



Skylar Analytics

Version 1.1.0

Table of Contents

Introduction to Skylar Analytics	1
What is Skylar AI?	3
Features of Skylar AI	3
Components of Skylar AI	3
Data Analyzed by Skylar AI	4
What is Skylar Analytics?	5
Getting Started with Skylar Analytics	6
Running the Skylar SL1 Management Script	6
Enabling Skylar Analytics for One or More SL1 Organizations	6
Skylar Analytics: Data Visualization and Data Exploration	8
What are Data Visualization and Data Exploration?	9
Data Visualization	9
Data Exploration	10
Viewing Dashboards and Charts in Skylar Analytics	10
Viewing and Customizing Skylar Analytics Dashboards	11
Default Skylar Analytics Dashboards	12
Working with Skylar Analytics Dashboards	13
Viewing and Customizing Skylar Analytics Charts	15
Working with Skylar Analytics Charts	16
Additional Tips for Creating and Customizing Charts	16
Viewing Skylar Analytics Datasets	17
Data Exploration: Exporting Data from Skylar AI	17
Additional Resources for Skylar Analytics (Apache Superset Training)	17
Skylar Analytics: Anomaly Detection	18
What is Anomaly Detection?	19
Enabling Anomaly Detection Events for Specific Metrics	19
Enabling Anomaly Detection Events for a Metric on the Device Investigator Page	19
Enabling Anomaly Detection Events for a Metric on the Service Investigator Page	20
Viewing Graphs and Data for Anomaly Detection	21
Creating an Event Policy for Anomalies	23
Using Anomaly-related Events to Trigger Automated Run Book Actions	25

Skylar Analytics: Predictive Alerting	27
What is Predictive Alerting?	28
Viewing Predictive Alerts in SL1	28

Chapter

1

Introduction to Skylar Analytics

Overview

Skylar Analytics contains a set of tools that lets you view, analyze, and use the data that SL1 gathers and sends to the Skylar AI engine. Skylar Analytics insights are presented in the SL1 user interface, in a ScienceLogic-hosted instance of Apache Superset, and in the Skylar AI API.

Skylar Analytics includes the following components:

- **Data Visualization.** Displays dashboards and charts based on data gathered by Skylar AI and SL1.
- **Data Exploration.** Enables third-party tools that use the Microsoft Open Database Connectivity (ODBC) interface to access the metric data from Skylar AI.
- **Anomaly Detection.** Uses always-on, unsupervised machine learning to identify unusual patterns that do not conform to expected behavior.
- **Predictive Alerting.** Generates events in SL1 that forecast when a future event could happen, instead of reporting on an event that has already occurred.

IMPORTANT: Skylar Analytics requires SL1 12.3.0 or later.

The following 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:

What is Skylar AI?	3
What is Skylar Analytics?	5
Getting Started with Skylar Analytics	6

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 soon Skylar Advisor, helps customers with their journey towards Autonomic IT.

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

SL1 collects data and leverages Skylar AI to learn the patterns for a particular device metric over a period of time. Skylar uses the resulting data to build a device metric-specific model that is used to define a scope of expected behavior as well as anomalous data points.

Features of Skylar AI

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

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

Skylar AI integrates seamlessly with the SL1 platform and other IT management tools. You can interact with Skylar AI through these familiar environments, where it enhances existing workflows with AI-driven insights and automation capabilities. Skylar AI can send you alerts and notifications, which can be customized to suit individual preferences or organizational needs. These alerts help you stay informed about potential issues, ongoing incidents, or opportunities for optimization.

Components of Skylar AI

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

- **Skylar Automated Root Cause Analysis (RCA)**, a log-based, root cause identification and analysis service powered by unsupervised AI.
- **Skylar Analytics**, an advanced reporting and custom analytics service that combines AI-powered analytics with deep data exploration and visualization.
- **Skylar Advisor**, a proactive IT problem-solving advisory service powered by human-centered AI.

Data Analyzed by Skylar AI

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




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

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

What is Skylar Analytics?

The Skylar Analytics suite of services uses data gathered by SL1 to explore data, generate visualizations, and monitor IT infrastructure metrics. Skylar Analytics can also use Skylar AI to predict alerts before they happen, and detect anomalies that could become events that might disrupt your IT infrastructure and functionality.



