com
```

2. Navigate to the Skylar AI user interface using your provided FQDN, such as <https://skylar.<yourdomain>.com>.
3. Log in for the first time with the default email of skylar@sciencelogic.com.
4. Set a password for your first login.

Monitoring Integration (Optional)

Prometheus Integration

If you have an existing Prometheus setup, you can configure it to scrape metrics from Skylar AI services:

```
# In your override.yaml - enables metrics endpoints
global:
  enableMonitoring: true
```

This will expose Prometheus metrics endpoints on Skylar services that can be scraped by your existing Prometheus deployment. Configure your Prometheus to discover and scrape these endpoints based on your service discovery method, such as Kubernetes service discovery or static configurations.

You can ingest your scraped metrics from Prometheus into Skylar One leveraging the "SL1 Prometheus" PowerPack.

Troubleshooting

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Common Issues

Image Pull Failures

- Verify registry secret is correctly configured.
- Check CLI secret is still valid in Harbor.
- Ensure namespace has access to the registry secret.
- Verify Helm registry login was successful.

TLS Certificate Issues

- Verify the TLS secret contains valid certificate and key.
- Check certificate matches the FQDN.
- Ensure certificate is not expired.

DNS Resolution Issues

- Verify FQDN points to the correct load balancer IP.
- Check DNS propagation if recently updated.

Ingress Controller Issues

- Verify ingress controller is running and healthy.
- Check ingress controller logs for errors.
- Ensure ingress class name matches your controller.
- Verify ingress annotations are compatible with your controller.

Pod Startup Failures

- Check resource constraints and storage availability.

Database Connection Issues

- Ensure ClickHouse and PostgreSQL pods are healthy.

Large Upload/Download Issues

- Verify ingress controller supports large request bodies (256 MB and up).
- Check proxy timeout configurations.
- Ensure proper buffering settings are applied.

Cloud-Specific Issues

AWS

- Check IAM permissions for EBS/EFS access

GCP

- Verify service account permissions for persistent disks

Azure

- Ensure proper RBAC for storage resources

Useful Commands

Deployment Status Commands

```
# Check deployment status
helm status skylar-prod -n skylar-production
```

```
# View pod logs for troubleshooting
kubectl logs <pod-name> -n skylar-production

# Describe pod for detailed information
kubectl describe pod <pod-name> -n skylar-production
```

Secret Verification Commands

```
# Verify registry secret
kubectl get secret harbor-creds -n skylar-production -o yaml

# Verify TLS secret
kubectl get secret skylar-tls-secret -n skylar-production -o yaml

# Check certificate details
kubectl get secret skylar-tls-secret -n skylar-production -o jsonpath='{.data.tls\.crt}' |
base64 -d | openssl x509 -text -noout
```

NOTE: Line breaks were added to some of the lines of code, above, to allow the code sample to display properly.

Infrastructure Verification Commands

```
# Check ingress controller status
kubectl get pods -n ingress-nginx # Adjust namespace based on your setup

# Check persistent volume claims
kubectl get pvc -n skylar-production

# Test registry connectivity
kubectl run test-registry --image=registry.scilo.tools/skyilar/skyilar-
charts:latest \
--dry-run=client -o yaml | kubectl apply -f -
```

Registry Authentication Issues

If you encounter image pull errors:

- **Verify CLI Secret:** Ensure the CLI secret has not expired in Harbor.
- **Test Registry Login:** Re-authenticate with helm registry login.
- **Check Secret Format:** Verify the Kubernetes secret was created correctly.
- **Namespace Access:** Ensure the secret exists in the correct namespace.

TLS Certificate Issues

If you encounter TLS-related problems:

- **Certificate Validation:** Verify the certificate is valid for your FQDN.
- **Certificate Format:** Ensure the certificate is in PEM format.
- **Certificate Chain:** Include intermediate certificates if required.
- **Private Key Match:** Verify the private key matches the certificate.

Ingress Controller Issues

If ingress resources are not working:

- **Controller Status:** Verify the ingress controller is running.
- **Class Name:** Ensure the ingress class name matches your controller.
- **Controller Logs:** Check ingress controller logs for configuration errors.
- **Service Endpoints:** Verify backend services are healthy and have endpoints.
- **Annotation Compatibility:** Ensure ingress annotations are supported by your controller.

Security Considerations

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Security Best Practices

- **Registry Credentials:** Store CLI secrets securely and rotate them regularly.
- **TLS Certificates:** Ensure certificates are from trusted CAs and renewed before expiration.
- **Network Policies:** Consider implementing Kubernetes network policies to restrict inter-pod communication.
- **RBAC:** Configure appropriate role-based access controls for the Skylar AI namespace.
- **Ingress Security:** Configure appropriate security headers and rate limiting on your ingress controller.
- **Secrets Management:** Use external secret management solutions, such as HashiCorp Vault, or AWS Secrets Manager for production.

- Cloud Security: Follow cloud provider security best practices.
- Third-Party Components: Monitor security advisories for included third-party components.

Third-Party Security Responsibilities

While ScienceLogic validates and updates third-party chart versions with each Skylar AI release, customers should:

- Monitor Security Advisories: Stay informed about security issues in included third-party components.
- Plan for Updates: Be prepared to upgrade Skylar AI when security patches are available.
- Vulnerability Assessment: Include third-party components in security scanning and assessment processes.
- Risk Management: Understand the security posture of all included dependencies.

Backup and Recovery

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Critical Data Components

Skylar AI stores critical data in the following persistent volumes, which must be backed up regularly.

ClickHouse Data Volumes

- Contains analytics data, metrics, and time-series information.
- Recommended backup frequency: Daily with 30-day retention.

PostgreSQL Data Volumes

- Contains application metadata, user data, and configuration.
- Recommended backup frequency: Daily with 30-day retention.

Backup Strategy Options

Volume-Level Backups

- Use cloud provider snapshot capabilities, such as AWS EBS, GCP Persistent Disks, or Azure Disks.
- Implement Kubernetes volume snapshots using CSI drivers.
- Consider third-party backup solutions like Velero for comprehensive cluster backup.

Application-Level Backups

- Database-native backup tools for ClickHouse and PostgreSQL.
- Export application configuration and secrets.
- Backup Helm chart values and deployment configurations.

Hybrid Approach

- Combine volume snapshots for fast recovery with application-level backups for granular restore options.
- Implement both local and off-site backup storage for disaster recovery.

Backup Verification

Regular Testing Requirements

- Monthly: Test backup restoration procedures in non-production environment.
- Quarterly: Full disaster recovery simulation.
- Document and validate recovery time objectives (RTO) and recovery point objectives (RPO).

Monitoring and Alerting

- Monitor backup job completion and success rates.
- Alert on backup failures or missing backups.
- Verify backup accessibility and integrity.

Configuration Backup

Helm and Kubernetes Configuration

- Export and store Helm values files.
- Backup Kubernetes secrets and ConfigMaps.
- Maintain version-controlled infrastructure as code.

CAUTION: Store configuration backups containing secrets in encrypted storage with restricted access. Never commit secrets to version control systems.

For detailed backup implementation guidance, consult your cloud provider's backup documentation and the Kubernetes Backup Best Practices.

Support

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

For deployment assistance, configuration guidance, or troubleshooting support, contact Skylar AI Enablement. They can provide:

- Environment-specific override file templates.
- Scaling recommendations based on your requirements.
- Custom configuration for enterprise integrations.
- Cloud-specific deployment guidance.
- Registry access troubleshooting.
- TLS certificate configuration assistance.
- Ingress controller configuration guidance.
- Post-deployment optimization.
- Third-party component guidance (limited to integration aspects only).

NOTE: Support for third-party components, such as Bitnami charts, is limited to integration and configuration guidance. For issues specific to these components, consult their respective documentation and support channels.

Appendix

B

Service Provider Administration for Skylar AI

Overview

This chapter explains the different tasks that a user with the **Service Provider** (MSP) role can perform in Skylar AI. A **Service Provider** user can provision new accounts and add an ODBC user for the Data Exploration component of Skylar Analytics.

IMPORTANT: This appendix is intended only for Skylar AI users with a role of **Service Provider**.

This chapter covers the following topics:

<i>First Login as a Service Provider User</i>	54
<i>Provisioning an Account</i>	54

First Login as a Service Provider User

When you first log in to your Skylar AI system, you will use the default service provider name of **provider@sciencelogic.com**. The user interface will prompt you to set the ScienceLogic user password before your first login can continue.

After you log in for the first time, you will see a link for just the **Skylar Settings** page on the **Skylar AI** home page. Click that link to start setting up new accounts.

After you set up a licensed version of Skylar Analytics in Skylar Settings, you will see an **Analytics** link on this home page. After you configure Skylar Advisor, you will see an **Advisor** link here.

Provisioning an Account

To provision a new customer account:

1. Log into Skylar AI as a Super User.
2. Click the **All Accounts** drop-down at the top of the page and select the customer account you want to provision. Skylar AI updates the contents of the page to show data for that account.
3. Click the **All Accounts** drop-down again and click **[Provision New Account]** at the bottom of the drop-down. The **Provision a New Account** wizard appears.
4. On the **Account Name** page, enter the name of the account in the **Account Name** field using only lower-case alphanumeric characters and type a user-friendly name for the account in the **Display Name** field.
5. Click **[Next]**.
6. On the **Account Owner** page, specify the **First Name**, **Last Name**, and **Email** for the first user of the new account.
7. When you enter an email address, Skylar AI automatically adds the domain name from that email to the **Claim Email Domain** field. If you are not using single sign-on (SSO) through Security Assertion Markup Language (SAML), or if you do not want to claim that email domain, you can remove the text from this field. This field is optional.

NOTE: If you claim this domain, when SSO through SAML is enabled, users that log in with this domain will be redirected to the SAML provider for this account.

8. Click **[Next]**.
9. On the **Licenses** page, select **Skylar Advisor** if you want to enable Skylar Advisor for this account. You have two additional options:
 - **Enable Advisories.** Pairs Skylar Advisor to the data lake to mine telemetry.
 - **Legacy Instance ID.** For data lake connections, use the legacy instance ID format.

10. Select **Skylar Analytics** to enable Skylar Analytics for this account. You have two additional options:
 - **Consolidate Instance Reporting.** Enables reporting from multiple Skylar One systems on this Skylar AI system. When this option is enabled, data from all instances under the account will be merged into a single database called **reporting**.

IMPORTANT: You must enable this option during the account provisioning wizard, as you cannot change it or enable it after the wizard is complete. When enabled, a note indicating that consolidated instance reporting is active appears on the **Dashboards** page ( > Analytics Admin > Dashboards) in Skylar Settings.

- **Enable ODBC.** Creates an ODBC connection for the Data Exploration component of Skylar Analytics. You will need to add the IP addresses for your ODBC client in the **ODBC Client IP Ranges** field. Be sure to add the public-facing IP address for the ODBC client to the "allow list" for Skylar AI. Click **[Next]**.

NOTE: You will need to add any ODBC users after you complete this procedure. For more information, see [Adding an ODBC User](#).

11. On the **Instances** page, enter a name for the instance associated with this account using only lowercase alphanumeric characters. Examples include a production, QA, or testing instance. You can also use *default* as the instance name. Click **[Next]**.
12. On the **Summary** page, review your settings and click **[Begin Provisioning]** to continue setting up the account. The provisioning process starts, and Skylar AI switches to the new account.

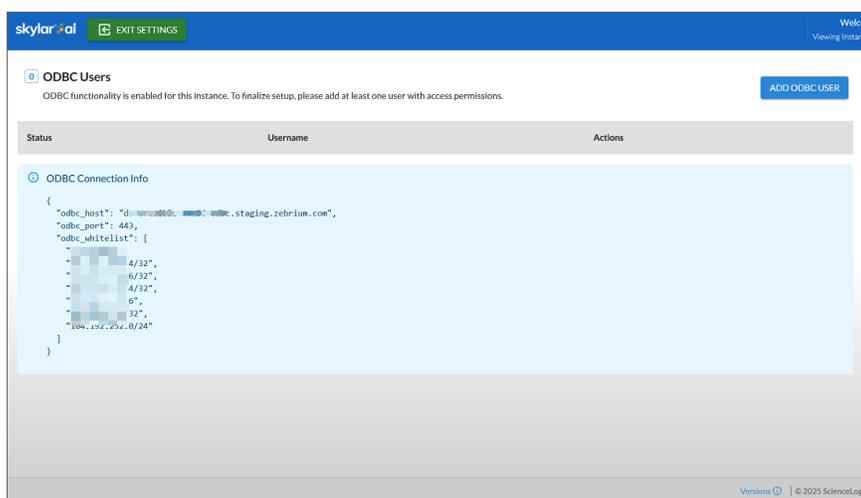
NOTE: When the account is set up, you will need to give the email address you used in step 7 to the first user. On first login, the new user will be prompted to change their password.

13. You can use the following links on the Super User toolbar to further configure this account:
 - **Account Info.** Lets you change the display name for this account.
 - **Licenses.** Lets you change the Skylar Advisor and Skylar Analytics licensing options. These options are the same options in steps 11-12, above.
 - **Instances.** Displays the current instances for this account, and lets you add another instance to the account.
 - **Authentication.** Lets you select the type of authentication you want to use for this site: email/password or SAML. For more information, see [Configuring SSO Authentication with SAML](#).
14. To set up other security and authentication settings, go to the **Authentication** page ( > Account Settings > Account Access > Authentication). For more information, see [Configuring Skylar AI System Settings](#).

Adding an ODBC User

When you create a new ODBC connection for the Data Exploration component of Skylar Analytics, you will need to create the ODBC user or users and set their password from the **[ODBC Users]** tab on the **Skylar Settings** page. You can add, edit, disable, and delete ODBC users through the **Skylar Settings** page.

1. In Skylar Settings, go to the **ODBC Users** page ( > Analytics Admin > ODBC Users). This page displays the ODBC connection information for the Skylar AI system:



2. Click the **[Add ODBC User]** button. The **Add ODBC User** window appears.
3. In the **Username** field, type a name after the "odbc_" prefix, and then type the password in the two **Password** fields.
4. Click the **[Add]** button. The ODBC user is added to the **ODBC Users** page.



Skylar Advisor

Version 2.1.0

Chapter

1

Introduction to 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 Skylar AI, see [Introduction to Skylar AI](#).

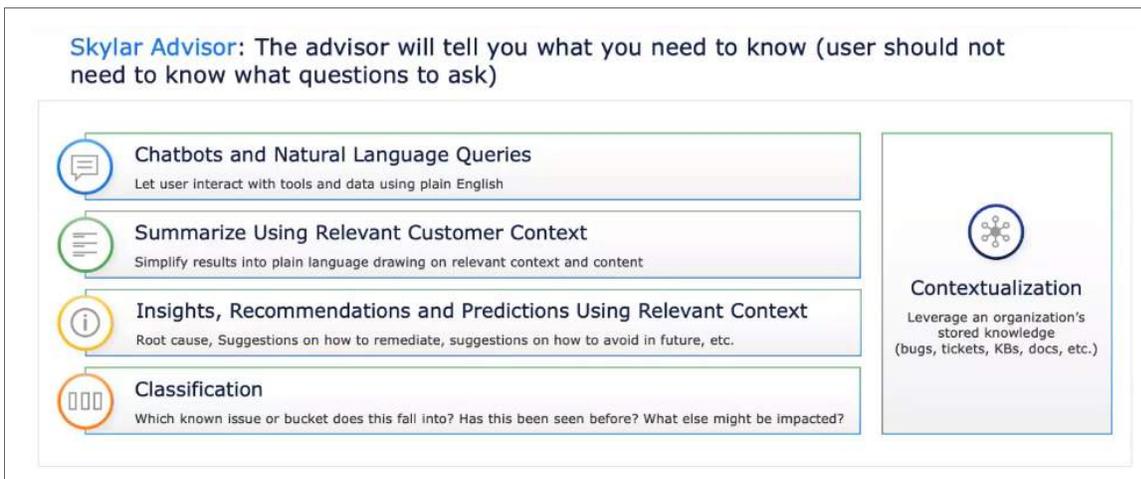
This chapter covers the following topics:

What is Skylar Advisor?	59
Logging In to the Skylar AI User Interface	59

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, Claude, or Microsoft Copilot, you can also use the Ask Skylar feature in Skylar Advisor to start a conversation with Skylar AI to get more information.

Logging In to the Skylar AI User Interface

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

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

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.

Clicking the **[Visit]** button for Skylar Predictive Alerting and Skylar Anomaly Detection opens new pages for those components in Skylar One.

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**.

Chapter

2

Interacting with Skylar Advisor

Overview

This chapter covers how to interact with Skylar Advisor to monitor diagnostic data for your Skylar One system using Skylar AI. You can upload files to Skylar Advisor to train it on your company workflows, policies, and processes.

You can also interact with Skylar Advisor by reviewing and following its proactive suggestions, called "Advisories." You can also 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.

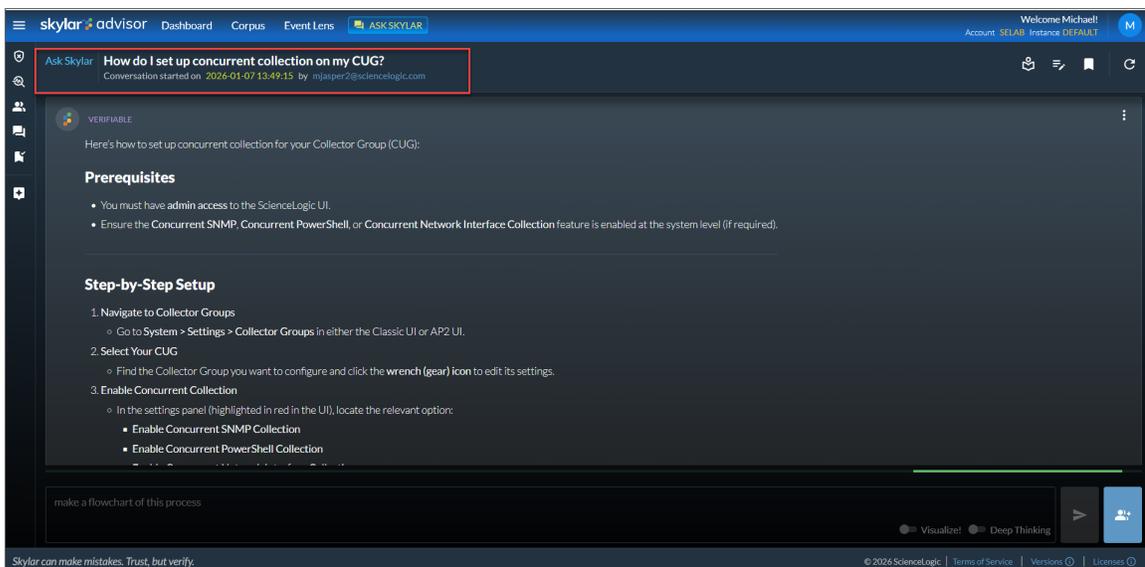
This chapter covers the following topics:

<i>The Skylar Advisor User Interface</i>	62
<i>Defining Your Persona Information</i>	63
<i>Managing Activities and Events on the Dashboards Page</i>	65
<i>Loading and Reviewing Your Documents on the Corpus Page</i>	73
<i>Managing Events on the Event Lens Page</i>	77
<i>Using the Ask Skylar Feature</i>	79
<i>Viewing Team Analytics</i>	82

The Skylar Advisor User Interface

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

- The **Dashboards** page is the default page that displays when you open Skylar Advisor. This page contains a list of all of your interactions with Skylar AI, and it also provides access to Skylar One events, organized by device group, issue type, or organization. For more information, see [Managing Activities and Events on the Dashboards Page](#).
- The **Corpus** page lets you upload your own documents to further train Skylar AI on your company's specific workflows, policies, and processes. You can also track how Skylar AI and other users are using your uploaded documents. For more information, see [Loading and Reviewing Your Documents on the Corpus Page](#).
- The **Event Lens** page lets you view details about the Skylar One events that have been received by Skylar AI and perform further investigations on them. For more information, see [Managing Events on the Event Lens Page](#).
- The **Ask Skylar** feature lets you query Skylar AI for tailored guidance based on real-time data sent from Skylar One, which empowers you to make informed decisions quickly and efficiently. The following image displays the results of a question a user asked Skylar AI. For more information, see [Using the Ask Skylar Feature](#).



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

- The Skylar AI menu (☰) gives you access to Skylar Advisor pages and settings, and the Skylar AI **User Settings**, **Account Settings**, and **Persona** pages. The lab icon (🏭) displays next to features that are currently experimental.
- If you have multiple instances of Skylar AI 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 also click the **[Sign Out]** button to sign out of this session.
- The enter selection mode icon (☑️) adds check boxes to items in a 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.
- All 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** panel 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.
- The notifications icon (🔔) at the bottom left of the window lets you view all of 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 **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 Information

One of the first tasks you should work on when you start working with Skylar Advisor is to set up your Persona information on the **Skylar Persona** page (☰ > **[Persona]** button).

Your persona information lets Skylar AI know what your role is in your company, the language you will use for interacting with Skylar AI, and how you want Skylar AI to structure your feedback, including the level of detail and the tone of the feedback.

