Skylar Analytics

Version 1.8.0

Table of Contents

Chapter 1

Introduction to Skylar Analytics

Overview

Skylar Analytics includes the following components:

- Data Visualization
- Data Exploration
- Anomaly Detection
- Predictive Alerting

For an overview of Skylar AI, including how to set up Skylar AI with Skylar One, see *Getting Started with Skylar AI*. To view the latest release notes, see the *Skylar AI Release Notes*.

IMPORTANT: While ScienceLogic recommends that you use the most recent version of Skylar One (formerly SL1), 12.3.2 is the minimum Skylar One version you can use with this release.

ScienceLogic recommends that you always use the most recent Skylar One 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 *Skylar One Platform and AP2 Release Notes*.

This chapter covers the following topics:

What is Skylar Analytics?

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

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

Skylar Analytics includes the following components:

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

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

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 Skylar One to enable the export of data from Skylar One to Skylar:

- *Create a Service Connection*
- *Enable Skylar Analytics for one or more organizations*

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

For information about setting up users, user groups, and user roles, see *Configuring Skylar AI System Settings*.

IMPORTANT: ScienceLogic strongly recommends that you always use the most recent Skylar One 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 *Skylar One Platform and AP2 Release Notes*.

Creating a Service Connection

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

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

To create a Skylar AI Engine service connection:

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

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

Enabling Skylar AI for One or More Organizations

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

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

To enable Skylar AI with Skylar One organizations:

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

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

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

For Older Versions: Running the Skylar Management Tool

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

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

To run the Skylar SL1 Management Tool:

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

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

   where:

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

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

   You can also use the following configuration options if needed:

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

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

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

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

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

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

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

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

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

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

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

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

You should look for the following messages in the output:

```
---------- checking feature toggles

SL_EXPORT_EVENTS = False

SL_EXPORT_METRICS = True

SL_EXPORT_CONFIG = True

---------- checking services

sl-otelcol is enabled and running

---------- checking connectivity

checking: Skylar endpoint is healthy

checking: local OTELCOL endpoint is healthy
```

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

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

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

Enabling Skylar AI Event Policies

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

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

To install the "Skylar Analytics Event Policies" PowerPack:

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

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

- *Skylar Analytics: Anomaly Detection*
- *Skylar Analytics: Predictive Alerting*

Mapping Skylar One Dynamic Application Object Names to Skylar Columns

When data from Skylar One Dynamic Applications is exported to Skylar AI, the names of collection and presentation objects are automatically converted into clean, standardized column names for the Skylar data lake.

The following rules ensure that all Skylar column names are consistent, machine-friendly, and easy to work with. If you are not sure how a name will be converted, use these common word replacements and clean-up rules as a guide.

The conversion process follows several steps:

1. **Standardize Special Characters**
    - If a letter is followed by a non-word character and an "a", replace it with the letter plus "A".
    - 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	Procssr
Processes	Procs

3. **Shorten Common Technical Terms**

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

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

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

4. **Clean Up the Name**

   - Remove all non-alphanumeric characters (like spaces, slashes, parentheses, etc.).
   - Replace common terms:
     - Average → Avg
     - QueueLength → QLen
     - sISl → SL
     - SL1Skylar → SL1Sky
     - Exporter → Exptr
     - Receiver → Rcvr

5. **Add Unit, if Applicable**

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

Chapter 2

Skylar Analytics: Data Visualization and Data Exploration

Overview

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

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

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

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

This chapter covers the following topics:

What is Data Visualization?

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

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

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

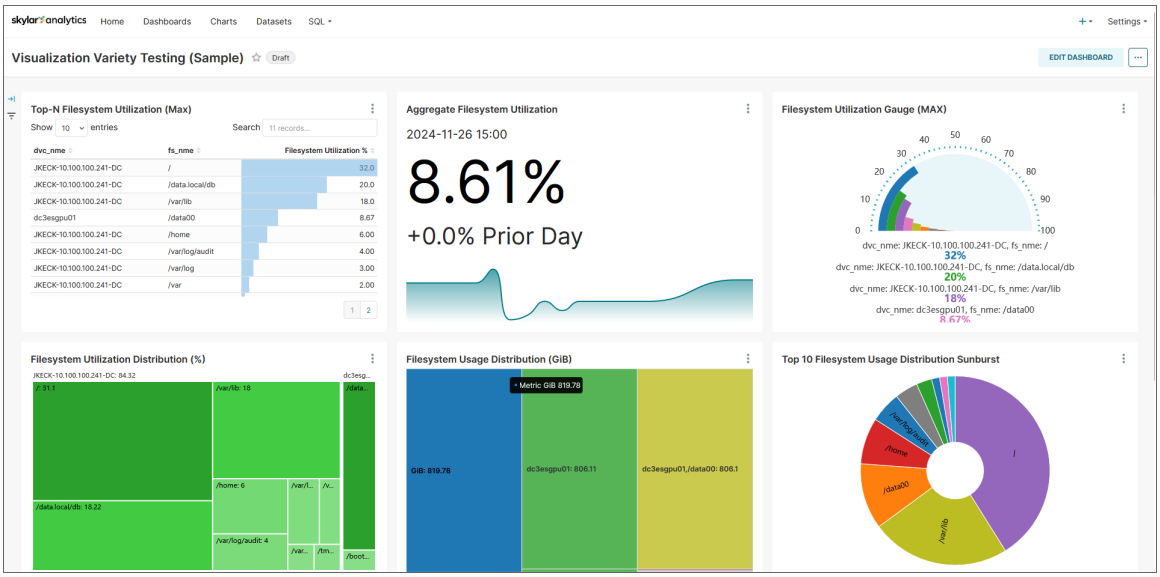

For more information, see *Viewing Dashboards and Charts in Data Visualization*.

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

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

Working with Datasets in Data Visualization

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

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

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

Components of a Dataset

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

- **One performance dataset**. In Skylar Analytics, these datasets use the naming convention of "Metric*<PowerPackName>*", such as "MetricMSWinSvr" for the "Microsoft: Windows Server" PowerPack.

- **Two configuration datasets**:
    - The "Current" dataset contains only the last recorded configuration change. You will typically use this dataset to quickly retrieve configuration details, unless you need to retrieve historical configuration changes.

      In Skylar Analytics, these datasets use the naming convention of "Conf*<PowerPackName>*Current", such as "ConfMSWinSvrCurrent" for the "Microsoft: Windows Server" PowerPack.

    - The other dataset contains configuration Dynamic Applications that have timestamps for each configuration snapshot taken by Skylar One. This dataset uses the naming convention of "Conf*<PowerPackName>*", such as "ConfMSWinSvr".

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

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

Viewing the List of Datasets

To view a list of the datasets in Skylar Analytics:

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

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

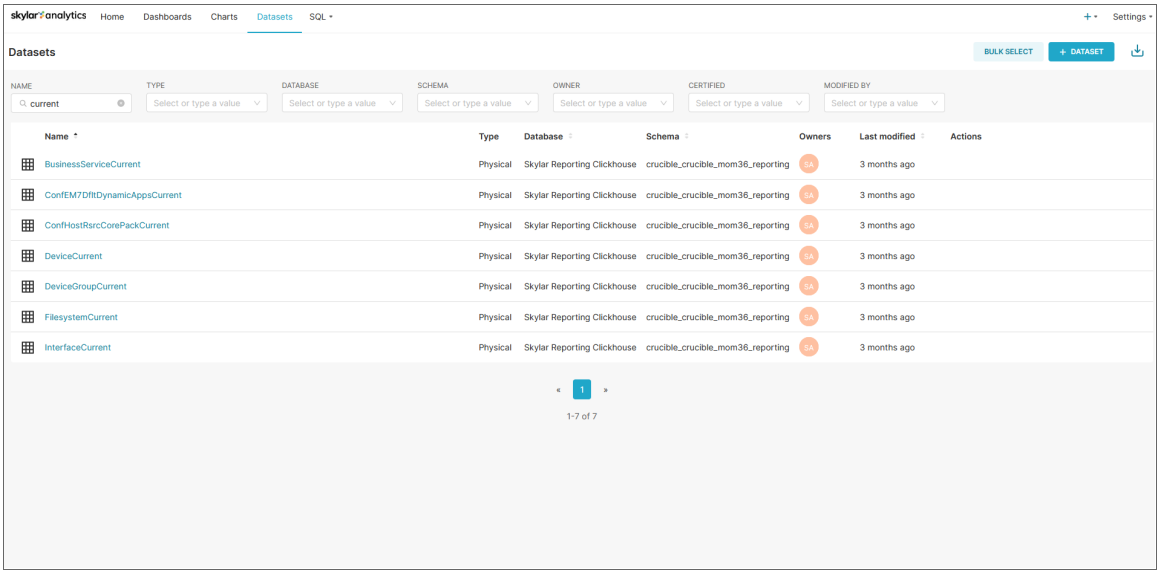

> **TIP:** You can filter the list of datasets by typing some or all of a dataset name in the ***Name*** field at top left.

