Machine Learning-based Anomaly Detection

SL1 version 12.1.0

Table of Contents

Chapter 1

Introduction

Overview

This chapter describes how to enable machine learning-based anomaly detection in SL1, as well as how to view recent anomalies for devices and services.

> **NOTE:** To use machine learning-based anomaly detection on the SL1 Extended Architecture, you must enable the Collector Pipeline to collect data from Performance Dynamic Applications. For more information, see the ***Using SL1 Publisher*** manual.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ( ).
- To view a page containing all of the menu options, click the Advanced menu icon (•••).

For more information about machine learning-based anomaly detection and business services, watch the video at https://www.youtube.com/watch?v=_GXArDFQ3dc.

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. SL1 does this by collecting data for a particular metric over a period of time, learning the patterns of that particular device metric, and then choosing the best possible algorithm to analyze that data.

SL1 uses the resulting combination of collected data and the auto-selected algorithm to build a model that is unique to that specific device and metric. That model is then used to anticipate the expected behavior for that device metric. Anomalies are detected when the actual collected data value falls outside the boundaries of the expected value range. SL1 then continuously refines the model as it collects more data.

TIP: Anomalies do not necessarily represent problems or events to be concerned about; rather, they represent unexpected behavior that you might want to investigate.

Anomaly Detection Terminology

The following are some terms that are used when discussing machine learning-based anomaly detection, and their definitions.

- ***Algorithm***. A mathematical formula for data analysis. SL1 currently uses period-based and multiple clustering algorithms to perform anomaly detection, with the ability to easily add more algorithms in the future.
- ***Model***. The combination of collected data and algorithm that SL1 uses to anticipate expected behavior and discover anomalies for a specific metric on a specific device. SL1 constantly refines these models.
- ***Model Selector***. An automatic model selector included in SL1 that examines the historical data for the selected metric and ensures there is enough data for successful analysis, applies all possible algorithms to the data, then determines the algorithm that is best able to distinguish anomalies in the data in order to build the model for the selected metric.

How is Anomaly Detection Different from Standard Deviation?

In SL1, you can use the deviation function to examine values collected by Dynamic Applications. The deviation function compares each collected value to the mean value for that hour and that day of the week. Deviation triggers an alert only when values fall outside the historical range of data, but will not trigger an alert when something abnormal happens within that range.

In contrast, anomaly detection learns the behavioral shapes and patterns of a data point and triggers an alert when values for that data point fall outside the behavioral shape. For example, anomaly detection could generate an alert when it discovers an unexpected flatline, a spike during a "low usage" period, or when collected values should repeat a pattern but do not. All these behaviors could occur within a standard deviation from the mean value so would not be discovered with the deviation function.

How Does SL1 Detect Anomalies?

The following steps describe the basic process SL1 uses to detect anomalies:

1. SL1 observes the behavior of a single metric on a single device, using historical and current time-series data.
2. Based on the observed behavior, SL1 performs calculations and builds a model that is specific to that single metric on that single device.
3. SL1 detects behavior that is abnormal compared to the model. This abnormal behavior is considered an anomaly.
4. SL1 then regularly refines its original model as more data is collected about the single metric on the single device.
5. Every two weeks, SL1 will rebuild the model. This enables it to adapt to new data patterns that have emerged since the last time the model was built.

> **NOTE:** Because the anomaly detection model is constantly being refined, you might experience a greater number of anomalies after you initially enable anomaly detection than you would after it has been enabled for a longer period of time. This is simply because there is less collected data to "train" the model after anomaly detection has initially been enabled, and it will begin to better understand longer-term behavior patterns the longer it collects and analyzes data.

Anomaly detection in SL1 can examine vitals data and any performance data collected by a Dynamic Application. When it discovers an anomaly, SL1 generates an alert. Optionally, you can choose to create events based on these alerts. For more information, see *Creating an Event Policy for Anomalies*.

What Can Anomaly Detection Do?

After you enable machine learning-based anomaly detection for a metric on a device, SL1 requires a certain amount of historical data in order to select the model it will use to detect anomalies. Depending on the configured polling frequency and the amount of historical data that is available for the device metric, it might take SL1 under an hour or up to several days to begin detecting anomalies.

