



---

## Skylar Analytics

Version 1.5.0

---

# Table of Contents

<b>Introduction to Skylar Analytics</b> .....	<b>1</b>
What is Skylar AI? .....	2
Features of Skylar AI .....	2
Components of Skylar AI .....	2
Data Analyzed by Skylar AI .....	3
What is Skylar Analytics? .....	4
Getting Started with Skylar Analytics .....	5
Running the Skylar SL1 Management Tool .....	5
Enabling Skylar Analytics for One or More SL1 Organizations .....	6
Mapping SL1 Dynamic Application Object Names to Skylar Columns .....	7
<b>Configuring Access Control in Skylar Analytics</b> .....	<b>10</b>
Overview of Authentication in Skylar AI .....	11
Role-Based Access Control in Skylar AI .....	11
Navigating the Skylar Settings User Interface .....	12
Elements of Role-based User Accounts in Skylar AI .....	12
Configuring SSO Authentication with SAML .....	13
Using Access Tokens for Users .....	15
<b>Skylar Analytics: Data Visualization and Data Exploration</b> .....	<b>16</b>
What is Data Visualization? .....	17
Viewing Dashboards and Charts in Data Visualization .....	18
Viewing and Customizing Skylar Analytics Dashboards .....	19
Default Skylar Analytics Dashboards .....	19
Working with Skylar Analytics Dashboards .....	20
Viewing and Customizing Skylar Analytics Charts .....	22
Working with Skylar Analytics Charts .....	23
Additional Tips for Creating and Customizing Charts .....	23
Viewing Skylar Analytics Datasets .....	23
Data Exploration: Exporting Data from Skylar AI .....	24
Configuring Skylar Analytics Data Exploration with Power BI .....	24
Additional Resources for Skylar Analytics (Apache Superset Training) .....	27
<b>Skylar Analytics: Anomaly Detection</b> .....	<b>29</b>

What is Anomaly Detection? .....	30
Viewing Graphs and Data for Anomaly Detection .....	31
Enabling Alerts and Thresholds for the Anomaly Chart .....	32
Enabling Anomaly Detection Events for Specific Metrics .....	33
Enabling Anomaly Detection Events for a Metric on the Device Investigator Page .....	34
Enabling Anomaly Detection Events for a Metric on the Service Investigator Page .....	34
Creating an Event Policy for Anomalies .....	35
Using Anomaly-related Events to Trigger Automated Run Book Actions .....	36
<b>Skylar Analytics: Predictive Alerting .....</b>	<b>38</b>
What is Predictive Alerting? .....	39
Viewing Predictive Alerts in SL1 .....	39
<b>Appendix: Service Provider Administration for Skylar AI .....</b>	<b>42</b>
First Login as a Service Provider User .....	43
Provisioning a New Account .....	43

---

# Chapter

# 1

## Introduction to Skylar Analytics

---

### Overview

Skylar Analytics includes the following components:

- **Data Visualization.** Enables SQL-based dashboards and charts based on data gathered by Skylar AI and SL1. 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 by visiting the **[Anomaly Detection]** tab on the **Device Investigator** page for a device.
- **Predictive Alerting.** Helps to avoid problems such as file systems running out of space, hosts running out of memory, or issues with network reliability due to oversubscription. The alerts appear as enriched events within SL1.

**IMPORTANT:** Skylar Analytics requires SL1 12.3.1 or later. ScienceLogic strongly recommends that you always use the most recent SL1 and AP2 releases in conjunction with the most recent Skylar AI release. 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 [SL1 Platform and AP2 Release Notes](#).

This video provides an overview of the different features of Skylar Analytics:

<https://player.vimeo.com/video/990317575?h=74e1aca2bf>

To view the latest Skylar Analytics release notes, see the [Skylar Analytics Release Notes](#).

This chapter covers the following topics:

<i>What is Skylar AI?</i> .....	2
<i>What is Skylar Analytics?</i> .....	4
<i>Getting Started with Skylar Analytics</i> .....	5
<i>Mapping SL1 Dynamic Application Object Names to Skylar Columns</i> .....	7

---

## 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 ScienceLogic AI Platform, which includes Skylar Automated RCA, Skylar Analytics, and Skylar Advisor (coming soon), 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.

SL1 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 SL1 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

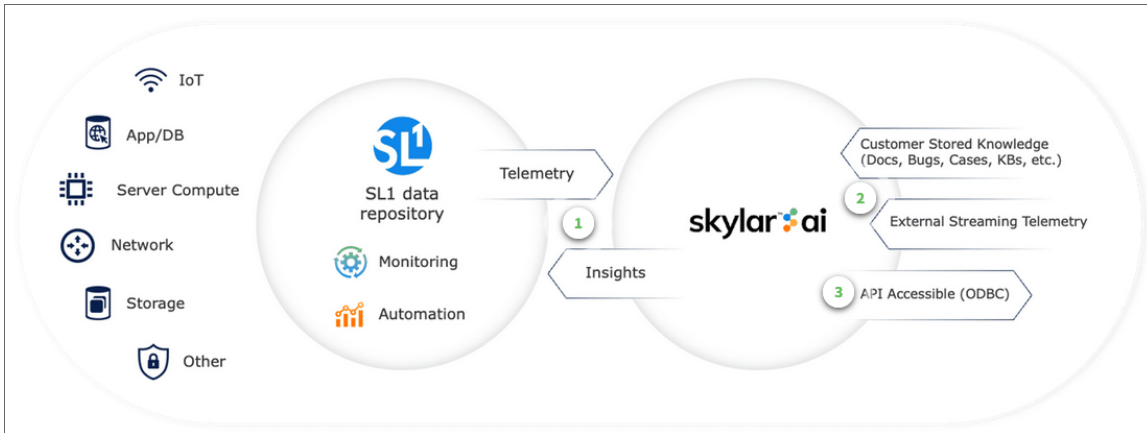
The Skylar AI family of services currently includes the three following components:

- **Skylar Automated Root Cause Analysis (RCA)**, a log-based, root cause identification and analysis service powered by unsupervised AI.

- **Skylar Analytics**, an advanced reporting and custom analytics service that combines AI-powered analytics with deep data exploration and visualization.
- **Skylar Advisor**, a proactive IT problem-solving advisory service powered by human-centered AI.

## Data Analyzed by Skylar AI

The following image shows the flow of data into and out of SL1 and the Skylar AI Engine:

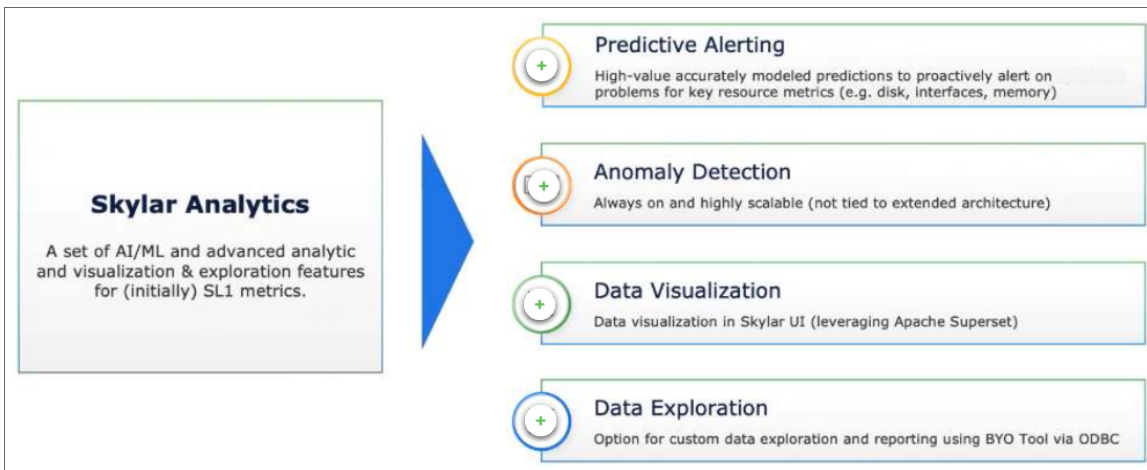


The following list contains some of the types of data that SL1 can send to the Skylar AI engine, where the data is analyzed and used by Skylar Automated RCA, Skylar Analytics, and Skylar Advisor:

- Alert and event logs
- Availability data collected by SL1
- Business Service health, availability, and risk metrics from SL1
- Class-Based Quality-of-Service (CBQoS) metadata and CBQoS time series data
- Data from Gen 1 SL1 agents, which use the SL1 Distributed Environment
- Data from Gen 3 SL1 agents, which use the SL1 Extended Architecture
- Dynamic Application mapping and performance data
- Topology data for L2, L3, CDP, LLDP, and ad-hoc relationships between devices
- DCM(+R) relationships
- Metadata for web content, SOAP/XML transaction, and domain name monitors
- Process and service data

# What is Skylar Analytics?

The Skylar Analytics suite of services uses data gathered by SL1 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.



Skylar Analytics includes the following components:

- **Data Visualization**. Enables SQL-based dashboards and charts based on data gathered by Skylar AI and SL1. 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 by visiting the **[Anomaly Detection]** tab on the **Device Investigator** page for a device.
- **Predictive Alerting**. Helps to avoid problems such as file systems running out of space, hosts running out of memory, or issues with network reliability due to oversubscription. The alerts appear as enriched events within SL1.