3. Click the name of a dataset on the **Datasets** page to access the **Charts** page, where you can create a chart based on the columns (metrics) in that dataset. For more information, see *Creating and Customizing Dashboards and Charts*.

Viewing the Contents of a Dataset

To view the contents of a dataset:

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

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

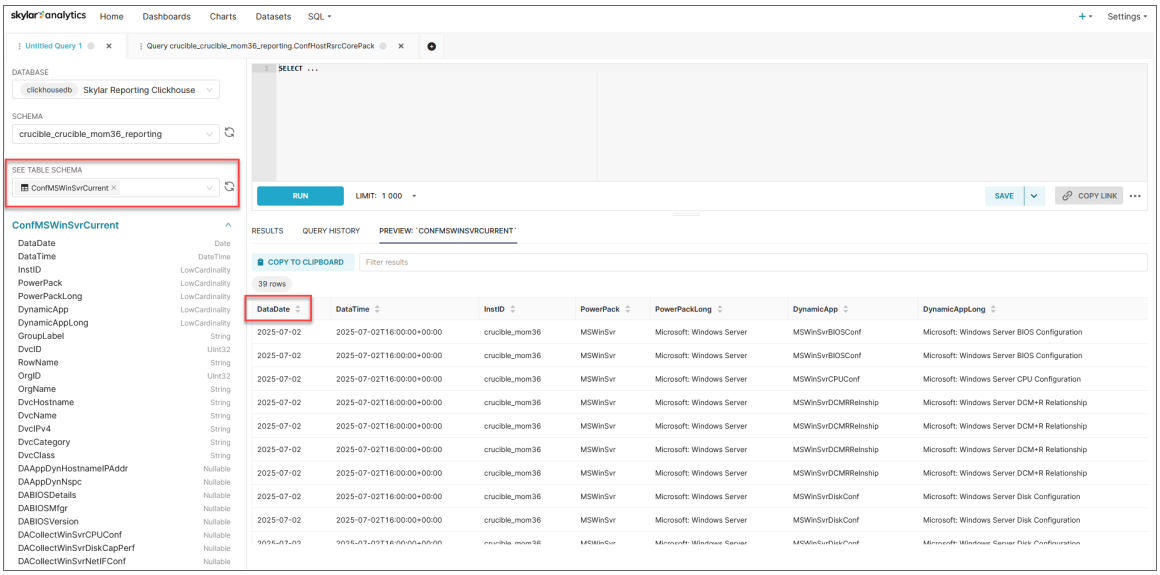

3. In the ***See Table Schema*** field, type the name of the dataset and press **[Enter]**. A list of metrics from that dataset are added below that field, and a preview of the table is displayed in a table to the right of that column. Each row in the preview table at the right reflects a set of data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

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

Viewing Dashboards and Charts in Data Visualization

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

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

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

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

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

Logging In to the Data Visualization Component

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

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

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:

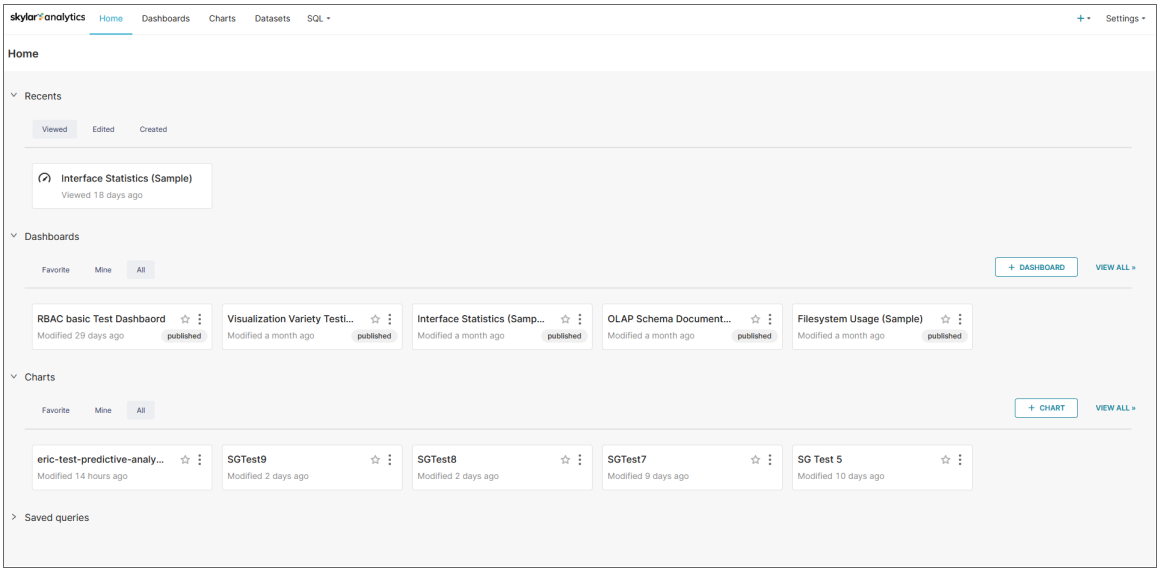

The **Home** page contains links to the dashboards and charts that you have used the most, including those that you have marked as favorites (⭐). You can create a dashboard or a chart from this page, and you can view all dashboards or charts by clicking the corresponding **View All** link.

3. Click a dashboard or chart from the **Home** page, or click the **Dashboards** page or the **Charts** page 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.

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

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

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

- **Filesystem Overview + Exploration (Sample)**.
    - Displays 95th percentile data, file system utilization distribution (as a percentage and 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)**.

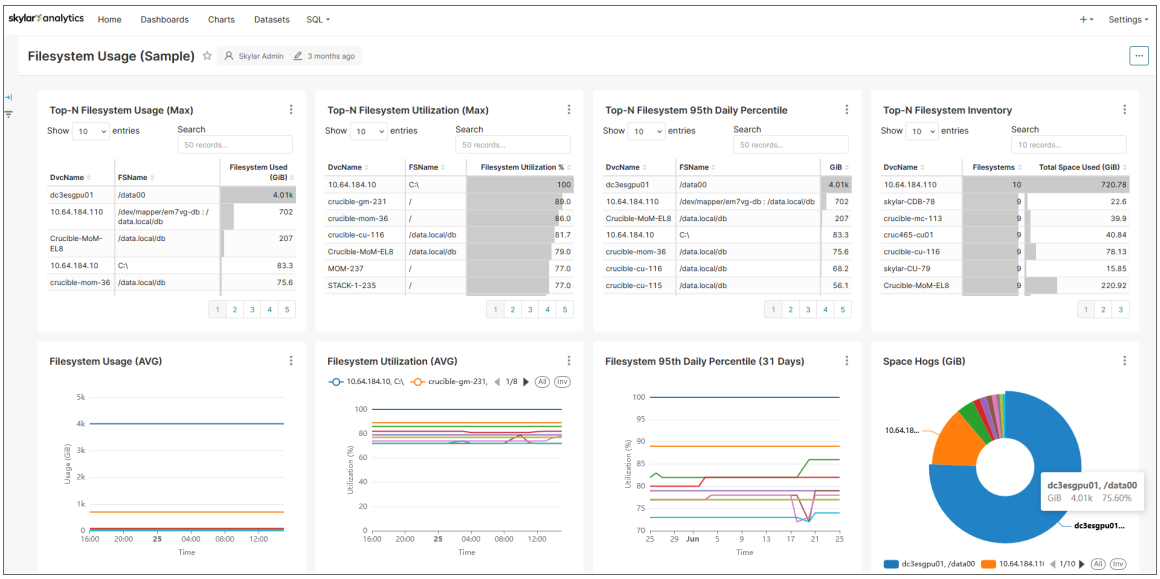

    - Displays a set of file system usage, utilization, 95th percentile and Top-N inventory charts for all devices, including a pie chart of "Space Hogs" (the devices using the most file system space).
    - You can click a device name on the "Space Hogs" pie chart to display chart details specifically for 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.
- **OLAP Schema Documentation (Skylar)**. Contains additional information about the structure of the charts in the Data Visualization component of Skylar Analytics. If you do not have this dashboard on the **Dashboards** page in Skylar Analytics, you can add it from the **Dashboards** page (Analytics Admin > Dashboards) in Skylar Settings.
- **Visualization Variety Testing (Sample)**. Contains a variety of chart visualizations related to file system utilization, including a table, a "big number" with a line graph, a gauge, a set of tree maps, and a sunburst map. You can use this dashboard to see how these different types of charts might work for your data.

