



---

## Skylar Analytics

Version 1.7.0

---

# 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	3
Data Analyzed by Skylar AI	3
What is Skylar Analytics?	4
Getting Started with Skylar Analytics	4
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	14
Using Access Tokens for Users	16
Adding and Upgrading Dashboards	16
<b>Skylar Analytics: Data Visualization and Data Exploration</b>	<b>18</b>
What is Data Visualization?	19
Working with Datasets in Data Visualization	20
Components of a Dataset	20
Viewing the List of Datasets	21
Viewing the Contents of a Dataset	22
Viewing Dashboards and Charts in Data Visualization	22
Logging In to the Data Visualization Component	23
Default Skylar Analytics Dashboards	24
Viewing Skylar Analytics Dashboards	25
Creating and Customizing Dashboards and Charts	27
Creating a Dashboard	27
Adding Contextual Cross-filtering to a Dashboard	32

Customizing a Dashboard .....	34
Icons for Chart Metrics .....	35
Customizing the Default Column Names for Charts .....	35
Data Exploration: Exporting Data to Skylar AI from Third-party Tools .....	37
Configuring Data Exploration with Power BI .....	38
Additional Resources for Skylar Analytics (Apache Superset Training) .....	41
<b>Skylar Analytics: Anomaly Detection .....</b>	<b>43</b>
What is Anomaly Detection? .....	44
How Anomaly Detection Works .....	44
Viewing Graphs and Data for Anomaly Detection .....	45
Enabling Thresholds and Alerts for the Anomaly Chart .....	46
Enabling Anomaly Detection Events for Specific Metrics .....	48
Creating an Event Policy for Anomalies .....	49
Using Anomaly-related Events to Trigger Automated Run Book Actions .....	50
<b>Skylar Analytics: Predictive Alerting .....</b>	<b>52</b>
What is Predictive Alerting? .....	53
How Predictive Alerting Works .....	53
Viewing Predictive Alerts in SL1 .....	53
Using Predictive Alerts to Trigger Automated Run Book Actions .....	56
<b>Service Provider Administration for Skylar AI .....</b>	<b>58</b>
First Login as a Service Provider User .....	59
Provisioning a New Account .....	59
Adding an ODBC User .....	60

---

# 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 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, hosts running out of memory, or issues with network reliability due to oversubscription. The alerts appear as enriched events within SL1.

**IMPORTANT:** While ScienceLogic recommends that you use SL1 version 12.3.6 or later, 12.3.2 is the minimum SL1 version you can use with this release. ScienceLogic recommends that you always use the most recent SL1 and AP2 releases in conjunction with the most recent Skylar AI release to 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:

<a href="#">What is Skylar AI?</a>	2
<a href="#">What is Skylar Analytics?</a>	4
<a href="#">Getting Started with Skylar Analytics</a>	4
<a href="#">Mapping SL1 Dynamic Application Object Names to Skylar Columns</a>	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

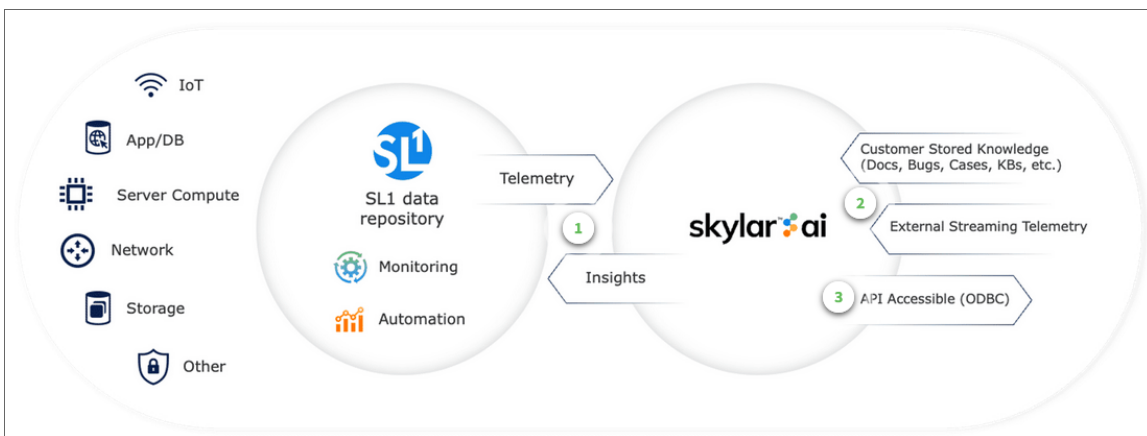
The Skylar AI family of services currently includes the following components:

- **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.

**NOTE:** This manual covers Skylar Analytics. At the Product Documentation site, you can find documentation for [Skylar Automated RCA](#).

## 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
- DCM(+R) relationships
- Dynamic Application mapping and performance data
- Metadata for web content, SOAP/XML transaction, and domain name monitors
- Process and service data

- SL1 Agent data, including Gen 1 (SL1 Distributed Environment) and Gen 3 (SL1 Extended Architecture) agents
- Topology data for L2, L3, CDP, LLDP, and ad-hoc relationships between devices

---

## 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.

**NOTE:** SL1 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 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 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, hosts running out of memory, or issues with network reliability due to oversubscription. The alerts appear as enriched events within SL1.

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


---

## Getting Started with Skylar Analytics

**NOTE:** These instructions are only for on-premises configurations of Skylar AI. The ScienceLogic SRE team performs these steps for SaaS configurations of Skylar AI.

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:** 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
- `<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](#).

**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 Skylar AI is installed.

This command will configure the OpenTelemetry Collector, restart services that export data, and check that connectivity to the supplied endpoints is healthy.

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.



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

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

**TIP:** To check to make sure you have connected Skylar AI to SL1, go to SL1 and look for the **Skylar AI** page (🔗). 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. After a few minutes, the Data Visualization charts will populate with data if the connection was successful.