For more information about these components, see the following chapters.

---

## Getting Started with Skylar Analytics

Before you can start using Skylar Analytics, you will need to perform the following configurations in SL1 to enable the export of data from SL1 to Skylar:

- [Run the Skylar SL1 Management Script](#)
- [Enable Skylar Analytics for one or more organizations](#)

After you perform these configurations, you can access Skylar Analytics and other key Skylar AI components from the **Skylar AI** page (🔗) in SL1.

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

**IMPORTANT:** Skylar Analytics requires SL1 12.3.1 or later. ScienceLogic strongly recommends that you always use the most recent SL1 and AP2 releases in conjunction with the most recent Skylar AI release. 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 [SL1 Platform and AP2 Release Notes](#).

## Running the Skylar SL1 Management Tool

The Skylar SL1 Management Tool configures SL1 data and SL1 processes, and it starts monitoring the Skylar connection and configuration. The script is named `sl-otelcol-mgmt.py`, and it is included with Skylar Analytics in the `sl-otelcol` RPM package.

To run the Skylar SL1 Management Tool:

1. Use the following command to run the Skylar SL1 Management script on the Database Server (an SL1 Central Database or an SL1 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, and `<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 [Using Access Tokens for Users](#).

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, they will populate with data. Please note that this process might take several minutes.



**TIP:** To check to make sure you have connected Skylar AI to SL1, go to the **Skylar AI** page in SL1. If the page loads, then the connection was successful. You can also go to the **Service Connections** page (Manage > Service Connections) and look for a service connection with a **Type** of "Skylar AI Engine" to verify that the connection was successful.