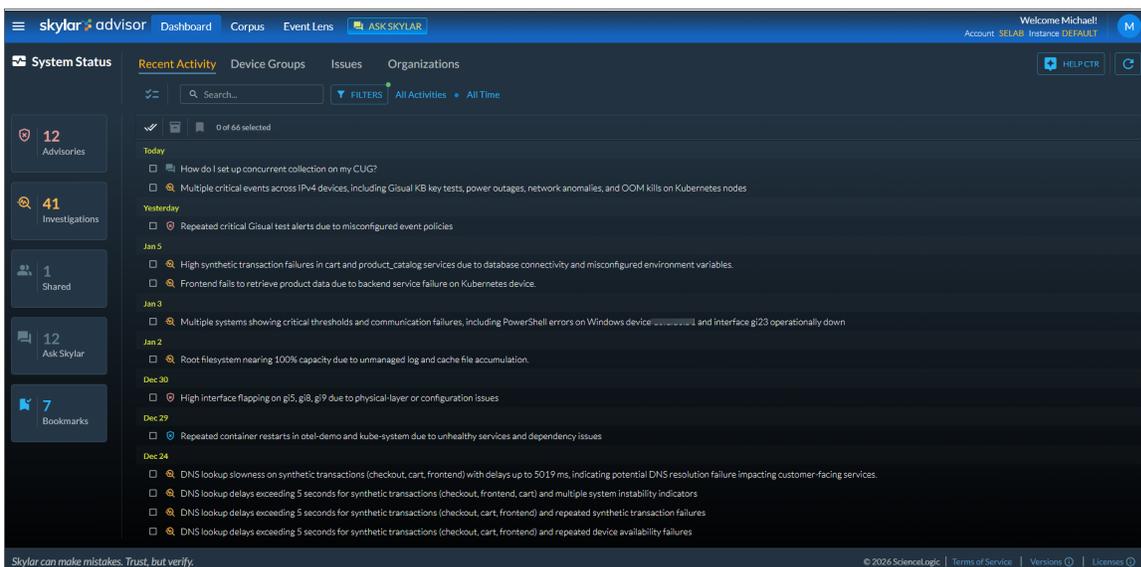
To define your Skylar Persona:

1. Log in to Skylar Advisor and click the Skylar AI menu (☰).
2. Click the **[Persona]** button. The **Skylar Persona** page appears:
3. On the **[Role]** tab, select a role that best matches what you do. A basic description of your role appears in the ***Customize this description*** text box, which you can edit to better fit your role.
4. If you do not find a good match, click **[Other]**. You can add more information about your role in the ***Customize this description*** text box.
5. Click **[Save & Preview]**. Based on your feedback on this tab, Skylar AI displays a sample response using your current settings. You can toggle between the **[Generic]** and **[Personalized]** tabs to see the difference.
6. On the **[Language]** tab, select a language for your responses from Skylar AI, and if needed, add more language information to the ***Customize Language Preference*** text box. Click **[Save & Preview]** to see a sample response based on your choice.
7. On the **[Format]** tab, select how you want Skylar AI to format your responses. You can let Skylar AI choose the format for you, or have Skylar AI format responses in paragraph format or in a structured format like bulleted or numbered lists. Click **[Save & Preview]** to see a sample response based on your choice.
8. On the **[Detail Level]** tab, select the depth of detail you want for your responses. You can let Skylar AI choose the depth or granularity, or have Skylar AI give either detailed or brief responses. Click **[Save & Preview]** to see a sample response based on your choice.
9. On the **[Tone]** tab, select the tone, or how formal you want the responses to be. You can let Skylar AI choose the tone, or have Skylar AI give either more formal or casual responses. Click **[Save & Preview]** to see a sample response based on your choice.
10. To return to the Skylar Advisor user interface, click the Skylar AI menu (☰) and select the page you want from the **Advisor** section.

TIP: You can return to this page at any time to update your persona information if you are not satisfied with the responses you are currently getting with Skylar AI.

Managing Activities and Events on the Dashboards Page

The **Dashboards** page is the default page that displays when you open Skylar Advisor. This page contains a list of all of your interactions and notices related to Skylar AI, including "Advisories," "Investigations," and conversations with Skylar AI. This page also provides access to filtered lists of alerts and events from Skylar One, organized by device group, issue type, or organization.

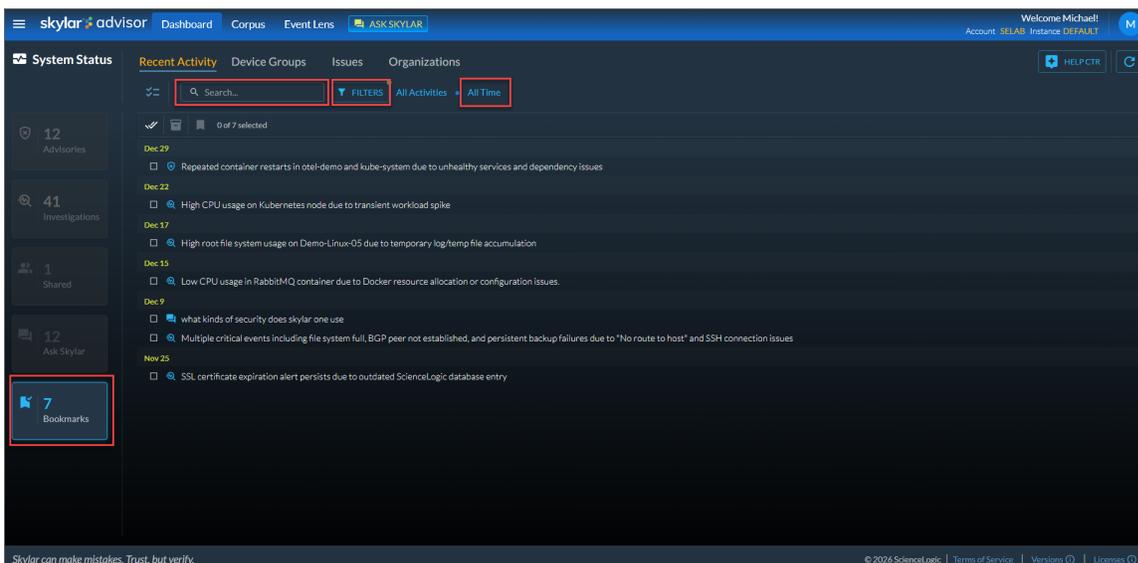


You can use the filters at the top of each tab on the **Dashboards** page to change the type of information displayed or update the time frame of the data display. Type activity information into the **Search** text box to display activities that match your search text, or click the **[Filters]** button to select additional filters.

If a page has a refresh icon (🔄), you can click that icon to update the content on the page, if needed.

The Recent Activity Tab

The **[Recent Activity]** tab on the **Dashboards** page displays a list of all Advisories, Investigations, and conversations with the **Ask Skylar** feature for a specific time frame. The following image shows only the activities you bookmarked, with a time frame of "All Time":



You can filter and sort this list, and you can also change the time frame if needed. If you do not see the activities that you need, click the **Filters** button at the top of the list of activities and adjust the time frame and any other filters as needed. You can view data from up to 90 days ago. To view data beyond 90 days, use the calendar feature to specify a start date and an end date, but the time range cannot be longer than 90 days.

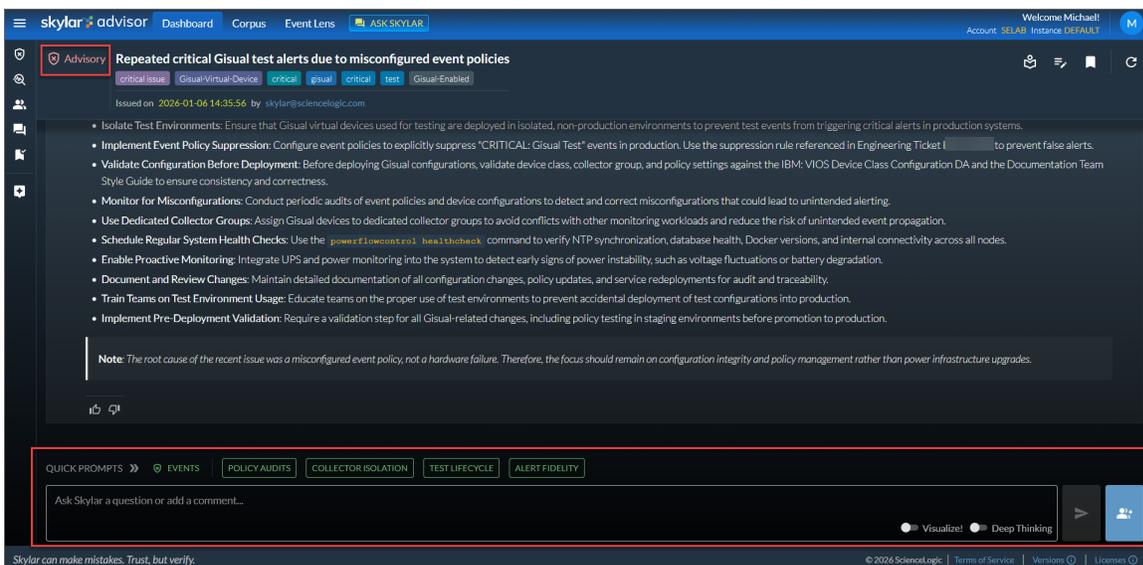
To search for a specific activity or issue, you can start typing in the **Search** text box at the top and the list of activities will get filtered by the text you type, such as "err" to show activities with the word "error" in their title.

The **[Recent Activity]** tab displays the following types of activities:

Icon	Activity Type	Description
	Advisories	<p>Advisories are events that Skylar AI 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 AI, and they directly correlate to alerts and events from Skylar One.</p> <p>When you click on an Advisory and open it, you can display additional details about it, including a summary of the issue, root cause analysis, troubleshooting information, and more.</p> <p>You can also use the Ask Skylar feature when you have an Advisory open. Skylar AI will generate a set of Quick Prompts, which are additional prompts related to the event. Skylar AI also displays an Events link that displays the Skylar One events that were used to generate this Advisory. A Quick</p>

Icon	Activity Type	Description
		Prompts Tasks link lets you generate a list of suggested tasks, which you can add to the Scratch Pad pane to edit and optionally insert them into the Advisory.
	Investigations	Investigations are ad-hoc event analyses that you can choose to generate instead of waiting for Skylar AI to create an Advisory for you. You create an Investigation by clicking the [Start Investigation] button for an event or event group on the Event Lens page. Skylar AI will then run an analysis on that event or event group. The result of this analysis is called an Investigation.
	Shared Conversations	Shared conversations are a list of interactions with Skylar AI using the Ask Skylar feature that have been shared with you by other Skylar Advisor users.
	Ask Skylar	The Ask Skylar section includes the questions you asked Skylar AI using the Ask Skylar feature. You can ask Skylar AI "plain English" questions just like you would for any other large-language model (LLM) like ChatGPT or Microsoft Copilot. To continue a conversation from the list, click a question.
	Bookmarks	A list of any of the above activities that you selected with the enter selection mode icon () enabled to save to this section. If you bookmark an Advisory or an Investigation, that icon will display in blue to show it has been bookmarked.
	Help Center	This icon is not specific to an activity type, but you can click it to open the Help Center pane.

To view more information about an activity on the **[Recent Activity]** tab, click the name of the activity. A detail page with the history of that activity displays. You can ask Skylar AI more questions on this detail page:



The screenshot displays the Skylar Advisor interface. At the top, the navigation bar includes 'skylar advisor', 'Dashboard', 'Corpus', 'Event Lens', and 'ASK SKYLAR'. The main content area shows an advisory titled 'Repeated critical Gisual test alerts due to misconfigured event policies'. Below the title, there are tags for 'critical issue', 'Gisual-Virtual-Device', 'critical', 'gisual', 'critical', 'test', and 'Gisual-Enabled'. The advisory text provides several remediation steps:

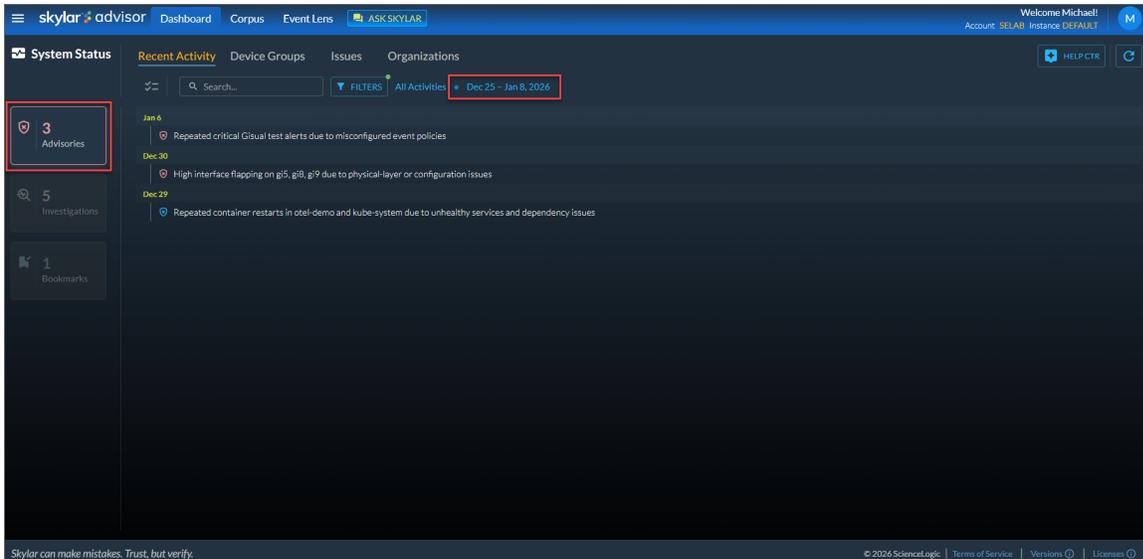
- Isolate Test Environments:** Ensure that Gisual virtual devices used for testing are deployed in isolated, non-production environments to prevent test events from triggering critical alerts in production systems.
- Implement Event Policy Suppression:** Configure event policies to explicitly suppress "CRITICAL: Gisual Test" events in production. Use the suppression rule referenced in Engineering Ticket [redacted] to prevent false alerts.
- Validate Configuration Before Deployment:** Before deploying Gisual configurations, validate device class, collector group, and policy settings against the IBM VIOS Device Class Configuration DA and the Documentation Team Style Guide to ensure consistency and correctness.
- Monitor for Misconfigurations:** Conduct periodic audits of event policies and device configurations to detect and correct misconfigurations that could lead to unintended alerting.
- Use Dedicated Collector Groups:** Assign Gisual devices to dedicated collector groups to avoid conflicts with other monitoring workloads and reduce the risk of unintended event propagation.
- Schedule Regular System Health Checks:** Use the `powerFlowControl healthcheck` command to verify NTP synchronization, database health, Docker versions, and internal connectivity across all nodes.
- Enable Proactive Monitoring:** Integrate UPS and power monitoring into the system to detect early signs of power instability, such as voltage fluctuations or battery degradation.
- Document and Review Changes:** Maintain detailed documentation of all configuration changes, policy updates, and service redeployments for audit and traceability.
- Train Teams on Test Environment Usage:** Educate teams on the proper use of test environments to prevent accidental deployment of test configurations into production.
- Implement Pre-Deployment Validation:** Require a validation step for all Gisual-related changes, including policy testing in staging environments before promotion to production.

A note at the bottom of the advisory states: "The root cause of the recent issue was a misconfigured event policy, not a hardware failure. Therefore, the focus should remain on configuration integrity and policy management rather than power infrastructure upgrades." At the bottom of the screen, there is a 'QUICK PROMPTS' section with buttons for 'EVENTS', 'POLICY AUDITS', 'COLLECTOR ISOLATION', 'TEST LIFECYCLE', and 'ALERT FIDELITY'. Below these buttons is an 'Ask Skylar' input field with the placeholder text 'Ask Skylar a question or add a comment...'. The bottom of the interface includes the text 'Skylar can make mistakes. Trust, but verify.' and copyright information for 2026 Sciencelogic.

Viewing Content by Activity Type

In the **System Status** panel, you can filter the activities on the **[Recent Activity]** tab to only show one type of activity. If there are no activities that meet the filter requirements, the filter button does not appear on the left in the **System Status** panel.

For example, if you clicked the **[Advisories]** filter in the **System Status** panel, you will only see the set of Advisories that occurred in the specified time frame of the past 14 days:



Click the **[Bookmarks]** button again to view all of the activity types.

TIP: On the **[Corpus]** tab, **[Event Lens]** tab, and the **Ask Skylar** page, a smaller version of the **System Status** buttons appear on the left of the screen. You can click one of these buttons to open a pop-out pane with a list of activities matching the button you clicked, such as **[Advisories]**.

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.

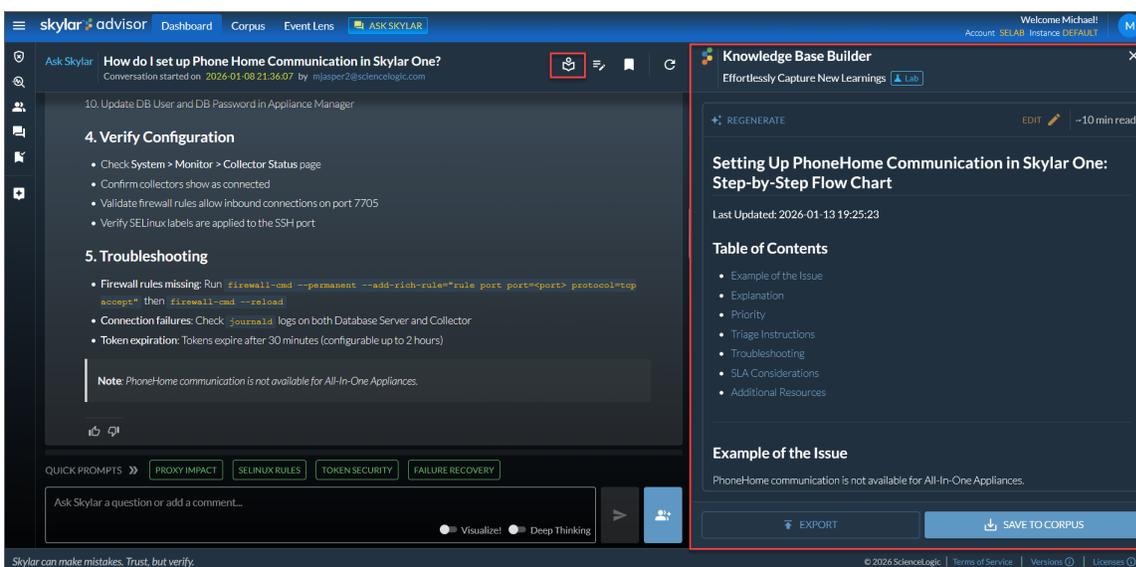
To bookmark one or more activities that you want to add to the **Bookmarks** section:

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, or a Conversation, you can click the knowledge base builder icon (📄) to have Skylar AI create a draft Knowledge Base (KB) article based on that activity.

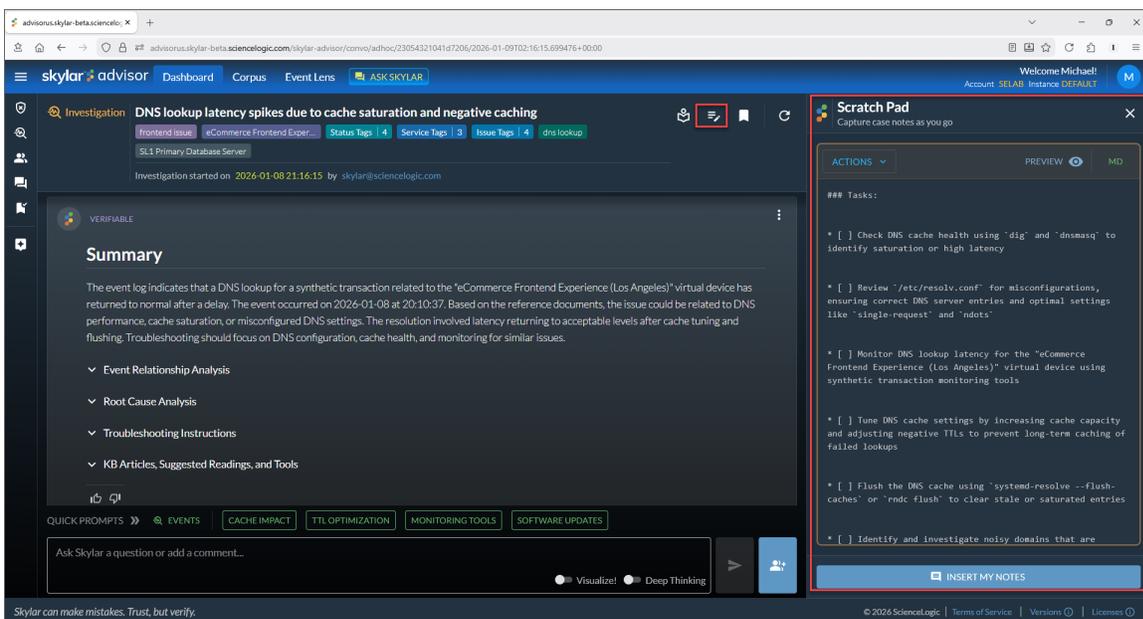


After Skylar AI creates the KB article, you can:

- Click **[Regenerate]** to create the article again. You might want to regenerate an article if you asked Skylar AI 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.

Using the Scratch Pad Feature

On an activity detail page for an Advisory, Investigation, or a Conversation, you can click the scratch pad icon (📝) 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.

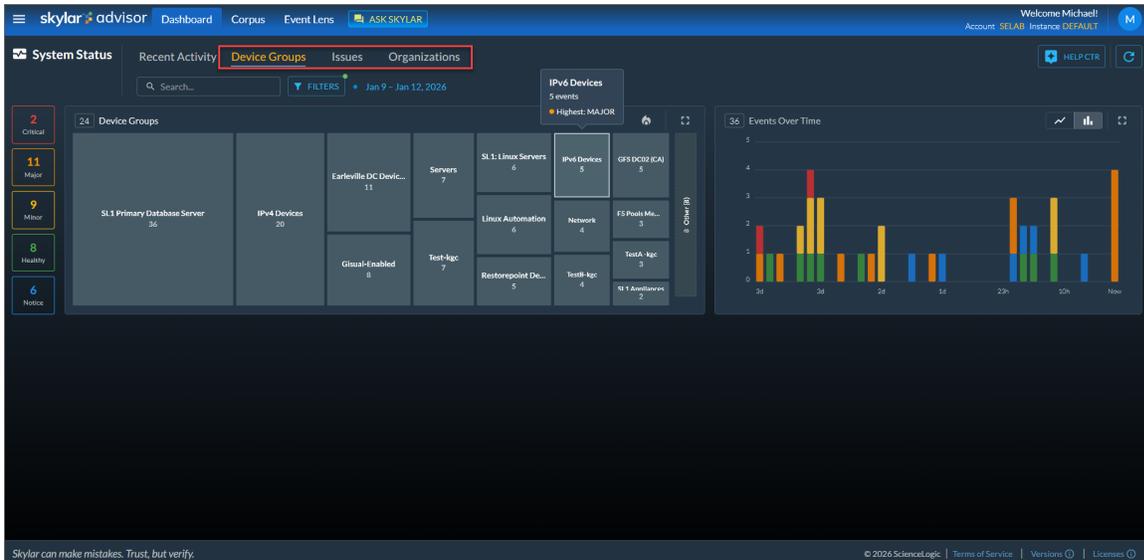
Click **[Preview]** to see how your text looks with formatting. Click **[MD]** for a list of basic Markdown commands.