3. 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
```

**NOTE:** 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 procedure 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*.

For information about how to use these components, see the following chapters:

- [Skylar Analytics: Anomaly Detection](#)
- [Skylar Analytics: Predictive Alerting](#)

For information about setting up authentication and access control for users in Skylar AI, see [Configuring Access Control in Skylar Analytics](#).

---

## 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
  - 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

---

# 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>	14
<i>Using Access Tokens for Users</i>	16
<i>Adding and Upgrading Dashboards</i>	16

---

## 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 Security Assertion Markup Language (SAML) single sign-on (SSO) authorization configured as needed.
- Access tokens for integration with external tools.

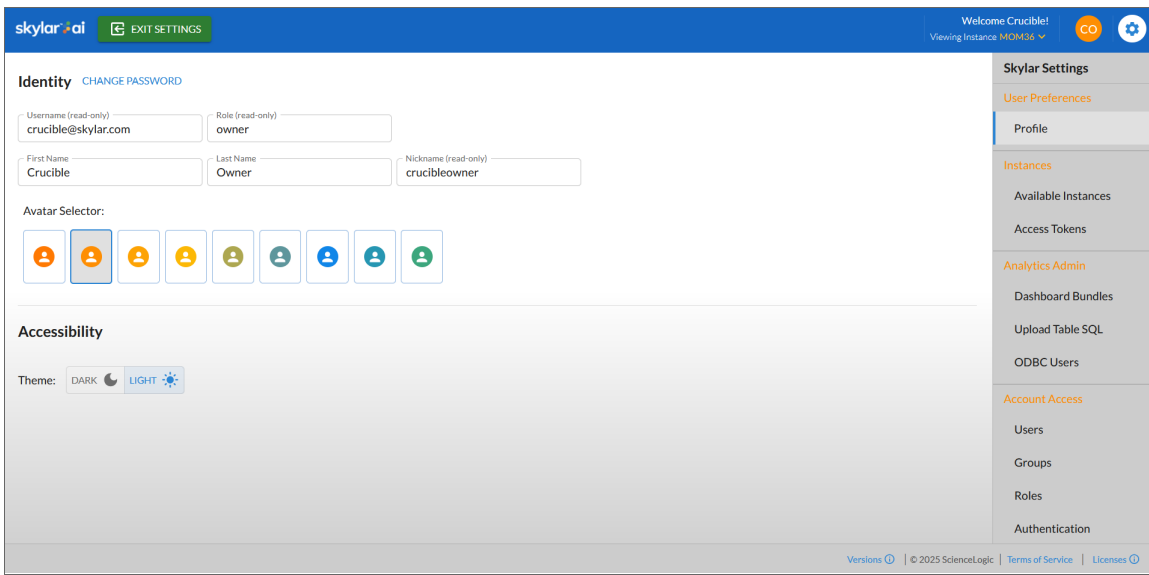
When a user is logged in to Skylar AI, that session uses the following rules:

- 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.
- 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]** tab of the **Skylar Settings** page.
- 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]** tab in **Skylar Settings**.
- User accounts that have not been active for 35 days are automatically locked.

---

## 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 following image displays a **[Profile]** tab for a user with the "Owner" role:



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 update new dashboards and 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 or the **[Exit Settings]** button 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.
- To view version numbers for the **Skylar Settings** user interface, click the **Versions** link in the footer of any page. From the footer, you can also click links to view the Terms of Service and information about licenses and open-source packages.

## 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** 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. If only one instance is available, you will use the instance labeled "default".

On the **[Available Instances]** tab of the **Skylar Settings** page, you can view a list of instances for the current user. An "Admin" user can also access the "Analytics Secrets" for an instance, which contains the Microsoft Open Database Connectivity (ODBC) host, password, port, and user information for Data Exploration using ODBC. Also, 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 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:
  - **Super User.** Assigned to the single **admin** user to 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. You can also add, edit, and delete users.

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

- **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 his or her 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.

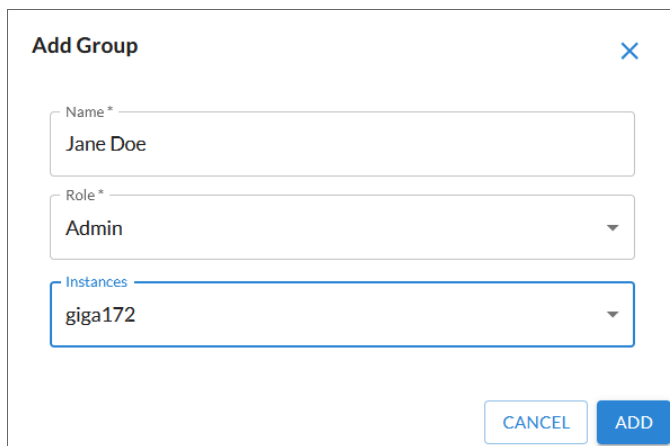
In case of an issue with authenticating, you can contact ScienceLogic to disable SAML for the account and potentially reset the owner's local (non-SAML) password if needed.

**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.

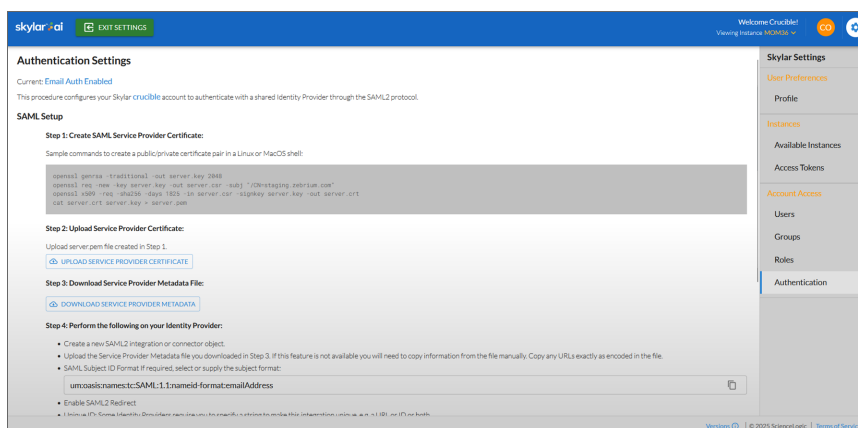
To set up SSO Authentication with SAML in the Skylar AI user interface:

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 shown. It has a title bar with 'Add Group' and a close button (X). Inside, there are three input fields: 'Name \*' with the value '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 is shown. It has a header with 'skylar AI' and 'EXIT SETTINGS'. The main content area is titled 'Authentication Settings' and shows 'Current: Email Auth Enabled'. Below this, there is a section for 'SAML Setup' with four steps: 1. Create SAML Service Provider Certificate, 2. Upload Service Provider Certificate, 3. Download Service Provider Metadata File, and 4. Perform the following on your Identity Provider. The page also has a sidebar with 'Skylar Settings' and various tabs like 'User Preferences', 'Profile', 'Instances', 'Available Instances', 'Access Tokens', 'Account Access', 'Users', 'Groups', 'Roles', and 'Authentication'.