2. 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 OTELCOL endpoint is healthy
```

3. If you need to turn off the Skylar connection, run the following command:

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

4. Continue to the next step to specify the organizations you want to use for exporting data to Skylar.

## Enabling Skylar Analytics for One or More SL1 Organizations

In SL1, if you want to use Anomaly Detection and Predictive Alerting, you will need to select one or more organizations 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 Anomaly Detection and Predictive Alerting:

1. In SL1, 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*.

---

## Mapping SL1 Dynamic Application Object Names to Skylar Columns

When data from SL1 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".
  - For example: ba\$ → bA
  - This ensures that column names are valid and avoid special symbols.

## 2. Replace Common Words

Certain words are automatically shortened to standard abbreviations. Here are the most common ones:

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 (like spaces, slashes, parentheses, etc.).
- Replace common terms:
  - Average → Avg
  - QueueLength → QLen
  - sISl → 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

---

# Chapter

# 2

## Configuring Access Control in Skylar Analytics

---

### Overview

This chapter explains the authentication and role-based access control used by Skylar AI, including how to use the **Skylar Settings** page in the Skylar AI user interface.

**IMPORTANT:** This chapter is intended for Skylar AI administrators only.

This chapter covers the following topics:

<i>Overview of Authentication in Skylar AI</i> .....	11
<i>Role-Based Access Control in Skylar AI</i> .....	11
<i>Configuring SSO Authentication with SAML</i> .....	13
<i>Using Access Tokens for Users</i> .....	15

---

# 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 (tenant ).
- Predefined roles for access control.
- Email and password (local accounts) authorization by default, and SAML single-sign-on (SSO) authorization configured as needed.
- Access tokens for integration with external tools.


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

- The first login for any new user will start with a prompt to create a new password.
- A session expires after 18 hours or three hours of inactivity.
- Logging into a session requires either an email and password combination or a successful SAML2 redirect workflow.
- Email domains and SAML are configured per account (tenant).
- Ten failed password validations within ten minutes disables the user and requires administrator intervention.

---

## 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**. The **[Profile]** tab for the **Skylar Settings** page appears:



On the different tabs of the **Skylar Settings** page, you can edit your user profile, add users and groups to your account, assign roles to groups, and create access tokens. Depending on your user role, you can also set up additional forms of authentication.

## Navigating the Skylar Settings User Interface

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

- To return to the Skylar AI login page, click the Skylar AI icon at top left.
- To view the email address and role for the current user in the Skylar AI user interface, click the user icon  at top right. You can also click the **[Sign Out]** button to sign out of this session.
- To return to the **[Profile]** tab for the current user, click the profile icon  at top right.

**TIP:** To return to the Skylar AI login page, click the Skylar AI icon at top left.

## 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 accounts in a single Skylar AI system. Another way of thinking of an account is that an account is a "tenant", as in "multi-tenant software".

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.

On the **[Available Instances]** tab of the **Skylar Settings** page, you can view a list of instances for the current user. You can also access the "Analytics Secrets" for an instance, which contains the ODBC host, password, port, and user information for Data Exploration using ODBC.

- **Access Tokens.** You can add access tokens to connect Skylar AI with SL1 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]** tab of the **Skylar Settings** page, you can view and add access tokens. For more information, see [Using Access Tokens for Users](#).

- **Users.** Each person that uses Skylar AI should have his or her own user account. A user must belong to at least one **group**.

On the **[Users]** tab of the **Skylar Settings** page, 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]** tab of the **Skylar Settings** page, 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:
  - **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. Please note that the **Admin**, **Editor**, and **Viewer** roles are the same for the current release of Skylar AI.
  - **Editor.** For a future release, this role will let a user edit (create, update, and delete) objects, particularly incident type metadata.
  - **Viewer.** For a future release, this role will give a user read-only access to Skylar AI. A **Viewer** user can edit their own profile.

On the **[Roles]** tab of the **Skylar Settings** page, you can view your assigned roles for this account.

- **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]** tab of the **Skylar Settings** page, you can configure SAML for this account. For more information, see [Configuring SSO Authentication with SAML](#).

---

## 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.

In case of 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.

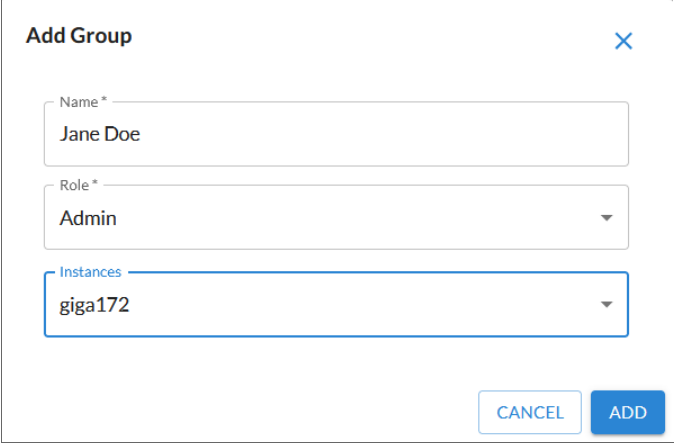
**IMPORTANT:** Before you can set up SSO authentication with SAML in Skylar AI, you will first need to create your user groups with your SAML identity provider if you do not already have them set up. Be sure to use the same names for your user groups with your SAML provider and with Skylar AI.



**IMPORTANT:** 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.

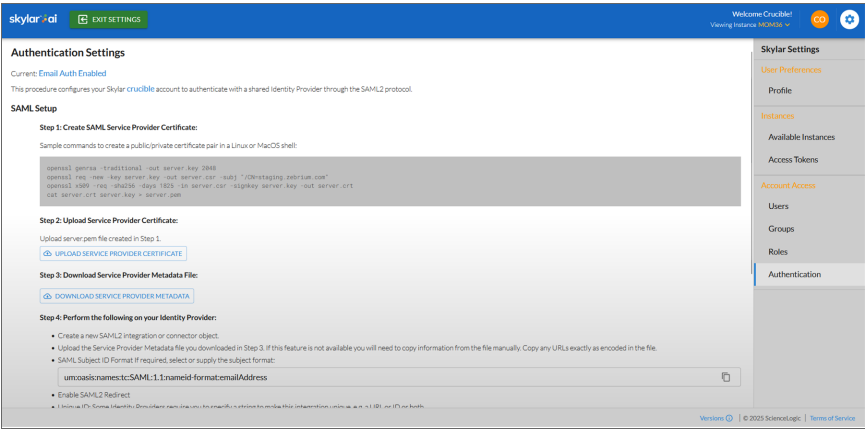
To set up SSO Authentication with SAML in the Skylar AI user interface:

1. On the **Skylar Settings** page, click the **[Groups]** tab and click **[Add Group]**. The **Add Group** dialog appears:



The 'Add Group' dialog box is a modal window with a title bar containing 'Add Group' and a close button (X). It contains three input fields: 'Name \*' with the text 'Jane Doe', 'Role \*' with a dropdown menu showing 'Admin', and 'Instances' with a dropdown menu showing 'giga172'. At the bottom right, there are two buttons: 'CANCEL' and 'ADD'.