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 Device Groups, Issues, and Organizations Tabs

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

The **[Device Groups]** tab displays a list of events organized by Skylar One device group, the **[Issues]** tab displays a list of issues organized by Skylar One events, and the **[Organizations]** tab of the **Dashboards** page displays a list of events organized by Skylar One organizations.



Viewing Basic Event Details on a Tab

In the **System Status** panel on the left, you can filter the list of events on these three 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 of the events again.

The show severity heatmap icon (📊) lets you view the events by severity within the organization 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 return to the default view.

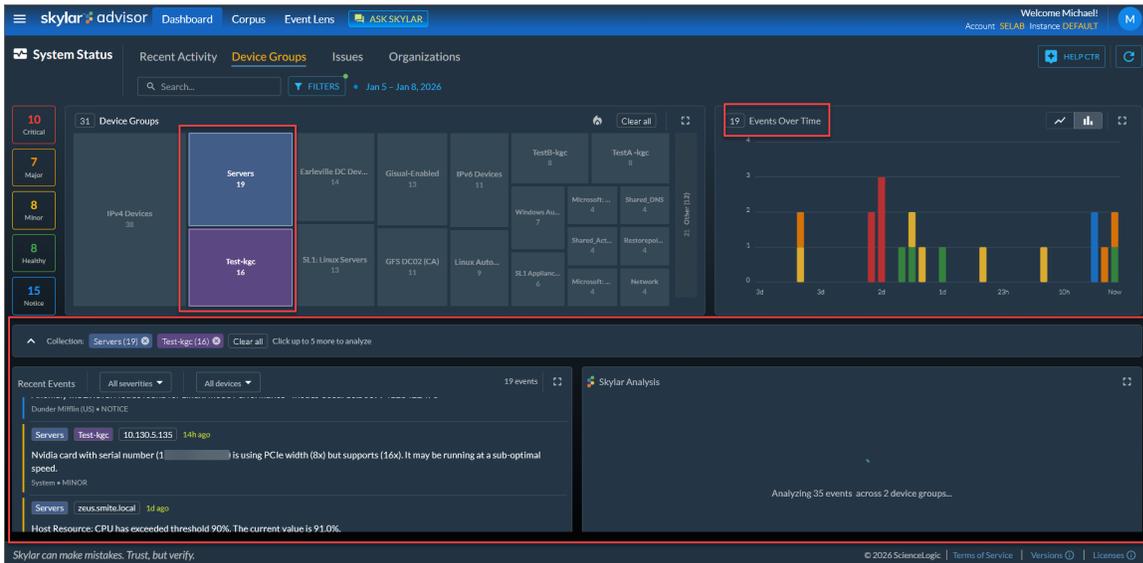
When you mouse over a set of events organized by device group, issue, or organization, a pop-up displays a summary of the events for that set of events.

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 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 data only for the selected items.



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

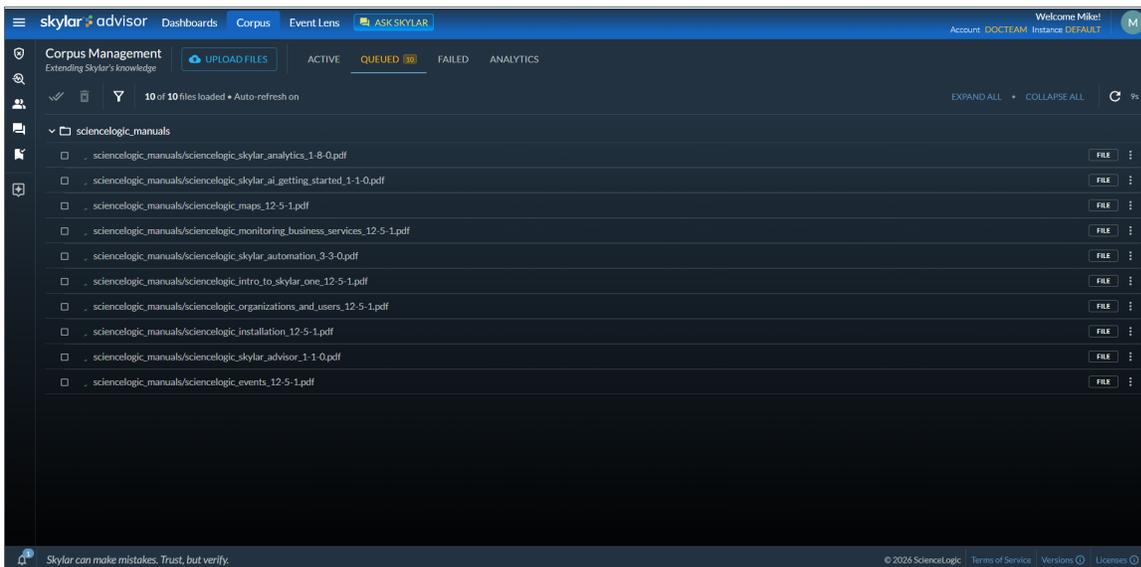
- Review the events from the item you selected above (in this example all of the "Server" events and the "Test-kgc" events).
- Click a device IP or device name to show only events for that device.

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

- Click **[Generate Analysis]** to have Skylar AI 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 activity that you can find on the **[Recent Activity]** tab of the **Dashboards** page.

Loading and Reviewing Your Documents on the Corpus Page

The **Corpus** page lets you upload and manage your own documents to further train Skylar AI on your company-specific workflows, policies, and processes. Skylar AI uses the documents in the Corpus to generate Advisories and answer your questions. You can also track how Skylar AI and other users are using your uploaded documents.



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

The Main Corpus Page

The **[Active]** tab on the **Corpus** 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 first three tabs on the **Corpus** page:

- **[Active]**. These documents are completely uploaded and can be used by Skylar AI.
- **[Queued]**. These documents are in the process of being uploaded and ingested by Skylar AI. The countdown timer on the right side of the tab displays how long until the list will be updated again, and it starts over every 15 seconds, or if the timer is clicked.
- **[Failed]**. 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.

You can upload .pdf, .txt, .doc, .ppt, .ppx, .xls, and .xlsx files to the **Corpus** page. The maximum file size for a document is 31 MB.

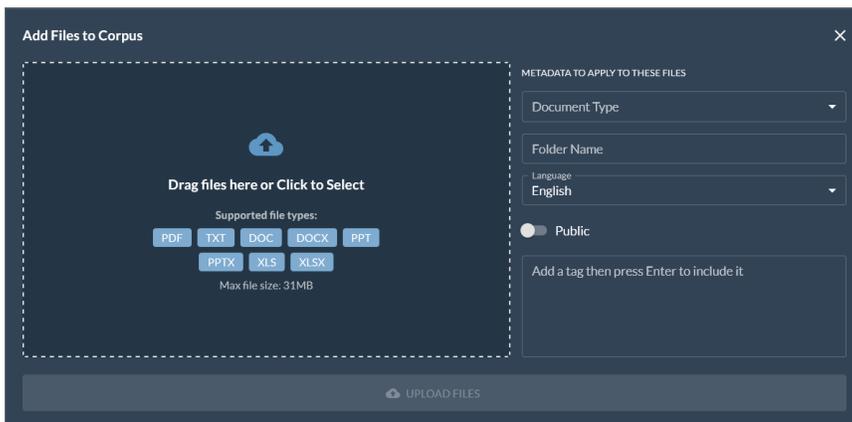
NOTE: Ingesting documents can take time time-consuming, as Skylar AI creates an "embedding" that lets it search the document. It also adds tokens to the document. A rough estimate for ingesting documents while onboarding is about 2,500 documents per hour.

On the **[Upload Files]** tab, you can:

- Click the batch actions icon (☰) to add check boxes next to each document. You can click a check box for one or more documents, and then click the delete icon (🗑️) to delete those documents. To delete a folder and all of its contents, check all of the documents in that folder and then click the delete icon.
- Click the filter corpus icon (🔍) to filter the list of documents by titles, tag names, regular expressions, documentation type, or last modified dates.
- If you have your documents organized in folders, you can click **Expand All** to open all the folders or **Collapse All** to close all the folders.
- Click the refresh icon (🔄) to update the content on the **Corpus** page.
- 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 browser tab or window. Please note that all uploaded documents currently display as PDF files after you upload them to the Corpus.
- Click the **Show Older Files** link to load files that were added to this page earlier.
- Click the **[Add Files to Corpus]** button to upload a PDF or text document. You can add metadata and make folders, language, and tags for additional search context.

To add one or more files to the **Corpus** page:

1. Click the **[Add Files to Corpus]** button. The **Add Files to Corpus** window appears:



2. Drag and drop one or more files onto the window, or select the files to add them.

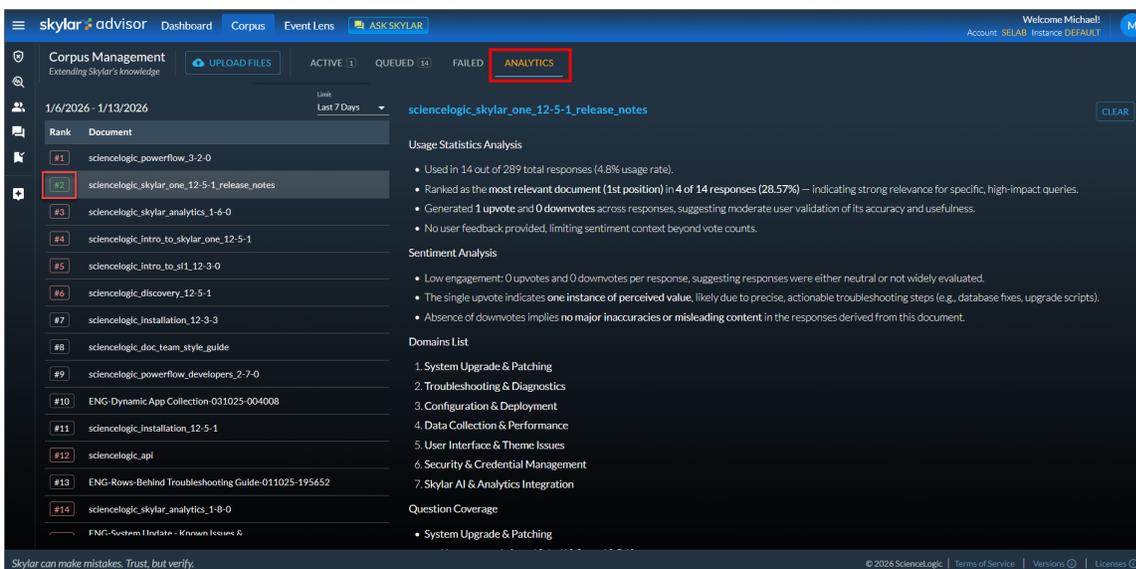
3. To help Skylar AI locate and learn each document, select the document type, type the folder name, and select the language for that document.
4. Click the **Public** toggle to turn it blue to make the document available to all users in all organizations using this instance, or turn it white to make 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.
6. Add more files as needed.
7. Click Upload Files. It will take Skylar AI a few moments to process the document you just uploaded, and when the document is processed, it will appear on the **Corpus** page.
8. Click the refresh icon () to update the content on the **Corpus** page as needed.

NOTE: If the same document is added the **Corpus** page, Skylar AI will highlight that file in yellow, and you can click the ellipsis icon () to delete the duplicate document.

The Analytics Tab

The **[Analytics]** tab on the **Corpus** page lets you view a list of the uploaded documents that have been used the most by Skylar AI to answer questions.

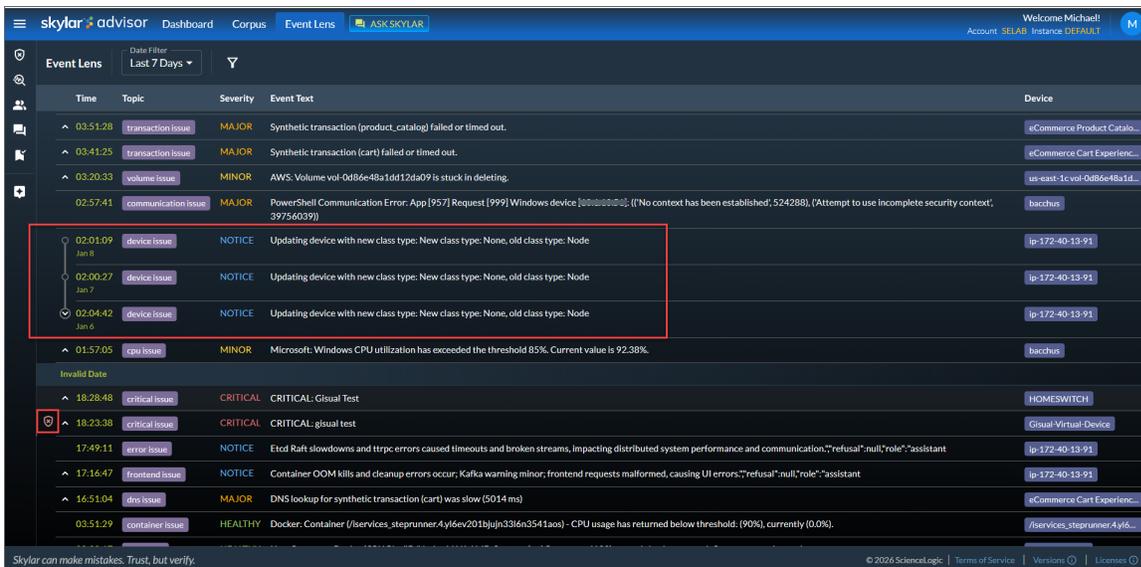
When you click a document in the list on the left, detailed information appears on the right, including usage statistics, sentiment analysis, domains list, question coverage, and potential document deficiencies.



The document analysis requires multiple question-and-answer interactions to provide meaningful insights. In addition, the numbers for the documents might have a colored # ranking based on the number of upvotes (thumbs up) or downvotes (thumbs down). For example, a document that has a green #2 received more upvotes than downvotes, while a pink #1 document has been downloaded the most but received more downvotes than upvotes by users. Based on the user voting, the green #2 document would appear to be a more reliable source of information and the pink #1 document.

Managing Events on the Event Lens Page

The **Event Lens** page lets you view details about the Skylar One events that have been received by Skylar AI. Skylar AI performs semantic event correlation to group together related events based on how Skylar AI understood the events. Skylar AI also considers related attributes like device groups.



Events are listed on the **Event 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 Event Lens Page](#).

Investigating Advisories and Events

The shield icon (🛡️) in the list indicates that the event or event group is an Advisory, an event that Skylar AI has determined to be important. If you select an Advisory from the list, you can view more information about the Advisory in a new pane to the right of the list. Clicking the **[Investigate Further]** button under the summary title will open the Advisory in a new tab. You can also find this Advisory in the **Advisories** section on the **Dashboards** 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.

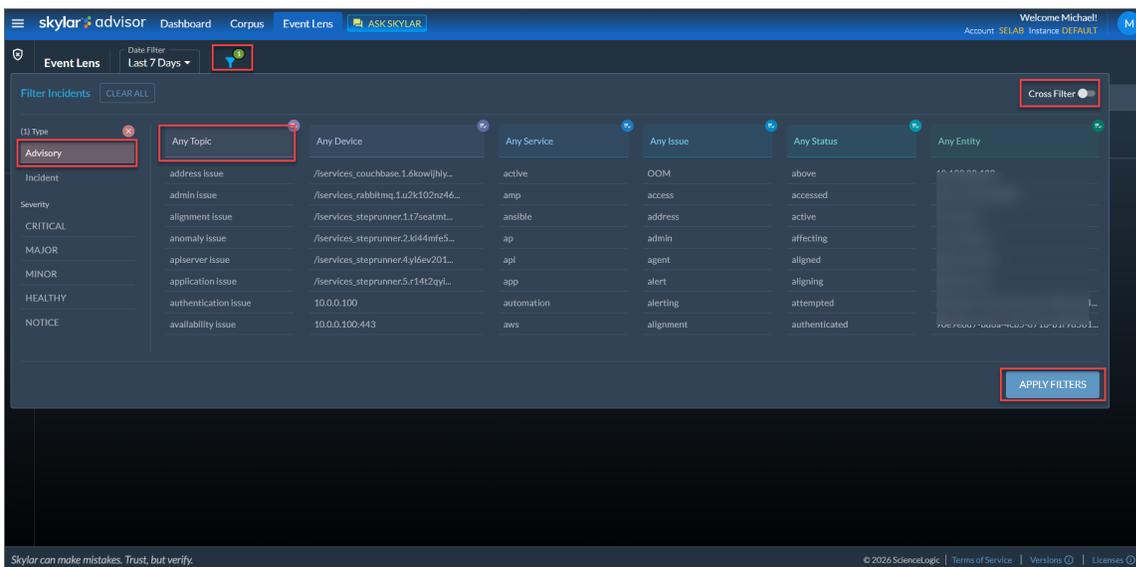
If you select an event or event group that has *not* been labeled as an Advisory, you can:

- Click the **[Start Investigation]** button to have Skylar AI run an analysis on this event or event group. The result of this analysis is called an **Investigation**.
- Open this Investigation in another tab by clicking the **[Investigate Further]** button. You will then be able to interact with and share the Investigation.

You can also find this Investigation in the **Investigations** section on the **Dashboards** page.

Filtering the Event Lens Page

On the **Event Lens** page, events are shown with the most recent at the top. You can change the filter settings to view more or less information on this page. By default, this page is filtered to just show **Advisories**.

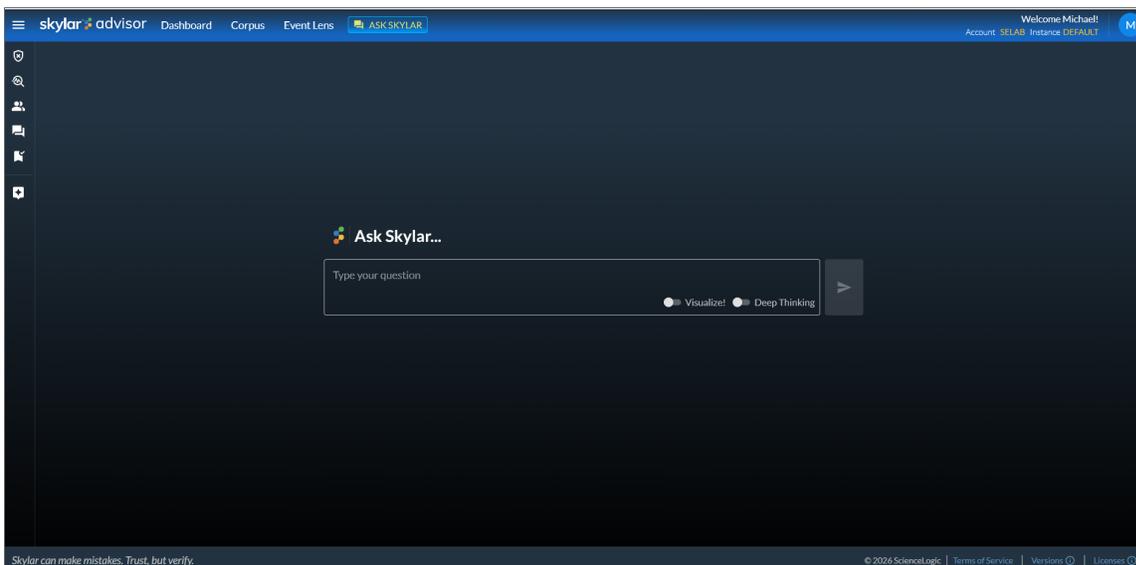


To filter this page, click the filter icon (🔍) to access the **Filter Incidents** window, where you can:

- Under **Type**, select *Advisory*, *Incident*, or both.
- Under **Severity**, select one or more severities to show on the **Event 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 Filters]** to save your filters.

Using the Ask Skylar Feature

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



You can ask Skylar AI any question, and Skylar AI 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.

TIP: The icons in the navigation pane on the left match the *icons on the System Status pane* on the **Dashboards** page.

Features of Ask Skylar

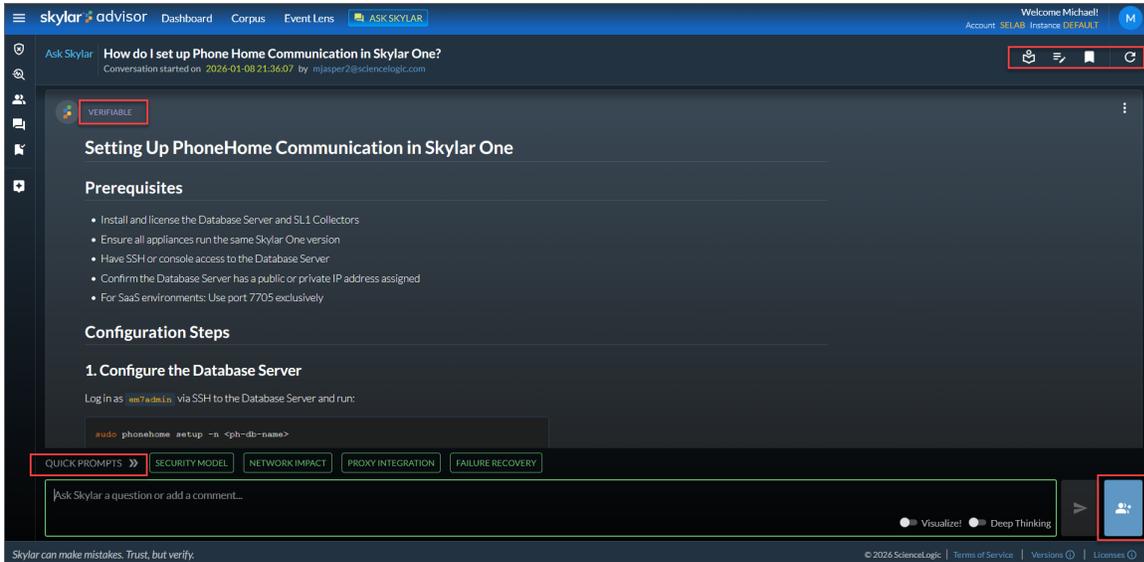
To access the **Ask Skylar** text box, click the **[Ask Skylar]** button at the top of the Skylar Advisor user interface or at the bottom of an existing activity from the **[Recent Activity]** tab of the **Dashboards** page.