Recommended Datasets

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

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

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

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

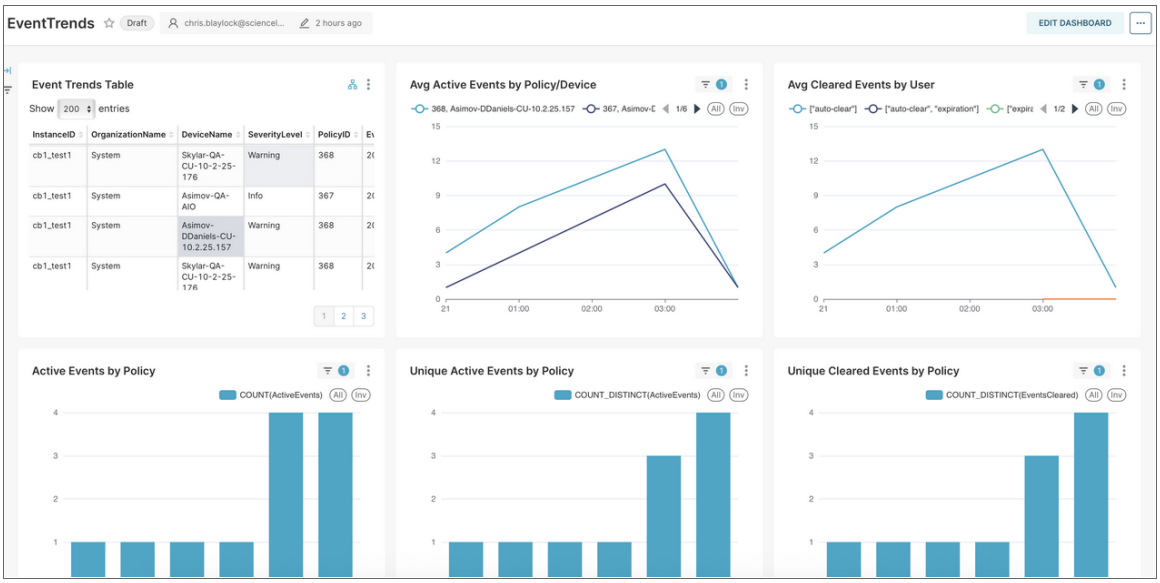

NOTE: These datasets contain raw data and do not have user-friendly column names.

These datasets include the following columns:

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

Viewing Skylar Analytics Dashboards

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

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

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

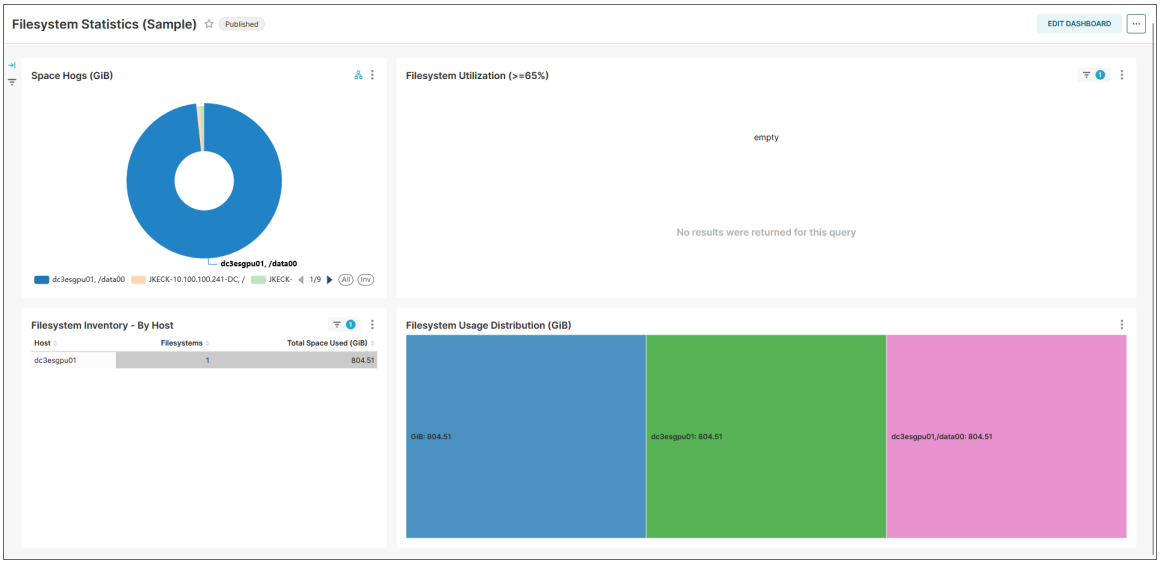

When viewing a dashboard, you can click the ellipsis button (⋯) at the top right of the **Dashboard** page to open a menu with the following dashboard options:

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

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

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

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

Creating and Customizing Dashboards and Charts

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

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

Creating a Dashboard

To create a dashboard:

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

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

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

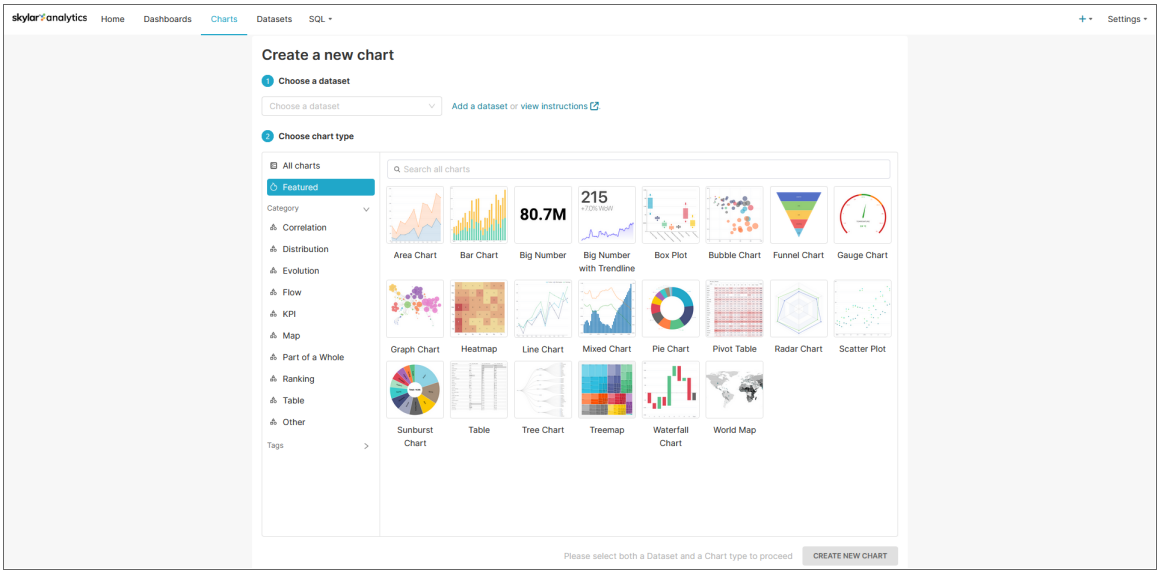

8. In the ***Choose a Dataset*** field, click to choose a dataset with the data you want to view in your new dashboard. In Skylar Analytics, a ***dataset*** contains a set of related metrics pulled from Dynamic Applications in Skylar One, such as server reports or Skylar One business service statistics.

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

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

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

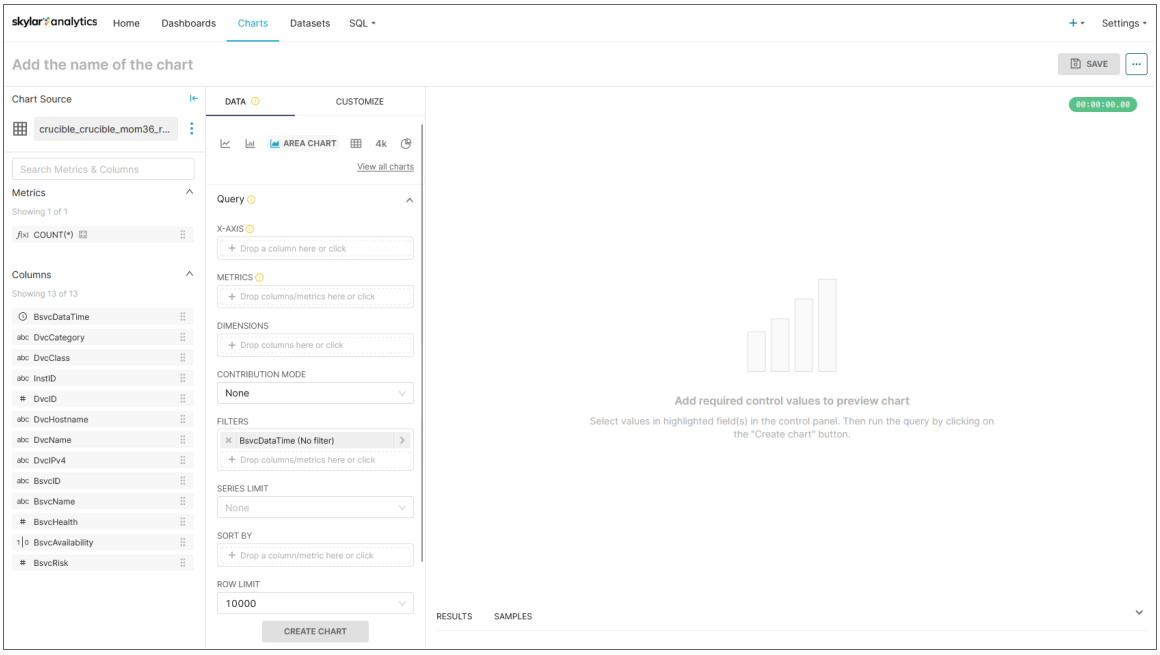

In the first column, the ***Chart Source*** field displays the dataset you selected (for this overview, it is the *BusinessServiceStatistics* chart source). Below that field, you can access the metrics and columns that you can add to the chart. You can drag a metric or column from the first column into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