Skylar Analytics includes the following components:

- **Data Visualization** contains dashboards and charts based on data gathered by Skylar AI and SL1. Currently, this data includes server-focused metrics and basic network interface metrics, with more metrics planned for future Skylar updates. Please note that the dashboards in Business Intelligence (BI) tools are independent of SL1 dashboards or reports. Data Visualization is achieved using a ScienceLogic-hosted instance of Apache Superset or with your own third party tool.
- **Data Exploration** enables third-party tools that use the Microsoft Open Database Connectivity (ODBC) interface to access the metric data from Skylar AI. This component lets you use ODBC to connect Skylar AI data with applications like Tableau, Microsoft Power BI, or other business intelligence tools. For Skylar Beta, this feature is not yet available.
- **Anomaly Detection** 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 unexpected behavior that might require further investigation.
- **Predictive Alerting** generates events in SL1 that forecast when a future event could happen, instead of reporting on an event that has already occurred. SL1 will display capacity predictions and even automate run book automations based on anticipation of out-of-capacity events before they become a production issue. Predictive alerts are based on real-time data analyzed by Skylar AI against expected device metrics.

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

Getting Started with Skylar Analytics

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

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

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

Running the Skylar SL1 Management Script

The Skylar SL1 Management Script lets you set up your SL1 connectors and SL1 services for exporting data to Skylar. The script is named `sl-otelcol-mgmt.py`, and it is included with Skylar Analytics in the `sl-otelcol` package.

To run the Skylar SL1 Management Script:

1. 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-metrics --skylar-config -  
-skylar-endpoint "https://skylar.com" --skylar-api-key "<Skylar-API-  
Key>" --ap2-feature-flags
```

where `<Skylar-API-Key>` is the API key for Skylar AI. Ask your ScienceLogic contact for this value.

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

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

Enabling Skylar Analytics for One or More SL1 Organizations

In SL1, if you want to use Anomaly Detection and Predictive Alerting, you will need to select one or more organizations that will share data with Skylar AI. This data will come from all of the devices in a selected organization. By default, the Skylar AI features are disabled.

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

To enable Anomaly Detection and Predictive Alerting:

1. In SL1, go to the **Organizations** page (Registry > Accounts > Organizations) and click the check box for one or more organizations.
2. In the **Select Action** drop-down, select *Send Data from Selected Orgs to Skylar AI* and click **[Go]** to start sending data about the selected organizations to Skylar AI. The **Skylar AI Status** column for the selected organizations changes to *Enabled*.

Chapter

2

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. You can use the Data Exploration component to drill down into specific data points, analyze IT infrastructure metrics, and generate new visualizations. Currently, this data includes server-focused metrics and basic network interface metrics, with more metrics planned for future Skylar updates.

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

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

This chapter will provide a general overview of how to view the charts, graphs, and other reports in the Skylar Analytics user interface, along with tips and best practices for users of SL1 and Skylar AI.

This chapter covers the following topics:

<i>What are Data Visualization and Data Exploration?</i>	9
<i>Viewing Dashboards and Charts in Skylar Analytics</i>	10
<i>Data Exploration: Exporting Data from Skylar AI</i>	17
<i>Additional Resources for Skylar Analytics (Apache Superset Training)</i>	17

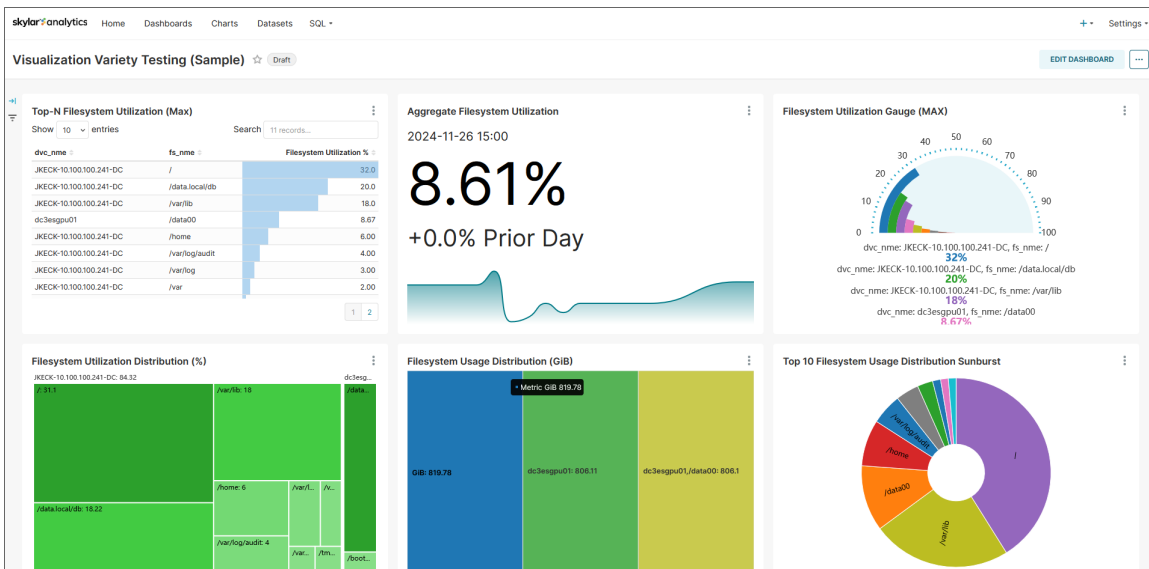
What are Data Visualization and Data Exploration?

Before the initial release of Skylar Analytics, SL1 stored data in a proprietary format that was not easily exported to other third-party applications for further research and insight. Skylar Analytics takes the data gathered by SL1 and Skylar AI, normalizes it, and makes it available in standard ODBC database format.

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

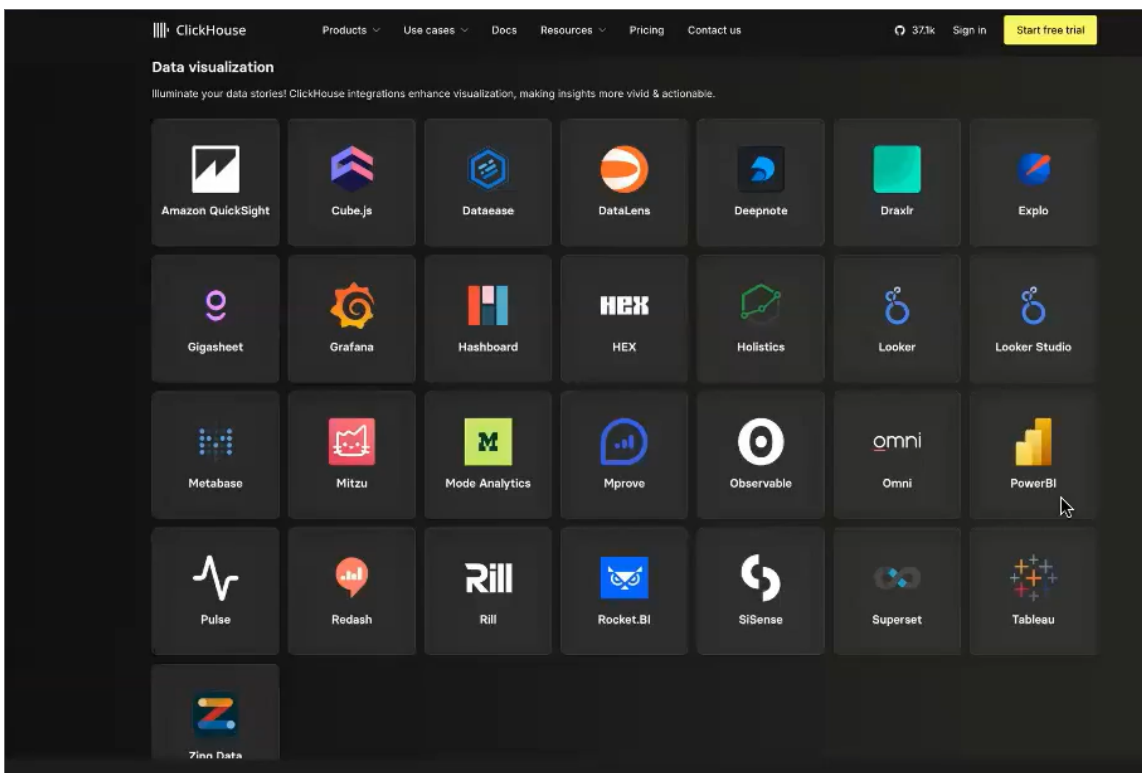
Data Visualization

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



Data Exploration

The **Data Exploration** component of Skylar Analytics lets you export this data to a third-party tool used by your company. This component lets you use ODBC to connect Skylar AI data with Grafana, Microsoft BI, and other business intelligence tools.



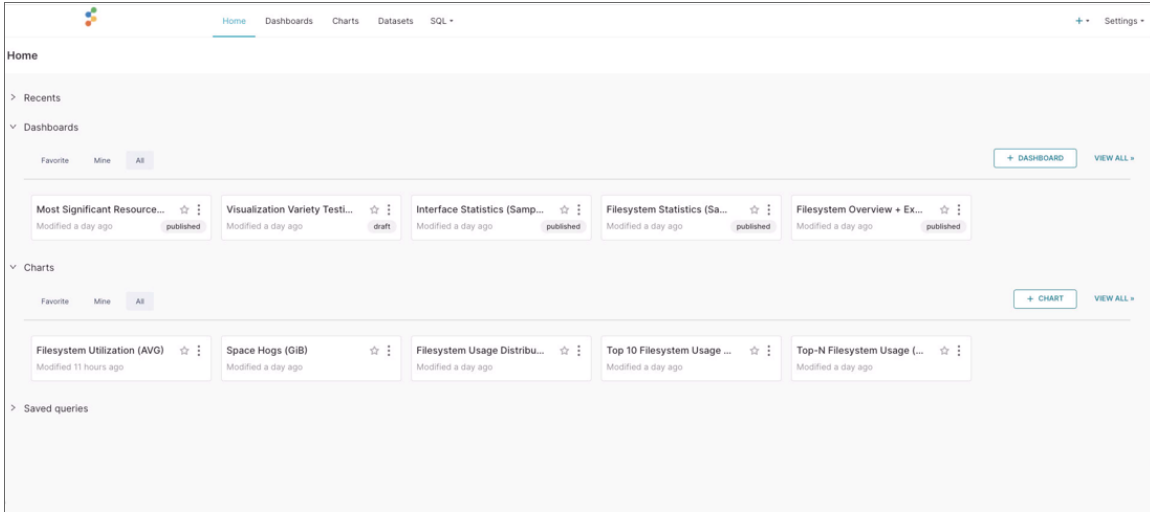
NOTE: Because ScienceLogic does not own the underlying framework for the Data Visualization and Data Exploration components, ScienceLogic is not responsible for maintaining or updating documentation for third-party open-source software, including Apache Superset. For the most current and accurate information,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

Viewing Dashboards and Charts in Skylar Analytics