In the **Ask Skylar** text box, type your question for Skylar AI. Before you click the send icon (➤), you can also enable one or both of the following options:

- **[Visualize!]**: Lets you choose to have Skylar AI create a flow chart, relationship map, or other relevant image based on your query. You can enlarge the image by clicking on it. You can also click on an element to highlight it and then click the **[Space Bar]** to view a "Quick Take" pane that defines that element.
- **[Deep Thinking]**: Tells Skylar AI to perform the same kind of thinking that a human would do, which takes longer than a typical question as Skylar AI gathers its resources. This option is useful for more complex questions.

After you click the send icon (➤), Skylar AI displays its progress and a timer, along with a **[Stop]** button if you would like to cancel the query.

After a few moments, Skylar AI displays the results of the query:



TIP: You can highlight and right-click any text in an answer from Skylar AI to open a **Quick Take** pane that provides more information about that text. If you are viewing a visualization, you can select an item in the visualization the press the **[Space Bar]** to view the **Quick Take** pane.

Additional Options

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

- **[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** panel, 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 AI, this tag displays as **[Unverifiable]**.

- Knowledge base builder icon (📄). Tells Skylar AI to create a draft Knowledge Base article based on that conversation. For more information, see [Creating a Knowledge Base Article Based on an Activity](#).
- Scratch pad icon (📝). 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](#).
- Add bookmark icon (🔖). Adds the conversation to the **Bookmarks** section on the **[Recent Activity]** tab of the **Dashboards** page.

- Refresh icon (). Updates the content on the current page.
- Thumbs up and thumbs down icons ( ). Lets you vote up or down on the quality of the response from Skylar AI, and you can use the text 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 AI smarter.

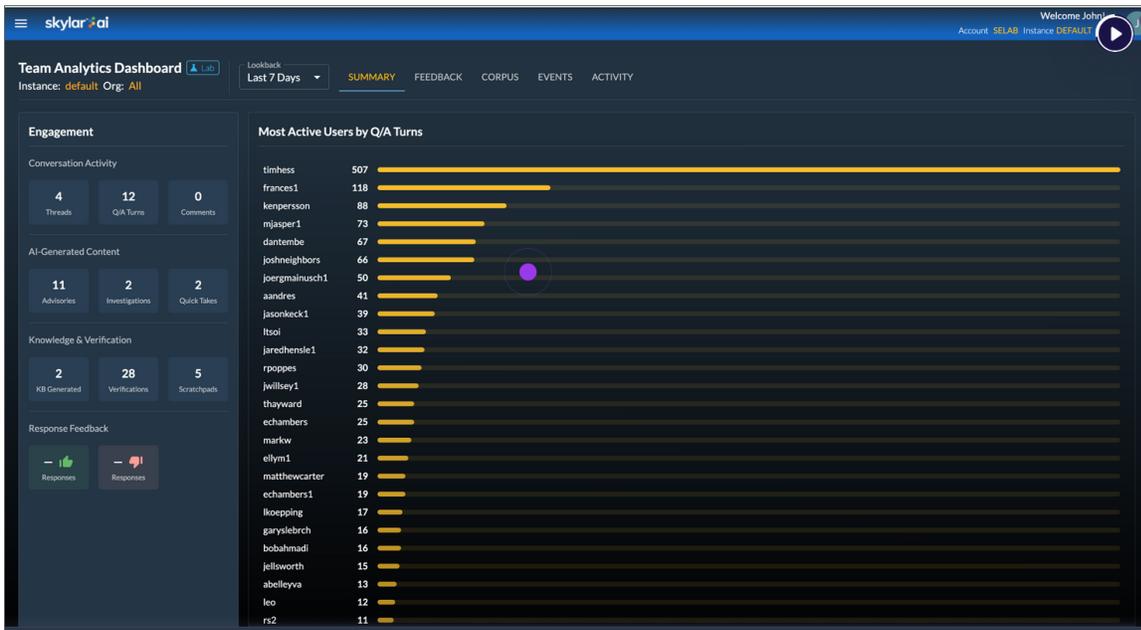
- **Quick Prompts.** Based on the original question, Skylar Advisor displays additional prompts in green boxes that you can click for additional information. You can hover over a Quick Prompt to see more information about it.

If you are using the **Ask Skylar** feature in an Advisory or Investigation detail record, this section will display an **Events** link that you can click to view the Skylar One events that were used to generate this Advisory. For Advisories or Investigations, a **Tasks** link also appears in this section, which you can click to view a list of suggested tasks in the **Scratch Pad** pane where you can edit them and optionally insert them into the Advisory.

- **Share this conversation** icon (). Lets you share your conversation. Skylar Advisor generates a URL you can send via slack or email, and this conversation is 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.

Viewing Team Analytics

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



NOTE: This page is a beta feature, so it has the lab (🔬 Lab) icon. 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.
- **[Feedback].** Displays the specific feedback users added when they click 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.
- **[Events].** Shows the number of events, Advisories, and Investigations over time. This page also shows how Skylar Advisor provided noise reduction, and how much Skylar Advisor saved you in money and time. You can adjust the sliders under *est. savings* and *analyst time saved* to view

savings over time.

- **[Activity]**. Charts the various activities by users over time.



Skylar Analytics

Version 2.1.0

Chapter

1

Introduction to Skylar Analytics

Overview

Skylar Analytics includes Data Visualization, Data Exploration, Anomaly Detection, and Predictive Alerting. This manual will explain all of these components, and how to use them.

For an overview of Skylar AI, see [Introduction to Skylar AI](#).

This chapter covers the following topics:

What is Skylar Analytics?	85
Mapping Skylar One Dynamic Application Object Names to Skylar AI Columns	86

What is Skylar Analytics?

The Skylar Analytics suite of services uses data gathered by Skylar One to explore data, generate visualizations, and monitor IT infrastructure metrics. Skylar Analytics can also use Skylar AI to predict alerts before they happen, and detect anomalies that could become events that might disrupt your IT infrastructure and functionality.

NOTE: Skylar One uses port 443 to communicate with your Skylar Analytics system. Skylar AI does not require a port.

Skylar Analytics includes the following components:

- **Data Visualization.** Enables SQL-based dashboards and charts based on data gathered by Skylar AI and Skylar One. Data Visualization is achieved using a ScienceLogic-hosted instance of Apache Superset.
- **Data Exploration.** Enables third-party tools that use the Open Database Connectivity (ODBC) interface to access the metric data from Skylar AI. This component lets you use ODBC to connect Skylar AI data with applications like Tableau, Microsoft Power BI, or other business intelligence tools.
- **Anomaly Detection.** Uses always-on anomaly detection to find metric outliers in Dynamic Application time series data. It also computes an anomaly score that characterizes the significance of each anomaly. You can view anomalies for all Dynamic Application metrics for a device by visiting the **[Anomaly Detection]** tab on the **Device Investigator** page for that device.
- **Predictive Alerting.** Helps to avoid problems such as file systems running out of space. The alerts appear as enriched events within Skylar One.

The other chapters in this manual cover each Skylar Analytics component in detail.

Mapping Skylar One Dynamic Application Object Names to Skylar AI Columns

When data from Skylar One Dynamic Applications is exported to Skylar AI, the names of collection and presentation objects are automatically converted into clean, standardized column names for the Skylar data lake. The following rules ensure that all Skylar column names are consistent, machine-friendly, and easy to work with. If you are not sure how a name will be converted, use these common word replacements and clean-up rules as a guide.

The conversion process follows several steps:

1. **Standardize Special Characters**

- If a letter is followed by a non-word character and an "a", replace it with the letter plus "A". This format ensures that column names are valid and avoid special symbols.
- Example: ba\$ → bA

2. Replace Common Words

Certain words are automatically shortened to standard abbreviations. The following table contains the most common abbreviations:

Original Word	Becomes
ScienceLogic	SL
Microsoft	MS
Server	Svr
Database	DB
FileSystem	FS
Interface	IF
Resource	Rsrc
Worker	Wrkr
Service	Svc
Relationship	Relnship
Total	Ttl
Interval	Ival
Baseboard	Basebrd
Num Of	Num
Distribution	Distro
Level	Lvl
Hardware	HW
Software	SW
Default	Dflt
Namespace	Nspc
Virtual Machine	VM
Kilobytes	KB
Megabytes	MB
Gigabytes	GB
Terabytes	TB
Backup	Bkup
Successful	Good
Expiration	Expiry
Manufacturer	Mfgr
Device	Dvc
Sockets	Socks
Command	Cmd

VMware Open	Open
Processor	Procscr
Processes	Procs

3. Shorten Common Technical Terms

Some longer technical words are shortened to their first few letters. Examples:

- Physical → P
- Utilization → U
- Capacity → C
- Configuration → C
- Discovery → D
- Storage → S
- Limit → L
- Network → N
- Address → Addr

Only the beginning of the word is kept for these cases.

4. Clean Up the Name

- Remove all non-alphanumeric characters, such as spaces, slashes, parentheses.
- Replace common terms:
 - Average → Avg
 - QueueLength → QLen
 - sISI → SL
 - SL1Skylar → SL1Sky
 - Exporter → Exptr
 - Receiver → Rcvr

5. Add Unit, if Applicable

- If the original name included a unit, like MB, GB, %, and so on, add it at the end after an underscore.
- Format: *columnname_unit*
- Example: MemoryUtilization (Gigabytes) → MemU_GB

Skylar Analytics: Data Visualization and Data Exploration

Overview

The **Data Visualization** component of Skylar Analytics contains dashboards and charts based on data gathered by Skylar AI and Skylar One. To display these dashboards and charts, Data Visualization uses a ScienceLogic-hosted instance of Apache Superset. The data for the dashboards and charts includes metrics for file systems, network interfaces, and all Dynamic Applications, with more metrics planned for future Skylar and Skylar One updates.

IMPORTANT: The dashboards and charts in the Data Visualization component of Skylar Analytics are *not* compatible with Skylar One dashboards, widgets, or reports.

The optional **Data Exploration** component of Skylar Analytics enables third-party tools that use the Open Database Connectivity (ODBC) interface to access the metric data from Skylar AI. This component lets you use ODBC to export Skylar AI data to Tableau, Microsoft BI, and other business intelligence tools.

This chapter provides a general overview of how to view the charts, graphs, and other reports in the Skylar Analytics user interface, along with tips and best practices for users of Skylar One and Skylar AI.

This chapter covers the following topics:

What is Data Visualization?	90
Working with Datasets in Data Visualization	91
Viewing Dashboards and Charts in Data Visualization	93

[Creating and Customizing Dashboards and Charts](#) 100

[Adding and Upgrading Dashboards and Datasets](#) 111

[Data Exploration: Exporting Data to Skylar AI from Third-party Tools](#) 112

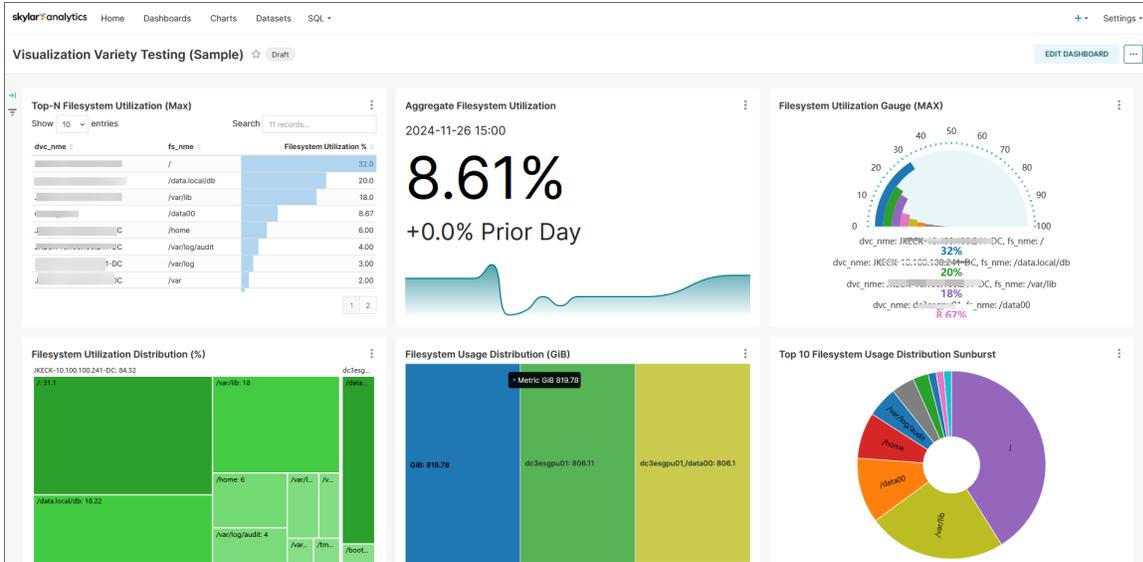
[Additional Resources for Skylar Analytics \(Apache Superset Training\)](#) 115

What is Data Visualization?

Before the initial release of Skylar Analytics, Skylar One stored data in a proprietary format that was not easily exported to other third-party applications for further research and insight. Skylar Analytics takes the data gathered by Skylar One and Skylar AI, normalizes it, and makes it available in standard ODBC database format.

The data originates from Skylar One data collectors, undergoes processing, and is then simultaneously transmitted to Skylar using the API. This data is stored in Skylar Analytics as **datasets**, which are curated representations of the data in your database gathered by Dynamic Application presentation objects in Skylar One. You can use the data in a dataset to populate dashboards and charts in Skylar Analytics. For more information, see [Working with Datasets in Data Visualization](#).

ScienceLogic hosts an instance of Apache Superset as an option for **Data Visualization** that lets you explore and view your data using business intelligence (BI) dashboards. Below is an example of one of the default dashboards in Skylar Analytics:



For more information, see [Viewing Dashboards and Charts in Data Visualization](#).

You can also use the Data Visualization component with your existing BI tools for your company that support ODBC; this option is called **Data Exploration**. For more information, see [Data Exploration: Exporting Data from Skylar AI](#).

NOTE: Because ScienceLogic does not own the underlying framework for the Data Visualization and Data Exploration components, ScienceLogic is not responsible for maintaining or updating documentation for third-party open-source software, including Apache Superset. For a list of the most current and accurate information, see [Additional Resources for Skylar Analytics](#).

Working with Datasets in Data Visualization

The data imported from Skylar One is stored in Skylar Analytics as **datasets**, which are curated representations of the data gathered by the Dynamic Application presentation objects from a PowerPack in Skylar One. A presentation object for a Dynamic Application defines how Skylar One uses the collected data to define and generate graphs.

You can use the data in a dataset to populate dashboards and charts in Skylar Analytics.

IMPORTANT: To update all of your datasets based on Skylar One PowerPacks, go to Skylar Settings and click the **[Sync Skylar Datasets]** button on the **[Customizations]** tab on the **Dashboards** page (Analytics Admin > Dashboards > Customizations) in Skylar Settings. If all datasets have been updated, the button does not appear, and the text "Datasets are current" appears instead. This button is only available to owner users in Skylar AI.

Components of a Dataset

In Skylar Analytics, each set of Dynamic Applications from a PowerPack is represented by three datasets:

- **One performance dataset.** In Skylar Analytics, these datasets use the naming convention of "Metric<PowerPackName>", such as "MetricMSWinSvr" for the "Microsoft: Windows Server" PowerPack.
- **Two configuration datasets:**
 - The "Current" dataset contains only the last recorded configuration change. You will typically use this dataset to quickly retrieve configuration details, unless you need to retrieve historical configuration changes. In Skylar Analytics, these datasets use the naming convention of "Conf<PowerPackName>Current", such as "ConfMSWinSvrCurrent" for the "Microsoft: Windows Server" PowerPack.
 - The other dataset contains configuration Dynamic Applications that have timestamps for each configuration snapshot taken by Skylar One. This dataset uses the naming convention of "Conf<PowerPackName>", such as "ConfMSWinSvr".

NOTE: You do not need to know the name of the Dynamic Application or the Dynamic Application structure to select data from one of these datasets. You just need to know the name of the PowerPack.

For database query purposes, tables in Skylar Analytics are abbreviated to be as short as possible. For example, "Microsoft" is shortened to "MS", "Windows" is "Win", and "Server" is "Svr". This results in the name "MSWinSvr". For more information about the abbreviations used for the metric names, see [Mapping Skylar One Dynamic Application Object Names to Skylar Columns](#).

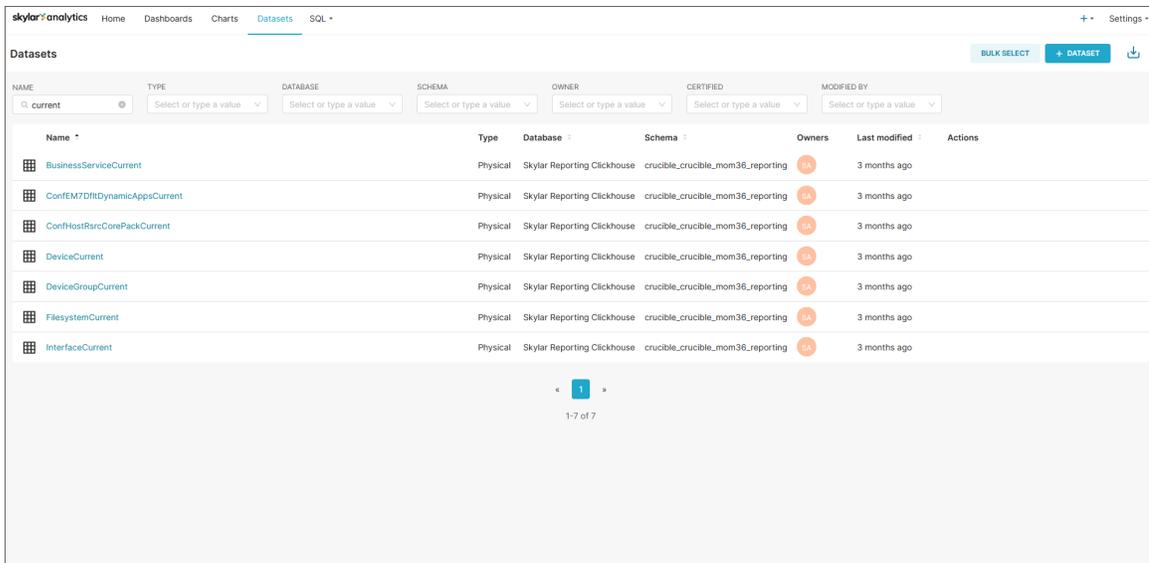
Viewing the List of Datasets

To view a list of the datasets in Skylar Analytics:

1. From Skylar One, go to the **Skylar AI** page () and click the **[Visit]** button next to **Skylar Data Visualization**. If you are not currently logged into Skylar AI, the Skylar AI login page appears. If not, log in and click **Analytics**.

TIP: If you know the URL of your Skylar AI system, you can go to that location instead of using Skylar One.

2. In the Skylar Analytics user interface, go to the **Datasets** page:



Name	Type	Database	Schema	Owner	Certified	Modified By
BusinessServiceCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago
ConfEM7DfllDynamicAppsCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago
ConfHostRsrcCorePackCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago
DeviceCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago
DeviceGroupCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago
FilesystemCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago
InterfaceCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting			3 months ago

TIP: You can filter the list of datasets by typing some or all of a dataset name in the **Name** field at top left.

3. Click the name of a dataset on the **Datasets** page to access the **Charts** page, where you can create a chart based on the columns (metrics) in that dataset. For more information, see [Creating and Customizing Dashboards and Charts](#).
4. In the list of datasets on the **Datasets** page, you can hover over an icon in the **Actions** column to delete, export, or edit (if you are an owner user) a specific dataset.

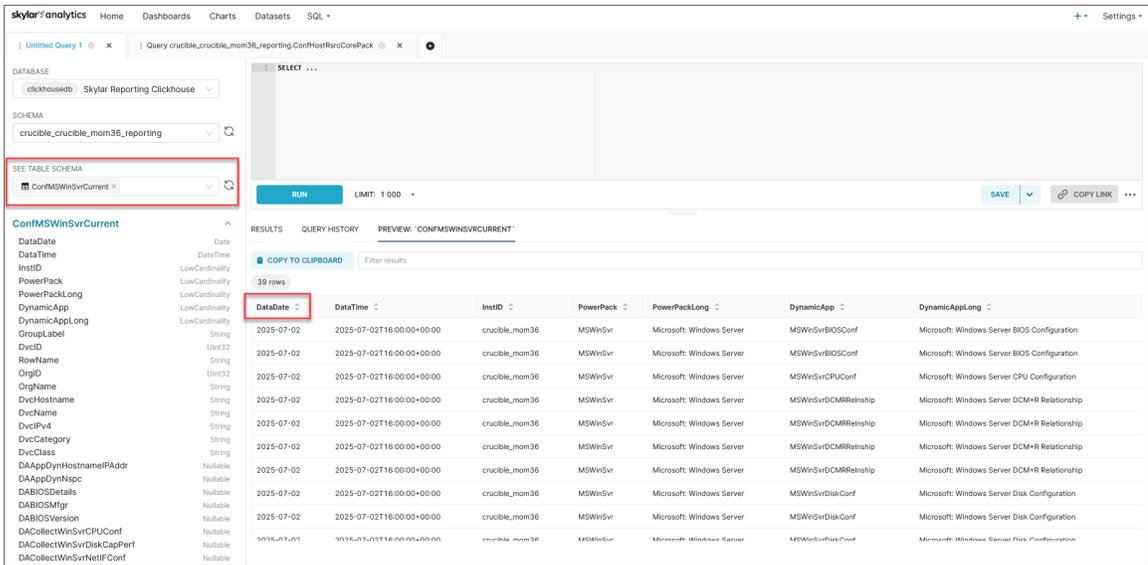
Viewing the Contents of a Dataset

To view the contents of a dataset:

1. From Skylar One, go to the **Skylar AI** page () and click the **[Visit]** button next to **Skylar Data Visualization**. If you are not currently logged into Skylar AI, the Skylar AI login page appears. If not, log in and click **Analytics**.