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

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:

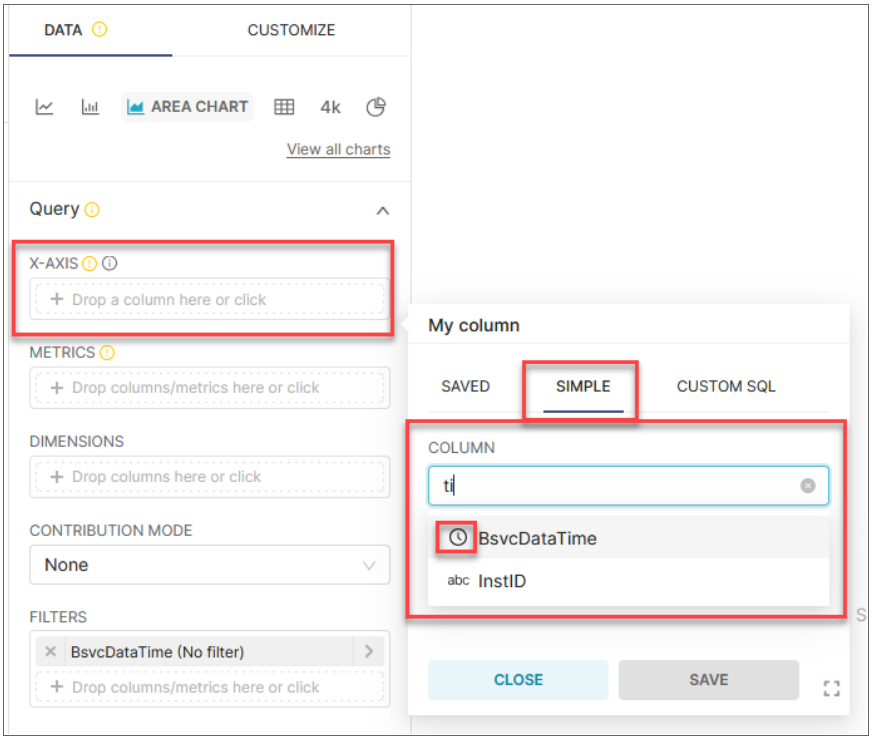

11. On the **[Simple]** tab of the modal, click the ***Column*** field to get a list of data. You can pre-filter the data by typing a column name or label, such as "time".

12. Select a column with a calendar icon (🕓) next to it to display a time range on the X-axis, such as *BsvcDataTime* from the *BusinessServiceStatistic* chart source, and click **[Save]**.

> **TIP:** For more information about the abbreviations used for the metric names, see *Mapping Skylar One Dynamic Application Object Names to Skylar Columns*.

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:

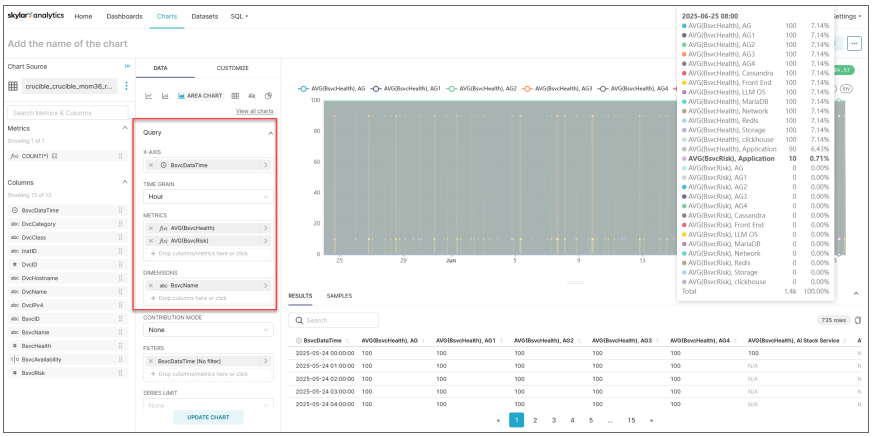

    The chart legend displays the average health and average risk in the legend at the top, and also in the columns at the bottom of the section. If you mouse over a line in the chart, you can see the specific data for those values.

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:

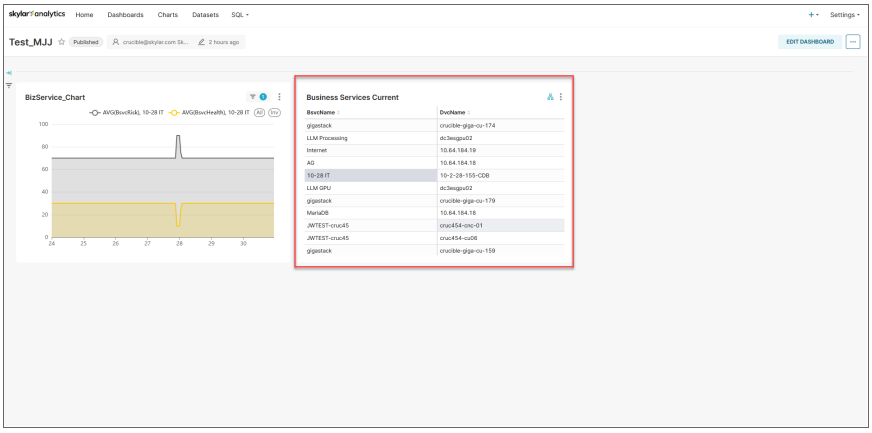

12. Typically, you would want to place the chart that drives the "context" (or the contextual cross-filtering) first, so click **[Edit Dashboard]** to enter Edit mode again.
13. Click and drag the next chart to the left of the first chart until a vertical rectangle appears. Drop the new chart onto that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

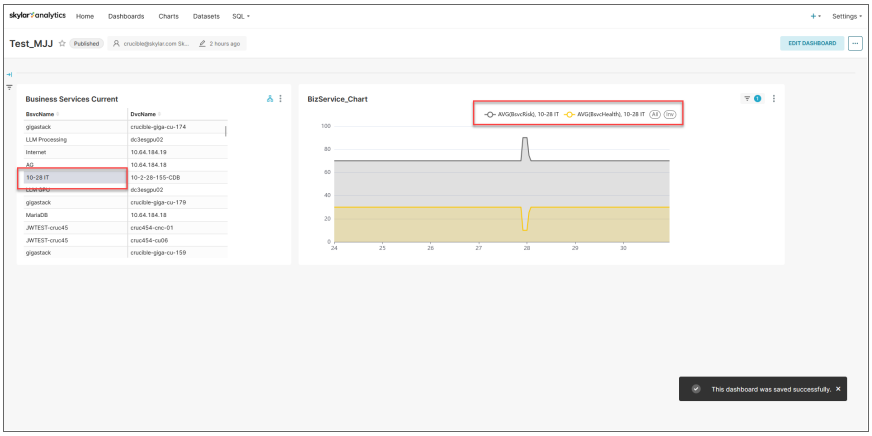

NOTE: If the second chart displays "No results were returned for this query", you might need to click the vertical ellipsis icon (⋮) for the second chart and select *Edit chart* to address the issue.

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:

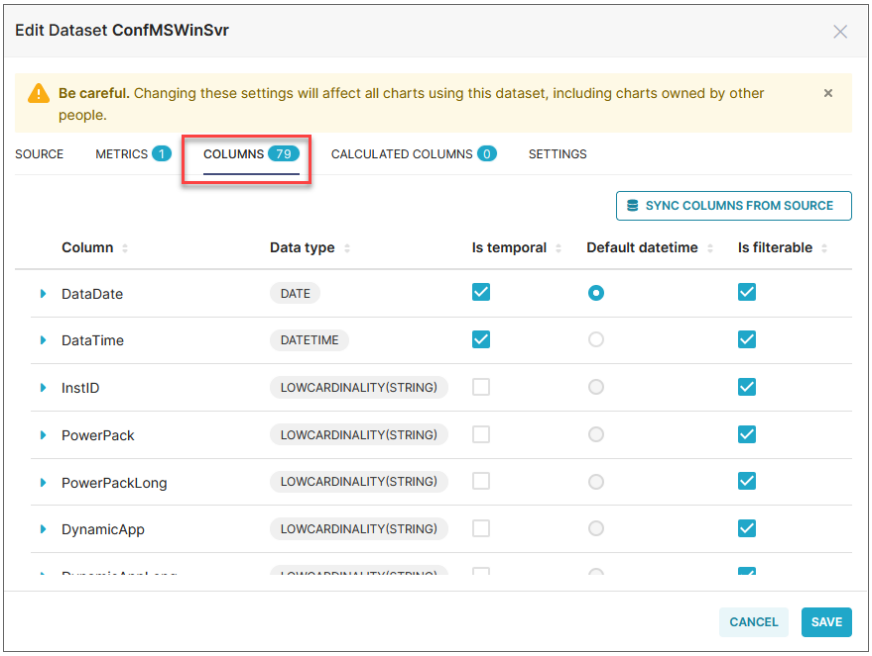

NOTE: You must be the dataset owner to edit a dataset.