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. 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.

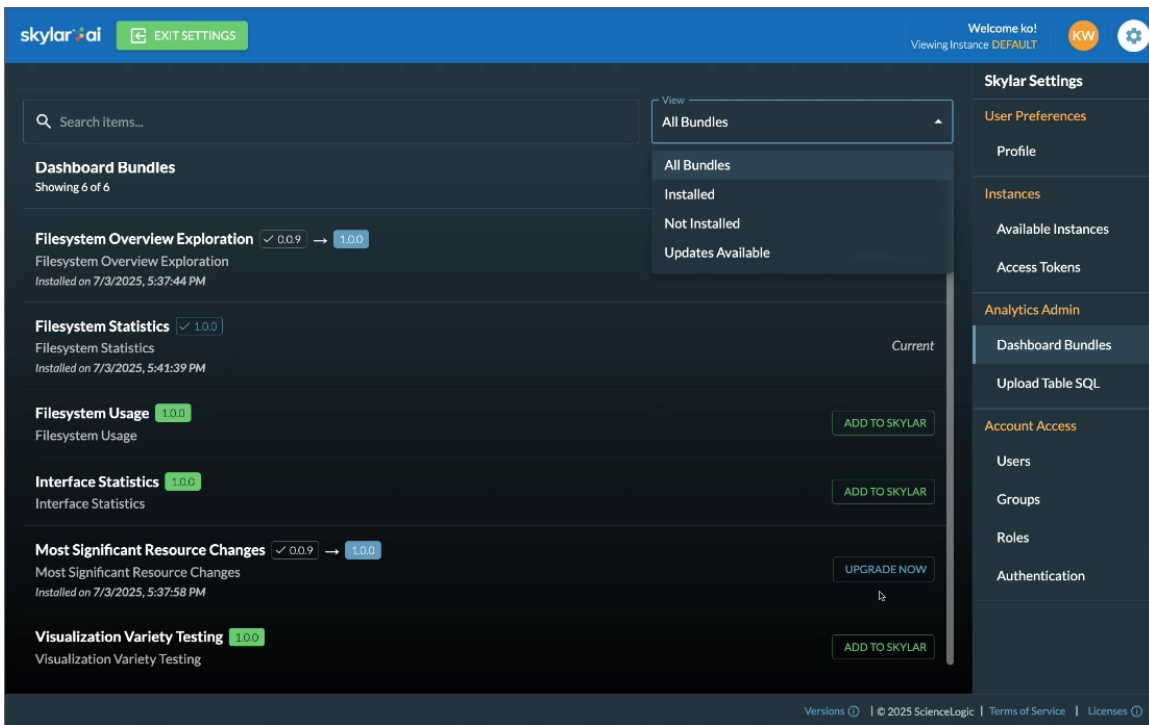
To create an access token:

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.

---

## Adding and Upgrading Dashboards

A user with an Owner role can add new dashboards and upgrade existing dashboards on the **[Dashboard Bundles]** tab of the **Skylar Settings** page:



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

These dashboards include "(Sample)" at the end of each dashboard name.

The options on the **[Dashboard Bundles]** tab include:

- **Current**. Shows that you are running the most recent version of a dashboard.
- **[Add to Skylar]**. Click this button to install a new dashboard for Skylar Analytics
- **[Upgrade Now]**. Click this button to upgrade an existing dashboard.

---

# 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 and SL1. 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 SL1 updates.

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

This chapter covers the following topics:

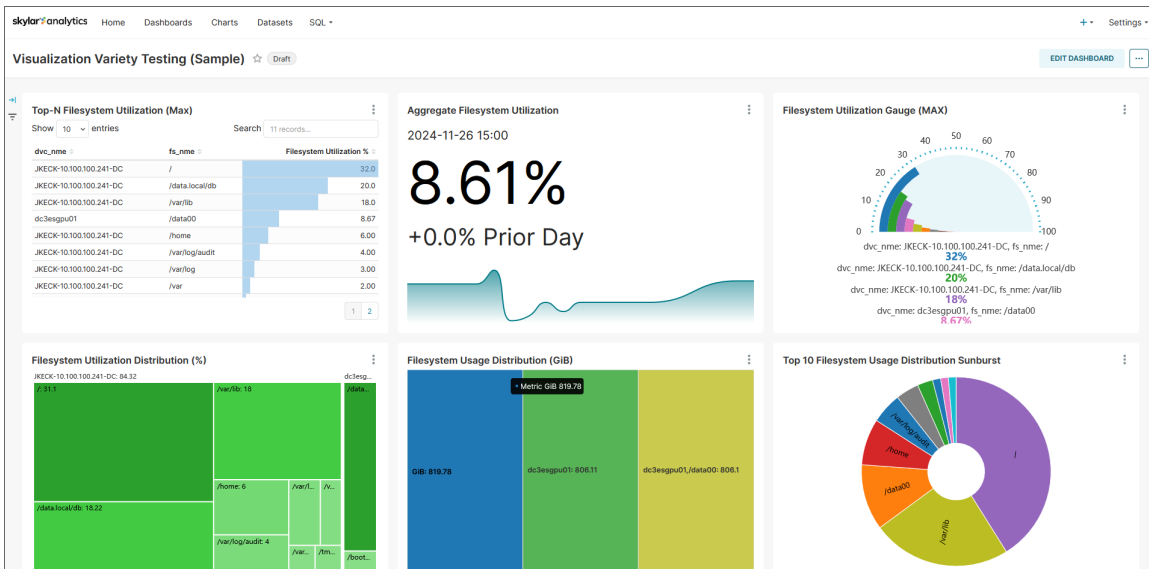
<i>What is Data Visualization?</i> .....	19
<i>Working with Datasets in Data Visualization</i> .....	20
<i>Viewing Dashboards and Charts in Data Visualization</i> .....	22
<i>Creating and Customizing Dashboards and Charts</i> .....	27
<i>Data Exploration: Exporting Data to Skylar AI from Third-party Tools</i> .....	37
<i>Additional Resources for Skylar Analytics (Apache Superset Training)</i> .....	41

# 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 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 SL1. 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 SL1 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 SL1. A presentation object for a Dynamic Application defines how SL1 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.

### Components of a Dataset

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

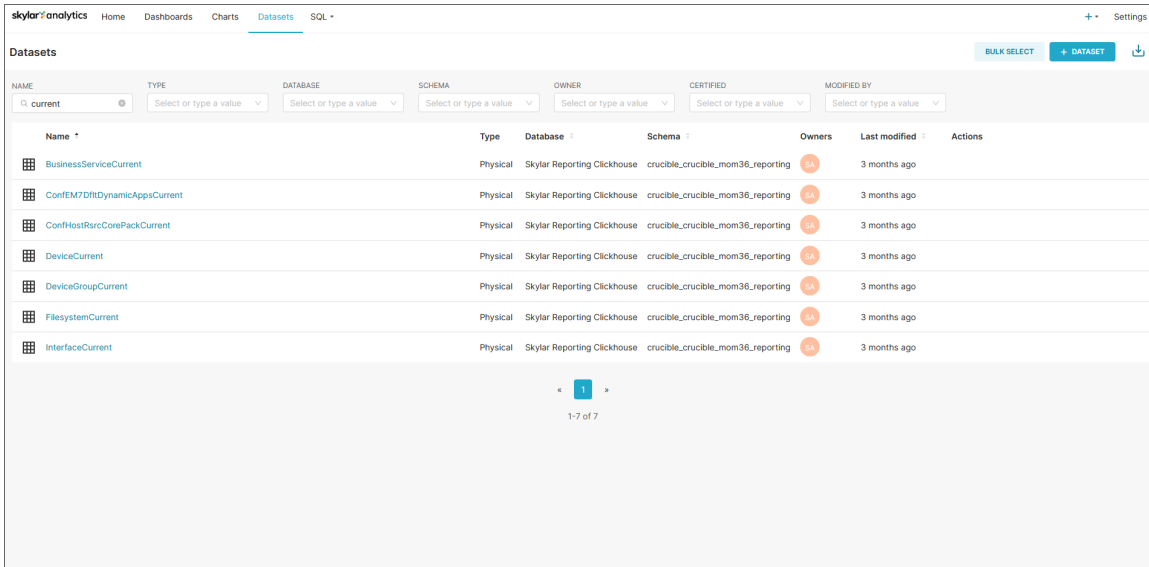
- **One performance dataset.** In Skylar Analytics, these datasets use the naming convention of "Metric<PowerPackName>", such as "MetricMSWinSvr" for the "Microsoft: Windows Server" PowerPack.
- **Two configuration datasets:**
  - The "Current" dataset contains only the last recorded configuration change. 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 SL1. 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 SL1 Dynamic Application Object Names to Skylar Columns](#).

## Viewing the List of Datasets

You can view a list of the datasets in Skylar Analytics on the **Datasets** page:



The screenshot shows the 'Datasets' page in Skylar Analytics. At the top, there's a navigation bar with 'Home', 'Dashboards', 'Charts', 'Datasets' (selected), and 'SQL'. Below the navigation bar, there's a search bar with 'current' entered. To the right of the search bar are buttons for 'BULK SELECT', '+ DATASET', and a download icon. Below the search bar is a table with columns: NAME, TYPE, DATABASE, SCHEMA, OWNER, CERTIFIED, and MODIFIED BY. The table lists seven datasets, all of which are 'Physical' type, owned by 'Skylar Reporting Clickhouse', and have a schema of 'crucible\_crucible\_mom36\_reporting'. The datasets are: BusinessServiceCurrent, ConfEM7DfnsDynamicAppsCurrent, ConfHostRarcCorePackCurrent, DeviceCurrent, DeviceGroupCurrent, FilesystemCurrent, and InterfaceCurrent. All datasets were last modified '3 months ago'. At the bottom of the table, there's a pagination control showing '1' of 7.

NAME	TYPE	DATABASE	SCHEMA	OWNER	CERTIFIED	MODIFIED BY
BusinessServiceCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	3 months ago	
ConfEM7DfnsDynamicAppsCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	3 months ago	
ConfHostRarcCorePackCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	3 months ago	
DeviceCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	3 months ago	
DeviceGroupCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	3 months ago	
FilesystemCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	3 months ago	
InterfaceCurrent	Physical	Skylar Reporting Clickhouse	crucible_crucible_mom36_reporting	SA	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.

Clicking the name of a dataset on the **Datasets** page takes you to 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](#).



## Viewing the Contents of a Dataset

To view the contents of a dataset:

1. From SL1, go to the **Skylar AI** page (🔗) and click the **[Visit]** button for **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 SL1.

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. On the left, the 'DATABASE' dropdown is set to 'clickhousedb' and the 'SCHEMA' dropdown is set to 'crucible\_crucible\_mon36\_reporting'. The 'SEE TABLE SCHEMA' dropdown is set to 'ConfMSWinSvrCurrent'. The 'RESULTS' tab is active, showing a table with 39 rows. The table has columns: DataDate, DataTime, InstID, PowerPack, PowerPackLong, DynamicApp, and DynamicAppLong. The first row of data is highlighted.

DataDate	DataTime	InstID	PowerPack	PowerPackLong	DynamicApp	DynamicAppLong
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrBIOSConf	Microsoft: Windows Server BIOS Configuration
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrBIOSConf	Microsoft: Windows Server BIOS Configuration
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrCPUConf	Microsoft: Windows Server CPU Configuration
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDCMRelationship	Microsoft: Windows Server DCM+R Relationship
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDCMRelationship	Microsoft: Windows Server DCM+R Relationship
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDCMRelationship	Microsoft: Windows Server DCM+R Relationship
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDCMRelationship	Microsoft: Windows Server DCM+R Relationship
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDiskConf	Microsoft: Windows Server Disk Configuration
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDiskConf	Microsoft: Windows Server Disk Configuration
2025-07-02	2025-07-02T16:00:00+00:00	crucible_mon36	MSWinSvr	Microsoft: Windows Server	MSWinSvrDiskConf	Microsoft: Windows Server Disk 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 SL1.

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. 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.

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

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.

**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.

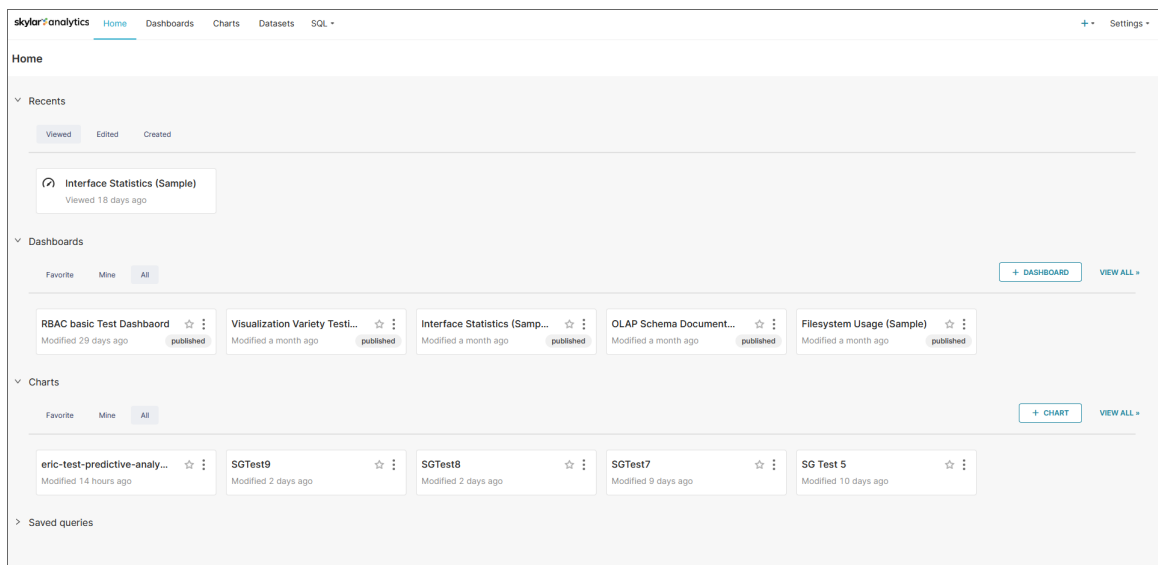
## Logging In to the Data Visualization Component

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

1. From SL1, go to the **Skylar AI** page (🔗) and click the **[Visit]** button for **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 SL1.

2. Click **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]** tab or the **[Charts]** tab to view a list of all dashboards or charts.

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

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.

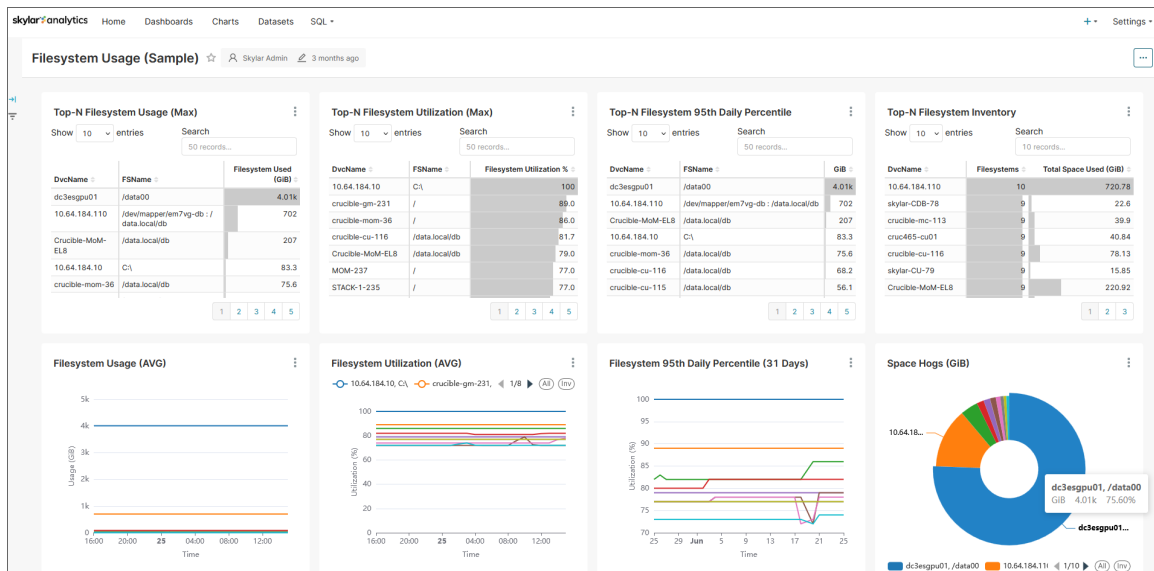
**IMPORTANT:** A user with an Owner role can add new dashboards and upgrade existing dashboards. For more information, see [Adding and Upgrading Dashboards](#).

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.

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.

## 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 **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 SL1, such as server reports or SL1 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 SL1 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:

The screenshot shows the 'skylar analytics' interface for creating a new chart. The top navigation bar includes 'Home', 'Dashboards', 'Charts', 'Datasets', and 'SQL'. The main area is titled 'Add the name of the chart' and has a 'SAVE' button. The interface is split into three columns:

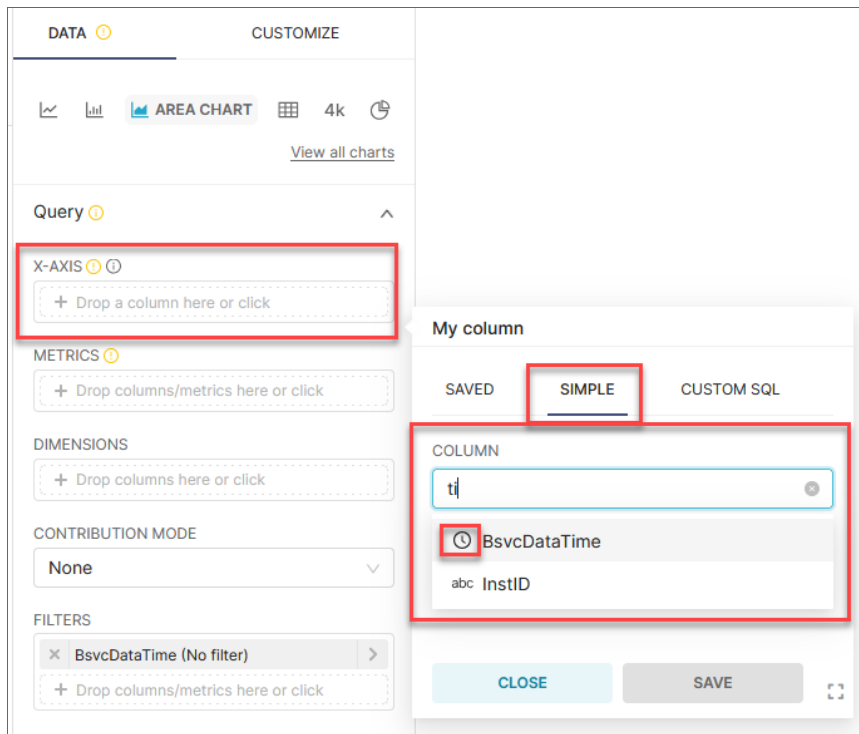
- Chart Source:** Contains a search bar and two lists: 'Metrics' (showing 1 of 1, with 'COUNT(\*)') and 'Columns' (showing 13 of 13, including 'BsvcDateTime', 'DvcCategory', 'DvcClass', 'InstID', 'DvcID', 'DvcHostname', 'DvcName', 'DvcIPv4', 'BsvcID', 'BsvcName', 'BsvcHealth', 'BsvcAvailability', and 'BsvcRisk').
- DATA:** Contains configuration options for the chart. It includes a 'Query' section with an 'X-AXIS' field, a 'METRICS' section, a 'DIMENSIONS' section, a 'CONTRIBUTION MODE' dropdown (set to 'None'), a 'FILTERS' section (with 'BsvcDateTime (No filter)' selected), a 'SERIES LIMIT' dropdown (set to 'None'), a 'SORT BY' field, and a 'ROW LIMIT' dropdown (set to '10000'). A 'CREATE CHART' button is at the bottom.
- CUSTOMIZE:** Contains a preview of the chart, which is currently a bar chart. It includes a message: 'Add required control values to preview chart. Select values in highlighted field(s) in the control panel. Then run the query by clicking on the "Create chart" button.'

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.



10. 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:



11. 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".
12. 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]**.

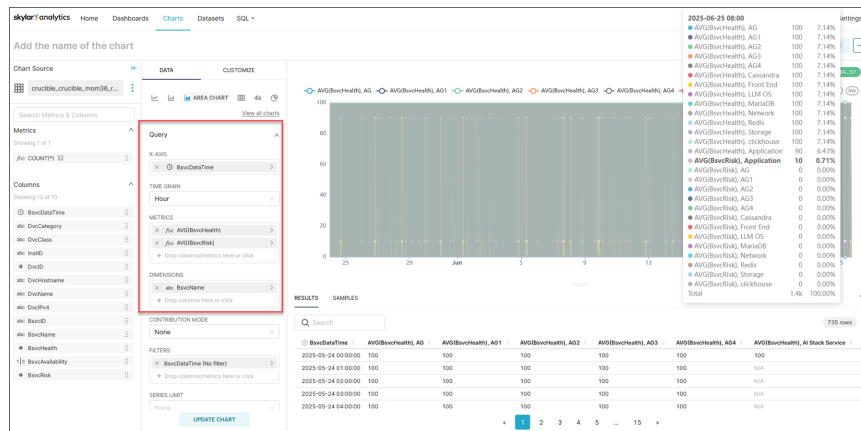
**TIP:** For more information about the abbreviations used for the metric names, see [Mapping SL1 Dynamic Application Object Names to Skylar Columns](#).

13. 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*.
14. 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.

15. Click **[Save]** to save the metrics.
16. 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.
17. 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.

18. 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.

19. 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.

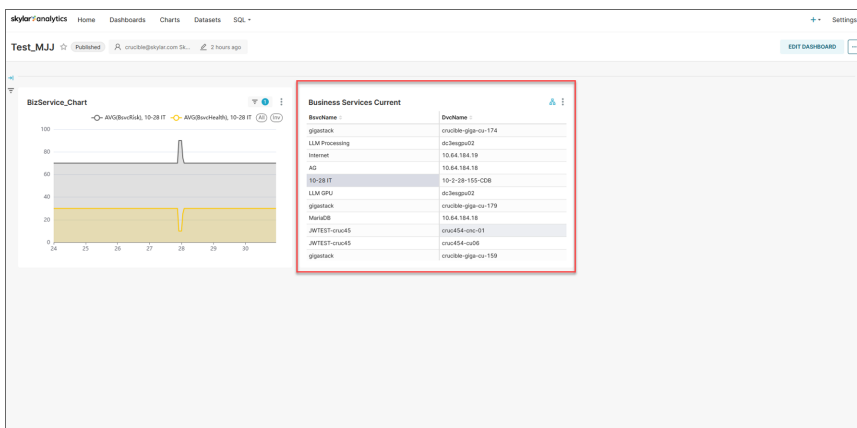
5. Select a chart type of *Table*, which will make it easy for users of the dashboard to select a business service to view more information.
6. Click **[Create New Chart]**. A new chart window appears.
7. Click and drag the *BsvcName* column onto the **Dimensions** field on the **[Data]** tab. This will 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*.

8. 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.

9. Click **[Create Chart]**. The chart displays in the preview area.
10. Click **[Save]** to save the metrics. The **Save chart** modal appears.
11. 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:

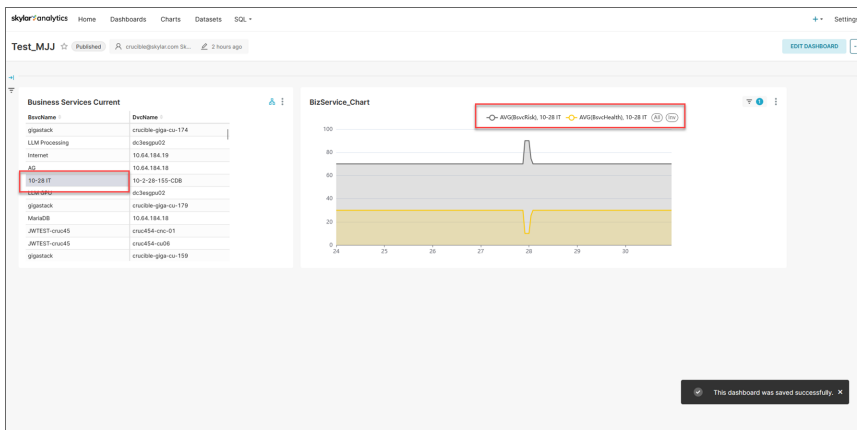



12. Typically, you would want to place the chart that drives the "context" (or the contextual cross-filtering) first, so click **[Edit Dashboard]** to enter Edit mode again.
13. Click and drag the next chart to the left of the first chart until a vertical rectangle appears. Drop the new chart onto that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

14. 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 (  ) 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 (  ) 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.

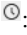
4. If you want to add extra elements to your dashboard, like a header, additional text, a divider, or other items, drag and drop the items onto the dashboard from the **[Layout Elements]** tab. Click **[Save]** when you are done, and click **[Edit Dashboard]** again to keep editing.
5. To edit a chart in the dashboard, click the vertical ellipsis button (  ) at the top right of the chart on the dashboard and select **Edit chart**. 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:

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
DataDate	DATE	<input checked="" type="checkbox"/>	<input checked="" type="radio"/>	<input checked="" type="checkbox"/>