TIP: If you know the URL of your Skylar AI system, you can go to that location instead of using Skylar One.

2. In the Skylar Analytics user interface, go to the **SQL Lab** page (SQL > SQL Lab):



The screenshot shows the Skylar Analytics SQL Lab interface. The top navigation bar includes 'Home', 'Dashboards', 'Charts', 'Datasets', and 'SQL'. The main area is divided into several sections:

- Query Editor:** Shows a query: `Query crucible_crucible_mon36_reporting.ConfMsrRrcCorePack`. Below the query is a 'RUN' button and a 'LIMIT: 1 000' dropdown.
- Database and Schema:** A dropdown menu shows 'clickhousedb' and 'Skylar Reporting Clickhouse'. Below it, the schema is 'crucible_crucible_mon36_reporting'. A 'SEE TABLE SCHEMA' button is visible.
- Table Schema:** A list of columns for the 'ConfMsrWinSvrCurrent' table is shown, including 'DataDate', 'DateTime', 'InstID', 'PowerPack', 'PowerPackLong', 'DynamicApp', 'DynamicAppLong', 'GroupLabel', 'DvcID', 'RowName', 'OrgID', 'OrgName', 'DvcInstName', 'DvcName', 'DvcIPV4', 'DvcCategory', 'DvcClass', 'DAAppDynHostnameIPAddr', 'DAAppDynHasc', 'DABIOSDetails', 'DABIOSMfrgr', 'DABIOSVersion', 'DACollectWinSvrCPUConf', 'DACollectWinSvrDiskCapPerf', and 'DACollectWinSvrNetIPConf'.
- Table Preview:** A table with 39 rows is displayed. The first row is highlighted. The columns are: DataDate, DateTime, InstID, PowerPack, PowerPackLong, DynamicApp, and DynamicAppLong. The data in the first row is: 2025-07-02, 2025-07-02T16:00:00+00:00, crucible_mon36, MSWinSvr, Microsoft: Windows Server, MSWinSvrBIOSConf, Microsoft: Windows Server BIOS Configuration.

3. In the **See Table Schema** field, type the name of the dataset and press **[Enter]**. A list of metrics from that dataset are added below that field, and a preview of the table is displayed in a table to the right of that column. Each row in the preview table at the right reflects a set of data representing all of the presentation objects in the PowerPack, along with the device information.

In the preview table, you can check to make sure that this dataset contains the data you need. You can also see how current the data is by checking the timestamp in the **DataDate** column.

Viewing Dashboards and Charts in Data Visualization

The Data Visualization component of Skylar Analytics contains dashboards and charts based on data gathered by Skylar AI and Skylar One.

A **dashboard** in Skylar Analytics is similar to a dashboard in Skylar One, in that they both contain a number of graphical "widgets" that display data in a variety of ways, such as pie charts, line graphs, maps, bar charts, and other visualizations. A **chart** in Skylar Analytics works much like a "widget" in Skylar One, in that a chart in Skylar Analytics is a building block that makes up a dashboard, and a dashboard can contain many charts.

IMPORTANT: The dashboards and charts in the Data Visualization component of Skylar Analytics are *not* compatible with Skylar One dashboards, widgets, or reports.

Unlike dashboards in Skylar One, a dashboard in Skylar Analytics is used only for laying out the various charts that make up that dashboard. You can use charts to customize the data. One significant difference is that a chart, when modified, impacts all dashboards using that chart definition.

TIP: As a best practice, you should make a copy of a chart if you want to modify that chart for different analyses on different dashboards.

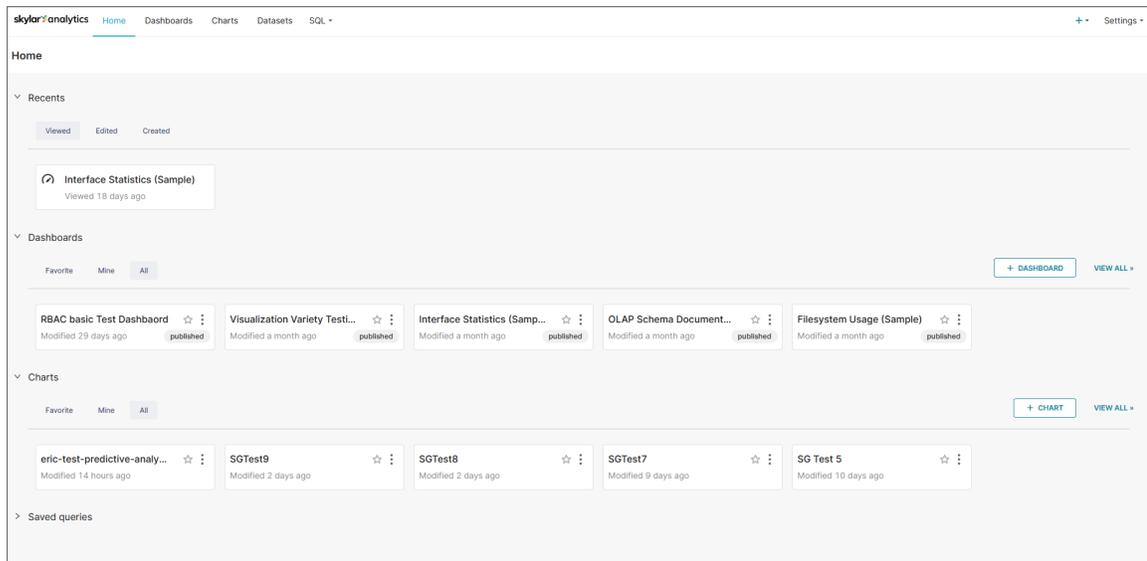
Logging In to the Data Visualization Component

To log in to the Data Visualization component of Skylar Analytics:

1. From Skylar One, go to the **Skylar AI** page (🚩) and click the **[Visit]** button next to **Skylar Data Visualization**. If you are not currently logged into Skylar AI, the Skylar AI login page appears. If not, log in and click the link for **Analytics**.

TIP: If you know the URL of your Skylar AI system, you can go to that location instead of using Skylar One.

2. Click the link for **Analytics** and, if needed, type in your user name and password. The **Home** page for the Data Visualization component of Skylar Analytics appears:



The **Home** page contains links to the dashboards and charts that you have used the most, including those that you have marked as favorites (★). You can create a dashboard or a chart from this page, and you can view all dashboards or charts by clicking the corresponding **View All** link.

3. Click a dashboard or chart from the **Home** page, or click the **Dashboards** page or the **Charts** page to view a list of all dashboards or charts.

TIP: In the list of dashboards on the **Dashboards** page, you can hover over an icon in the **Actions** column to delete, export, or edit (if you are an owner user) a specific dashboard. You can also perform the same actions with charts on the **Charts** page.

4. For more information about viewing existing dashboards, see [Viewing Skylar Analytics Dashboards](#).
5. For more information about creating or customizing dashboards, see [Creating and Customizing Dashboards and Charts](#).

Default Skylar Analytics Dashboards

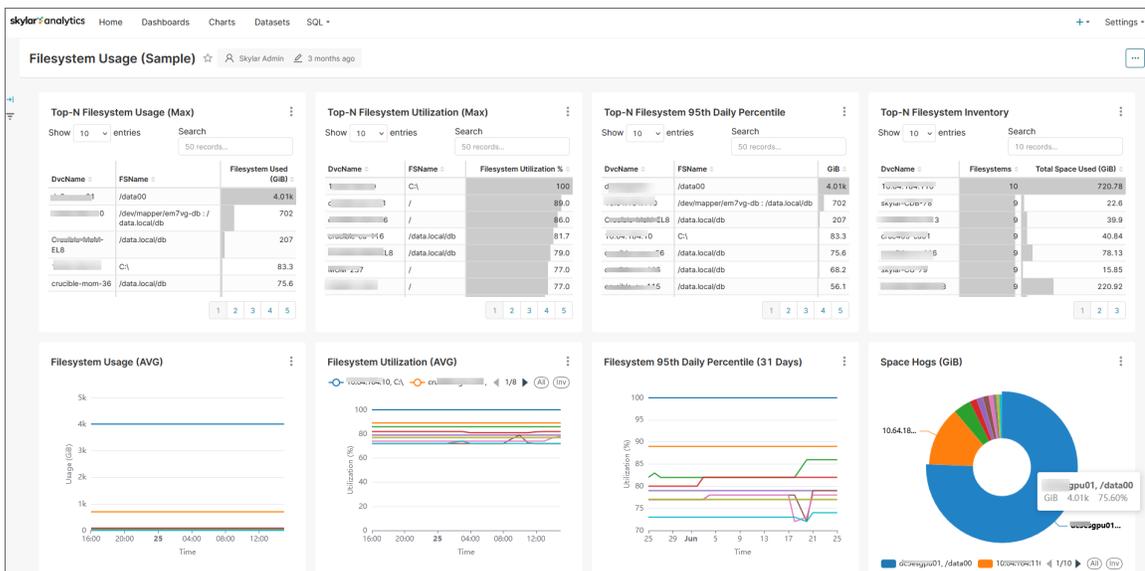
Skylar Analytics includes a set of default dashboards created by ScienceLogic that you can use to view data. You can also customize these dashboards as needed.

On the **Dashboards** page ( > Analytics > Settings > Dashboard Mgmt) in Skylar Settings, owner users can also install dashboards with **(Sample)** in their names by clicking the **[Add to Skylar]** button. These dashboards display a variety of visualization or chart configurations to show users different ways to display data, and users can reference these dashboards as examples. Many of the **(Sample)** dashboard layouts display multiple visualizations of the same raw data that would not typically be used at the same time on a production dashboard. These charts can also be copied and modified as needed, saving

development time when building new dashboards. ScienceLogic recommends that you keep the original versions of these (**Sample**) dashboard as unpublished draft dashboards and use them only for reference. For more information, see [Adding and Upgrading Dashboards](#).

The **Dashboards** page for Skylar Analytics contains the following default dashboards:

- **Filesystem Overview + Exploration (Sample).**
 - Displays 95th percentile data, file system utilization distribution (as a percentage and Gigabit or GiB), and "Space Hogs" (the devices using the most file system space).
 - You can click a device name on the "Space Hogs" pie chart to display chart details specifically for that device.
 - Also includes the **[Ad-Hoc Comparative Analysis]** tab, which displays additional file system charts for all devices or selected devices from the **[Overview]** tab.
- **Filesystem Statistics (Sample).** Displays a pie chart of "Space Hogs" (the devices using the most file system space), file system utilization as a percentage, file system inventory by host, and file system usage distribution.
- **Filesystem Usage (Sample).**
 - Displays a set of file system usage, utilization, 95th percentile and Top-N inventory charts for all devices, including a pie chart of "Space Hogs" (the devices using the most file system space).
 - You can click a device name on the "Space Hogs" pie chart to display chart details specifically for that device.



- **Interface Statistics (Sample).** Displays interface traffic in a variety of charts, including active hosts, active interfaces, dropped packets, and 95th percentile for the last 30 days (as a percentage and MIBPs).

- **Most Significant Resource Changes (Sample).**
 - Displays devices with the highest delta of file system usage, along with average file system usage, Top-N interface usage delta, and interface traffic in the past seven days.
 - You can click a device name on the "Top-N Filesystem Usage" or the "Top-N Interface Usage" tables to display chart details specifically for that device.
- **OLAP Schema Documentation (Skylar).** Contains additional information about the structure of the charts in the Data Visualization component of Skylar Analytics. If you do not have this dashboard on the **Dashboards** page in Skylar Analytics, you can add it from the **Dashboards** page (Analytics Admin > Dashboards) in Skylar Settings.
- **Visualization Variety Testing (Sample).** Contains a variety of chart visualizations related to file system utilization, including a table, a "big number" with a line graph, a gauge, a set of tree maps, and a sunburst map. You can use this dashboard to see how these different types of charts might work for your data.

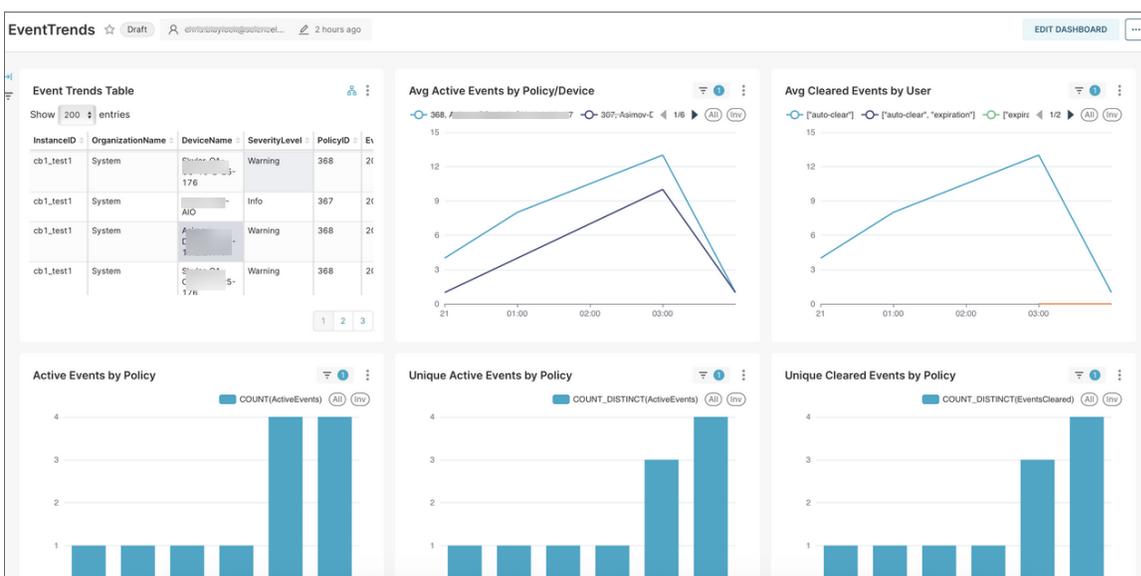
Recommended Datasets

If you want to view Skylar One event data in Skylar Analytics, ScienceLogic recommends the following datasets and their related dashboards and charts:

- **evtstat**
- **EventTrendsDevice**
- **EventTrendsDeviceGroup**

These datasets and their related dashboards and charts let you view Skylar One event data in the Data Visualization component of Skylar Analytics. The data is similar to the data on the **Events** page in Skylar One, but the data is focused more on historical trends as opposed to what is happening right now in your environment.

The following Data Visualization dashboard includes data from the **EventTrendsDevice** dataset:



NOTE: These datasets contain raw data and do not have user-friendly column names.

These datasets include the following columns:

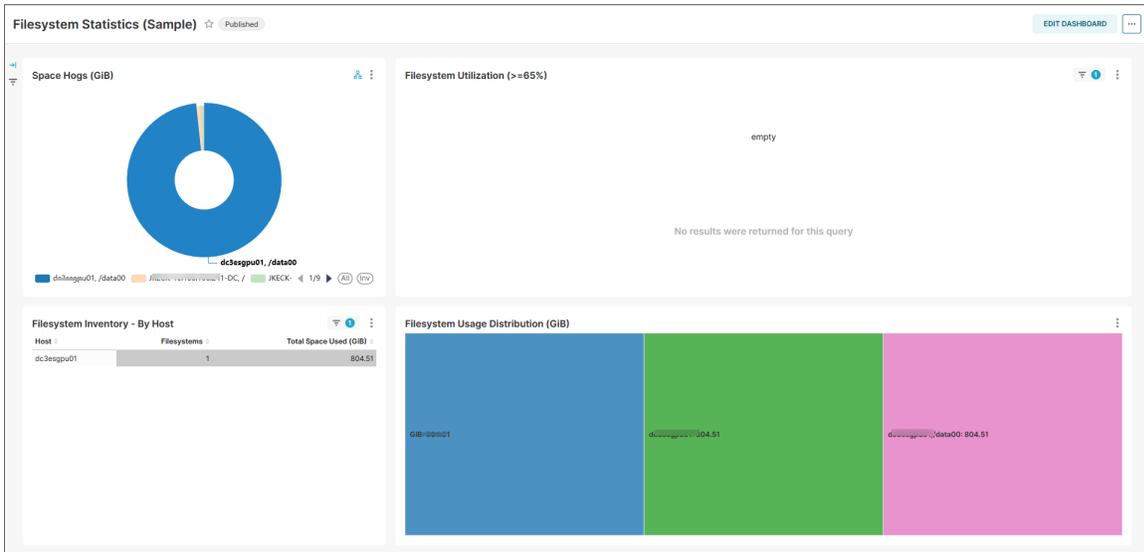
- **EventsCreated.** Events opened this hour.
- **LowLifeSpanEvents.** Events opened and then closed this hour, shows life span.
- **RunningTotalEventsResolved.** All events closed this hour, including events created before this hour.
- **EventsWithTickets.** Number of tickets for this grouping of events.
- **EventsWithDuplicates.** Number of events with duplicates in this grouping.
- **DuplicateNotificationCount.** Number of duplicates of events in this grouping.
- **PolicyName.** Name of the event policy.
- **ClearedBy.** User or action of clearing the event.
- **AutoClearedCount.** Number of events cleared by auto-clear.
- **ManualClearedCount.** Number of events cleared by users.
- **AvgResolutionHours.** Average length of time the events in this grouping stayed open before being closed this hour.
- **AvgActiveEventAgeHours.** Average time events in this grouping have stayed open and are still not closed.
- **MaxEventAgeHours.** Longest running event in this grouping.

Viewing Skylar Analytics Dashboards

You can use the following tips to get more data from your Skylar Analytics dashboards:

- You can hover over a graphical element in a chart, such as a piece of a pie chart or a colored metric in a tree map, to view a pop-up with more information about that element.
- If a dashboard is editable, you can click **[Edit Dashboard]** to make changes to the dashboard and the charts that make up the dashboard. For more information, see [Creating and Customizing Dashboards and Charts](#).
- For most dashboards, you can click a single device or item in the first chart at the top left of the **Dashboard** page (or any "Top-N" chart types) to view data specific to just that device. Click the device a second time to clear the filter. For more information, see [Adding Contextual Cross-filtering to a Dashboard](#).

The following image displays a dashboard with a device selected in the "Space Hogs" graph that forces the other graphs to only display data for that device:



When viewing a dashboard, you can click the ellipsis button (⋮) at the top right of the **Dashboard** page to open a menu with the following dashboard options:

- *Refresh dashboard.* Updates all of the charts in the dashboard to account for any changes you might have made.
- *Enter fullscreen.* Displays the browser window containing the dashboard display as full screen. Select *Exit fullscreen* from the menu to return to the previous setting.
- *Save as.* If a dashboard is editable, lets you save a copy of the dashboard, with the option of overwriting the existing dashboard or changing the name to make a new dashboard (if you have appropriate permissions).
- *Download.* Lets you export the dashboard as a PDF or download the dashboard as an image.
- *Share.* Lets you copy a link to the chart to the clipboard of your computer, and also lets you share a link to a chart using email.
- *Set auto-refresh interval.* Lets you choose how often you want Skylar Analytics to update the data for the dashboard. The default is *Don't refresh*.

On a **Dashboard** page, you can also click the vertical ellipsis button (⋮) at the top right of a *chart* on the dashboard to open a menu with the following chart options:

- *Enter fullscreen.* Displays the browser window containing just this chart display as full screen. Click the exit fullscreen icon (⌵) or select *Exit fullscreen* from the menu to return to the previous setting.
- *Edit chart.* Opens the **Edit Chart** page so you can add metrics, edit queries, and make other updates to this chart. Click **[Save]** to keep your changes (if you have appropriate permissions).
- *Cross-filtering scoping.* Lets you add **cross-filtering**, where you apply a data element from a chart (like a table row or a slice from a pie chart) and then apply it as a filter across all eligible charts in the dashboard. For more information, see [Adding Contextual Cross-filtering to a Dashboard](#).

- *View query.* Displays the SQL query for that chart. You can use this option to determine which data from the dataset is being used in this chart, and how the data is being used.
- *View as table.* Displays the chart in table format.
- *Drill to detail.* Displays all the data that makes up a chart.
- *Share.* Lets you copy a shareable chart link to your system's clipboard, or launches your system's default email client and composes a new message featuring the chart URL.
- *Download.* Lets you export the chart to .CSV or Excel, or you can download the chart as an image.

Creating and Customizing Dashboards and Charts

You can create a new dashboard in Skylar Analytics, or you can customize any of the default dashboards and save them with a new name. You can also create and customize the charts that make up the various dashboards.

TIP: To optimize dashboard speed, you should always try to use the smallest table needed for the dashboard.