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:

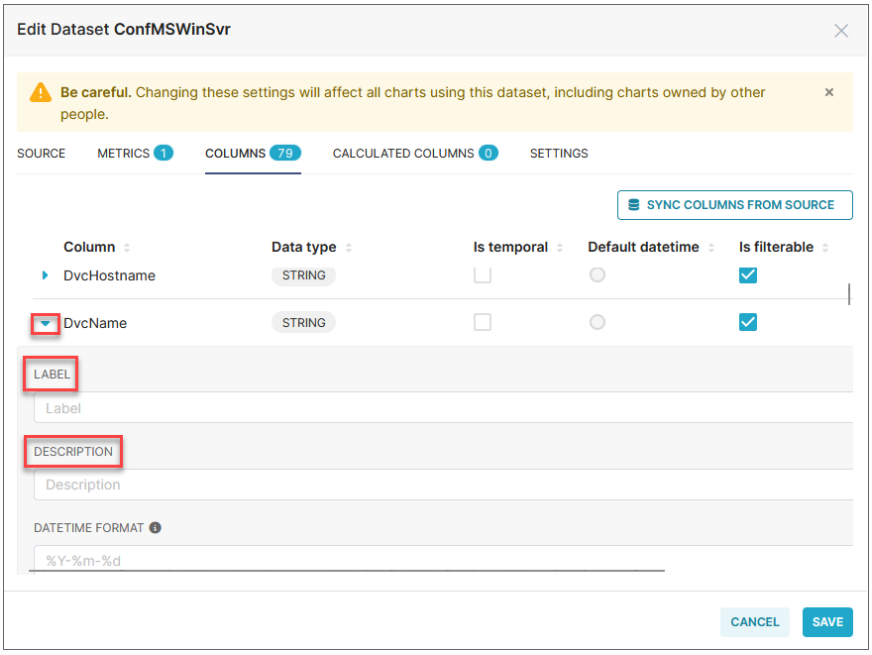

6. In the ***Label*** field, type the new column name you want to display in all charts that use this dataset.
7. In the ***Description*** field, type a short description of the column. This text, along with the **Label** text, appears when you hover over that column in the **Columns** section when you 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.

Adding and Upgrading Dashboards and Datasets

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

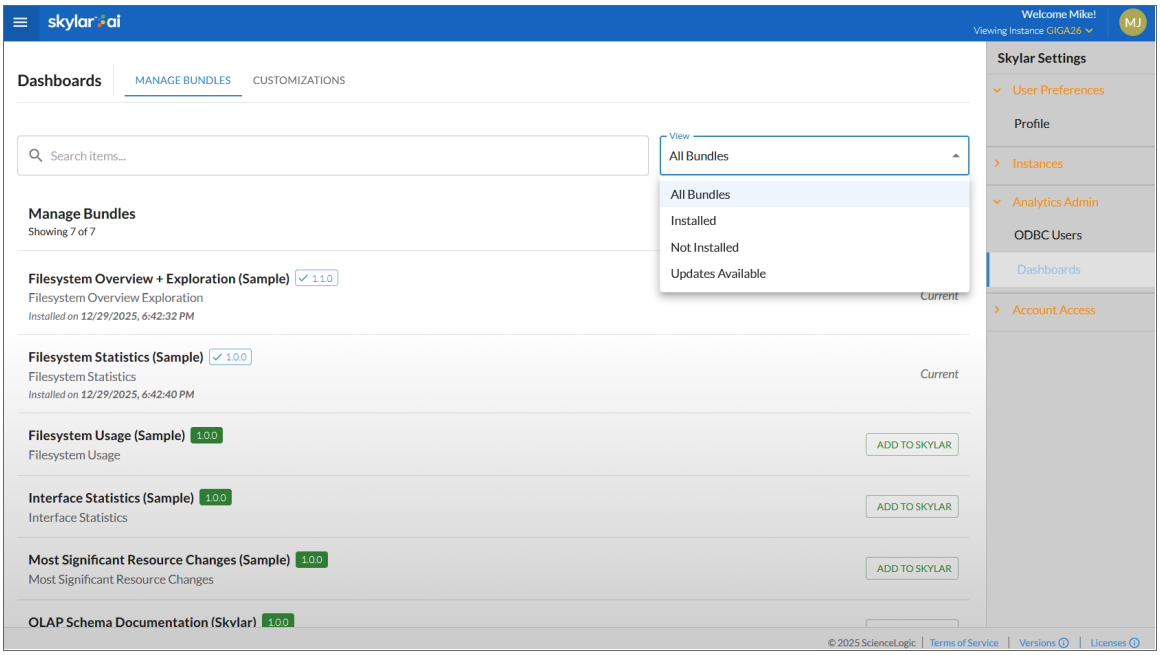

You can search for dashboard bundles and sort the list of bundles by *All Bundles*, *Installed*, *Not Installed*, and *Updates Available*.

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

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

The options on the **Dashboards** page 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.

In addition, you can use the **[Sync Skylar Datasets]** button on the **[Customizations]** tab on the **Dashboards** page to update all of your datasets based on Skylar One PowerPacks, including PowerPacks that have been updated in Skylar One. If all datasets have been updated, the button does not appear, and the text "Datasets are current" appears instead. This button is only available to owner users in Skylar AI.

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

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

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

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

Configuring Data Exploration with Power BI

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

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

To install and configure the ODBC connection:

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

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

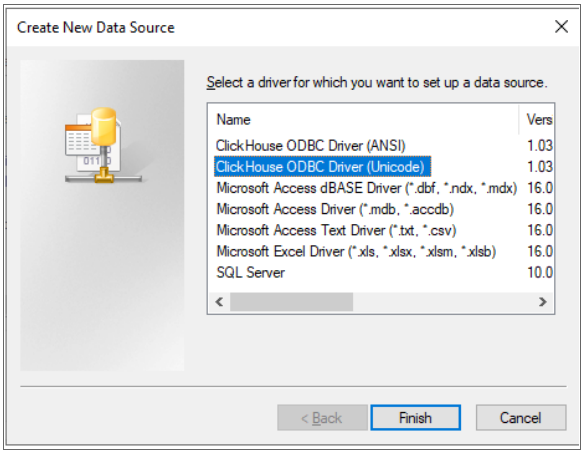

5. Select **ClickHouse ODBC (Unicode)** and click **[Finish]**. The **Create data source for Clickhouse** dialog appears:

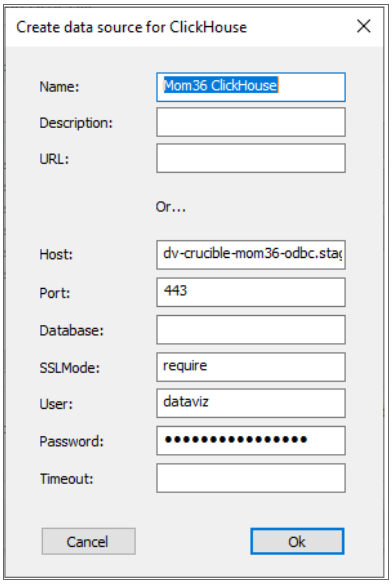

6. Complete the following fields with ODBC connection details from the ScienceLogic Site Reliability Engineering (SRE) team:
   - ***Name***: Add a name to identify this connection. This will be used later in the BI tools.
   - ***Host***: Specify the host URL, provided by SRE.
   - ***Port***: 443.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

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

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

> **TIP:** For an example of how to connect Power BI, see How to connect Power BI to Skylar AI via ODBC.