The Data Visualization component of Skylar Analytics contains dashboards and charts based on data gathered by SL1. The visualizations currently contain server-focused metrics and basic network interface metrics, with more metrics planned for future Skylar Analytics updates.

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

When you log into the Data Visualization component of Skylar Analytics, the **Home** page appears:



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

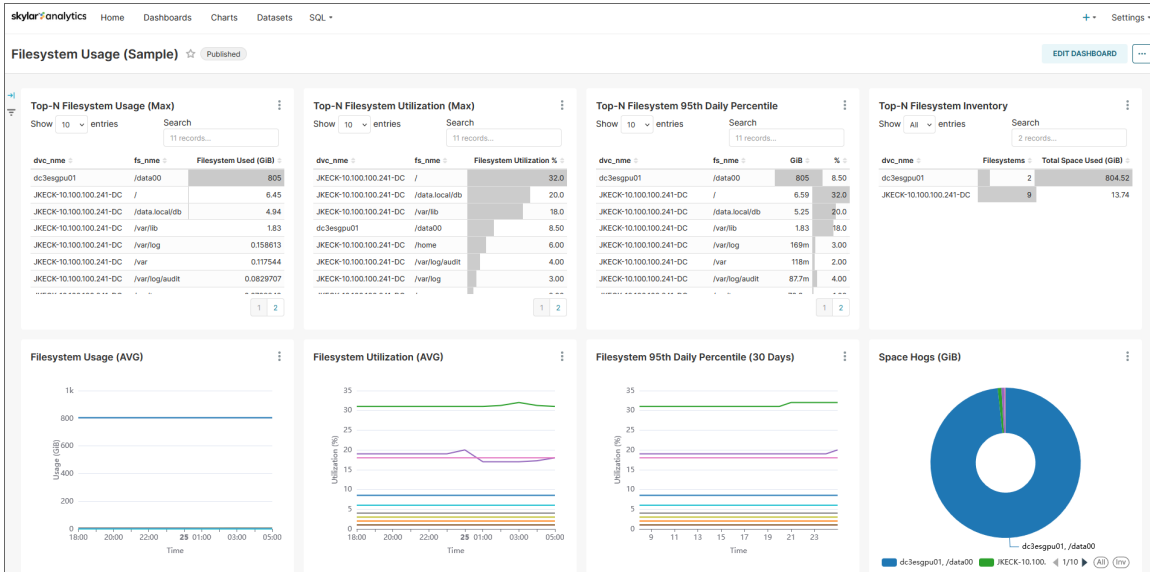
For Skylar Analytics, you will mainly use the following tabs to view SL1 and Skylar AI data visualizations:

- [Dashboards](#)
- [Charts](#)
- [Datasets](#) (for administrators)

Viewing and Customizing Skylar Analytics Dashboards

A **dashboard** in Skylar Analytics is similar to a dashboard in SL1, in that they both contain a number of graphical "widgets" that display data in a variety of ways, such as pie charts, line graphs, maps, bar charts, and other visualizations. In Skylar Analytics, a widget is called a "chart".

NOTE: Unlike dashboards in SL1, a dashboard in Skylar Analytics is used only for laying out the various charts that make up that dashboard. You can use charts to customize the data. One significant difference is that a chart, when modified, impacts all dashboards using that chart definition. Charts can be duplicated to be modified for different analyses on different dashboards.



Default Skylar Analytics Dashboards

Skylar Analytics contains the following default dashboards:

- **Filesystem Overview + Exploration (Sample).**
 - Displays 95th percentile data, filesystem utilization distribution (as a percentage and Gigibit or GiB), and "Space Hogs" (the devices using the most filesystem 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 filesystem 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 filesystem space), filesystem utilization as a percentage, filesystem inventory by host, and filesystem usage distribution.
- **Filesystem Usage (Sample).**
 - Displays a set of filesystem usage, utilization, 95th percentile and Top-N inventory charts for all devices, including a pie chart of "Space Hogs" (the devices using the most filesystem 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 filesystem usage, along with average filesystem usage, Top-N interface usage delta, and interface traffic in the past seven days.
 - You can click a device name on the "Top-N Filesystem Usage" or the "Top-N Interface Usage" tables to display chart details specifically for that device.
- **Visualization Variety Testing (Sample).**
 - Displays a variety of chart visualizations related to filesystem utilization, including a table, a "big number" with a line graph, a gauge, a set of tree maps, and a sunburst map.
 - This table is not meant to be informational so much as an example of the types of visualizations you can use with Skylar Analytics


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

Working with Skylar Analytics Dashboards

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

- For most dashboards, you can click a single device or item in the first chart at the top left of the Dashboard page to view data specific to just that device. Click the device a second time to clear the filter.
- Hover over a graphical element in a chart, such as a piece of a pie chart or a colored metric in a tree map to view a pop-up with more information about that element.
- Click **[Edit Dashboard]** to make changes to the dashboard and the charts that comprise the dashboard. For more information, see <https://docs.preset.io/docs/creating-a-dashboard>.