Viewing Graphs 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 **Machine Learning** (🧠) page.
- The **[Machine Learning]** tab in the **Device Investigator**.
- The **Anomalies** widget in the **Service Investigator** for a business, IT, or device service.

You can view these graphs by clicking the **Expand** icon ( ) next to the device or the metric for the device. The **Anomaly Chart** modal appears, displaying the "Anomaly Index" chart above the chart for the specified metric you are monitoring.

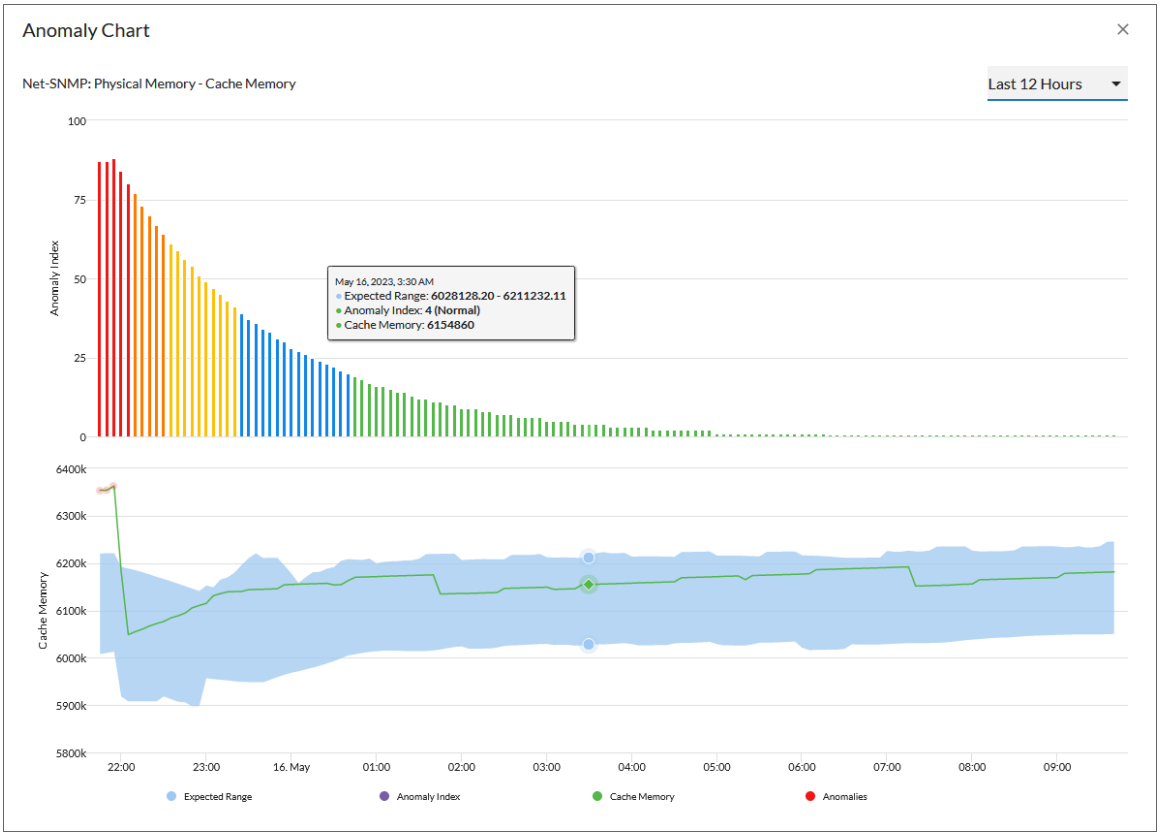

The "Anomaly Index" chart displays a graph of values from 0 to 100 that represent how far the real data for a metric diverges from its normal patterns. The lines in the chart are color-coded by the level of event that gets triggered as the data diverges further and further. You can define the thresholds for the Anomaly Index, and whether those values generate alerts, on the **Machine Learning Thresholds** page (Machine Learning > Thresholds). For more information, see *Enabling Alerts and Thresholds for the Anomaly Index*.