2. Type a name for the group, select a role of *Admin*, and select one or more instances. Click **[Add]**. The group is added to the **[Groups]** tab.
3. Go to the **[Authentication]** tab and review the instructions for SAML setup:



The 'Authentication Settings' page in the Skylar AI interface. The top bar shows 'skylar.ai' and 'EXIT SETTINGS'. The main content area is titled 'Authentication Settings' and includes a status 'Current: Email Auth Enabled'. Below this, it says 'This procedure configures your Skylar crucible account to authenticate with a shared Identity Provider through the SAML2 protocol.' The 'SAML Setup' section contains four steps: 1. Create SAML Service Provider Certificate (with sample commands), 2. Upload Service Provider Certificate (with an 'UPLOAD SERVICE PROVIDER CERTIFICATE' button), 3. Download Service Provider Metadata File (with a 'DOWNLOAD SERVICE PROVIDER METADATA' button), and 4. Perform the following on your Identity Provider (with a list of instructions and a text input field for 'umsoasname: SAML:1:1:nameid-format:emailAddress'). The right sidebar shows the 'Skylar Settings' menu with options like 'User Preferences', 'Profile', 'Instances', 'Available Instances', 'Access Tokens', 'Account Access', 'Users', 'Groups', 'Roles', and 'Authentication' (which is currently selected).

4. Follow steps 1-7 from the **[Authentication]** tab on the **Skylar Settings** page.

**TIP:** For step 7 on the **[Authentication]** 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.

---

## Using Access Tokens for Users

You can use the **[Access Tokens]** tab from the **Skylar Settings** page to add access tokens to connect Skylar AI with SL1 or a third-party application. Starting with Skylar Analytics version 1.4.0, 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 select **Skylar Settings**.
2. Click the **[Access Tokens]** tab.
3. Click the **[Add Access Token]** button. The **Add Access Token** window appears.
4. Complete the following fields:
  - **Name.** Type a name for the token, such as "SL1 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 use with the [Skylar SL1 Management Tool](#), select both *sl1\_connector* and *telemetry*.
  - **Expiration Date.** Select an expiration date.
5. Click the **[Add]** button. The access token is added to the **[Access Tokens]** tab.
6. Click the copy icon (📋) to copy the access token to the clipboard.

---

# Chapter

# 3

## 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. Data Visualization is achieved using a ScienceLogic-hosted instance of Apache Superset or with your own third party tool.

Currently, this data includes metrics for file systems, network interfaces, and all Dynamic Applications, with more metrics planned for future Skylar updates.

**IMPORTANT:** The dashboards, charts, and reports in the Data Visualization component of Skylar Analytics are *not* compatible with SL1 dashboards, widgets, or reports.

The optional **Data Exploration** component enables third-party tools that use the Microsoft 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 Tableau, Microsoft BI, and other business intelligence tools.

This chapter will provide 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 SL1 and Skylar AI.

This chapter covers the following topics:

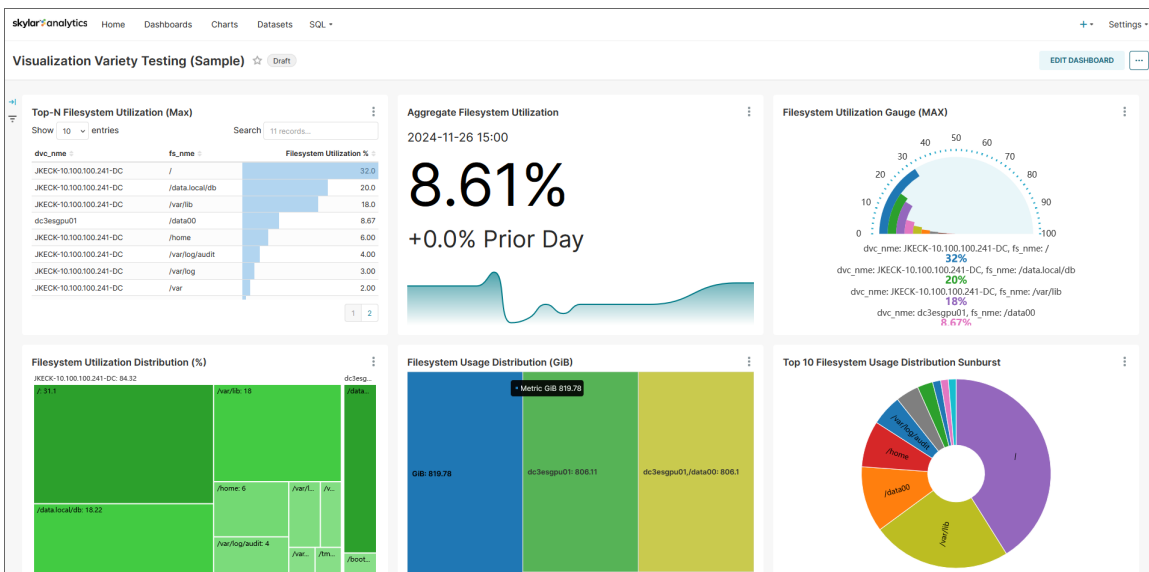
<a href="#">What is Data Visualization?</a> .....	17
<a href="#">Viewing Dashboards and Charts in Data Visualization</a> .....	18
<a href="#">Data Exploration: Exporting Data from Skylar AI</a> .....	24
<a href="#">Additional Resources for Skylar Analytics (Apache Superset Training)</a> .....	27

# What is Data Visualization?

Before the initial release of Skylar Analytics, SL1 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 SL1 and Skylar AI, normalizes it, and makes it available in standard ODBC database format.

The data originates from SL1 data collectors, undergoes processing, and is then simultaneously transmitted to Skylar via API.

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. You can also leverage the Data Visualization component with your existing BI tools for your company that support ODBC.



**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 the most current and accurate information, see [Additional Resources for Skylar Analytics](#).

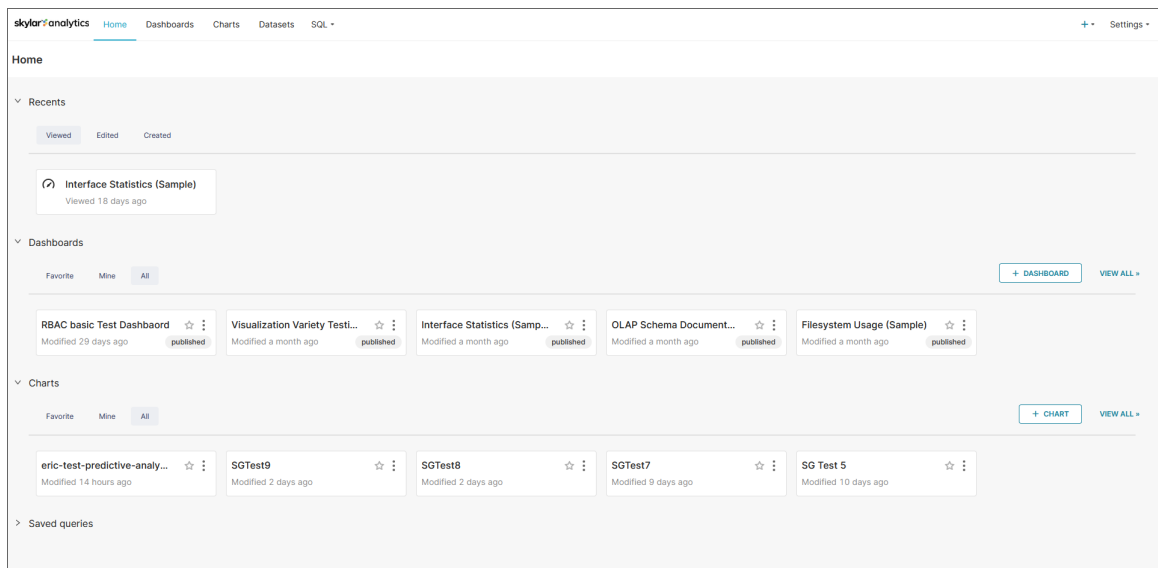
# Viewing Dashboards and Charts in Data Visualization

The Data Visualization component of Skylar Analytics contains dashboards and charts based on data gathered by Skylar AI and SL1.

**IMPORTANT:** The dashboards in the Data Visualization component of Skylar Analytics are *not* compatible with SL1 dashboards, widgets, or reports.

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

1. From SL1, go to the **Skylar AI** page (🔗) and click the **[Visit]** button for **Skylar Data Visualization**. The Skylar AI login page appears.
2. Click **Analytics** and type in your user name and password. The **Home** page for Data Visualization component of Skylar Analytics appears:



**TIP:** To return to the Skylar AI login page, click the Skylar Analytics icon at top left.

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 also create a dashboard or a chart from this page, and you can view all dashboards and charts by clicking the **View All** link.

For Skylar Analytics, you will mainly use the following tabs to view SL1 and Skylar AI data visualizations:

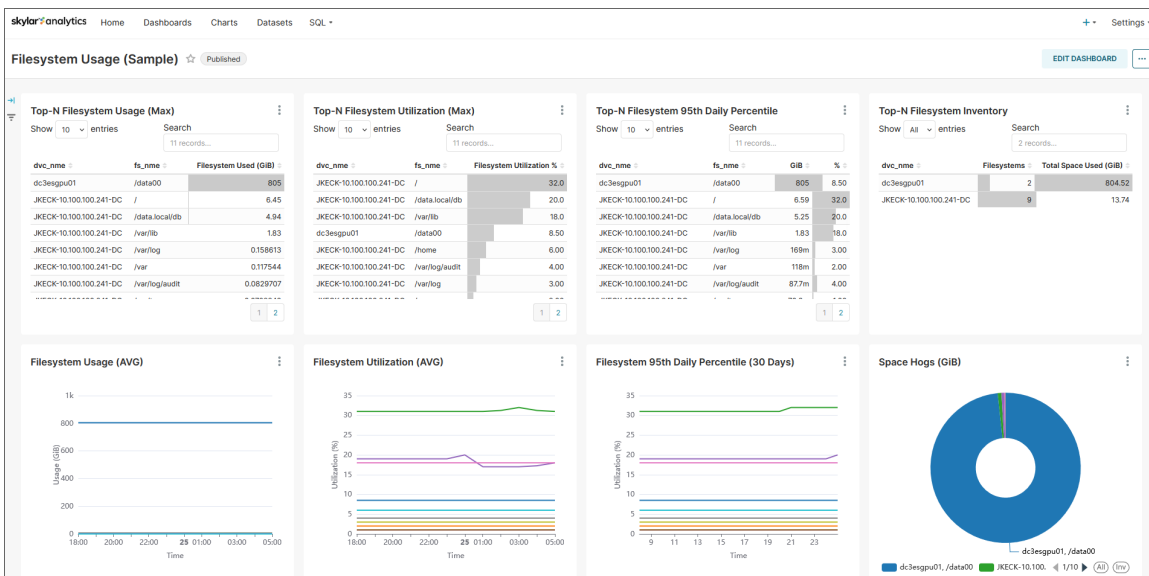
- [Dashboards](#)
- [Charts](#)

- [Datasets](#) (for administrators)

## Viewing and Customizing Skylar Analytics Dashboards

A **dashboard** in Skylar Analytics is similar to a dashboard in SL1, 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. In Skylar Analytics, a widget is called a "chart".

**NOTE:** Unlike dashboards in SL1, 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. Charts can be duplicated to be modified for different analyses on different dashboards.



## Default Skylar Analytics Dashboards

The **[Dashboards]** tab for Skylar Analytics contains the following default dashboards:

- **Filesystem Overview + Exploration (Sample).**
  - Displays 95th percentile data, file system utilization distribution (as a percentage and Gigibit 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.
- **Visualization Variety Testing (Sample).**
  - Displays 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.
  - This table is not meant to be informational so much as an example of the types of visualizations you can use with Skylar Analytics.

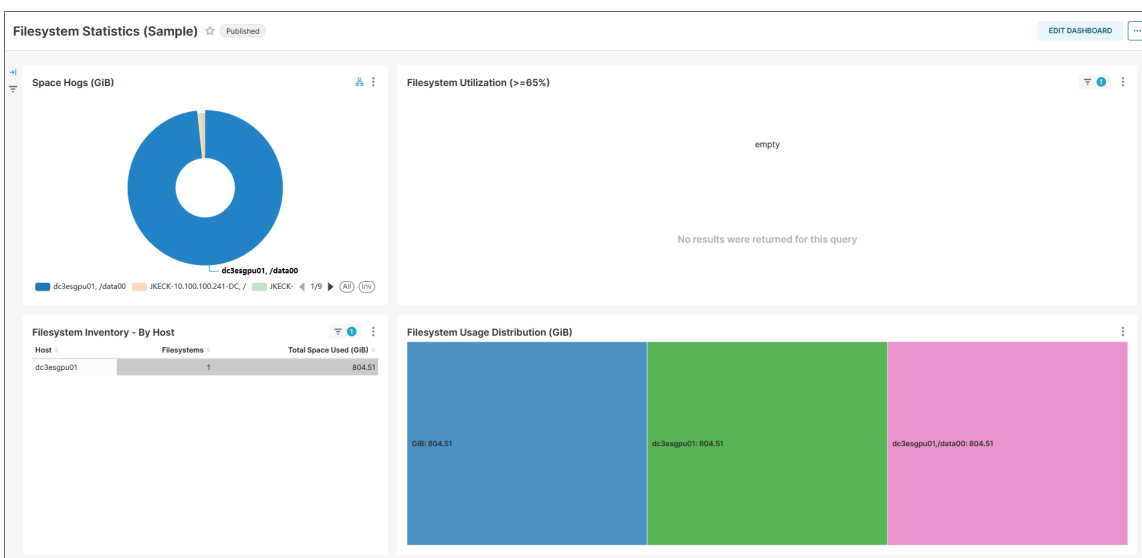