DateTime	DATETIME	<input checked="" type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>
InstID	LOWCARDINALITY(String)	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>
PowerPack	LOWCARDINALITY(String)	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>
PowerPackLong	LOWCARDINALITY(String)	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>
DynamicApp	LOWCARDINALITY(String)	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>
DynamicAppLong	LOWCARDINALITY(String)	<input type="checkbox"/>	<input type="radio"/>	<input checked="" type="checkbox"/>

CANCEL SAVE

**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:

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.

---

## 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.



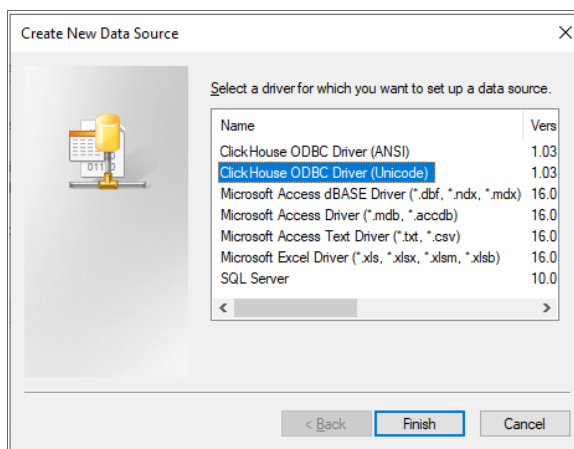
## Configuring Data Exploration with Power BI

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.

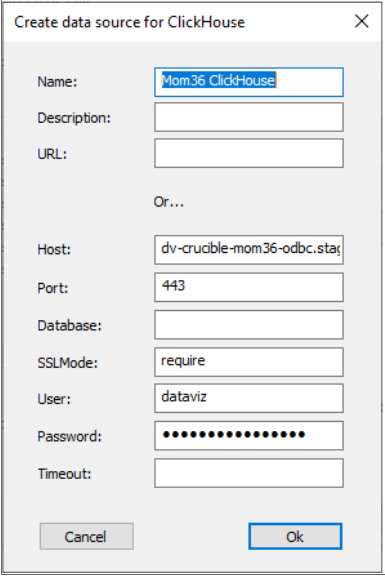
**TIP:** For an example of how you can connect DBeaver, another business intelligence tool, to Skylar Analytics, see the following video: [BYO Tool and Data: Connect via ODBC and Import Data](#).

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:



Create data source for ClickHouse

Name: Mom36 Click-house

Description:

URL:

Or...

Host: dv-crucible-mom36-odbc.stat

Port: 443

Database:

SSLMode: require

User: dataviz

Password: .....

Timeout:

Cancel Ok

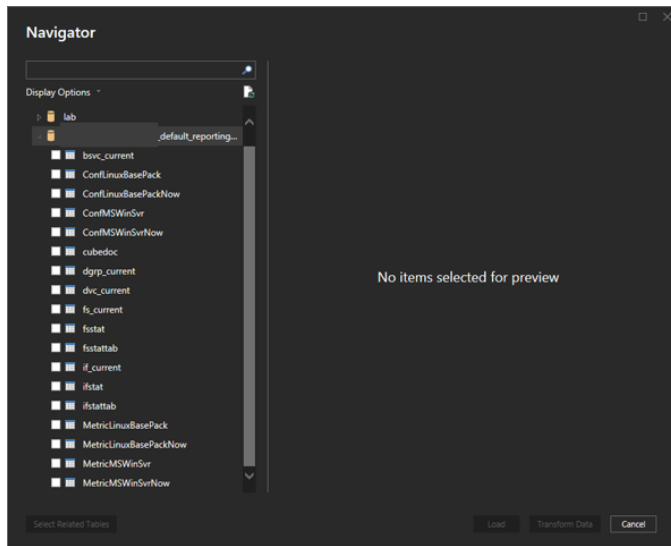
6. Complete the following fields with ODBC connection details from the ScienceLogic Site Reliability Engineering (SRE) team:
  - **Name:** Add a name to identify this connection. This will be used later in the BI tools.
  - **Host:** Specify the host URL, provided by SRE.
  - **Port:** 443.
  - **Database:** 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:



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

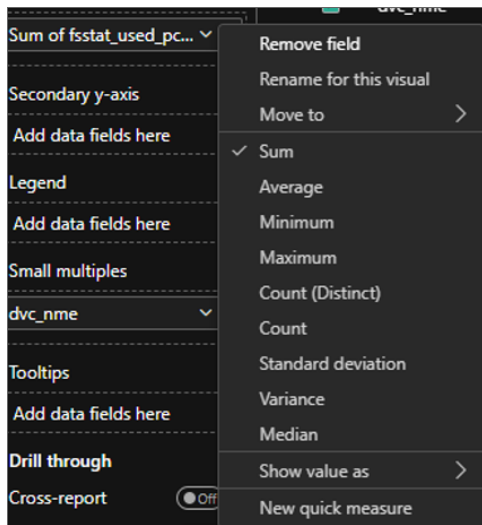
**TIP:** When selecting datasets to import, choose only the necessary tables to optimize performance. The following procedure creates a sample dashboard in Power BI.

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.

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 SL1 display their data in Skylar Analytics, how to verify if you have current data from SL1 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](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.

---

# 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 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.

**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> .....	44
<i>Viewing Graphs and Data for Anomaly Detection</i> .....	45
<i>Enabling Anomaly Detection Events for Specific Metrics</i> .....	48
<i>Creating an Event Policy for Anomalies</i> .....	49
<i>Using Anomaly-related Events to Trigger Automated Run Book Actions</i> .....	50

---

## 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 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 SL1 to Skylar AI.

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

## 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 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 SL1 **Events** page, filtered by "Anomaly messages (Skylar AI (🔗) > [Visit] button for Skylar Anomaly Detection).
- The **[Anomaly Detection]** tab in the **Device Investigator**.
- The **[Anomaly Detection]** tab in the **Service Investigator** for a business, IT, or device service.

**NOTE:** For Skylar Analytics version 1.7.0, the **Anomaly Detection** page in SL1 was replaced by the filtered SL1 **Events** page. If you need to access the Anomaly Detection page, add the following directories to your SL1 URL: **/skylar-ai/anomaly-detection**. For example: <https://12.0.0.79/skylar-ai/anomaly-detection>.

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 SL1.

For more information, see [Enabling Thresholds and Alerts for the Anomaly Chart](#).





The second graph displays the following data:

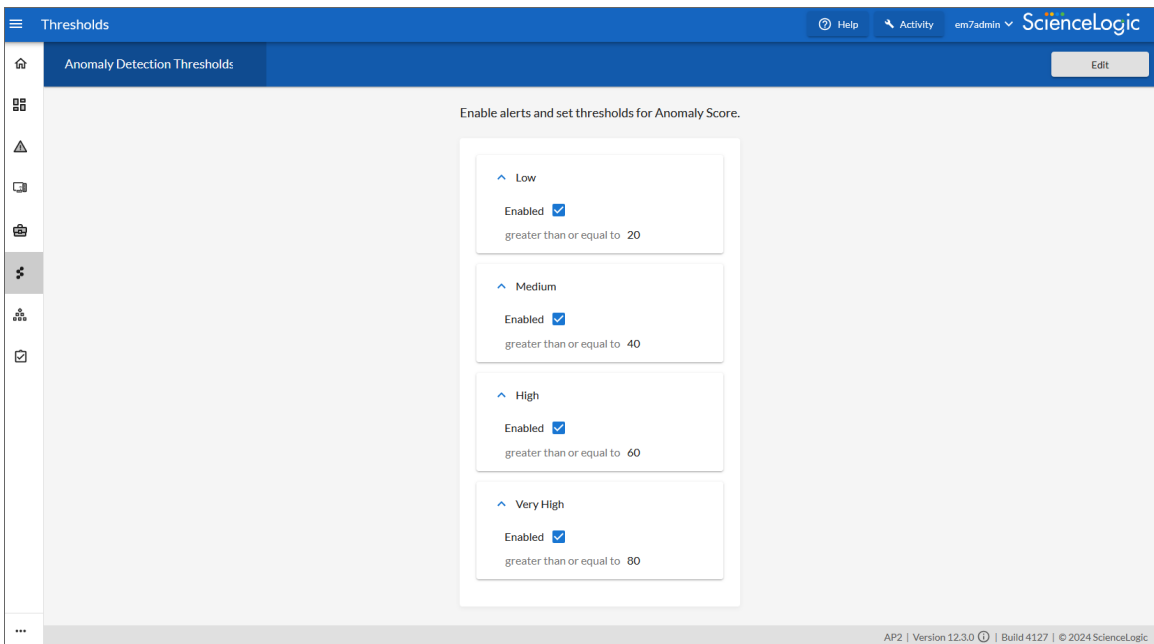