For more information about the charts and graphs, see *Viewing Device Anomalies* and *Viewing Business Service Anomalies.*

Using Anomaly Detection to Trigger Events

You can use anomaly detection to trigger an event or to add extra criteria to an event policy. For example, you could specify that if an anomaly occurs five times within 10 minutes, SL1 should trigger an event. For more information, see *Creating an Event Policy for Anomalies*.

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.

Using Anomaly Detection to Trigger Run Book Actions

You can also use events based on anomaly detection to trigger run book automation actions that perform further diagnostics or send notifications. For more information, see *Using Anomaly-related Events to Trigger Automated Run Book Actions*.

Chapter 2

Enabling Machine Learning-based Anomaly Detection

Overview

This chapter describes how to enable machine learning-based anomaly detection in SL1, as well as how to view recent anomalies for devices and business services.

> **NOTE:** To use machine learning-based anomaly detection on the SL1 Extended Architecture, you must enable the Collector Pipeline to collect data from Performance Dynamic Applications. For more information, see the ***Using SL1 Publisher*** manual.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ( ).
- To view a page containing all of the menu options, click the Advanced menu icon (•••).

This chapter covers the following topics:

Viewing the List of Devices that Have Anomaly Detection Enabled

The **Machine Learning** page displays a list of devices that are currently using machine learning for anomaly detection, as well as devices for which you can enable anomaly detection if it is not enabled already.

To navigate to the **Machine Learning** page, click the **Machine Learning** icon ():

TIP: To filter the devices that appear on the page based on whether anomaly detection is enabled or disabled, type "MachineLearningPolicy.enabled" in the ***Search*** field. A "MachineLearningPolicy.enabled" pill appears below the ***Search*** field. Click that pill and then select *True* to filter the page to display only those devices on which anomaly detection is enabled, or select *False* to filter the page to display only those devices on which anomaly detection is disabled.

TIP: You can filter the items on this inventory page by typing filter text or selecting filter options in one or more of the filters found above the columns on the page. For more information, see "Filtering Inventory Pages" in the ***Introduction to SL1*** manual.

For each device in the list, the **Machine Learning** page displays the following information:

- ***Device Name.*** Displays the name of the device. Click the hyperlink to go to the **[Machine Learning]** tab of the **Device Investigator** page for that device. Each row in the list represents a specific device and metric; therefore, a device might appear in the list multiple times if anomaly detection is enabled for multiple metrics on that device.
- ***Anomaly Detection.*** Indicates the build status for the for the metric that SL1 is evaluating for anomalies on the device. Possible values include:
    - *Disabled*. Anomaly detection is disabled for the metric.
    - *Enabled*. Anomaly detection is enabled for the metric.
    - *Queued*. The metric has been selected for anomaly detection, but SL1 has not yet begun building the anomaly detection model for that metric.
    - *Building*. SL1 is building the anomaly detection model that is specific to the selected device and metric.
    - *Failed*. The anomaly detection model build process failed.
    - *Failed (building)*. SL1 could not find any winners for a predictor for this metric, and SL1 will continue to monitor for anomalies on that device.
    - *Waiting for Data*. Anomaly detection for this metric lacks sufficient data, either because detection needs at least one day of monitoring or the data for this metric is irregular.

- ***Metric Type.*** Indicates the metric that SL1 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.
- ***Class.*** Displays the Device Class for the device.
- ***Category.*** Displays the device's Device Category.
- ***Anomaly Count***. Displays the number of anomalies detected by SL1.

Enabling Machine Learning on a Device

For SL1 to collect and analyze data for the sake of detecting anomalies for a specific metric on a particular device, you must first enable machine learning on that device. You can do that from several different places within SL1.

The following sections describe each of these methods.

Enabling Machine Learning from the Machine Learning Page

To enable machine learning for a device from the **Machine Learning** page:

1. Click the **Machine Learning** icon (). The **Machine Learning** page displays.
2. Locate the device on which you want to enable machine learning.

> **TIP:** To sort by devices that do not have anomaly detection enabled, select *Disabled* in the **Anomaly Detection** column.