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



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

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


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

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

- *View as table*. Displays the chart in table format.
- *Drill to detail*. Displays all the data that makes up a chart. For more information, see <https://docs.preset.io/docs/drilling-to-chart-details>.
- *Share*. Lets you copy a shareable chart link to your system's clipboard, or launches your system's default email client and composes a new message featuring the chart URL.
- *Download*. Lets you export the chart to .CSV or Excel, or you can download the chart as an image.

To customize a dashboard:

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



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

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

Viewing and Customizing Skylar Analytics Charts

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

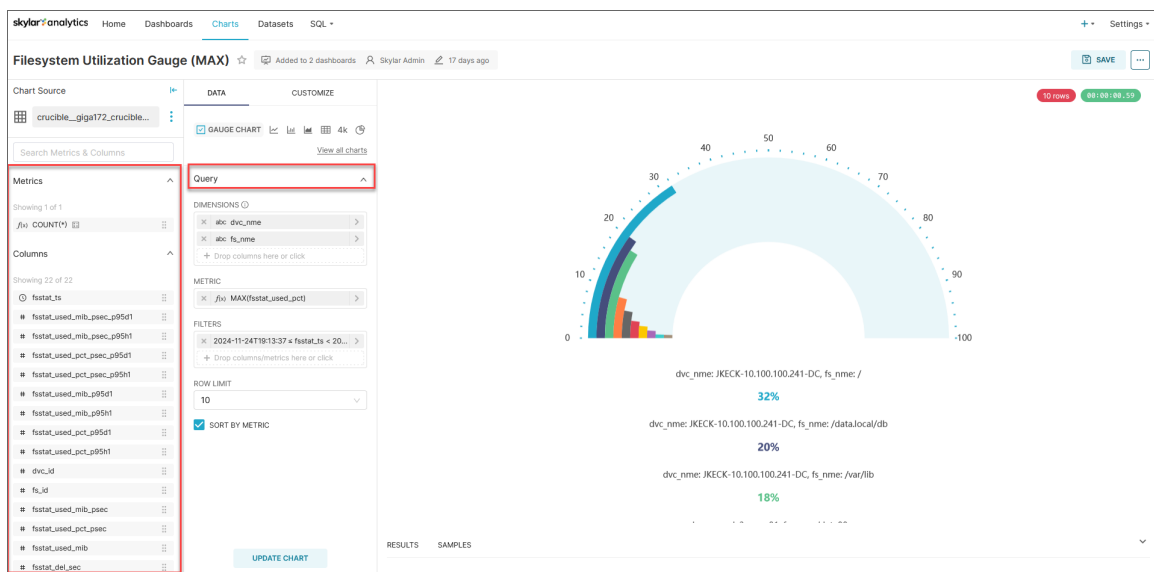
TIP: On the **Dashboards** tab, the "Visualization Variety Testing (Sample)" dashboard contains a variety of chart visualizations related to filesystem utilization, including a table, a "big number" with a line graph, a gauge, a set of tree maps, and a sunburst map. You can use this dashboard to see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

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

Working with Skylar Analytics Charts

To create or customize a chart:

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



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

Additional Tips for Creating and Customizing Charts

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

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

For more information, see <https://docs.preset.io/v1/docs/using-preset-explore>.

Viewing Skylar Analytics Datasets

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

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

Data Exploration: Exporting Data from Skylar AI

You can use the optional Data Exploration component of Skylar Analytics to enable Microsoft Open Database Connectivity (ODBC) to export Skylar AI data with third-party tools like Grafana, Power BI, Tableau, Cognos, Crystal Reports, SAP, Excel, and other Business Intelligence applications.

For Skylar Beta, this feature is not yet available.

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 UI.

The following links go to Apache Superset-related documentation:

- <https://superset.apache.org/docs/intro>
- <https://docs.preset.io/docs/getting-started>

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

- <https://www.udemy.com/course/apache-superset-for-data-engineers-hands-on/>
- https://www.youtube.com/watch?v=znmco3eK-M&list=PLzRV_ObEwmNhRjhMNcvcDP7ZDjOXtodd
- <https://superset.apache.org/community>

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.

You can view a list of all devices that have metrics being monitored for anomalies on the corresponding **Device Investigator** or **Service Investigator** pages.

NOTE: Unlike the Data Visualization and Exploration and Predictive Alerting components of Skylar Analytics, this release of Anomaly Detection with Skylar Analytics works on all Dynamic Applications in SL1. Data Visualization and Exploration as well as Predictive Alerting currently monitor server and network metrics, with more metrics planned for future SL1 releases.

This chapter covers the following topics:

What is Anomaly Detection?	19
Enabling Anomaly Detection Events for Specific Metrics	19
Viewing Graphs and Data for Anomaly Detection	21
Creating an Event Policy for Anomalies	23
Using Anomaly-related Events to Trigger Automated Run Book Actions	25

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.

NOTE: Unlike the Data Visualization and Exploration and Predictive Alerting components of Skylar Analytics, this release of Anomaly Detection with Skylar Analytics works on all Dynamic Applications in SL1. Data Visualization and Exploration as well as Predictive Alerting currently monitor server and network interface metrics, with more metrics planned for future SL1 releases.

Anomaly detection is calculated and displayed in the SL1 user interface for all Performance Dynamic Applications. This detection is enabled by default and cannot be disabled. You can control which device data gets sent to Skylar for analysis based on the organization aligned with the device or devices. All devices in the selected organization will get anomaly detection analysis.


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

Enabling Anomaly Detection Events for Specific Metrics


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

Enabling Anomaly Detection Events for a Metric on the Device Investigator Page

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

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

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

2. On the **[Anomaly Detection]** tab, click the **Actions** icon () for any of the listed metrics and select **Enable**. The **Select Available Metrics** modal appears.

3. In the **Select Metric** drop-down, use the **Search** field to search for a specific metric or click one of the category names, such as "Dynamic Apps" or "Collection Labels", to view a list of available metrics for that metric category.
4. Click the name of the metric on which you want to enable anomaly detection events for the device.
5. For some metrics, a second drop-down field might display that enables you to specify the device directory. If this field appears, click the name of the directory on which you want to enable anomaly detection.
6. Click **[Enable]**. That metric is enabled for events for that device.

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

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

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

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

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

1. On the **Business Services** page (📁), select a service from the list of business, IT, and device services by clicking its name. The **Service Investigator** displays.
2. On the **Service Investigator** page, click the **[Anomaly Detection]** tab.
3. Click the **Actions** icon (⋮) for any of the listed metrics and select *Enable*. The **Select Available Metrics** modal appears.
4. In the **Select Metric** drop-down, use the **Search** field to search for a specific metric or click one of the category names, such as "Dynamic Apps" or "Collection Labels", to view a list of available metrics for that metric category.
5. Click the name of the metric on which you want to enable anomaly detection events for the device.
6. 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 .
7. Click **Enable**.

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

Viewing Graphs and Data for Anomaly Detection

After SL1 begins performing anomaly detection for a device, you can view graphs and data about each anomaly. Graphs for anomalies appear on the following pages in SL1:

- The **[Anomaly Detection]** tab in the **Device Investigator**.
- The **Anomalies** tab in the **Service Investigator** for a business, IT, or device service.

You can view the anomaly detection graphs for the metrics by clicking the **Open** icon (↗) next to the metric for the device. The **Anomaly Chart** modal appears, displaying the "Anomaly Score" chart above the chart for the specified metric you are monitoring.

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



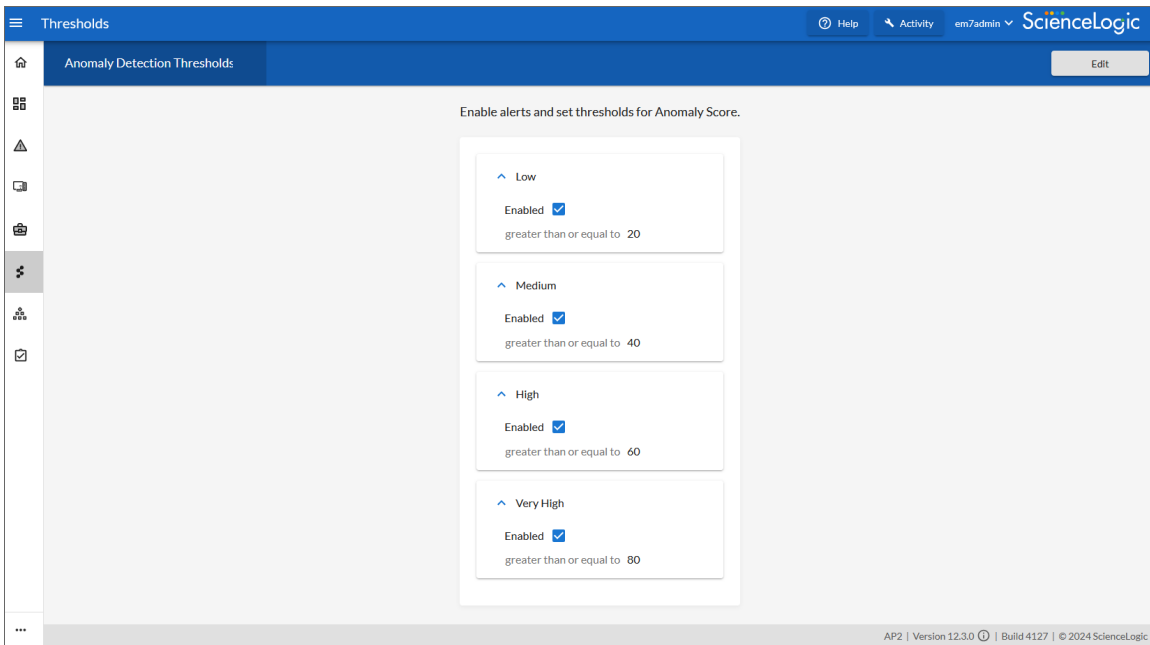
The second graph displays the following data:

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


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

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

You can define the thresholds for the "Anomaly Score" chart on the **Anomaly Chart** modal, and whether those values generate alerts, on the **Anomaly Detection Thresholds** page.



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



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

To edit the Anomaly Score thresholds:

1. On the **Anomaly Detection Thresholds** page, click **[Edit]**.
2. For each of the four severity levels, from Low to Very High, you can select **Enabled** to have SL1 generate an alert when the Anomaly value for a metric is equal to or greater than the threshold for that severity level.
3. You can edit the threshold value for each level if SL1 is generating too many (or not enough) anomalies of a certain severity level.
4. For example, if you want to enable a Low level alert when the Anomaly Score value is between 25 and 39, you would go to the **Low** panel, select **Enabled**, and update the value from "20" to "25".
5. Click **[Save]**.
6. You can then edit an event policy that uses alerts based on the settings on this page to generate events in SL1. For more information, see [Creating an Event Policy for Anomalies](#).

Creating an Event Policy for Anomalies


After you have enabled anomaly detection for devices, 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).
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:

The screenshot shows the 'Automation Policy Editor' interface for creating a new automation policy. The browser address bar indicates the URL: `https://10.128.88.92/em7/index.em7?exec=registry_policies_automation_editor&height=700&wid=`. The page title is 'Automation Policy Editor | Creating New Automation Policy'. A 'Reset' button is located in the top right corner.

The configuration fields are as follows:

- Policy Name:** Anomaly High
- Policy Type:** [Active Events]
- Policy State:** [Enabled] (highlighted with a red box)
- Policy Priority:** [Default]
- Organization:** Sample

The logic configuration section includes:

- Criteria Logic:** [Severity >=] [Minor,] [and 5 minutes has elapsed] [since the first occurrence,] [and event is NOT cleared] and all times are valid
- Match Logic:** [Text search]
- Match Syntax:** (empty field)
- Repeat Time:** [Only once]
- Align With:** [Devices]
- Include events for entities other than devices (organizations, assets, etc.)
- Trigger on Child Rollup

The interface is divided into four main sections for device and event management:

- Available Devices:** Lists 'Sample' (AWS: Service: test, ScienceLogic, Inc.: EM7 Data Collector: mrktng-dc1, ScienceLogic, Inc.: EM7 Data Collector: mrktng-dc2) and 'System'.
- Aligned Devices:** Currently contains '(All devices)'.
- Available Events:** Lists 'anom' with three entries: '[1768] Critical: Anomaly Score Critical - new york', '[18] Minor: Anomaly Score Minor', and '[17] Notice: Anomaly Score Notice'.