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

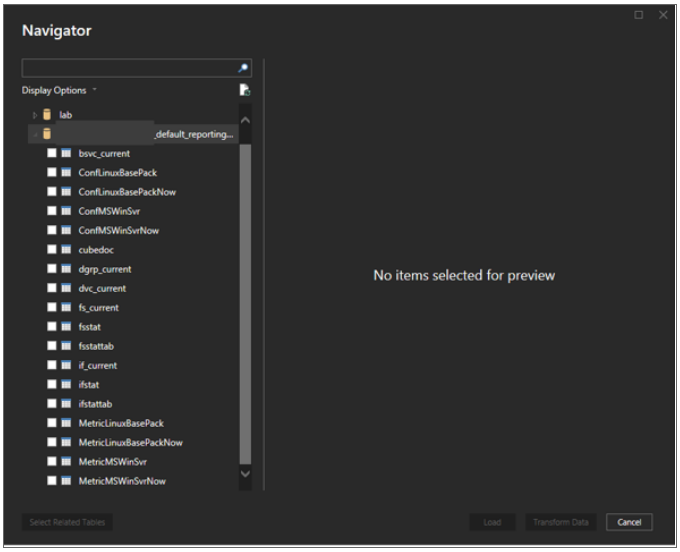

To import data and create a dashboard with Skylar AI data in Power BI:

> **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*:

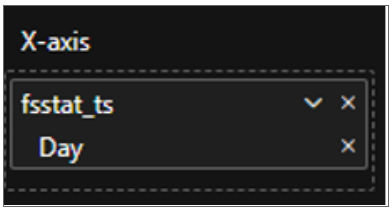

6. To configure the Y-Axis, drag **fsstat_used_pct_psec** (Used Percentage Per Second) to the Y-Axis.
7. To customize the data fields, click the drop-down arrow next to the selected data field. You can rename the field or modify how the value is calculated:

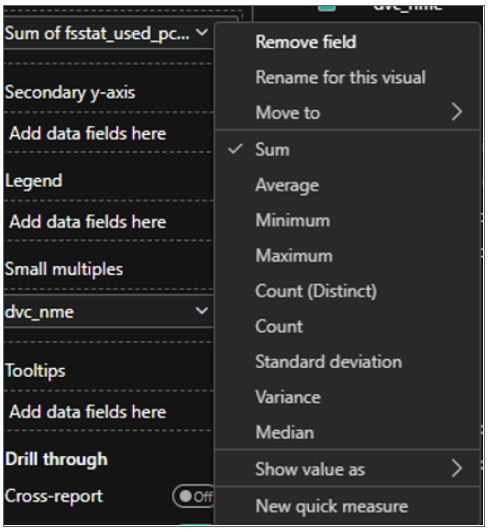

8. Continue adding additional charts and visuals as needed to finish up your dashboard.

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 Skylar One display their data in Skylar Analytics, how to verify if you have current data from Skylar One in Analytics, how to create a dataset (temporary), and a new workflow for creating charts from datasets.
- Tips & Tricks: Friendly Column Names: An explanation of the right way and the wrong way to rename columns in Skylar Analytics.
- BYO Tool and Data: Connect via ODBC and Import Data: How to connect with ODBC using an open-source database tool called DBeaver (for Mac or PC), and how to import data from an external source and store it in Skylar "local" database schema.
- How to connect Power BI to Skylar AI via ODBC: How to connect Power BI to Skylar AI.
- Writable Schema, Joins and Views: How to bring in external pricing data and use it in a Skylar Analytics chart, and how to create a database view by using a JOIN query.

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

- Apache Superset-related documentation: https://superset.apache.org/docs/intro
- Apache Superset Community: https://superset.apache.org/community
- Udemy course for Apache Superset: https://www.udemy.com/course/apache-superset-for-data-engineers-hands-on/
- What is Apache Superset - Quick Overview: https://www.youtube.com/watch?v=znnmco3eK-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 3

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 Skylar One PowerPacks.

This chapter covers the following topics:

What is Anomaly Detection?

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

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

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

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

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

You can view a list of all devices that are being monitored for anomalies on the **Anomaly Detection** page in Skylar One (Skylar AI ( ) > **[Advanced: Anomaly Alerting]** button):

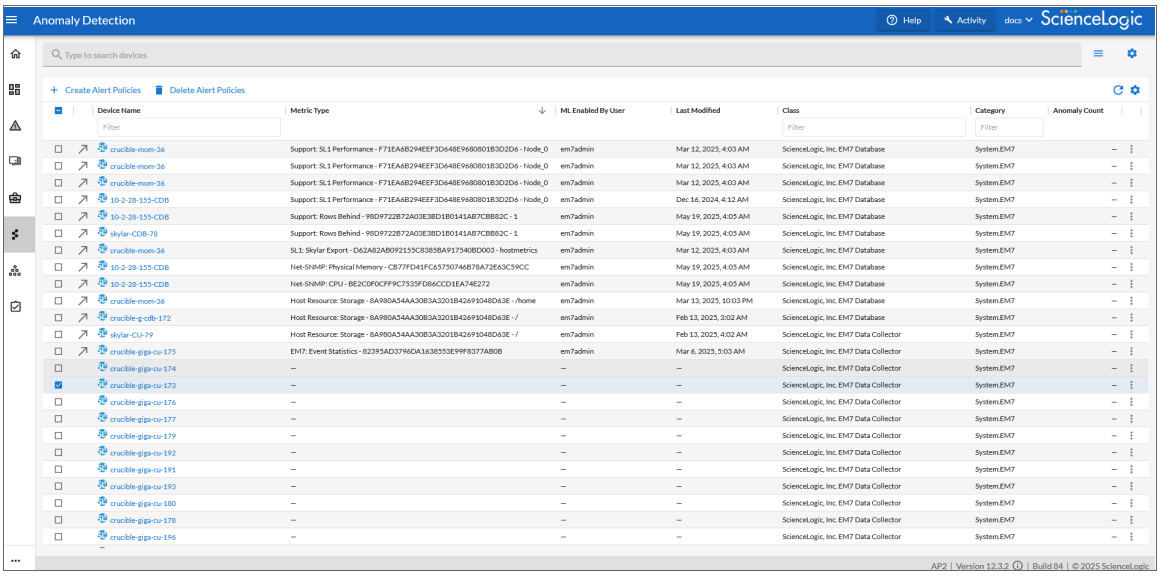

NOTE: The filtered list will appear blank until an Anomaly Score alert triggers an event.

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

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

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

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

How Anomaly Detection Works

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

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

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

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

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

Viewing Graphs and Data for Anomaly Detection

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

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

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

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

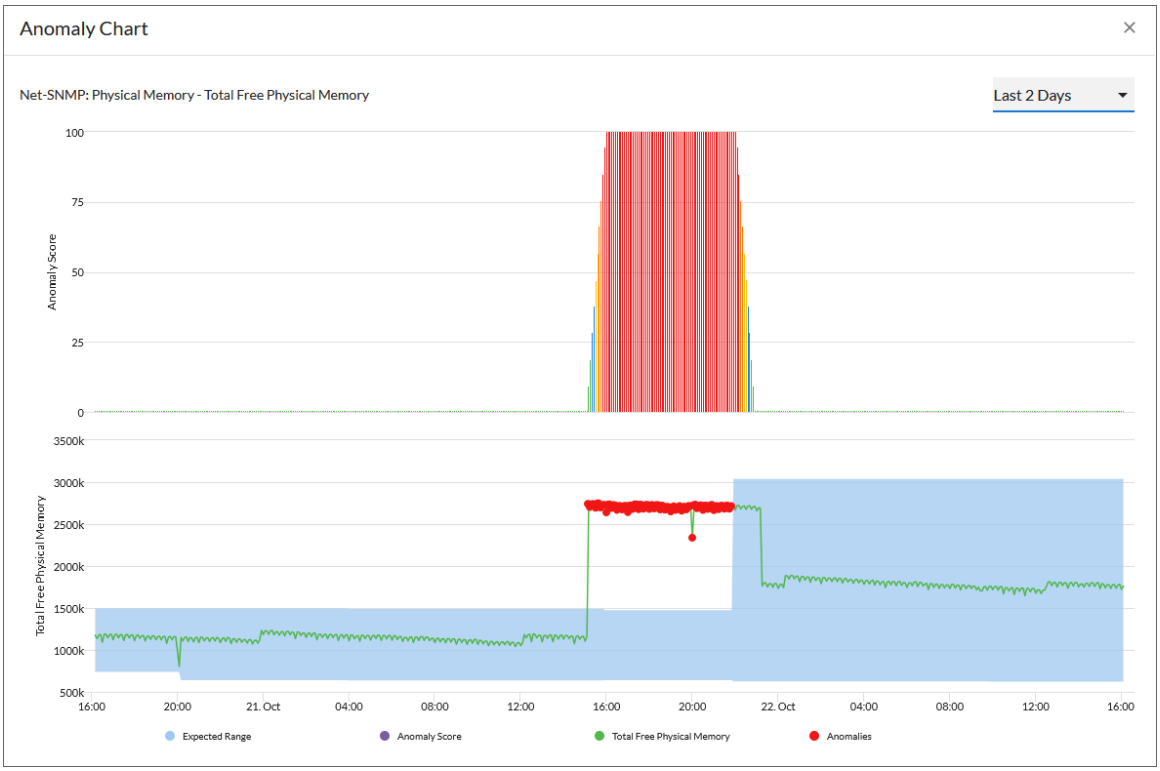

You can define the thresholds for the "Anomaly Score" chart on the **Anomaly Detection Thresholds** page (Skylar AI ( ) > **[Advanced: Adjust Thresholds]** button). You can also use this page to specify whether

the Anomaly Score values generate alerts in Skylar One. For more information, see *Enabling Thresholds and Alerts for the Anomaly Chart*.