3. Click the **[Actions]** icon () for that device and select *Enable*. The **Select Metric to Enable Machine Learning** modal page appears.

> **TIP:** Alternatively, you can select the checkbox for the device on which you want to enable machine learning and then click **[Enable]** at the top of the page.

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 machine learning 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 machine learning.
7. Click **[Enable]**. That metric is enabled for the device, and the metric is listed in the **Metric Type** column on the **Machine Learning** page.

TIP: To enable or disable machine learning for more than one device, select the checkboxes for each device for which you want to enable or disable machine learning and then click **[Enable]** or **[Disable]** at the top of the page.

Enabling Machine Learning in the Device Investigator

To enable machine learning for a device in the **Device Investigator**:

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

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

3. On the **[Machine Learning]** tab, click the **[Add ML Metric]** button or click the **Actions** icon () for any of the listed metrics and select *Enable*. The **Select Metric to Enable Machine Learning** modal page 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 machine learning 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 machine learning.
7. Click **[Enable]**. The metric appears on the **Machine Learning** tab.

TIP: To disable machine learning for a metric, click the **Actions** icon () for that metric and select *Disable*. The metric is removed from the **Machine Learning** tab.

Enabling Machine Learning in the Service Investigator

The **Anomalies** widget in the **Service Investigator** displays a list of devices within the selected business, IT, or device service that have anomaly detection enabled. From this widget, you can also enable machine learning for additional metrics or disable machine learning metrics on which it is currently enabled.

NOTE: The **Anomalies** widget appears only if you have at least one device in the selected service that has anomaly detection enabled.

To enable machine learning in the **Service Investigator**:

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 **Anomalies** widget.
3. Click the **Actions** icon () for any of the listed metrics and select *Enable*. The **Select Metric to Enable Machine Learning** modal page 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 machine learning 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 machine learning.
7. Click **[Enable Machine Learning]**. The metric appears in the **Anomalies** widget.

TIP: To disable machine learning for a metric, click the **Actions** icon () for that metric and select *Disable*. The metric is removed from the **Anomalies** widget.

Viewing Device Anomalies

On the **[Machine Learning]** tab of the **Device Investigator**, you can view a list of machine learning metrics that are enabled for the device:

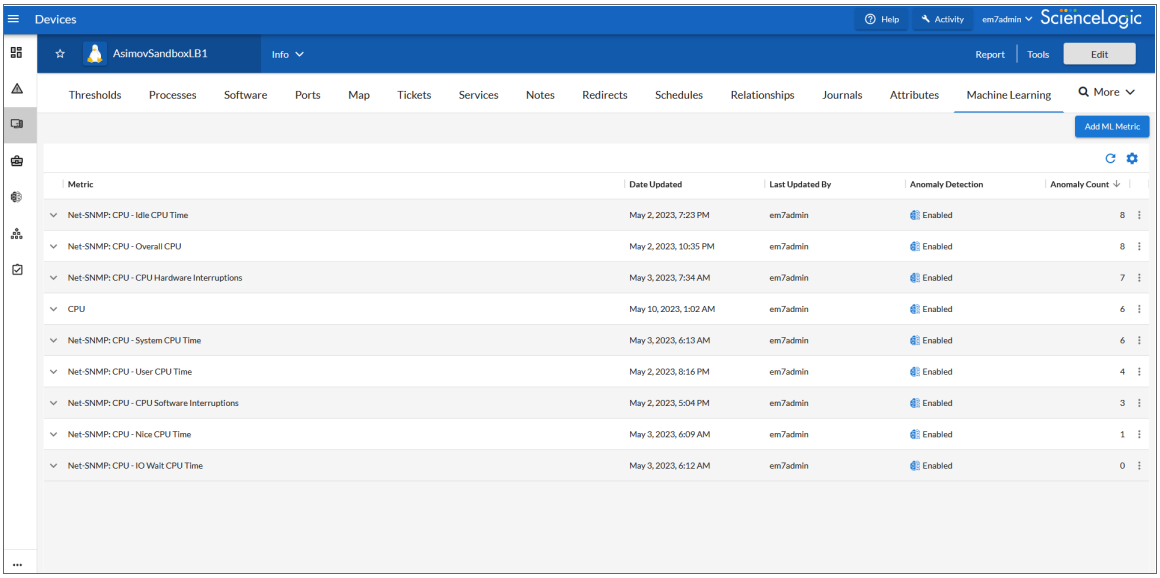

On this tab, you can view the Anomaly Detection graphs by clicking the **Expand** icon ( ) next to the metric for the device. The **Anomaly Chart** modal appears, displaying the "Anomaly Index" chart above the chart for the specified metric you are monitoring.

The "Anomaly Index" 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 level of event that gets triggered as the data diverges further.

You can view these graphs by clicking the **Expand** icon ( ) next to the device or the metric for the device. The **Anomaly Chart** modal appears, displaying the "Anomaly Index" chart above the chart for the specified metric you are monitoring.

The "Anomaly Index" 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 level of event that gets triggered as the data diverges further and further. You can define the thresholds for the Anomaly Index, and whether those values generate alerts, on the **Machine Learning Thresholds** page (Machine Learning > Thresholds). For more information, see *Enabling Alerts and Thresholds for the Anomaly Index*.

In the second graph, the blue shape represents the expected value range for the selected device metric over the given time period, the green line indicates the actual values that SL1 collected over that time period, and the small red dots at top left represent the anomalies where the actual value fell outside of the expected range.