**NOTE:** Each default dashboard has the word "(Sample)" or "(Skylar)" at the end of its name to show that it is a ScienceLogic dashboard, and also to remind you to duplicate any of these dashboards or charts if you wish to make modifications. They are also owned by the System Administrator ("SA") user. These SA-owned dashboards and charts might be updated by ScienceLogic periodically.

## Working with Skylar Analytics Dashboards

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

- For most dashboards, you can click a single device or item in the first chart at the top left of the Dashboard page to view data specific to just that device. Click the device a second time to clear the filter.
- 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.
- Click **[Edit Dashboard]** to make changes to the dashboard and the charts that comprise the dashboard. For more information, see <https://docs.preset.io/docs/creating-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 horizontal 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*. 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 permalink to the chart to the clipboard of your computer, and also lets you share 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:

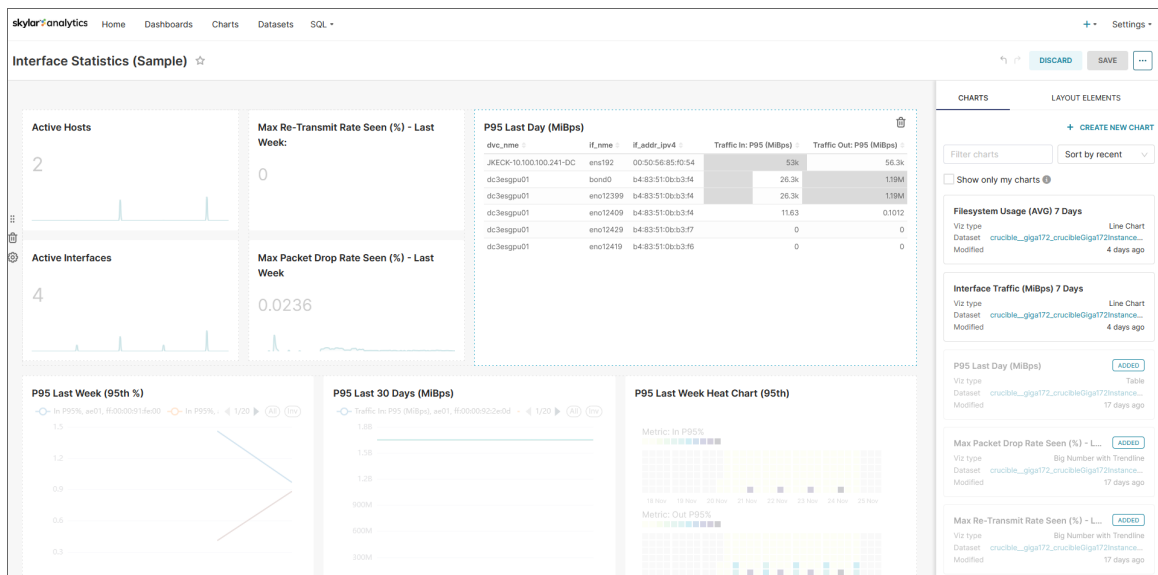
- *Enter fullscreen*. Displays the browser window containing just this chart display as full screen. Click the *Exit fullscreen* icon (⌵) or select *Exit fullscreen* from the menu to return to the previous setting.
- *Edit chart*. Opens the **Edit Chart** page so you can add metrics, edit queries, and 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**, which lets 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 <https://docs.preset.io/docs/cross-filtering#scoping-cross-filters>.
- *View query*. Displays the SQL query for that chart.
- *View as table*. Displays the chart in table format.



- *Drill to detail*. Displays all the data that makes up a chart. For more information, see <https://docs.preset.io/docs/drilling-to-chart-details>.
- *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.

To customize a dashboard:

1. Select the dashboard from the **Dashboards** page. You can also hover over the dashboard and click the Edit icon.
2. On the Dashboard page, click **[Edit Dashboard]**. The **Edit Dashboard** page appears:



3. For more information, see <https://docs.preset.io/docs/creating-a-dashboard>.

**TIP:** To watch a related video, see <https://superset.apache.org/docs/using-superset/creating-your-first-dashboard/>.

## Viewing and Customizing Skylar Analytics Charts

A **chart** in Skylar Analytics works much like a "widget" in SL1, in that a chart in Skylar Analytics is a building block that makes up a dashboard, and a dashboard can contain many charts.

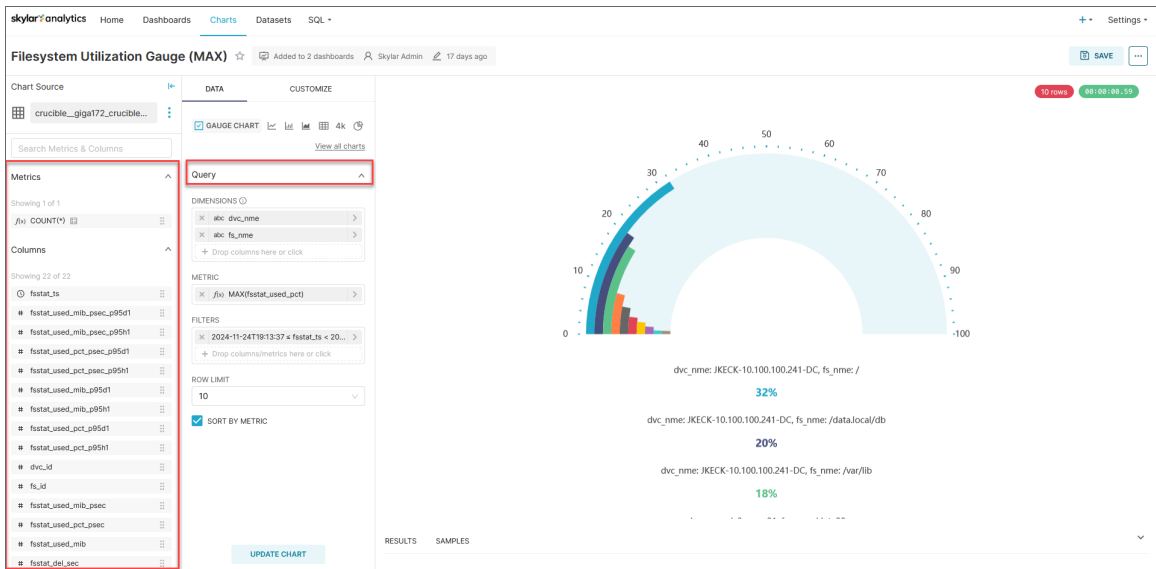
**TIP:** On the **Dashboards** tab in Skylar Analytics, the "Visualization Variety Testing (Sample)" dashboard contains a variety of chart visualizations related to file system utilization, including a table, a "big number" with a line graph, a gauge, a set of tree maps, and a sunburst map. You can use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

For more information about the types of charts you can use in a Skylar Analytics dashboard, see <https://docs.preset.io/docs/chart-walkthroughs>.

## Working with Skylar Analytics Charts

To create or customize a chart:

1. Select the chart from the **Charts** page, or edit the chart from an existing dashboard. If you are creating a new chart, click the **[+ Chart]** button on the **Charts** page.
2. On the Chart page, click **[Edit Chart]**. The **Edit Chart** page appears:



3. You can drag and drop **Metrics** and **Columns** into the **Query** panel to configure your visualization. For more information, see <https://docs.preset.io/docs/creating-a-chart>.

## Additional Tips for Creating and Customizing Charts

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

- **f**: Function used for metrics
- **Clock**: The time column for the data source
- **ABC**: Text data
- **#**: Numeric value data

## Viewing Skylar Analytics Datasets

**Datasets** are curated representations of the data in your database that let you quickly create dashboards and charts in Skylar Analytics. These dashboards and charts are based on the metrics stored in the datasets. In Skylar Analytics, each dataset contains a set of related metrics, such as server reports, which you can use to build a custom dashboard or chart or to enhance an existing dashboard or chart.

You will not need to create new datasets for Skylar Advisor.

---

## Data Exploration: Exporting Data from Skylar AI

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.

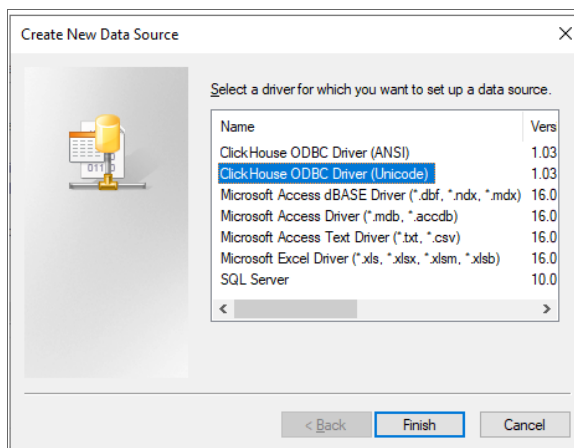
Data Exploration lets you view Skylar AI data alongside other business sources, offering a holistic perspective on your operations.

### Configuring Skylar Analytics 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.

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:

The screenshot shows a dialog box titled "Create data source for ClickHouse". It contains the following fields and values:

- Name: Mom36 ClickHouse
- Description: (empty)
- URL: (empty)
- Or...
- Host: dv-crucible-mom36-odbc.sta
- Port: 443
- Database: (empty)
- SSLMode: require
- User: dataviz
- Password: (masked with dots)
- Timeout: (empty)

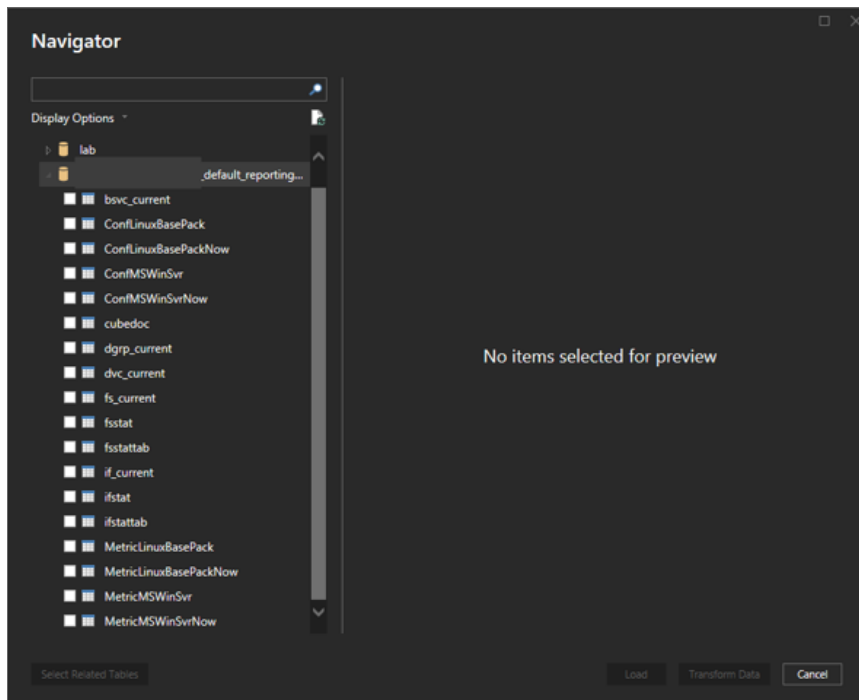
Buttons: Cancel, Ok

6. Complete the following fields with ODBC connection details from ScienceLogic SRE:
  - **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:

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:

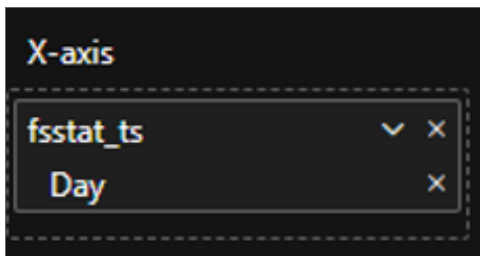


To import data and create a dashboard with Skylar AI data in Power BI:

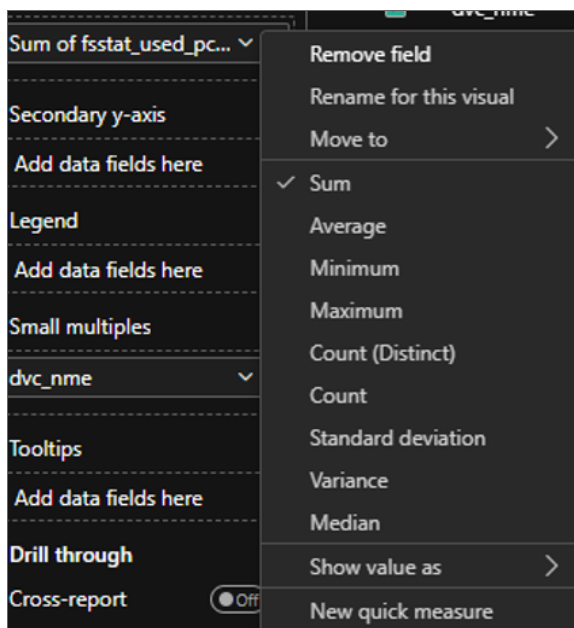
**TIP:** When selecting datasets to import, choose only the necessary tables to optimize performance. The following procedure creates a sample dashboard in Power BI.

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.

**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.

Apache Superset-related documentation: <https://superset.apache.org/docs/intro>

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

- <https://www.udemy.com/course/apache-superset-for-data-engineers-hands-on/>
- [https://www.youtube.com/watch?v=znnmco3eK-M&list=PLzRV\\_ObjEwmNhRjhMNcvcDP7ZDjOXtodd](https://www.youtube.com/watch?v=znnmco3eK-M&list=PLzRV_ObjEwmNhRjhMNcvcDP7ZDjOXtodd)
- <https://superset.apache.org/community>

---