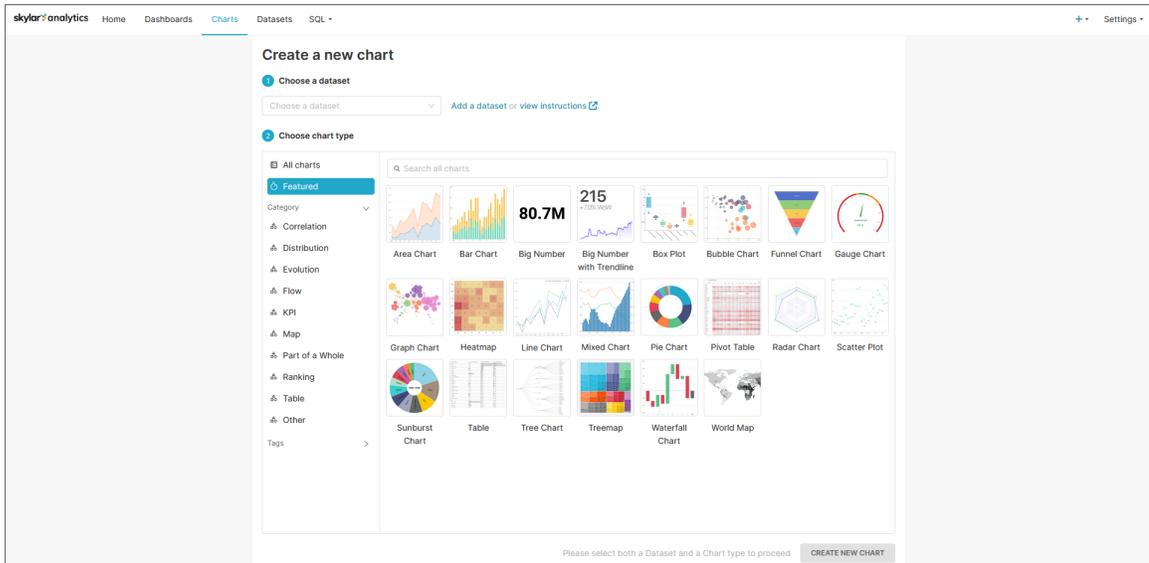
Creating a Dashboard

To create a dashboard:

1. Log in to the Skylar AI user interface and click the link for **Analytics**. The **Home** page for Skylar Analytics Data Visualization appears.
2. In the **Dashboards** section of the **Home** page, click the **[+Dashboard]** button. An **Untitled Dashboard** page appears.
3. Triple-click in the **[untitled dashboard]** field at the top right and type a name for the new dashboard. If you are using a shared system, you might want to add your initials to the end of the name.
4. Click **[Save]** in the upper right corner of the page.
5. Click **[Edit the Dashboard]**.
6. If there are existing charts that you want to add to this dashboard, click and drag each chart from the **[Charts]** tab on the right and drag the chart onto the dashboard. Click **[Save]** when you are done, and click **[Edit Dashboard]** again to keep editing.

TIP: If you want to see only the charts that you have created, check **Show only my charts**. If you want to see charts by all users, clear this option.

7. If you have not yet created any charts, or no charts exist on your system from other users, click **[Create New Chart]**. The **Create a new chart** window appears:

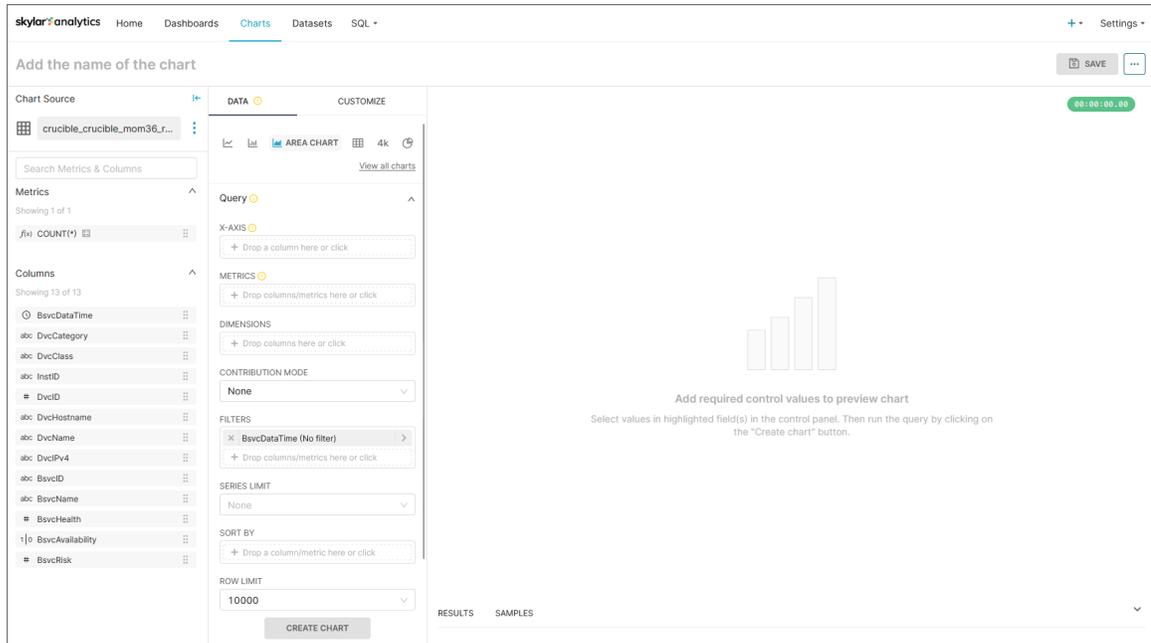


8. In the **Choose a Dataset** field, click to choose a dataset with the data you want to view in your new dashboard. In Skylar Analytics, a **dataset** contains a set of related metrics pulled from Dynamic Applications in Skylar One, such as server reports or Skylar One business service statistics.

A dataset that has "Current" at the end of its name contains the latest updated configuration data collected for that dataset. A dataset that has "Statistics" at the end of its name includes the time series metric data.

For this overview, we will select the *BusinessServiceStatistics* dataset, which contains data about Skylar One business services.

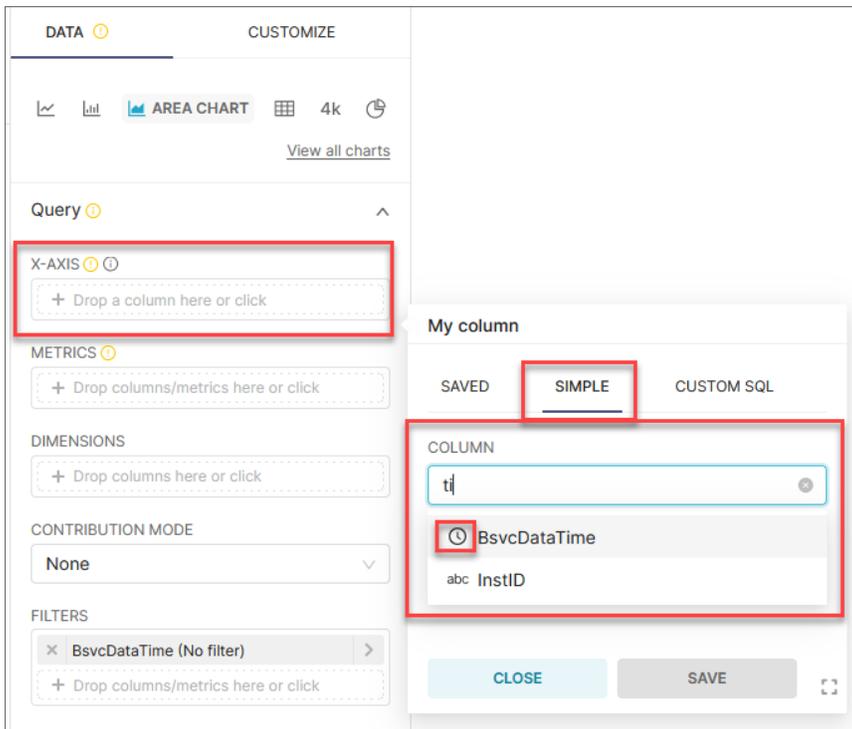
9. Select a chart type from the **Choose chart type** section and click **[Create New Chart]**. For this overview, we will select *Area Chart*. A new chart window appears:



In the first column, the **Chart Source** field displays the dataset you selected (for this overview, it is the *BusinessServiceStatistics* chart source). Below that field, you can access the metrics and columns that you can add to the chart. You can drag a metric or column from the first column into the second column to add it to the chart.

In the second column, you can select which data will appear in the chart, and how the data will be displayed in the chart. The large section to the right displays a preview of the chart as you build it after you click the **[Create Chart]** button to run the query.

- For example, with an Area Chart type, you could define the X-axis of the chart to display a time range by clicking in the **X-axis** field in the **Query** section. A modal appears:



- On the **[Simple]** tab of the modal, click the **Column** field to get a list of data. You can pre-filter the data by typing a column name or label, such as "time".
- Select a column with a calendar icon (📅) next to it to display a time range on the X-axis, such as *BsvcDateTime* from the *BusinessServiceStatistic* chart source, and click **[Save]**.

TIP: For more information about the abbreviations used for the metric names, see [Mapping Skylar One Dynamic Application Object Names to Skylar Columns](#).

- To set the granularity of the time frame to a shorter time frame, click in the **Time Grain** field and select *Hour* instead of the default of *Day*.

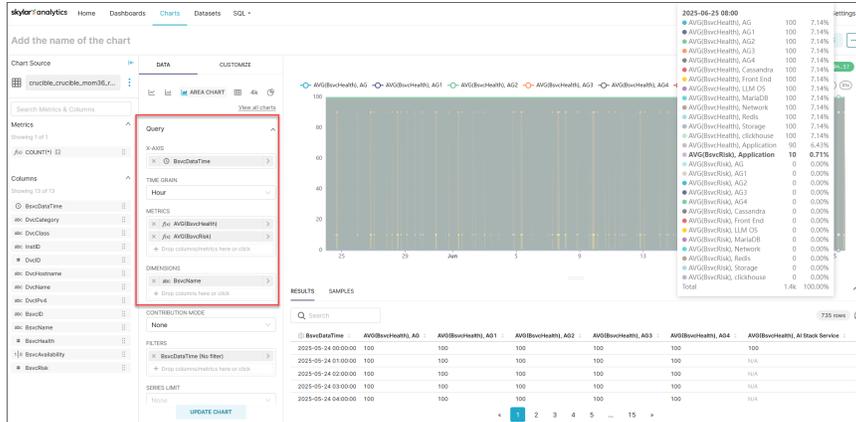
- Next, select the metrics you want to visualize in the chart by clicking in the **Metrics** field and clicking the **[Simple]** tab. You can also drag a metric from the first column and drop it on this field.

For this overview, we will select *BsvcHealth* (business service health) in the **Column** field and *AVG* in the **Aggregate** field. We will also select *BsvcRisk* (business service risk) in the **Column** field and *AVG* in the **Aggregate** field. These settings will show the average values for business service health and risk over time.

TIP: In the **Column** field on the **[Simple]** tab, type % to filter the list down to Utilization or Percentages.

- Click **[Save]** to save the metrics.
- To see a preview of the chart so far, click the **[Update Chart]** button or the **click here** link in the large section to the right. You will need to do this every time you make a change if you want to see the latest preview.
- You can use the **Dimensions** field to add descriptive elements to the chart that help users understand the data being visualized. For example, for line charts and area charts, the dimensions will appear in the legends and mouseover text. For tables, dimensions represent the columns to display.

For this overview, we will add *BsvcName* (business service name) to the **Dimensions** field. Click the **[Update Chart]** button to see an updated preview:



The chart legend displays the average health and average risk in the legend at the top, and also in the columns at the bottom of the section. If you mouse over a line in the chart, you can see the specific data for those values.

- In the **Filters** field, you can edit the existing filter by clicking on it and specifying what data to display on the new chart. The filter currently has no specific filter set right now.

For this overview, click on *BsvcDataTime* in the **Filters** field and then click in the **Time Range** field. An **Edit time range** modal appears.

- In the **Range Type** field, select a range, such as *Last*, as in *Last day*, *Last week*, and so on.

20. Select the *Last week* option and click **[Apply]**.
21. Click the **[Update Chart]** button to review your updates.
22. To finish the chart, be sure to give it a name in the top right of the window, and then click the **[Save]** button.
23. In the **Save chart** modal, click **[Save & Go to Dashboard]**. The new chart is added to your dashboard.
24. Continue adding charts to the dashboard as needed.
25. When your dashboard is complete, click the **[Draft]** button at top left to publish it. The button changes from **[Draft]** to **[Published]**.

TIP: For more information, see [Creating Your First Dashboard and registering a new table](#) in the Superset documentation.

Adding Contextual Cross-filtering to a Dashboard

You can add contextual cross-filtering (also called "context") to the widgets to create an interactive dashboard in Skylar Analytics. When you do, you can click on a device in one widget to make another widget in the dashboard display data specific to that device.

For this process, we will use the chart and the dashboard that we created in the previous procedure and add a new chart to that dashboard as an example.

Adding context to a dashboard:

1. Select the dashboard from the **Dashboards** page. You can also hover over the dashboard and click the edit icon (✎) in the **Actions** column.
2. Click **[Edit Dashboard]**. The **[Charts]** and **[Layout Elements]** tabs appear on the right.
3. On the **[Charts]** tab, click **[Create New Chart]**. The **Create a new chart** window appears.
4. In the **Choose a Dataset** field, click to choose a dataset with the data you want to view in your new dashboard.

A dataset that has "Current" at the end of its name contains the latest updated configuration data collected for that dataset. A dataset that has "Statistics" at the end of its name includes the time series metric data.

Because the other chart we just created used the *BusinessServiceStatistics* dataset, select the *BusinessServiceCurrent* dataset as a complement to the first chart.

5. Select a chart type of *Table*, which will make it easy for users of the dashboard to select a business service to view more information.
6. Click **[Create New Chart]**. A new chart window appears.

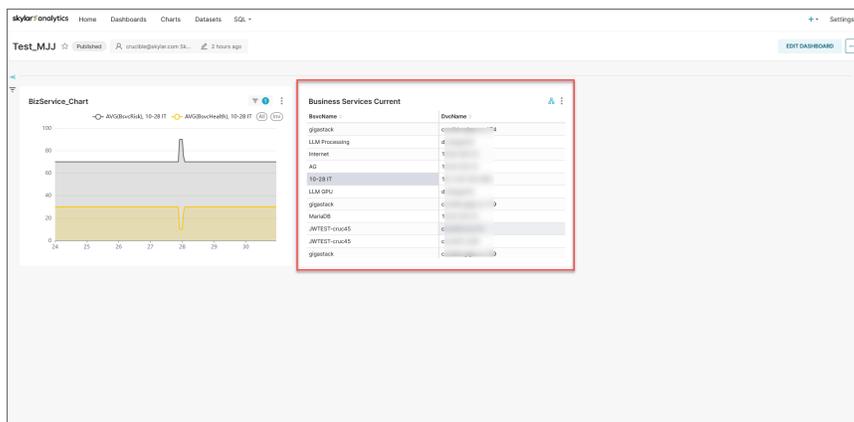
- Click and drag the *BsvcName* column onto the **Dimensions** field on the **[Data]** tab. This will add a list of business services as the first column in the new table.

TIP: You can also click in the **Dimensions** field on the **[Data]** tab, go to the **Column** field on the **[Simple]** tab, and then select *BsvcName*.

- Click and drag the *DvcName* column onto the **Dimensions** field, under *BsvcName*. This will add device names as the second column in the new table.

NOTE: Because you are not using this chart to display metrics, you do not need to add any values to the Metrics field.

- Click **[Create Chart]**. The chart displays in the preview area.
- Click **[Save]** to save the metrics. The **Save chart** modal appears.
- As needed, add a name and select a dashboard for the chart, and then click **[Save & Go to Dashboard]**. The new chart is added to your dashboard, to the right of the first dashboard:

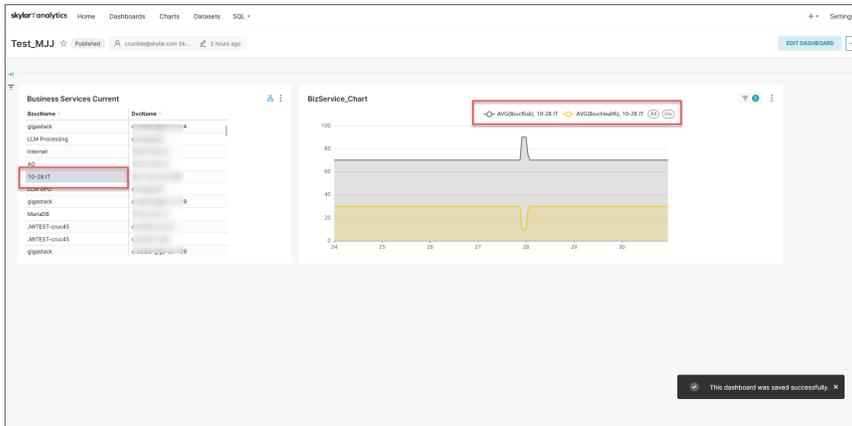


- Typically, you would want to place the chart that drives the "context" (or the contextual cross-filtering) first, so click **[Edit Dashboard]** to enter Edit mode again.
- Click and drag the next chart to the left of the first chart until a vertical rectangle appears. Drop the new chart onto that rectangle.

TIP: While the dashboard is in Edit mode, you can also resize a chart by hovering over a corner and then dragging the arrows to change the size.

- Click **[Save]**.

15. When you select a business service or a device in the first chart, the second chart updates with only the data for the selected item:



NOTE: If the second chart displays "No results were returned for this query", you might need to click the vertical ellipsis icon (\vdots) for the second chart and select *Edit chart* to address the issue.

Customizing a Dashboard

To customize a dashboard:

1. Select the dashboard from the **Dashboards** page. You can also hover over the dashboard and click the edit icon (\pencil) in the **Actions** column.
2. On the **Dashboard** page, click **Edit Dashboard**. The **Charts** and **Layout Elements** tabs appear.
3. If there are existing charts that you want to add to this dashboard, click and drag each chart from the **Charts** tab on the right and drag the chart onto the dashboard. Click **Save** when you are done, and click **Edit Dashboard** again to keep editing.

TIP: If you want to see only the charts that you have created, check **Show only my charts**. If you want to see charts by all users, clear this option.

4. If you want to add extra elements to your dashboard, like a header, additional text, a divider, or other items, drag and drop the items onto the dashboard from the **Layout Elements** tab. Click **Save** when you are done, and click **Edit Dashboard** again to keep editing.
5. To edit a chart in the dashboard, click the vertical ellipsis icon (\vdots) at the top right of the chart on the dashboard and click the **Edit chart** link. For more information, see [steps 8-21](#) in the "To create a dashboard" procedure.

6. When you are done updating the dashboard, you might need to click **[Edit Dashboard]** and rename the dashboard if the dashboard was created by ScienceLogic, with the word "(Sample)" or "(Skylar)" at the end of the name.

TIP: On the **Dashboards** tab in Skylar Analytics, the "Visualization Variety Testing (Sample)" dashboard contains a variety of chart visualizations related to file system utilization, including a table, a "big number" with a line graph, a gauge, a set of tree maps, and a sunburst map. You can use this dashboard to see how these different types of charts might work for your data.

Icons for Chart Metrics

Each data type includes a small icon that conveys its type:

- **abc**: Text data
- **#**: Numeric value data
- **🕒**: The time column for the data source
- **f(x)**: Function used for metrics

Customizing the Default Column Names for Charts

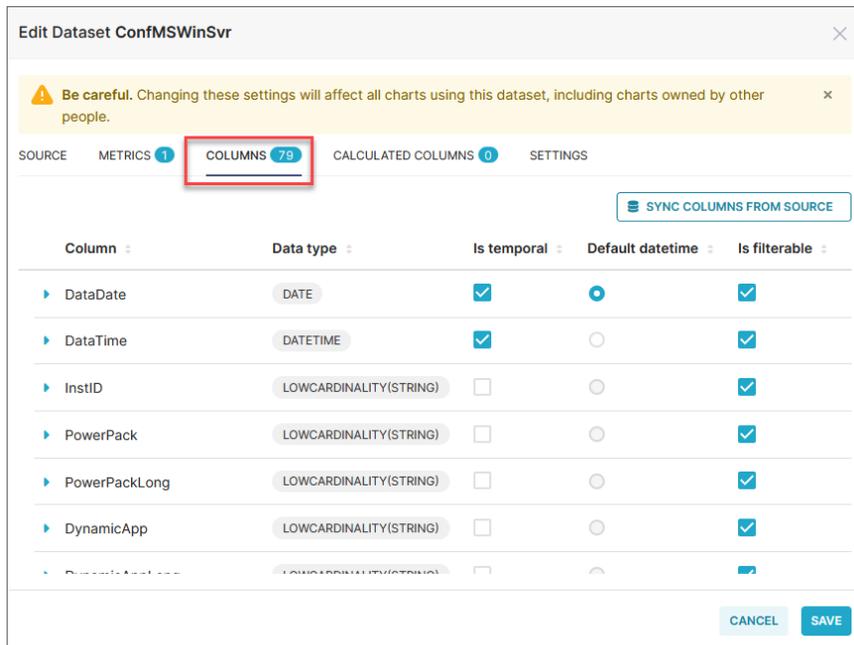
You can customize the default column names that Skylar Analytics created for the charts in your dashboards to make the names more useful for your users. You can rename columns by using the Label and Description parameters in a dataset.

IMPORTANT: Do not edit the columns at the **chart** or **dashboard** level in Skylar Analytics, as doing so will break the contextual cross-filtering. Instead, edit the columns at the **dataset** level, as discussed in the following procedure. Also, make sure that the icon to the left of the renamed column is **abc** and not **f(x)**, because string values can drive context, but functions cannot drive context.

To customize default column names:

1. In Skylar Analytics, open to the chart in *Edit mode* and make a note of the dataset name in the Chart Source field. The dataset name is at the very end of the chart source name, such as "BusinessServiceCurrent".
2. Go to the **Datasets** page and locate that dataset.

3. Hover over the dataset and click the edit icon (✎) in the **Actions** column. The **Edit Dataset** dialog appears:



NOTE: You must be the dataset owner to edit a dataset.

4. Click the **[Columns]** tab to view all of the columns in that dataset.

5. Locate the column you want to rename and click the expand icon (▸) next to the current column name:

Be careful. Changing these settings will affect all charts using this dataset, including charts owned by other people.