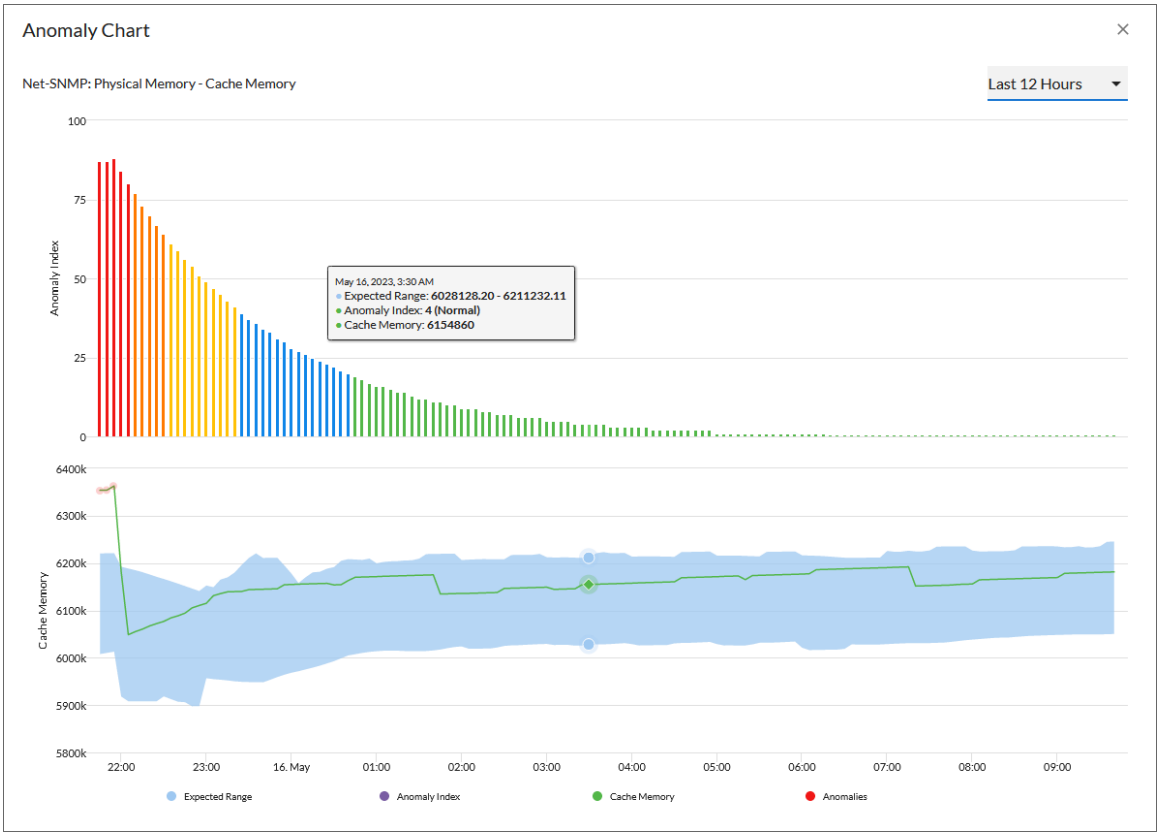

TIP: You can hover over a value in one of the charts to see a pop-up box with the **Expected Range** and the metric value. The **Anomaly Index** value also displays in the pop-up box, with the severity in parentheses: Normal, Low, Medium, High, or Very High.

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.

TIP: 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.

NOTE: For more information about devices, see the ***Device Management*** manual.

Viewing Business Service Anomalies

If one or more devices within a business, IT, or device service has anomaly detection enabled, the **Anomalies** widget will appear on the **[Overview]** tab of the **Service Investigator**. The **Anomalies** widget displays a list of all the devices within the selected service that have anomaly detection enabled.

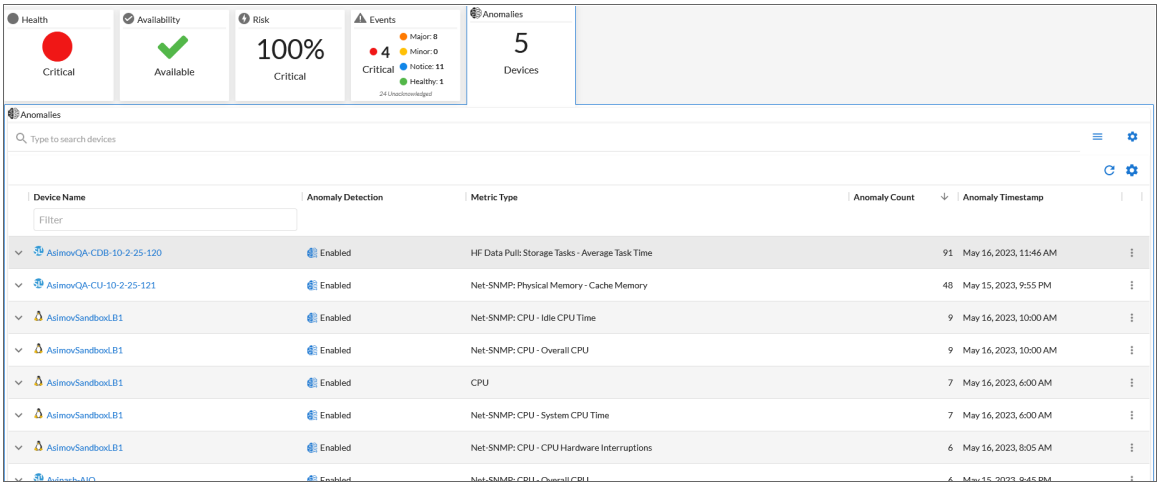

To view the **Service Investigator** page, select a service from the list of business, IT, and device services on the **Business Services** page ( ). The **[Overview]** tab opens by default. This tab provides a single-page view of the selected service, including key metrics, events, and anomalies that are impacting the service.

On the **[Anomalies]** tab of the **Device Investigator**, you can view a list of devices that are enabled for anomaly detection. Each device has a set of graphs that tracks the anomaly detection data for that device.

You can view these graphs by clicking the **Expand** icon ( ) next to the device or the metric for the device. The **Anomaly Chart** modal appears, displaying the "Anomaly Index" chart above the chart for the specified metric you are monitoring.

The "Anomaly Index" 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 level of event that gets triggered as the data diverges further and further. You can define the thresholds for the Anomaly Index, and whether those values generate alerts, on the **Machine Learning Thresholds** page (Machine Learning > Thresholds). For more information, see *Enabling Alerts and Thresholds for the Anomaly Index*.