# Chapter 4

## 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 anomalies on the **[Anomaly Detection]** tab on the **Device Investigator** page for each device, as well as corresponding **Service Investigator** pages.

**NOTE:** Anomaly Detection with Skylar Analytics works with all of the Performance Dynamic Applications in all SL1 PowerPacks.

This chapter covers the following topics:

<i>What is Anomaly Detection?</i> .....	30
<i>Viewing Graphs and Data for Anomaly Detection</i> .....	31
<i>Enabling Anomaly Detection Events for Specific Metrics</i> .....	33
<i>Creating an Event Policy for Anomalies</i> .....	35
<i>Using Anomaly-related Events to Trigger Automated Run Book Actions</i> .....	36



# What is Anomaly Detection?

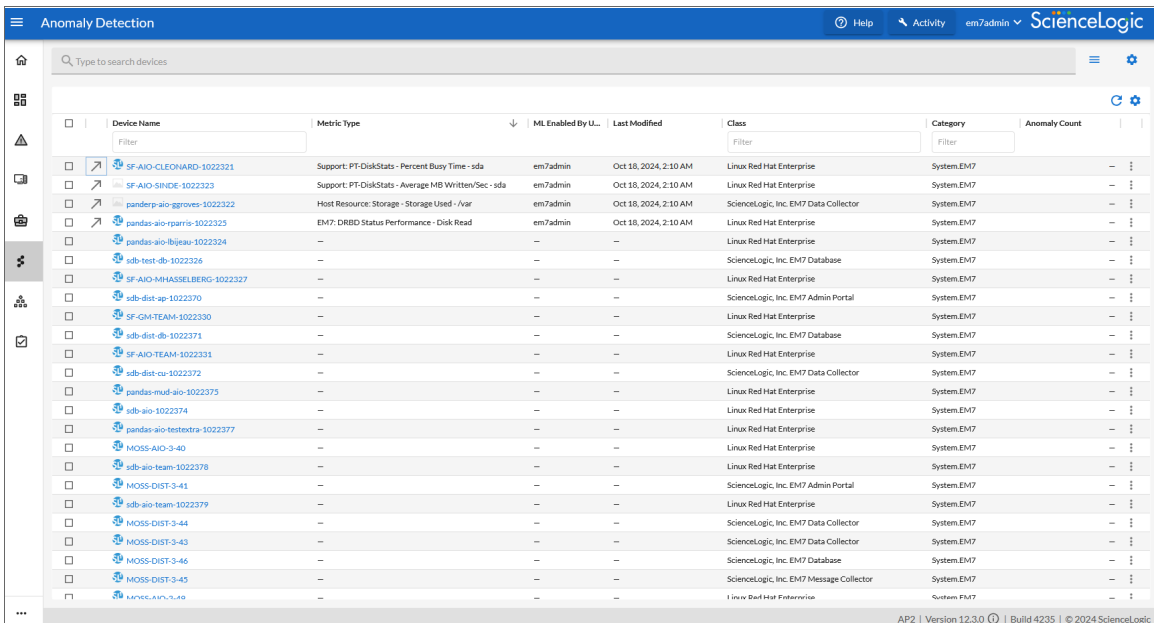
**Anomaly detection** is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. Anomaly detection provides always-on, unsupervised machine learning-based monitoring that automatically identifies unusual patterns in the real-time performance metrics 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 SL1 user interface for all Performance Dynamic Applications. 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.

For more information, see [Enabling Skylar Analytics for One or More SL1 Organizations](#).

You can view a list of all devices that have metrics being monitored for anomalies on the **Anomaly Detection** page in SL1 (Skylar AI (🔍) > **[Visit]** button for Skylar Anomaly Detection):



	Device Name	Metric Type	ML Enabled By U...	Last Modified	Class	Category	Anomaly Count
<input type="checkbox"/>	SF-AIO-CLEONARD-1022321	Support: PT-DiskStats - Percent Busy Time - sda	em7admin	Oct 18, 2024, 2:10 AM	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	SF-AIO-SINDE-1022323	Support: PT-DiskStats - Average MB Written/Sec - sda	em7admin	Oct 18, 2024, 2:10 AM	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	pandas-aio-ggroves-1022322	Host Resource: Storage - Storage Used - /var	em7admin	Oct 18, 2024, 2:10 AM	ScienceLogic, Inc. EM7 Data Collector	System.EM7	—
<input type="checkbox"/>	pandas-aio-rparris-1022325	EM7: DRBD Status Performance - Disk Read	em7admin	Oct 18, 2024, 2:10 AM	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	pandas-aio-lbjeau-1022324	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	sdb-test-db-1022326	—	—	—	ScienceLogic, Inc. EM7 Database	System.EM7	—
<input type="checkbox"/>	SF-AIO-MHASSELBERG-1022327	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	sdb-dist-ap-1022370	—	—	—	ScienceLogic, Inc. EM7 Admin Portal	System.EM7	—
<input type="checkbox"/>	SF-GM-TEAM-1022330	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	sdb-dist-db-1022371	—	—	—	ScienceLogic, Inc. EM7 Database	System.EM7	—
<input type="checkbox"/>	SF-AIO-TEAM-1022331	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	sdb-dist-cu-1022372	—	—	—	ScienceLogic, Inc. EM7 Data Collector	System.EM7	—
<input type="checkbox"/>	pandas-mad-aio-1022375	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	sdb-aio-1022374	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	pandas-aio-testextra-1022377	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	MOSS-AIO-3-40	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	sdb-aio-team-1022378	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	MOSS-DIST-3-41	—	—	—	ScienceLogic, Inc. EM7 Admin Portal	System.EM7	—
<input type="checkbox"/>	sdb-aio-team-1022379	—	—	—	Linux Red Hat Enterprise	System.EM7	—
<input type="checkbox"/>	MOSS-DIST-3-44	—	—	—	ScienceLogic, Inc. EM7 Data Collector	System.EM7	—
<input type="checkbox"/>	MOSS-DIST-3-43	—	—	—	ScienceLogic, Inc. EM7 Data Collector	System.EM7	—
<input type="checkbox"/>	MOSS-DIST-3-46	—	—	—	ScienceLogic, Inc. EM7 Database	System.EM7	—
<input type="checkbox"/>	MOSS-DIST-3-45	—	—	—	ScienceLogic, Inc. EM7 Message Collector	System.EM7	—
<input type="checkbox"/>	EM7-AIO-TEAM-1022379	—	—	—	Linux Red Hat Enterprise	System.EM7	—

## Viewing Graphs and Data for Anomaly Detection


After SL1 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 SL1:

- The **Anomaly Detection** page (Skylar AI (🔍) > [Visit] button for Skylar Anomaly Detection).
- The [Anomaly Detection] tab in the **Device Investigator**.
- The **Anomalies** tab in the **Service Investigator** for a business, IT, or device service.

You can view the anomaly detection graphs for the metrics by clicking the **Open** icon (↗) next to the metric for the device. 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 normal patterns. The lines in the chart are color-coded by the severity level of the event that gets triggered as the data diverges further. The anomaly score is basically a running sum over a small window of time, so after anomalies stop, the score will drop to zero over that time.

You can define the thresholds for the "Anomaly Score" chart, and whether those values generate alerts, on the **Anomaly Detection Thresholds** page (Skylar AI (🔍) > [Advanced: Adjust Thresholds] button). For more information, see [Enabling Alerts and Thresholds for the Anomaly Chart](#).




The second graph displays the following data:

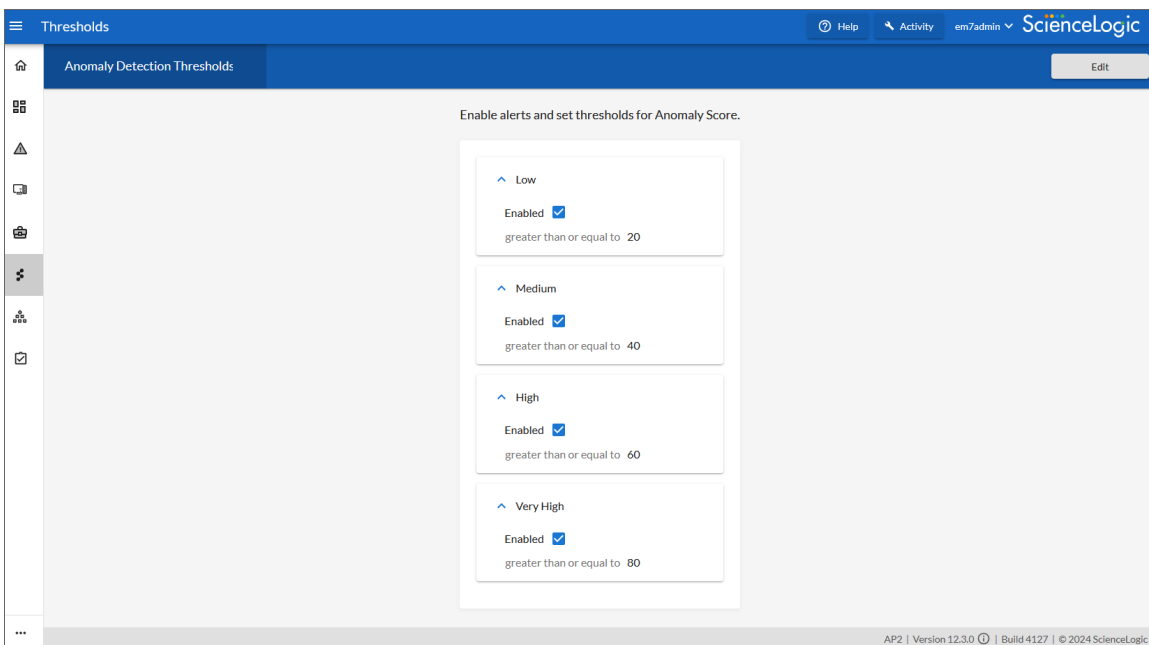
- A blue band representing the range of probable values that SL1 expected for the device metric.
- A green line representing the actual value for the device metric.
- A red dot indicating anomalies where the actual value appears outside of the expected value range.

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 Alerts and Thresholds for the Anomaly Chart

You can define the thresholds for the "Anomaly Score" chart on the **Anomaly Chart** modal, and whether those values generate alerts, 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, click **[Edit]**.
2. For each of the four severity levels, from Low to Very High, you can select **Enabled** to have SL1 generate an alert when the Anomaly value for a metric is equal to or greater than the threshold for that severity level.
3. You can edit the threshold value for each level if SL1 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 SL1. For more information, see [Creating an Event Policy for Anomalies](#).

## Enabling Anomaly Detection Events for Specific Metrics


You can set up anomaly detection events for specific metrics for devices and business services so that event policies are triggered when an anomaly is detected for that metric.


## Enabling Anomaly Detection Events for a Metric 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. The **[Anomaly Detection]** tab for **Device Investigator** displays.

**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.

2. On the **[Anomaly Detection]** tab, click the **Actions** icon () for any of the listed metrics and select *Enable*. The **Select Available Metrics** modal appears.