- A blue band representing the range of probable values that SL1 expected for the device metric.
- A green line representing the actual value for the device metric.
- A red dot indicating anomalies where the actual value appears outside of the expected value range. The number of the red dots are listed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

While anomaly detection is enabled automatically as soon as you [enable Skylar Analytics for one or more SL1 organizations](#), you can also set up anomaly detection events for specific Dynamic Application metrics on a device. When this is configured, an event policy is triggered when an anomaly is detected for that metric. Anomaly detection events display with an **Event Source** of *Skylar AI* on the **Events** page in SL1.


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 SL1 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**. 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*.

---

## Creating an Event Policy for Anomalies

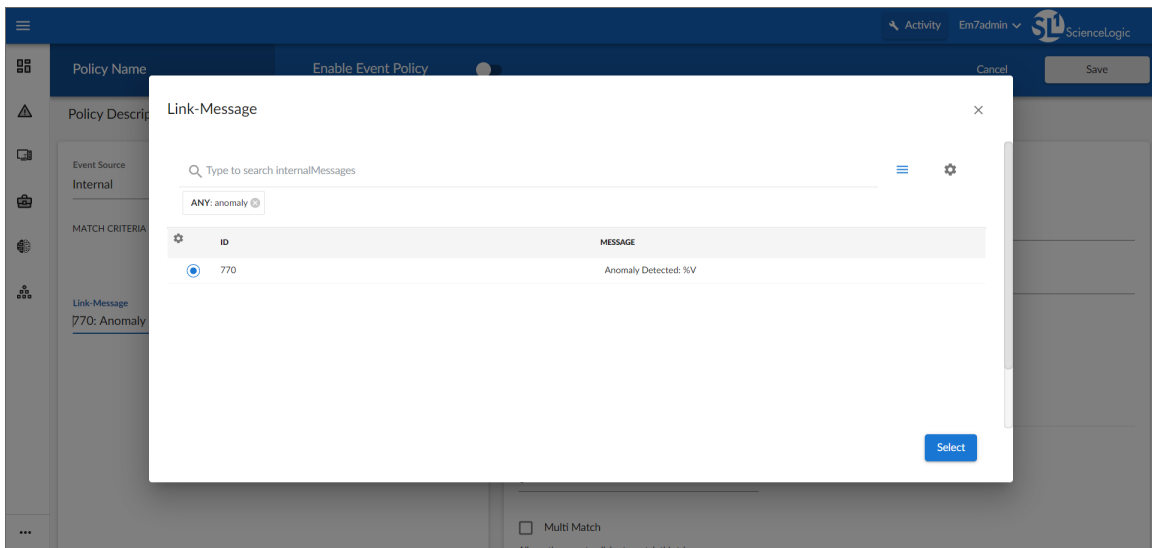
You can create additional event policies that will trigger events in SL1 when anomalies are detected for those devices.

**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:  Policy Type: [Active Events] Policy State: [Enabled] Policy Priority: [Default] Organization: [Robert and Sons]

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]

[and all times are valid]

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

☐ Trigger on Child Rollup

Available Devices:

- Linux: CentOS: 8.2.0-107.el8
- Linux: Oracle Linux Server release 8: smb\_svm6
- 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 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]**.

---

# 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 appear as enriched events in SL1, 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, SSH), network interfaces (utilization, errors, discards), and memory.

This chapter covers the following topics:

<i>What is Predictive Alerting?</i> .....	53
<i>Viewing Predictive Alerts in SL1</i> .....	53
<i>Using Predictive Alerts to Trigger Automated Run Book Actions</i> .....	56

---

## 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.

Skylar Analytics will start generating predictive alerts about 48 hours after data starts getting exported from SL1 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, and in the case of memory or network interfaces, Skylar AI looks at the daily **p95** value: the 95th percentile value, where 95 percent of the data in the past 30 days is lower than the p95 value and five percent of the data is higher than this 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.

Additionally, Skylar AI looks at the 99th percentile of daily differences and the **interquartile range** (also known as IQR, which is the spread of the data based on the difference between the 75th and 25th percentiles of the data) of daily differences. If today's difference exceeds the value of  $\text{<99th percentile>} + 3 * \text{IQR}$ , Skylar AI assumes there is a breakout and uses this new value to calculate a new slope to predict when a threshold will be reached.

Skylar AI also uses a number of other heuristics to prevent false positives, depending on the metric type. For example:

- at 70%, increasing consistently by 1% per day; will predict 100% in 30 days.
- at 50%, increasing consistently at 1% per day, but now increasing at 5% per day; will predict 100% in 10 days.

Additionally, for network interfaces, Skylar AI does not generate a predictive alert until at least one error or discard has been seen. The purpose of this is to weed out noise where no problems are likely to occur soon. In other words, Skylar AI is noticing some transients indicating network congestion; looking backward, does this seem to be the result of a recent trend of greater use?

---