- Aligned Events:** Contains two entries: '[20] Critical: Anomaly Score Critical' and '[19] Major: Anomaly Score Major'.

At the bottom, there are sections for 'Available Actions' (listing SNMP Trap, Snippet, and AWS actions) and 'Aligned Actions' (empty). A 'Save' button is located at the bottom center.

3. In the **Policy State** field, select *Enabled*.
4. In the **Available Events** field, search for and select an anomaly-related event policy, and then click the right-arrow icon to move it to the **Aligned Events** field. For more information about anomaly-related events, see [Creating an Event Policy for Anomalies](#).
5. Complete the remaining fields on the **Automation Policy Editor** page based on the specific parameters that you want to establish for the automation policy. For more information about the fields on the **Automation Policy Editor** page, see [Automation Policies](#).
6. When you are finished, click **[Save]**.

Skylar Analytics: Predictive Alerting

Overview

The Predictive Alerting component of Skylar Analytics generates events in SL1 that forecast when a future event could happen, instead of reporting on an event that has already occurred.

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

This chapter covers the following topics:

<i>What is Predictive Alerting?</i>	28
<i>Viewing Predictive Alerts in SL1</i>	28

What is Predictive Alerting?

A **predictive alert** is a warning about an event before it happens, based on analysis from Skylar AI.

The predictive alerting feature generates events in SL1 that forecast when a future event could happen, instead of reporting on an event that has already occurred.

Viewing Predictive Alerts in SL1

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


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

The screenshot shows the SL1 Events page with a table of events. The table has columns for Organization, Severity, Name, Message, Last Det., Age, Ticket ID, Count, Event Type, Event N., Masked Events, Event So., Acknowledge, and Clear. A red box highlights a group of events where the Message column contains predictive alerts. The first event in the highlighted group is: "Prediction: File System JKECK-10.100.100.241-D". Other events in the group include "Prediction: CPU Utilization will reach 100% in 5 d" and "Prediction: File System JKECK-10.100.100.241-D".

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

To view details about a predictive alert:

1. On the **Events** page, click the message for a predictive alert with the Skylar icon (🌟). The **Event Investigator** page for that alert appears.
2. On the **Event Investigator** page, the **Skylar Analytics Summary** panel displays a timeline of data from Skylar AI about a specific metric:



The dotted line on the graph on 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/usr/fig will reach 100% in 0 days.
Nov 4, 2024, 9:22 PM	AIEngine	94022	Major	-	Prediction: CPU Utilization will reach 100% in 11 days.
Nov 4, 2024, 7:35 PM	AIEngine	93939	Notice	-	Prediction: File System mrktng-dc2/ will reach 100% in 28 days.
Nov 3, 2024, 9:28 PM	AIEngine	93091	Notice	-	Prediction: CPU Utilization will reach 100% in 20 days.

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

© 2003 - 2024, 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. For more information, see <https://sciencelogic.com/company/legal>.

ScienceLogic

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

International: +1-703-354-1010