SOURCE METRICS 1 COLUMNS 79 CALCULATED COLUMNS 0 SETTINGS

SYNC COLUMNS FROM SOURCE

Column	Data type	Is temporal	Default datetime	Is filterable
▸ DvcHostname	STRING	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>
▾ DvcName	STRING	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>

LABEL
Label

DESCRIPTION
Description

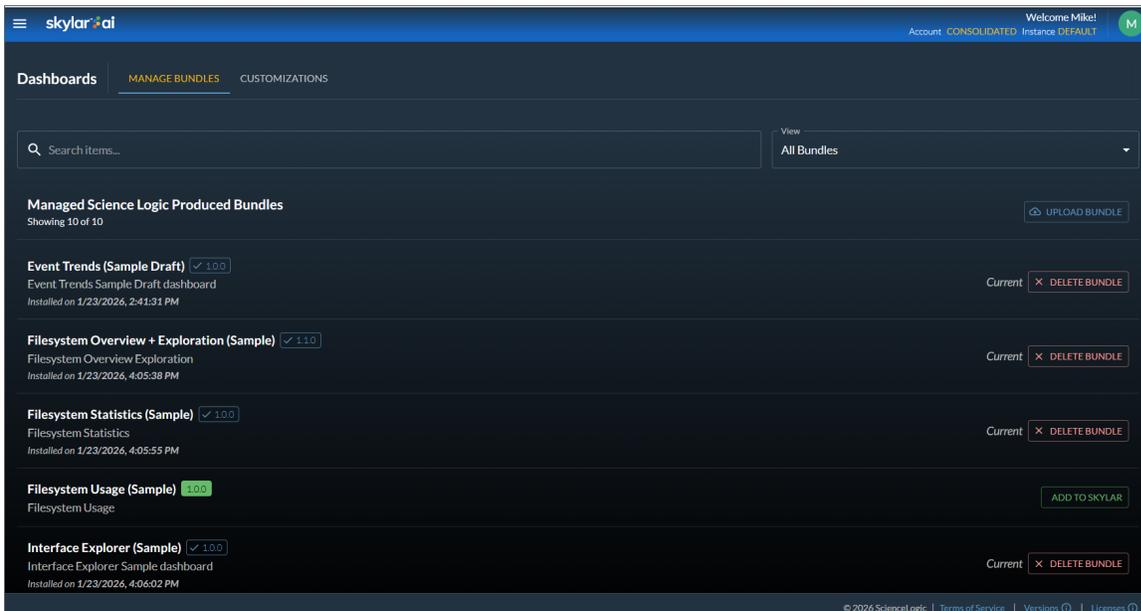
DATETIME FORMAT ⓘ
%Y-%m-%d

CANCEL SAVE

6. In the **Label** field, type the new column name you want to display in all charts that use this dataset.
7. In the **Description** field, type a short description of the column. This text, along with the **Label** text, appears when you hover over that column in the **Columns** section when you are editing the chart.
8. Click **[Save]**. A confirmation message appears to remind you that these edits will affect all charts that use this dataset.
9. Click **[OK]**. The column label and description are updated for all charts that use this dataset.
10. Repeat steps 5-9 for any other column names that you want to rename.

Adding and Upgrading Dashboards and Datasets

A user with an owner role can manage the Skylar Analytics dashboards on the **[Manage Bundles]** tab of the **Dashboards** page ( > Analytics > Settings > Dashboard Mgmt):



You can search for dashboard bundles and sort the list of bundles by *All Bundles*, *Installed*, *Not Installed*, and *Updates Available*.

Owner users can install dashboards with **(Sample)** in their names by clicking the **[Add to Skylar]** button. These dashboards display a variety of visualization or chart configurations to show users different ways to display data, and users can reference these dashboards as examples. Many of the **(Sample)** dashboard layouts display multiple visualizations of the same raw data that would not typically be used at the same time on a production dashboard. These charts can also be copied and modified as needed, saving development time when building new dashboards.

ScienceLogic recommends that you keep the original versions of these **(Sample)** dashboard as unpublished draft dashboards and use them only for reference. For more information, see [Adding and Upgrading Dashboards](#).

Other Dashboard Options

You have the following options on the **[Manage Bundles]** tab of the **Dashboards** page:

- **[Upload Bundle]**. Upload a .zip file containing a Skylar Analytics dashboard that has been exported from a Skylar Analytics system. This feature lets you share dashboards that were created and then exported by other users, for example.
- **[Add to Skylar]**. Click this button to install a new dashboard for Skylar Analytics

- **[Upgrade Now]**. Click this button to upgrade an existing dashboard. **Current** displays next to a dashboard to show that you are running the most recent version of that dashboard.
- **[Delete Bundle]**. Remove this bundle from the Skylar Analytics system.

Update All Datasets

In addition, you can use the **[Sync Skylar Datasets]** button on the **[Customizations]** tab on the **Dashboards** page to update all of your datasets based on Skylar One PowerPacks, including PowerPacks that have been updated in Skylar One.

If all datasets have been updated, the button does not appear, and the text "Datasets are current" appears instead. This button is only available to owner users in Skylar AI.

Data Exploration: Exporting Data to Skylar AI from Third-party Tools

You can use the optional Data Exploration component of Skylar Analytics to enable Open Database Connectivity (ODBC) to connect Skylar AI data with third-party tools like Grafana, Power BI, Tableau, Cognos, Crystal Reports, SAP, Excel, and other business intelligence applications. When the tool is connected using ODBC, you can export data from Skylar Analytics to that third-party tool.

You can also import data from third-party tools, such as billing data, environmental data, or service level objectives (SLOs), and then use that data in Skylar AI.

Data Exploration with ODBC lets you view Skylar AI data alongside other business sources, offering a holistic perspective on your operations.

Configuring Data Exploration with Power BI

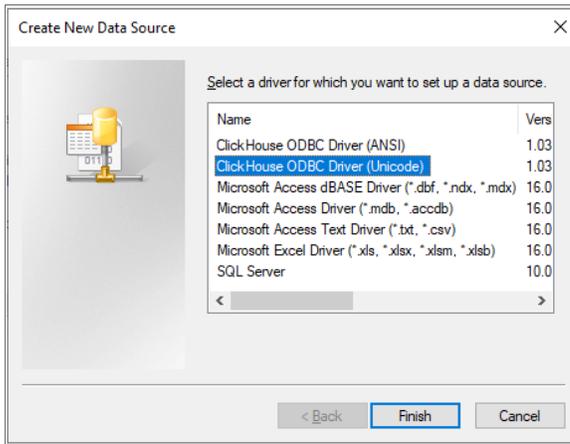
This section covers how to set up an ODBC connection for Skylar Analytics so you can use it with Power BI for data visualization. Other business intelligence applications will use a similar process to integrate with Skylar Analytics.

TIP: For an example of how you can connect DBeaver, another business intelligence tool, to Skylar Analytics, see the following video: [BYO Tool and Data: Connect via ODBC and Import Data](#).

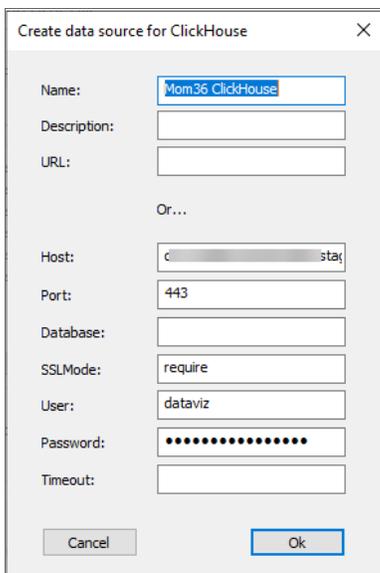
To install and configure the ODBC connection:

1. Go to the **ClickHouse ODBC driver releases** page at <https://github.com/ClickHouse/clickhouse-odbc/releases>.
2. Download the relevant version for your operating system.
3. Open the ODBC Data Source Administrator application.

4. On the [User DSN] tab, click [Add]. The **Create New Data Source** dialog appears:



5. Select **ClickHouse ODBC (Unicode)** and click [Finish]. The **Create data source for Clickhouse** dialog appears:



6. Complete the following fields with ODBC connection details from the ScienceLogic Site Reliability Engineering (SRE) team:

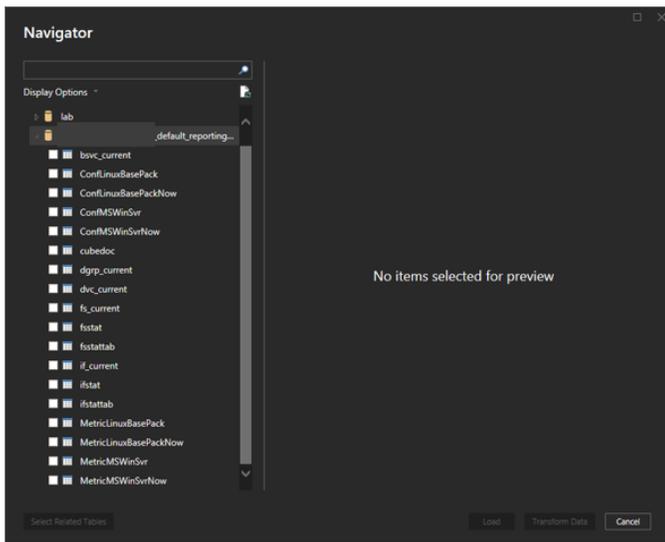
- **Name:** Add a name to identify this connection. This will be used later in the BI tools.
- **Host:** Specify the host URL, provided by SRE.
- **Port:** 443.
- **Database:** Leave blank.
- **SSLMode:** Type the word "require".

- **User:** dataviz
- **Password:** Specify the password, provided by SRE.

To connect your BI tool, such as the Power BI Desktop:

TIP: For an example of how to connect Power BI, see [How to connect Power BI to Skylar AI via ODBC](#).

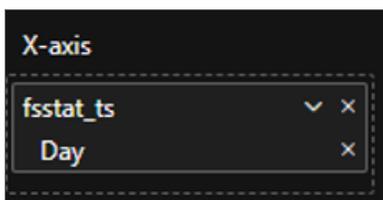
1. Launch the Power BI Desktop and click **[Blank Report]**.
2. Click **Get data from another source**, select **Other**, and then select **ODBC**.
3. Click **[Connect]**.
4. In the pop-up window, click the drop-down menu and select the ODBC connection you just created in the previous procedure.
5. Click **[OK]**.
6. If prompted, re-enter your username and password, and then click **[Connect]**.
7. After you are connected, a menu will appear displaying available datasets, which you can use to create dashboards in your BI tool:



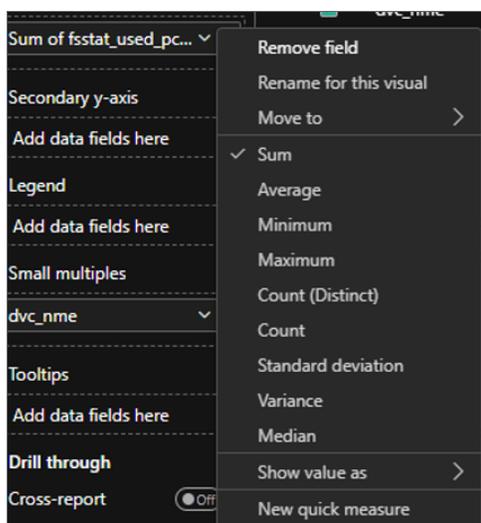
TIP: When selecting datasets to import, choose only the necessary tables to optimize performance. The following procedure creates a sample dashboard in Power BI.

To import data and create a dashboard with Skylar AI data in Power BI:

1. On the **Home** screen of the Power BI Desktop, click **[New Visual]**.
2. Select a Line Chart as an example.
3. To configure the X-Axis, expand the **fsstattab** dataset from the right-hand Data Column.
4. Drag **fsstat_ts** (Timestamp) to the X-Axis in the **Visualizations** panel.
5. Remove the options for *Year*, *Quarter*, and *Month*, keeping only *Day*.



6. To configure the Y-Axis, drag **fsstat_used_pct_psec** (Used Percentage Per Second) to the Y-Axis.
7. To customize the data fields, click the drop-down arrow next to the selected data field. You can rename the field or modify how the value is calculated:



8. Continue adding additional charts and visuals as needed to finish up your dashboard.

Additional Resources for Skylar Analytics (Apache Superset Training)

This section has been provided as an independent study guide to help you identify and develop basic knowledge and skills to build data visualizations within Skylar Analytics user interface.

The following videos from ScienceLogic cover some of the key features of Data Visualization and Exploration:

- [How to create your first Dashboard](#): How to create a dashboard and a chart, and how to configure the axis, time grain, dimensions, metrics, and filters.
- [How to add cross-filtering to a Dashboard](#): How to adding a second chart to a dashboard, how to correct a filtering issue, how to reconfigure a chart to another dataset, and how to update a dashboard with contextual cross filtering (context).
- [Datasets Overview](#): Definition of a dataset, dataset naming conventions, how Dynamic Applications and PowerPacks from Skylar One display their data in Skylar Analytics, how to verify if you have current data from Skylar One in Analytics, how to create a dataset (temporary), and a new workflow for creating charts from datasets.
- [Tips & Tricks: Friendly Column Names](#): An explanation of the right way and the wrong way to rename columns in Skylar Analytics.
- [BYO Tool and Data: Connect via ODBC and Import Data](#): How to connect with ODBC using an open-source database tool called DBeaver (for Mac or PC), and how to import data from an external source and store it in Skylar "local" database schema.
- [How to connect Power BI to Skylar AI via ODBC](#): How to connect Power BI to Skylar AI.
- [Writable Schema, Joins and Views](#): How to bring in external pricing data and use it in a Skylar Analytics chart, and how to create a database view by using a JOIN query.

ScienceLogic recommends the following resources for a deeper understanding of Apache Superset:

- Apache Superset-related documentation: <https://superset.apache.org/docs/intro>
- Apache Superset Community: <https://superset.apache.org/community>
- Udemy course for Apache Superset: <https://www.udemy.com/course/apache-superset-for-data-engineers-hands-on/>
- What is Apache Superset - Quick Overview: https://www.youtube.com/watch?v=znmco3eK-M&list=PLzRV_ObjEwmNhRjhMNcvcDP7ZDjOXtodd

NOTE: Because ScienceLogic does not own the underlying framework for the Data Visualization and Data Exploration components, ScienceLogic is not responsible for maintaining or updating documentation for third-party open-source software, including Apache Superset.

Skylar Analytics: Anomaly Detection

Overview

The Anomaly Detection component of *Skylar Analytics* uses Skylar AI to identify unusual patterns that do not conform to expected behavior. Anomaly Detection provides always-on, unsupervised, machine-learning-based monitoring that automatically identifies unusual patterns in the real-time performance metrics and resource data that it observes. Anomalies do not necessarily represent problems or events to be concerned about; rather, they represent anomalous behavior that might require further investigation.

You can view device anomalies for each Dynamic Application metric on the **[Anomaly Detection]** tab on the **Device Investigator** page for each device. Anomaly Detection also computes an Anomaly Score that characterizes the significance of each anomaly.

<p>NOTE: Anomaly Detection with Skylar Analytics works with all of the Performance Dynamic Applications in all Skylar OnePowerPacks.</p>

This chapter covers the following topics:

What is Anomaly Detection?	118
Viewing Graphs and Data for Anomaly Detection	119
Enabling Anomaly Detection Events for Specific Metrics	122
Creating an Event Policy for Anomalies	124
Using Anomaly-related Events to Trigger Automated Run Book Actions	126

What is Anomaly Detection?

Anomaly detection is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. Anomaly detection provides always-on, unsupervised machine learning-based monitoring that automatically identifies unusual patterns in the real-time performance metrics and resource data that it observes.

Anomalies do not necessarily represent problems or events to be concerned about; rather, they represent unexpected behavior that might require further investigation.

Anomaly detection is calculated and displayed in the Skylar One user interface for all Dynamic Application metrics. This detection is enabled by default and cannot be disabled.

You can control which device data gets sent to Skylar for analysis based on the organization aligned with the device or devices. All devices in the selected organization will get anomaly detection analysis.

Skylar Analytics starts generating anomaly detection charts and alerts about six to eight hours after data starts getting exported from Skylar One to Skylar AI.

How Anomaly Detection Works

Initially, a historic profile for anomaly detecting is based on 24 hours of data. These values include minimum and maximum values, median lag differences, and median absolute deviation of those lag values (capturing the variance of lag values from the median lag value.)

Skylar AI uses these statistics to create bands at prediction time that determine anomalous and non-anomalous behavior.

Skylar AI periodically re-calculates and blends these values with the previously calculated values. In general, if the recent period shows more extreme behavior, then Skylar AI uses these values to update the model. If the recent period is less extreme, then the model statistics will move in the direction of these less extreme values.

At prediction time, the bands also take into consideration recent behavior that was deemed non-anomalous, allowing for gradual trends that go outside the pre-computed bands.

With the final min/max expected values computed, Skylar AI considers anything outside of those values to be anomalous. Skylar AI calculates a score based on the distance outside of the band, normalized by a value based on typical point-by-point changes.

Viewing Graphs and Data for Anomaly Detection

After Skylar One begins performing anomaly detection for a device, you can view graphs and data about each anomaly. Graphs for anomalies appear on the following pages in Skylar One:

- The Skylar One **Events** page, filtered by "Anomaly messages" (Skylar AI (🚩) > **[Visit]** button for Skylar Anomaly Detection).
- The **Anomaly Detection** page (Skylar AI (🚩) > **[Advanced: Anomaly Alerting]** button).
- The **[Anomaly Detection]** tab in the **Device Investigator**.
- The **[Anomaly Detection]** tab in the **Service Investigator** for a business, IT, or device service.

You can view the anomaly detection graphs for devices by clicking the open icon (↗) in the first column of the table on the inventory page. The **Anomaly Chart** modal appears, displaying the "Anomaly Score" chart above the chart for the specified metric you are monitoring.

The "Anomaly Score" chart displays a graph of values from 0 to 100 that represent how far the real data for a metric diverges from its expected values. The anomaly score indicates the significance of an anomaly, with a greater severity as the number gets bigger. The lines in the chart are color-coded by the severity level of the event that gets triggered as the data diverges further. The score is basically a running sum over a small window of time, so after the anomalies stop, the score will drop to zero over that time.



You can define the thresholds for the "Anomaly Score" chart on the **Anomaly Detection Thresholds** page (Skylar AI (🚩) > **[Advanced: Adjust Thresholds]** button). You can also use this page to specify whether

the Anomaly Score values generate alerts in Skylar One. For more information, see [Enabling Thresholds and Alerts for the Anomaly Chart](#).

The second graph displays the following data:

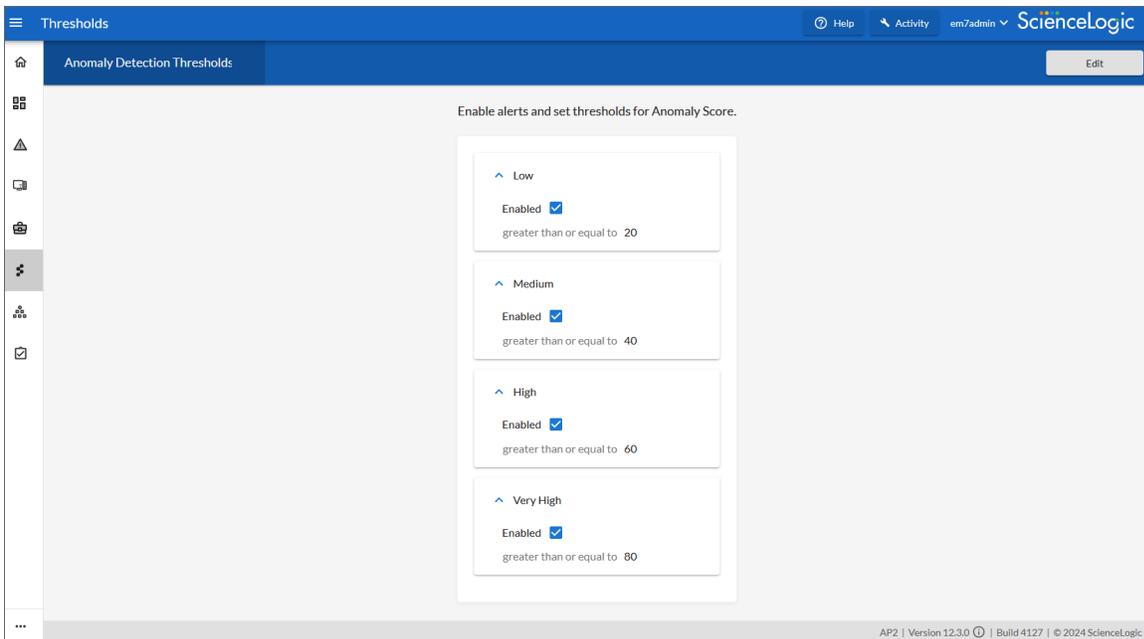
- A blue band representing the range of probable values that Skylar One expected for the device metric.
- A green line representing the actual value for the device metric.
- A red dot indicating anomalies where the actual value appears outside of the expected value range. The number of the red dots are listed in the **Anomaly Count** column on the **[Anomaly Detection]** tab of the **Device Investigator** page.

You can hover over a value in one of the charts to see a pop-up box with the **Expected Range** and the metric value. The **Anomaly Score** value also displays in the pop-up box, with the severity in parentheses: Normal, Low, Medium, High, or Very High.

You can zoom in on a shorter time frame by clicking and dragging your mouse over the part of the chart representing that time frame, and you can return to the original time span by clicking the **[Reset zoom]** button.

Enabling Thresholds and Alerts for the Anomaly Chart

You can define the thresholds for the "Anomaly Score" chart that displays on the **Anomaly Chart** modal, and whether those values generate alerts in Skylar One, on the **Anomaly Detection Thresholds** page (Skylar AI () > **[Advanced: Adjust Thresholds]** button).