3. In the **Select Metric** drop-down, click the name of the metric on which you want to enable anomaly detection events for the device.
4. 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.
5. 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*.

## Enabling Anomaly Detection Events for a Metric on the Service Investigator Page

On the **[Anomaly Detection]** tab on a **Service Investigator** page, you can enable anomaly detection events for additional metrics or disable anomaly detection metric events on which it is currently enabled.

**NOTE:** The **[Anomaly Detection]** tab appears only if you have at least one device in the selected service that has anomaly detection enabled.

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

1. On the **Business Services** page () , select a service from the list of business, IT, and device services by clicking its name. The **Service Investigator** displays.
2. On the **Service Investigator** page, click the **[Anomaly Detection]** tab.
3. On the **[Anomaly Detection]** tab, click the **Actions** icon () for any of the listed metrics and select *Enable*. The **Select Available Metrics** modal appears.

4. In the **Select Metric** drop-down, click the name of the metric on which you want to enable anomaly detection events for the device.
5. 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 .
6. Click **[Enable]**.

**TIP:** To disable anomaly detection for a metric, click the **Actions** icon (⋮) for that metric and select *Disable*. The metric is removed from the **[Anomaly Detection]** tab.

---

## Creating an Event Policy for Anomalies

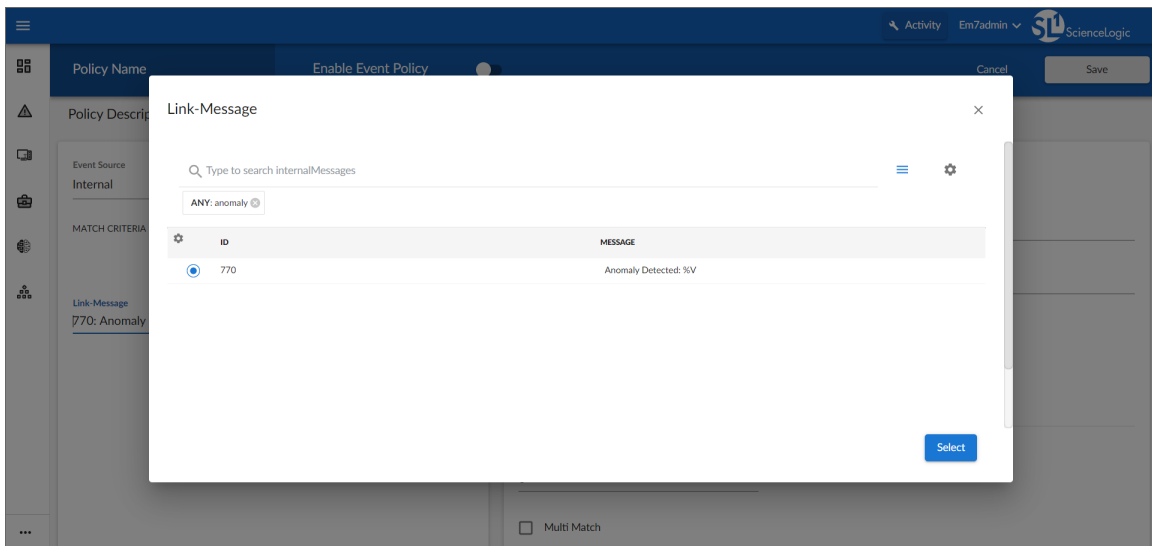
You can create additional event policies that will trigger events in SL1 when anomalies are detected for those devices.

**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 SL1 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 SL1 has had an opportunity to learn more about the device metric's data patterns.

To create an event policy for anomalies:

1. Go to the **Event Policies** page (Events > Event Policies, or Registry > Events > Event Manager in the classic SL1 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

SL1 includes automation features that allow you to define specific event conditions and the actions you want SL1 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 SL1.

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:

Automation Policy Editor | Creating New Automation Policy

Policy Name: Anomaly High

Policy Type: [Active Events]

Policy State: [Enabled]

Policy Priority: [Default]

Organization: Sample

Criteria Logic: [Severity >=] [Minor]

Match Logic: [Text search]

Match Syntax: [ ]

Repeat Time: [Only once]

Align With: [Devices]

Include events for entities other than devices (organizations, assets, etc.): ☐

Trigger on Child Rollup: ☐

Available Devices:

- Sample
- AWS: Service: test
- ScienceLogic, Inc.: EM7 Data Collector: mrktng-dc1
- ScienceLogic, Inc.: EM7 Data Collector: mrktng-dc2
- System

Aligned Devices:

(All devices)

Available Events:

- anom
- [1768] Critical: Anomaly Score Critical - new york
- [18] Minor: Anomaly Score Minor
- [17] Notice: Anomaly Score Notice

Aligned Events:

- [20] Critical: Anomaly Score Critical
- [19] Major: Anomaly Score Major

Available Actions:

- SNMP Trap [1]: SL1 Event Trap
- Snippet [5]: Automation Utilities: Calculate Memory Size for Ea
- Snippet [5]: AWS: Account Creation
- Snippet [5]: AWS: Account Write Back
- Snippet [5]: AWS: Disable Instance By Tag

Aligned Actions:

Save

3. In the **Policy State** field, select *Enabled*.
4. In the **Available Events** field, search for and select an anomaly-related event policy, and then click the right-arrow icon to move it to the **Aligned Events** field. For more information about anomaly-related events, see [Creating an Event Policy for Anomalies](#).
5. 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](#).
6. When you are finished, click **[Save]**.



---

# Chapter

# 5

## Skylar Analytics: Predictive Alerting

---

### Overview

The Predictive Alerting component of Skylar Analytics helps to avoid problems such as file systems running out of space, hosts running out of memory, or issues with network reliability due to oversubscription. The alerts are generated in advance of the problem and can provide days, weeks, or months of notice depending upon the conditions.

For this release, the Predictive Alerting component monitors file systems (SNMP, PowerShell, SSH) network interfaces (utilization, errors, discards), and memory.

This chapter covers the following topics:

<i>What is Predictive Alerting?</i> .....	39
<i>Viewing Predictive Alerts in SL1</i> .....	39

# What is Predictive Alerting?

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

## Viewing Predictive Alerts in SL1

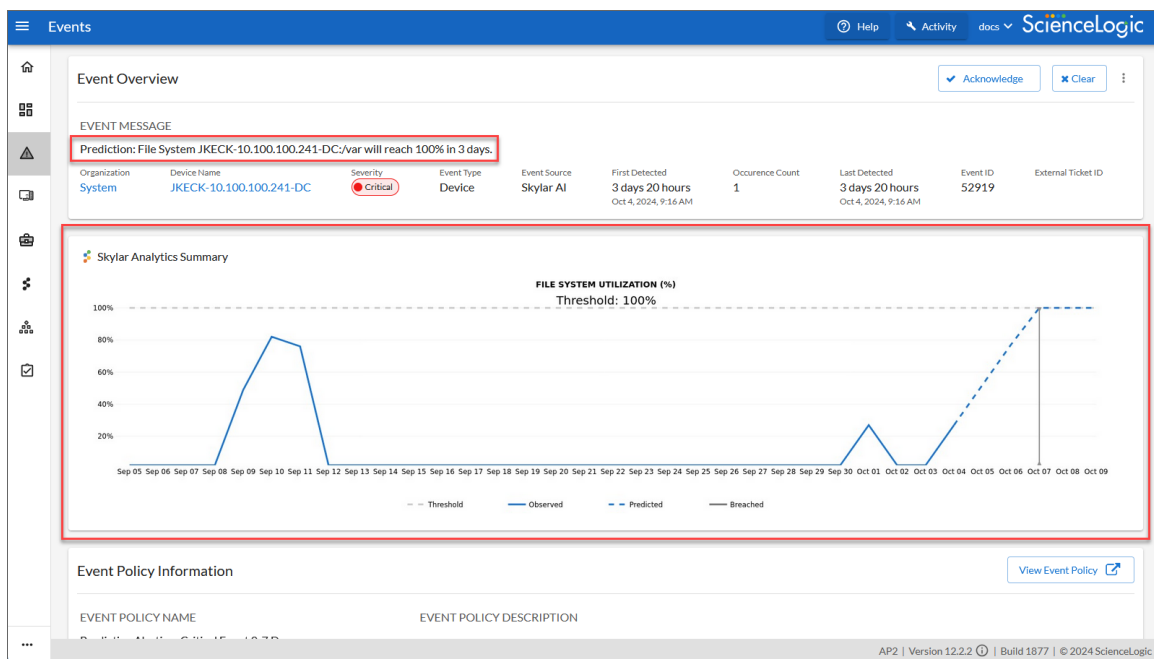
When your SL1 system is connected to Skylar AI, you can start viewing predictive alerts in SL1. No additional configuration is needed.

Predictive alerts display the Skylar icon (🧙) to the left of the event message in the **Message** column of the **Events** page, and the message starts with the word "Prediction":

	Organiz...	Severity	Name	Message	Last Det...	Age	Ticket ID	Count	Event Ty...	Event N...	Masked Events	Event So...	Acknowledge	Clear
<input type="checkbox"/>	Sample	Critical	mrktng-dc2	Host Resource: Storage Utilization (/var/log) of type H...	Oct 8, 2024,	1 month 18 d	—	14274	Device			Dynamic	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Host Resource: Storage Utilization (/var/tmp) of type H...	Oct 8, 2024,	1 day 22 hou	—	561	Device			Dynamic	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	ISR-4331-R1	Fan problem, Fan (Fan 1 Critical) state: shutdown	Oct 8, 2024,	2 months 7 c	—	3310	Device			Dynamic	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	4948-SW-01	Power supply problem, Power supply (Power Supply 1)	Oct 8, 2024,	2 months 2 c	—	3055	Device			Dynamic	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	4948-SW-01	Power supply problem, Power supply (Power Supply 1)	Oct 8, 2024,	2 months 2 c	—	3055	Device			Dynamic	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	Sample	Critical	mrktng-dc2	/var/log: File system usage exceeded critical threshold	Oct 8, 2024,	1 month 18 d	—	4732	Device			Internal	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	/var/tmp: File system usage exceeded critical threshold	Oct 8, 2024,	1 day 22 hou	—	184	Device			Internal	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Prediction: File System JKECK-10.100.100.241-D	Oct 5, 2024,	2 days 9 hou	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Prediction: File System JKECK-10.100.100.241-D	Oct 4, 2024,	3 days 19 ho	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	xdemo-vc1-	Prediction: CPU Utilization will reach 100% in 5 d	Oct 1, 2024,	6 days 9 hou	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Prediction: File System JKECK-10.100.100.241-D	Oct 1, 2024,	6 days 13 ho	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Prediction: File System JKECK-10.100.100.241-D	Oct 1, 2024,	6 days 13 ho	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	linux-web02	Prediction: CPU Utilization will reach 100% in 3 d	Sep 30, 2024,	7 days 9 hou	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	linux-web02	Prediction: CPU Utilization will reach 100% in 3 d	Sep 30, 2024,	7 days 11 ho	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	linux-web02	Prediction: CPU Utilization will reach 100% in 4 d	Sep 30, 2024,	7 days 15 ho	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	skylar-ai-de	Prediction: File System skylar-ai-demo/home will	Sep 27, 2024,	10 days 14 h	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Prediction: File System JKECK-10.100.100.241-D	Sep 27, 2024,	10 days 15 h	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	Sample	Critical	mrktng-dc2	Prediction: File System mrktng-dc2:/var/log will re	Sep 27, 2024,	10 days 15 h	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear
<input type="checkbox"/>	System	Critical	JKECK-10.1	Prediction: File System JKECK-10.100.100.241-D	Sep 23, 2024,	14 days 22 h	—	1	Device			Skylar AI	<input checked="" type="checkbox"/> Acknowledge	<input checked="" type="checkbox"/> Clear

To view details about a predictive alert:

1. On the **Events** page, click the message for a predictive alert with the Skylar icon (🌟). The **Event Investigator** page for that alert appears.
2. 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 SL1, and the gray dotted line represents the threshold set in SL1. 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 an SL1 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.

**TIP:** If the graph in the **Skylar Analytics Summary** panel does not load, open the `/opt/em7/nextui/nextui.conf` file and make sure the following line is included in the file:  
`GQL_USE_AI_EVENT_VISUALIZATIONS=enabled`

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

1. 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:

Date/Time	Source	Event ID	Severity	Syslog Severity	Message
Nov 17, 2024, 9:17 PM	AIEngine	89455	Minor	—	Prediction: CPU Utilization will reach 100% in 18 days.
Nov 14, 2024, 9:21 PM	AIEngine	89455	Minor	—	Prediction: CPU Utilization will reach 100% in 17 days.
Nov 13, 2024, 9:18 PM	AIEngine	89455	Minor	—	Prediction: CPU Utilization will reach 100% in 18 days.
Nov 12, 2024, 9:19 PM	AIEngine	89455	Minor	—	Prediction: CPU Utilization will reach 100% in 19 days.
Nov 11, 2024, 9:20 PM	AIEngine	89455	Minor	—	Prediction: CPU Utilization will reach 100% in 18 days.
Nov 9, 2024, 9:17 PM	AIEngine	93091	Notice	—	Prediction: CPU Utilization will reach 100% in 29 days.
Nov 8, 2024, 9:17 PM	AIEngine	93091	Notice	—	Prediction: CPU Utilization will reach 100% in 28 days.
Nov 7, 2024, 7:11 PM	AIEngine	94604	Critical	—	Prediction: File System mrktng-dc2-/var/log will reach 100% in 0 days.
Nov 4, 2024, 9:22 PM	AIEngine	94022	Major	—	Prediction: CPU Utilization will reach 100% in 11 days.
Nov 4, 2024, 7:35 PM	AIEngine	93939	Notice	—	Prediction: File System mrktng-dc2-/ will reach 100% in 28 days.
Nov 3, 2024, 9:28 PM	AIEngine	93091	Notice	—	Prediction: CPU Utilization will reach 100% in 20 days.

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

---

# Appendix

# A

## Appendix: Service Provider Administration for Skylar AI

---

### Overview

This chapter explains the different tasks that a user with the **Service Provider** role can perform in Skylar AI. A **Service Provider** user can provision new accounts.

**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> .....	43
<i>Provisioning a New Account</i> .....	43

---

## 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.

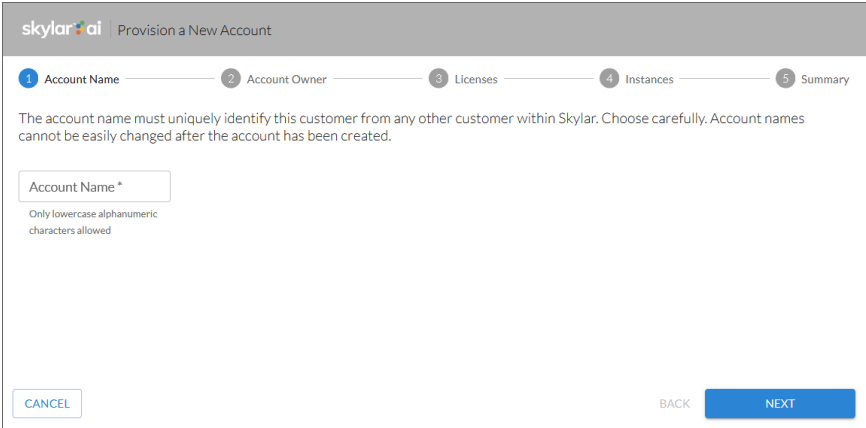
---

## Provisioning a New Account

You can have multiple accounts in a single Skylar AI system. To add a new account, you will need to provision that account in the **Skylar Settings** page.

To create a new account:

1. Create a new account by clicking the **All Accounts** drop-down at the top of the **Skylar Settings** page and clicking **[Provision New Account]**. The **Provision a New Account** wizard appears:



2. On the **Account Name** page, type the **Account Name** using only lower-case alphanumeric characters, and then click **[Next]**.
3. On the **Account Owner** page, specify the **First Name**, **Last Name**, and **Email** for the first user of the new account. When you type the email address, Skylar AI adds the domain name from that email into the **Claim Email Domain** field. Click **[Next]**.

**NOTE:** When Single-Sign-On (SSO) through SAML is enabled, users that log in with the domain used by SAML will be redirected to the SAML provider for this account.

4. On the **Licenses** page, select **Skylar Analytics** to enable Skylar Analytics for this account.
5. If you select **Enable ODBC**, you will need to add the IP addresses for your ODBC client in the **ODBC Client IP Ranges** field. Click **[Next]**.

6. On the **Instances** page, type the name of your instance for this account, using only lower-case alphanumeric characters. You can also use *default* as the instance name. Click **[Next]**.
7. On the **Summary** page, review your settings and click **[Begin Provisioning]** to continue setting up the account. The provisioning process begins, 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 3 to the first user. On first login, the new user will be prompted to change their password.

8. To set up single-sign-on (SSO) authentication with SAML for this new account, see [Configuring SSO Authentication with SAML](#).

© 2003 - 2025, ScienceLogic, Inc.

All rights reserved.

#### LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

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 on this Site, information on this Site 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 in this Site at any time without notice.

#### Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

#### Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: [legal@sciencelogic.com](mailto:legal@sciencelogic.com). For more information, see <https://sciencelogic.com/company/legal>.





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

International: +1-703-354-1010