The second graph displays the following data:

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

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

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

Enabling Thresholds and Alerts for the Anomaly Chart

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

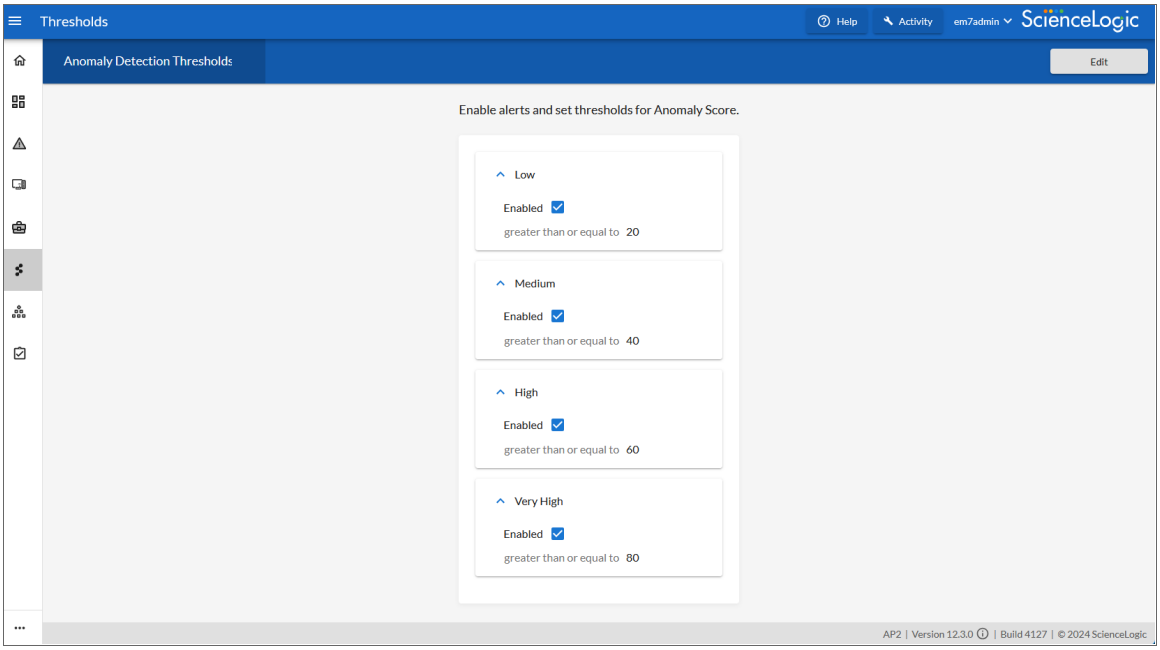

You can view the alert levels when you hover over a value in one of the charts on the **Anomaly Chart** modal. The Anomaly Score severity level displays after the index value, in parentheses: Normal, Low, Medium, High, or Very High:

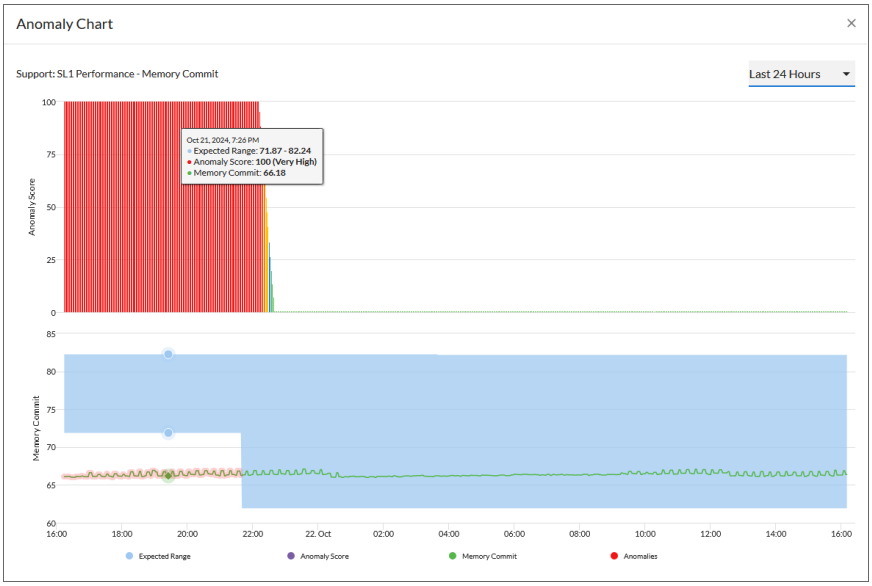

NOTE: An Anomaly Score severity level of ***Normal*** is assigned to a value in the chart that is *lower* than the lowest enabled alert level. For example, if the threshold for the Low severity is enabled and set to 20 or higher, an Anomaly Score of 16 would have a severity level of Normal.

To edit the Anomaly Score thresholds:

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

Enabling Anomaly Detection Events for Specific Metrics

While anomaly detection is enabled automatically as soon as you *enable Skylar Analytics for one or more Skylar One 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 Skylar One.

> **IMPORTANT:** If you are using only SNMP to monitor a device that does not have any Dynamic Applications aligned to it, you currently cannot enable anomaly detection events for that device.

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

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

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

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

2. On the **[Anomaly Detection]** tab, click the **Actions** icon ( ) for any of the listed devices and select *Enable Alerting*. You can also select multiple devices using the check box on the left and click the **[Create Alert Policies]** button at the top. The **Select Available Metrics** modal appears.
3. In the ***Select Metric*** drop-down of the **Select Available Metrics** modal, 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 Alerting]**. That metric is enabled for events for that device.

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

Creating an Event Policy for Anomalies

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

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

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

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

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

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

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

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

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

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

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

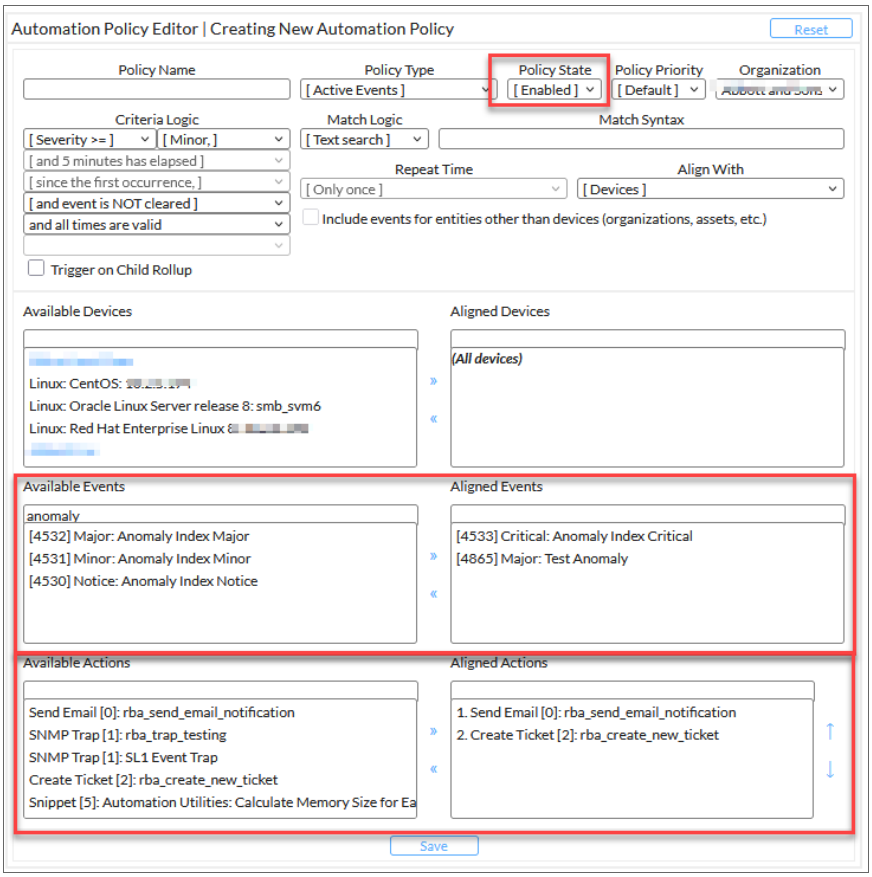

3. In the ***Policy State*** field, select *Enabled*.
4. In the ***Available Events*** field, search for and select one or more anomaly-related event policies, and then click the right-arrow icon to move each event to the ***Aligned Events*** field. For more information about anomaly-related events, 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 4

Skylar Analytics: Predictive Alerting

Overview

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

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

This chapter covers the following topics:

What is Predictive Alerting?