You can view the alert levels when you hover over a value in one of the charts on the **Anomaly Chart** modal. The Anomaly Score severity level displays after the index value, in parentheses: Normal, Low, Medium, High, or Very High:



NOTE: An Anomaly Score severity level of **Normal** is assigned to a value in the chart that is *lower* than the lowest enabled alert level. For example, if the threshold for the Low severity is enabled and set to 20 or higher, an Anomaly Score of 16 would have a severity level of Normal.

To edit the Anomaly Score thresholds:

1. On the **Anomaly Detection Thresholds** page (Skylar AI (⚙️) > **[Advanced: Adjust Thresholds]** button), click **[Edit]**.
2. For each of the four severity levels, from Low to Very High, you can click to check **Enabled** to have Skylar One generate an alert when the Anomaly Score is equal to or greater than the threshold for that severity level.
3. You can edit the threshold value for each level if Skylar One is generating too many (or not enough) anomalies of a certain severity level.
4. For example, if you want to enable a Low level alert when the Anomaly Score value is between 25 and 39, you would go to the **Low** panel, select **Enabled**, and update the value from "20" to "25".
5. Click **[Save]**.
6. You can then edit an event policy that uses alerts based on the settings on this page to generate events in Skylar One. For more information, see [Creating an Event Policy for Anomalies](#).

TIP: To filter the list of devices on this page by name, type some or all of a device name in the **Search** field at the top of the window, based on the device-naming convention you used for your devices.

For each device in the list, the **Anomaly Detection** page displays the following information:

- **Device Name.** Displays the name of the device. Click the hyperlink to go to the **[Anomaly Detection]** tab of the **Device Investigator** page for that device. Each row on the **Anomaly Detection** page represents a specific device and metric for that device. As a result, a device might appear in the list multiple times if anomaly detection is enabled for multiple metrics on that device.
- **Metric Type.** Indicates the metric that Skylar One is evaluating for anomalies on the device.
- **ML Enabled By User.** Indicates the username of the user that enabled anomaly detection for the device and metric.
- **Last Modified.** Date the metric was most recently updated.
- **Class.** Displays the Device Class for the device.
- **Category.** Displays the device's Device Category.
- **Anomaly Count.** Displays the number of anomalies detected by Skylar One.

NOTE: On the **Anomaly Detection** page, the **Anomaly Count** column does not currently display the number of anomalies. Go to the **[Anomaly Detection]** tab on the **Device Investigator** page for a device to see the correct anomaly count. You can sort the **Anomaly Count** column to see which anomalies are happening the most often.

Enabling Anomaly Detection Events on the Anomaly Detection Page

To enable anomaly detection events for a metric on the **Anomaly Detection** page:

1. In Skylar One, go to the **Anomaly Detection** page in Skylar One (Skylar AI (🔗) > **[Advanced: Anomaly Alerting]** button)
2. On the **Anomaly Detection** page, click the **[Enable Alert Policies]** button. The **Enable Alert Policies** modal appears.

TIP: Alternately, you can click the **Actions** icon (⋮) for any of the listed devices and select **Enable**. You can also select multiple devices using the check box on the left and click the **[Enable Alert Policies]** button at the top.

3. Select the devices, Dynamic Application metrics, and indexes on which you want to enable alerting. This modal automatically filters out any devices that are not aligned to organizations that use Skylar AI. Click **[Next]**.
4. Select a Dynamic Application and its corresponding metric. For some metrics, a second drop-down field might display that enables you to specify the device directory. If this field appears, click the name of the directory on which you want to enable anomaly detection. Click **[Next]**.
5. Select an index and click **[Enable]**.

TIP: To disable anomaly detection events for a metric, click the **Actions** icon (⋮) for that metric and select *Disable*.

Enabling Anomaly Detection Events on the Device Investigator Page

To enable anomaly detection events for a metric on the **Device Investigator** page:

1. On the **Devices** page (📁), click the **Device Name** for the device on which you want to enable anomaly detection events and click the **[Anomaly Detection]** tab on the **Device Investigator** page.

TIP: If the **[Anomaly Detection]** tab does not already appear on the **Device Investigator**, click the **More** drop-down menu and select it from the list of tab options.

TIP: If your Skylar One system does not have any Dynamic Applications enabled, you will see only dashes (–) listed in the table on the **[Anomaly Detection]** tab for a device.

2. On the **[Anomaly Detection]** tab, click the **Actions** icon (⋮) for any of the listed metrics and select *Enable Alert Policy*. The **Enable Alert Policy** modal appears.
3. Click **[Enable]**. That metric is enabled for events for that device.

TIP: To disable anomaly detection events for a metric, click the **Actions** icon (⋮) for that metric and select *Disable Alerting*.

Creating an Event Policy for Anomalies

You can create additional event policies that will trigger events in Skylar One when anomalies are detected for those devices.

TIP: Because anomalies do not always correspond to problems, ScienceLogic recommends creating an event policy only for scenarios where anomalies appear to be correlated with some other behavior that you cannot otherwise track using an event or alert.

NOTE: Because the anomaly detection model is constantly being refined as Skylar One collects more data, you might experience a larger number of anomaly-related events if you create an event policy for anomalies soon after enabling anomaly detection compared to if you were to do so after Skylar One has had an opportunity to learn more about the device metric's data patterns.

The **Event Policies** page in Skylar One was completely updated in version 12.5.1. Use the following procedure if you are on Skylar One 12.5.1 or later, or use the next procedure if you are on an older version of Skylar One.

To create an event policy for anomalies in Skylar One version 12.5.1 or later:

1. Go to the **Event Policies** page (Events > Event Policies) and click the **[Create Event Policy]** button. The **[Basic]** tab of the **Event Policy Editor** page appears.
2. In the **Event Policy Name** field, type a name for the new event policy.
3. Click to select the checkbox for **Enable Event Policy**.
4. In the **Event Source** field, select *Internal*.
5. Click the **[Select Link-Message]** button.
6. In the **Link-Message** modal page, search for "Anomaly" to locate the message "Anomaly Detected: %V":
7. Select the radio button for the message "Anomaly Detected: %V", and then click **[Select]**.
8. Complete the remaining fields and tabs in the **Event Policy Editor** based on the specific parameters that you want to establish for the event. For more information about the fields and tabs in the **Event Policy Editor**, see [Defining an Event Policy](#).
9. When you are finished entering all of the necessary information into the event policy, click **[Save]**.

To create an event policy for anomalies in versions of Skylar One before 12.5.1:

1. Go to the **Event Policies** page (Events > Event Policies, or Registry > Events > Event Manager in the classic user interface).
2. On the **Event Policies** page, click the **[Create Event Policy]** button. The **Event Policy Editor** page appears.
3. In the **Policy Name** field, type a name for the new event policy.
4. Click the **[Match Logic]** tab.
5. In the **Event Source** field, select *Internal*.
6. In the **Match Criteria** field, click the **[Select Link-Message]** button.
7. In the **Link-Message** modal page, search for "Anomaly" to locate the message "Anomaly Detected: %V".

8. Click the radio button for the message "Anomaly Detected: %V", and then click **[Select]**.
9. Complete the remaining fields and tabs in the **Event Policy Editor** based on the specific parameters that you want to establish for the event. For more information about the fields and tabs in the **Event Policy Editor**, see [Defining an Event Policy](#).
10. To enable the event policy, click the **Enable Event Policy** toggle so that it is in the "on" position.
11. When you are finished entering all of the necessary information into the event policy, click **[Save]**.

Using Anomaly-related Events to Trigger Automated Run Book Actions

Skylar One includes automation features that allow you to define specific event conditions and the actions you want Skylar One to execute when those event conditions are met. You can use these features to trigger automated run book actions whenever an anomaly-related event is generated in Skylar One.

To use anomaly-related events to trigger automated run book actions:

1. Go to the **Automation Policy Manager** page (Registry > Run Book > Automation).
2. Click the **[Create]** button. The **Automation Policy Editor** page appears:

The screenshot displays the 'Automation Policy Editor' interface for creating a new automation policy. The 'Policy State' is highlighted in red and set to 'Enabled'. The 'Available Events' section is also highlighted in red and contains the following items:

- anomaly
- [4532] Major: Anomaly Index Major
- [4531] Minor: Anomaly Index Minor
- [4530] Notice: Anomaly Index Notice

The 'Aligned Events' section contains the following items:

- [4533] Critical: Anomaly Index Critical
- [4865] Major: Test Anomaly

The 'Aligned Actions' section contains the following items:

1. Send Email [0]: rba_send_email_notification
2. Create Ticket [2]: rba_create_new_ticket

At the bottom of the form, there is a 'Save' button.

3. In the **Policy State** field, select *Enabled*.
4. In the **Available Events** field, search for and select one or more anomaly-related event policies, and then click the right-arrow icon to move each event to the **Aligned Events** field. For more information about anomaly-related events, see [Creating an Event Policy for Anomalies](#).
5. In the **Available Actions** field, search for and select one or more run book actions that you want to run when the anomaly event from step 4 occurs. Click the right-arrow icon to move each action to the **Aligned Actions** field. For example, you might want to send an email or create a ticket for that anomaly event.
6. Complete the remaining fields on the **Automation Policy Editor** page based on the specific parameters that you want to establish for the automation policy. For more information about the fields on the **Automation Policy Editor** page, see [Automation Policies](#).
7. When you are finished, click **[Save]**.

Skylar Analytics: Predictive Alerting

Overview

The Predictive Alerting component of Skylar Analytics helps to avoid problems such as file systems running out of space. The alerts appear as enriched events in Skylar One, and they are generated in advance of the problem and can provide days, weeks, or months of notice depending upon the conditions.

The Predictive Alerting component monitors file systems (SNMP, PowerShell, and SSH).

This chapter covers the following topics:

<i>What is Predictive Alerting?</i>	129
<i>Viewing Predictive Alerts in Skylar One</i>	129
<i>Using Predictive Alerts to Trigger Automated Run Book Actions</i>	132

What is Predictive Alerting?

Predictive alerts help to avoid problems such as file systems running out of space. The alerts are generated in advance of the problem and can provide days, weeks, or months of notice depending upon the conditions.

Skylar Analytics will start generating predictive alerts about 48 hours after data starts getting exported from Skylar One to Skylar AI.

NOTE: A prediction cannot be made less than three times of the observation window. In other words, if you have one day of information, Skylar AI will not generate a prediction more than three days in the future.

How Predictive Alerting Works

To generate predictive alerts, Skylar AI looks at utilization trends over the past 30 days. In the case of file systems, Skylar AI looks at maximum value. Skylar AI uses these values to compute a linear trend, which provides a very simple slope to predict when a threshold will be reached.

Starting with version 1.8.0, Skylar Analytics uses the same approach for both the 30-day trend and the "breakout" or 1-day trend, which is to calculate the slope of the data over that time period.

Then, to choose the best slope for generating the prediction on, Skylar AI calculates root mean square deviation (RMSE) and the "R squared" (R^2) error rates for both slopes as well as a flat slope, across all 30 days of data as well as the last day of data. Skylar AI finds the best match, weighted towards the daily slope then the flat slope.

If none of the predictions are above the threshold, or if the flat slope is determined to be the best, then Skylar AI will not generate a prediction. Otherwise, Skylar AI generates a prediction based on the slope that has the best fit against the data.

Viewing Predictive Alerts in Skylar One

When your Skylar One system is connected to Skylar AI, you can start viewing predictive alerts in Skylar One. The alerts appear as enriched events in Skylar One, and they are generated in advance of the problem. No additional configuration is needed.

Predictive alerts display the Skylar icon (🌟) to the left of the event message in the **Message** column of the **Events** page. The filter text in the **Message** column and the text of the message contains the word "Prediction":

Organiz...	Severity	Name	Message	Last Det...	Age	Ticket ID	Count	Event Ty...	Event N...	Masked Events	Event So...	Acknowledge	Clear
System	Critical	JKECK-10.1	Prediction: File System JKE	Oct 5, 2024	2 days 9 hou	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	JK	Prediction: File System JKE	Oct 4, 2024	3 days 19 ho	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	xd	Prediction: CPU Utilization	Oct 1, 2024	6 days 9 hou	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	JK	Prediction: File System JKE	Oct 1, 2024	6 days 13 ho	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	JK	Prediction: File System JKE	Oct 1, 2024	6 days 13 ho	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	lin	Prediction: CPU Utilization	Sep 30, 2024	7 days 9 hou	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	lin	Prediction: CPU Utilization	Sep 30, 2024	7 days 11 ho	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	lin	Prediction: CPU Utilization	Sep 30, 2024	7 days 15 ho	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	sk	Prediction: File System sky	Sep 27, 2024	10 days 14 h	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	JK	Prediction: File System JKE	Sep 27, 2024	10 days 15 h	-	1	Device			Skylar AI	Acknowledge	Clear
Sample	Critical	mv	Prediction: File System mv	Sep 27, 2024	10 days 15 h	-	1	Device			Skylar AI	Acknowledge	Clear
System	Critical	JK	Prediction: File System JKE	Sep 23, 2024	14 days 22 h	-	1	Device			Skylar AI	Acknowledge	Clear

NOTE: The filtered list will appear blank until an active predictive alert triggers an event.

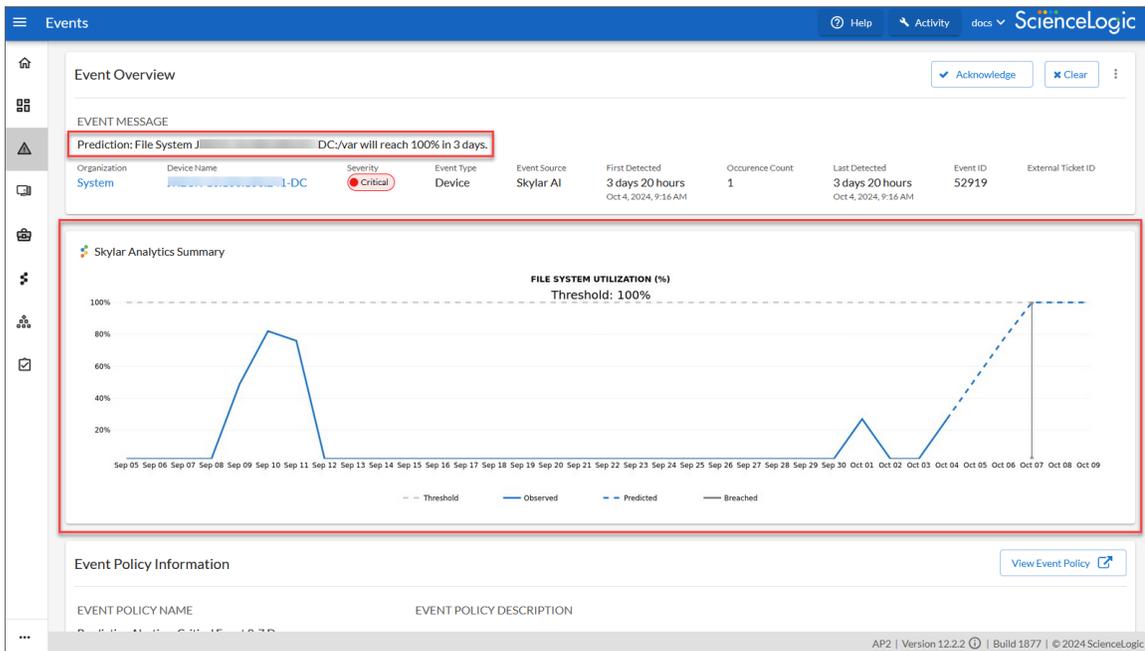
To view details about a predictive alert:

1. In Skylar One, go to the **Skylar AI** page (🌟) and click the **[Visit]** button for **Skylar Predictive Alerting**. A filtered **Events** page displays a list of predictive alerts.

TIP: The word "Prediction" appears in the filter field for the **Message** column. To clear the list of predictive alerts to view all events, click the X button in the filter.

2. On the **Events** page, click the message for a predictive alert with the Skylar icon (🌟). The **Event Investigator** page for that alert appears.

- On the **Event Investigator** page, the **Skylar Analytics Summary** panel displays a timeline of data from Skylar AI about a specific metric:



The dotted line on the graph in the **Skylar Analytics Summary** panel represents a time frame in the future that Skylar AI is forecasting, based on pattern recognition.

The blue line represents the activity observed so far by Skylar One, and the gray dotted line represents the threshold set in Skylar One. The blue dotted line represents where Skylar AI is predicting a potential alert in the future, with the gray line representing a potential problem in the future, also predicted by Skylar AI.

In the example above, Skylar AI predicts that the file system utilization will hit the threshold of 100% in three days, on October 7th. By tracking the timeline on the graph, you can see when a potential event might happen, and you can take action now to prevent it.

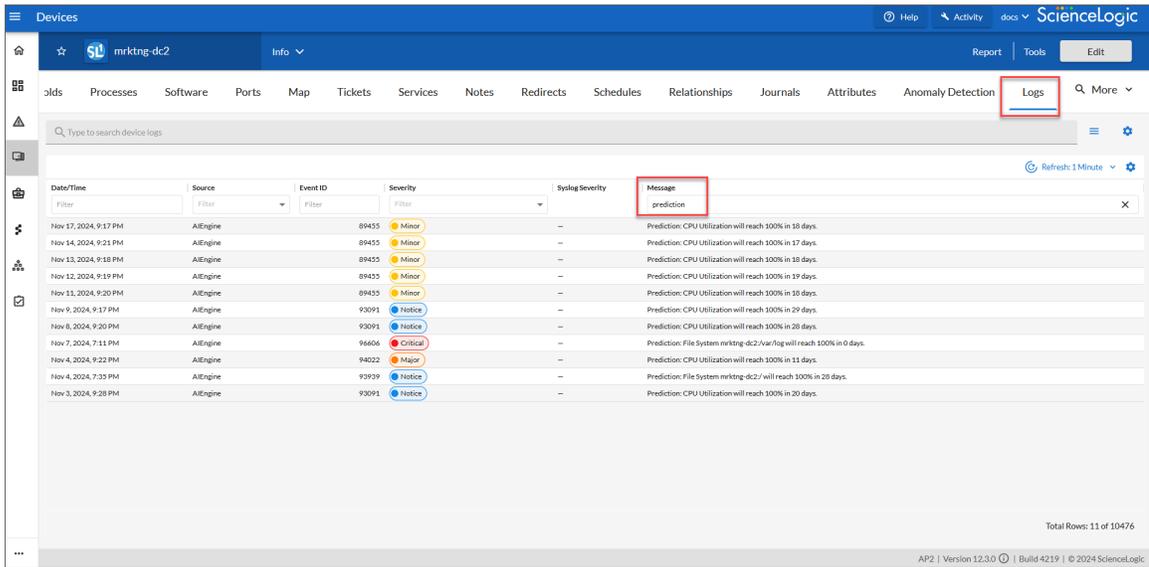
In addition, if you have an event policy monitoring a metric that is now being tracked by Predictive Alerting, you can disable that event policy.

NOTE: Because the data for the chart on the **Skylar Analytics Summary** panel is coming from Skylar AI, you will not be able to use that data in a Skylar One dashboard. Also, this chart is rendered at prediction time and is static, so that when opening an event, you can see the state and prediction at the time of prediction.

You can also review the logs for a specific device to view the history of the predictions:

- On the **Devices** page or the **Events** page, select the device with the predictive alerts. The Device Investigator page for that device appears.

2. Click the **[Logs]** tab. A list of recent logs displays:



3. If needed, type "prediction" in the **Message** column to view only the predictive alerts.

Using Predictive Alerts to Trigger Automated Run Book Actions

After Skylar AI creates a Skylar One event for a predictive alert, you can create a run book automation policy that runs one or more run book actions when a predictive alert is generated.

The predictive alert must have an **Event Type** of *Device* and an **Event Source** of *Skylar AI*.

To use predictive alerts to trigger automated run book actions:

1. Go to the **Automation Policy Manager** page (Registry > Run Book > Automation).
2. Click the **[Create]** button. The **Automation Policy Editor** page appears:

Automation Policy Editor | Creating New Automation Policy Reset

Policy Name: Policy Type: [Active Events] Policy State: [Enabled] Policy Priority: [Default] Organization: [Robert Anderson]

Criteria Logic: [Severity >=] [Minor.] Match Logic: [Text search] Match Syntax:

[and 5 minutes has elapsed] Repeat Time: [Only once] Align With: [Devices]

[since the first occurrence,] [and event is NOT cleared] Include events for entities other than devices (organizations, assets, etc.)

[and all times are valid] Trigger on Child Rollup

Available Devices: Linux: CentOS: 8.0.214, Linux: Oracle Linux Server release 5, Linux: Red Hat Enterprise Linux 8

Aligned Devices: (All devices)

Available Events: anomaly, [4532] Major: Anomaly Index Major, [4531] Minor: Anomaly Index Minor, [4530] Notice: Anomaly Index Notice

Aligned Events: [4533] Critical: Anomaly Index Critical, [4865] Major: Test Anomaly

Available Actions: Send Email [0]: rba_send_email_notification, SNMP Trap [1]: rba_trap_testing, SNMP Trap [1]: SL1 Event Trap, Create Ticket [2]: rba_create_new_ticket, Snippet [5]: Automation Utilities: Calculate Memory Size for Ea

Aligned Actions: 1. Send Email [0]: rba_send_email_notification, 2. Create Ticket [2]: rba_create_new_ticket

Save

3. In the **Policy State** field, select *Enabled*.
4. In the **Available Events** field, search for and select one or more event policies related to predictive alerts, and then click the right-arrow icon to move each event to the **Aligned Events** field.
5. In the **Available Actions** field, search for and select one or more run book actions that you want to run when the predictive alert event from step 4 occurs. Click the right-arrow icon to move each action to the **Aligned Actions** field. For example, you might want to send an email or create a ticket for that predictive alert.
6. Complete the remaining fields on the **Automation Policy Editor** page based on the specific parameters that you want to establish for the automation policy. For more information about the fields on the **Automation Policy Editor** page, see [Automation Policies](#).
7. When you are finished, click **[Save]**.

© 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.

ScienceLogic

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

International: +1-703-354-1010