## Viewing Predictive Alerts in SL1

When your SL1 system is connected to Skylar AI, you can start viewing predictive alerts in SL1. The alerts appear as enriched events in SL1, and they are generated in advance of the problem and can provide days, weeks, or months of notice depending upon the conditions. No additional configuration is needed.



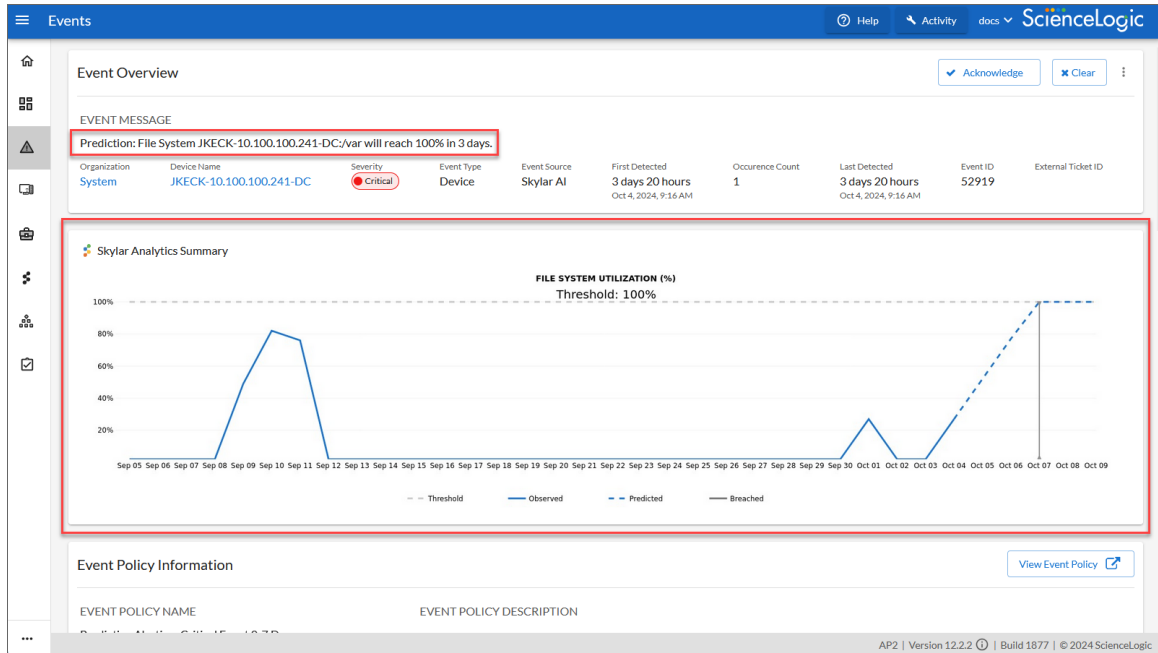
Predictive alerts display the Skylar icon (🤖) to the left of the event message in the **Message** column of the **Events** page. The text of the message contains the word "Prediction":

Events													
Total Events: 71 Critical: 23 Major: 16 Minor: 15 Notice: 16 Healthy: 1 View All													
Type to search events													
Refresh: 5 Minutes													
	Organiz...	Severity	Name	Message	Last Det...	Age	Ticket ID	Count	Event Ty...	Event N...	Masked Events	Event So...	Acknowledge
<input type="checkbox"/>	Filter	Filter	Filter	Filter	Filter				Filter			F...	
<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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="Clear"/>
<input type="checkbox"/>	System	Critical	linux-web02	Prediction: CPU Utilization will reach 100% in 4 d	Sep 30, 2024,	7 days 13 ho	—	1	Device			Skylar AI	<input type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 type="checkbox"/> Acknowledge <input type="button" value="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 SL1, 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.
2. On the **Events** page, click the message for a predictive alert with the Skylar icon (🔗). The **Event Investigator** page for that alert appears.
3. 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.

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, 9:20 PM	AIEngine	93091	Notice	—	Prediction: CPU Utilization will reach 100% in 28 days.
Nov 7, 2024, 7:11 PM	AIEngine	94606	Critical	—	Prediction: File System mrktng-dc2:/var/log will reach 100% in 0 days.
Nov 4, 2024, 9:23 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.

## Using Predictive Alerts to Trigger Automated Run Book Actions

After Skylar AI creates an SL1 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

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

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]

[and all times are valid]

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

☐ Trigger on Child Rollup

Available Devices:

- Linux: CentOS: 8.2.0-107.el8
- Linux: Oracle Linux Server release 8: smb\_svm6
- Linux: Red Hat Enterprise Linux 8: 8.1-1.el8

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]**.

---

# Appendix

# A

## 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> .....	59
<i>Provisioning a New Account</i> .....	59

---

## 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 the **Skylar Settings** page, you will see an **Analytics** link on this home page.

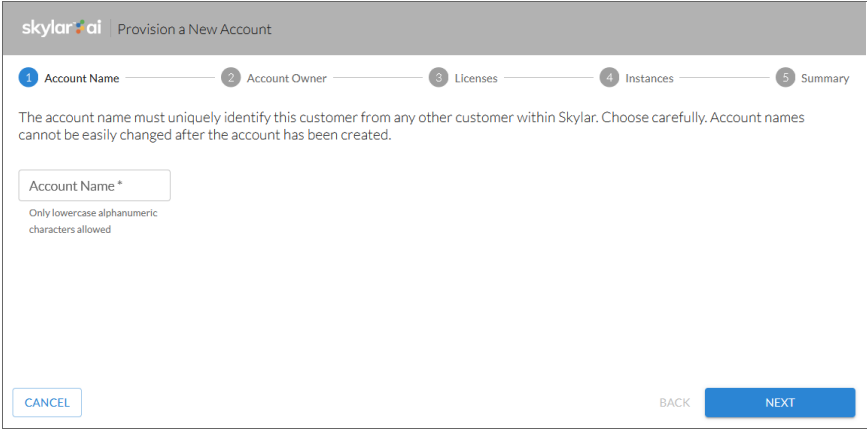
---

## 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. On the **Skylar Settings** page, create a new account by clicking the **All Accounts** drop-down at the top of the **Skylar Settings** page and clicking **[Provision New Account]** at the bottom of the drop-down. 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.

- On the **Licenses** page, select **Skylar Analytics** to enable Skylar Analytics for this account.
- If you select **Enable ODBC**, 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](#).

- 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]**.
- 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.

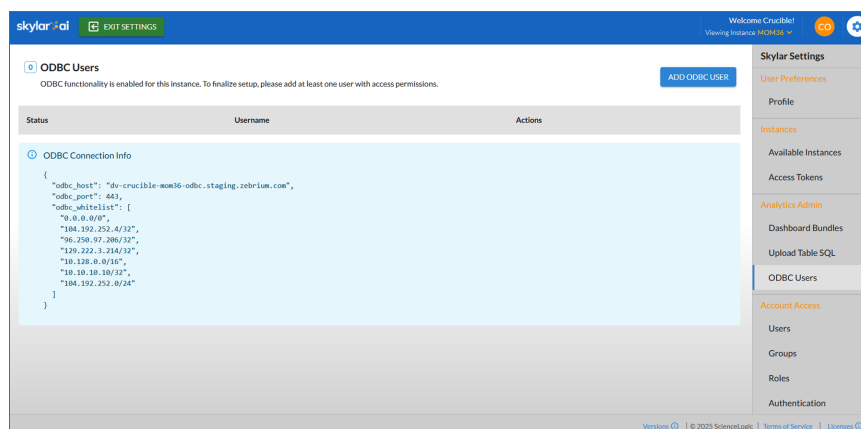
**NOTE:** When the account is set up, you will need to give the email address you used in step 3 to the first user. On first login, the new user will be prompted to change their password.

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

## 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.

- On the **Skylar Settings** page, click the **[ODBC Users]** tab. The tab displays the ODBC connection information for the Skylar AI system:



- 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]** tab.



© 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