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

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

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

How Predictive Alerting Works

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

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

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

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

Viewing Predictive Alerts in Skylar One

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

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

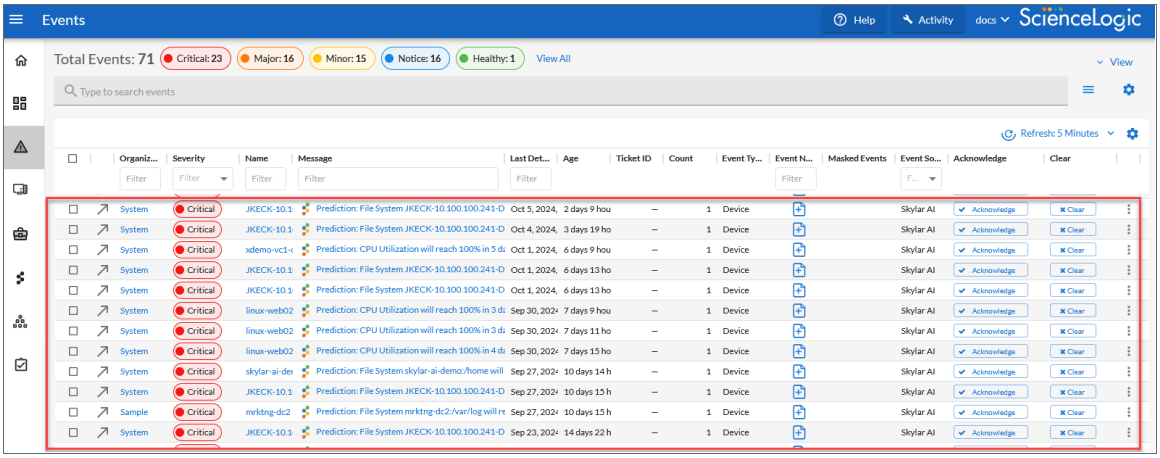

NOTE: The filtered list will appear blank until an active predictive alert triggers an event.

To view details about a predictive alert:

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

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

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

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

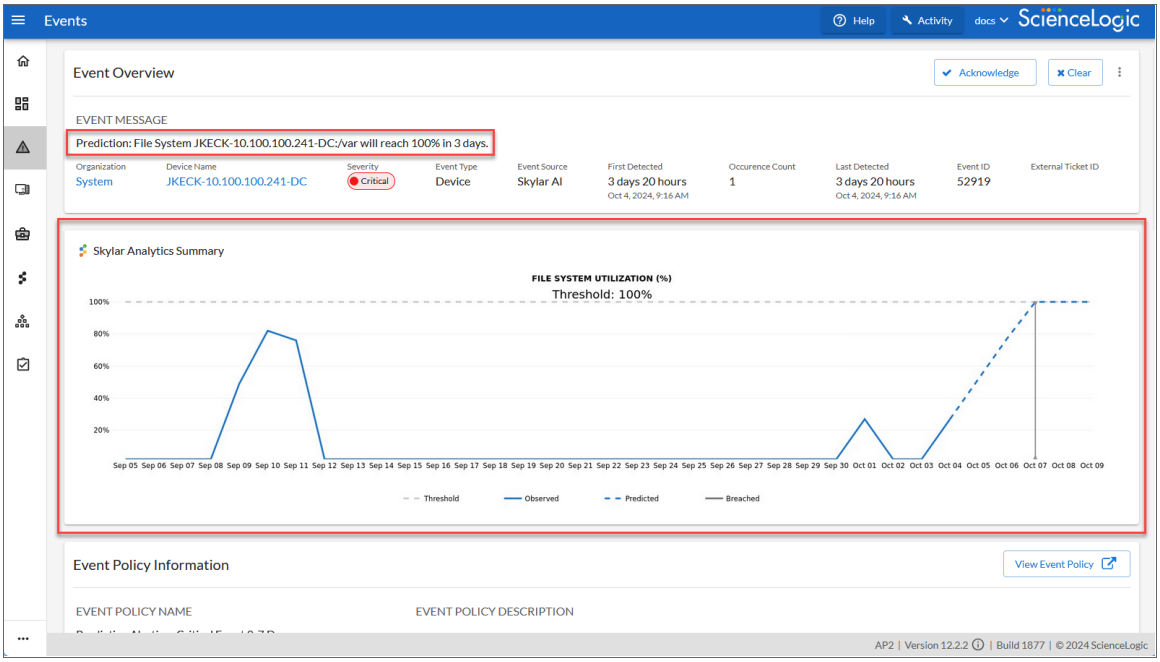

The dotted line on the graph in the **Skylar Analytics Summary** panel represents a time frame in the future that Skylar AI is forecasting, based on pattern recognition.

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

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

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

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

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

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:

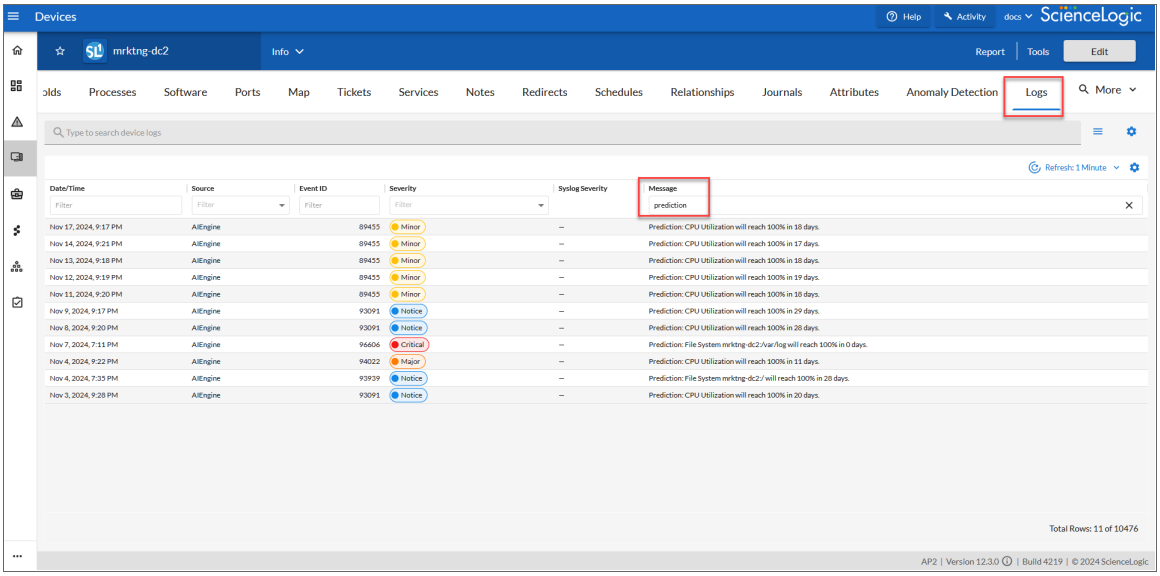

3. If needed, type "prediction" in the ***Message*** column to view only the predictive alerts.

Using Predictive Alerts to Trigger Automated Run Book Actions

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

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

To use predictive alerts to trigger automated run book actions:

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

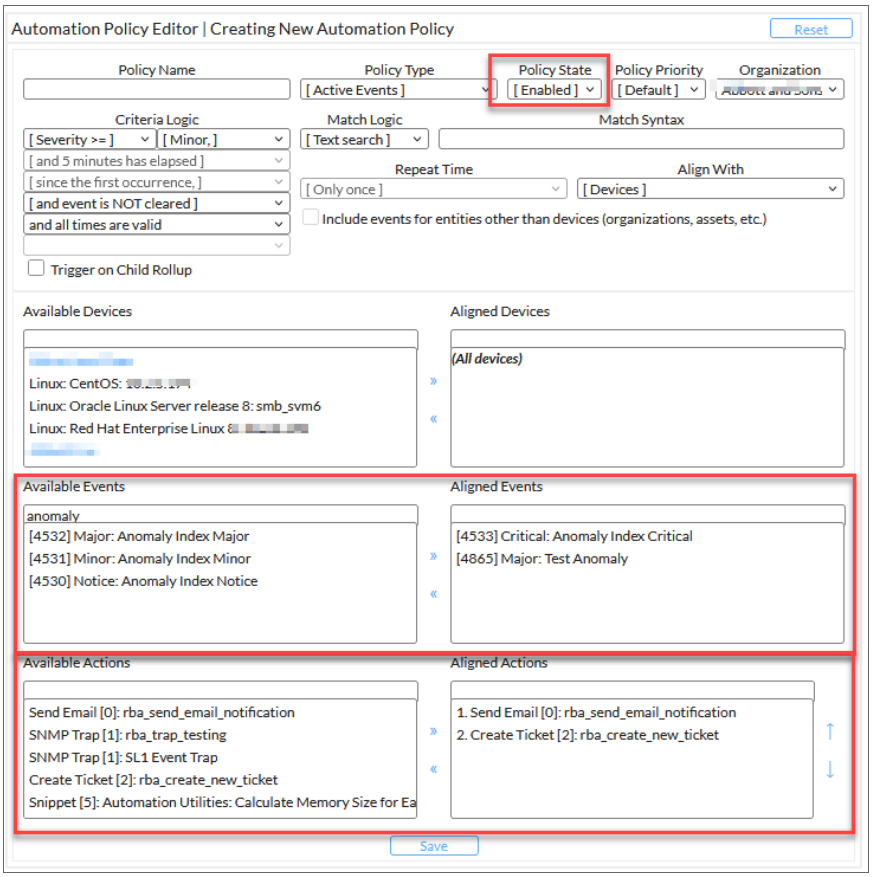

3. In the ***Policy State*** field, select *Enabled*.
4. In the ***Available Events*** field, search for and select one or more event policies related to predictive alerts, and then click the right-arrow icon to move each event to the ***Aligned Events*** field.
5. In the ***Available Actions*** field, 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]**.

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

International: +1-703-354-1010