In the second graph, the blue shape represents the expected value range for the selected device metric over the given time period, the green line indicates the actual values that SL1 collected over that time period, and the small red dots at top left represent the anomalies where the actual value fell outside of the expected range.

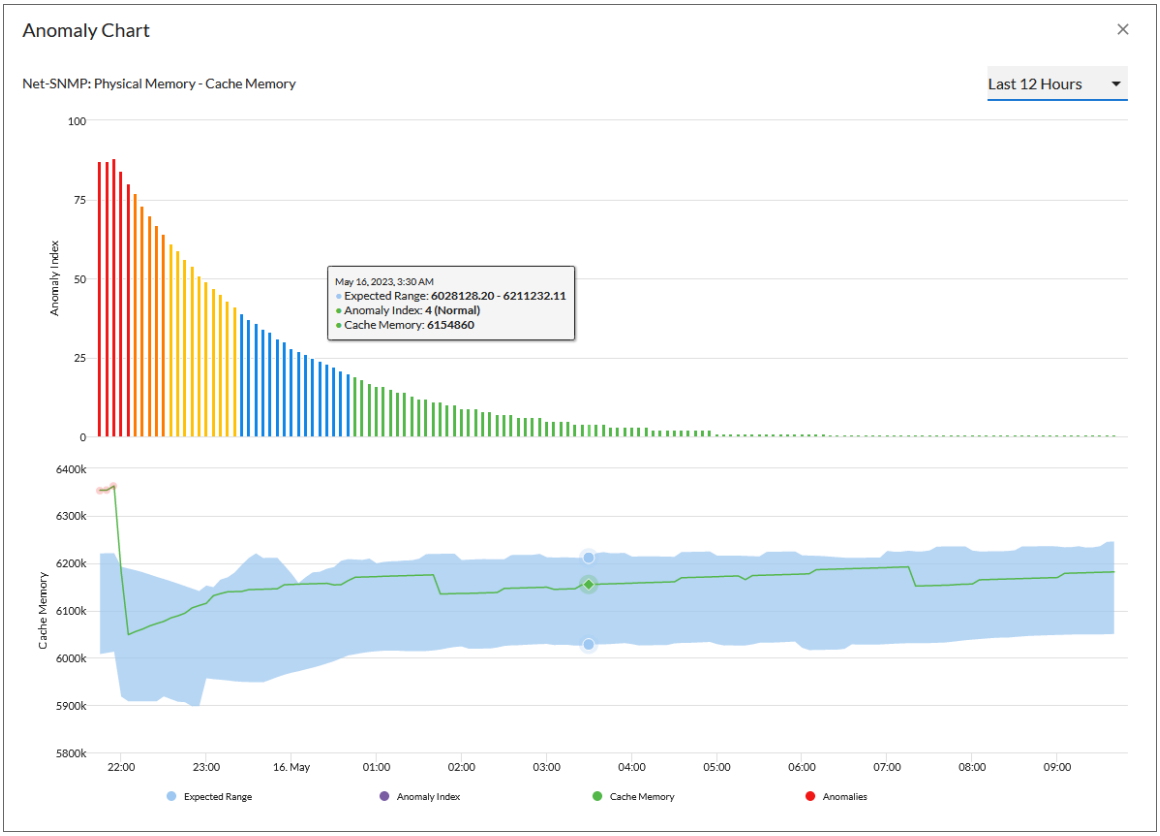

TIP: You can hover over a value in one of the charts to see a pop-up box with the **Expected Range** and the metric value. The **Anomaly Index** value also displays in the pop-up box, with the severity in parentheses: Normal, Low, Medium, High, or Very High.

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.

TIP: You can use the time span filter on the **Anomalies** widget to adjust the time span of anomalies that appears in the graph. The default filter is *Last 24 hours*, but you can select a time span ranging from *Last Hour* up to *Last 2 Years*. You can also 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.

NOTE: For more information about business services, see the ***Monitoring Business Services*** manual.

Enabling Alerts and Thresholds for the Anomaly Index

You can define the thresholds for the "Anomaly Index" chart, and whether those values generate alerts, on the **Machine Learning Thresholds** page (Machine Learning > Thresholds).

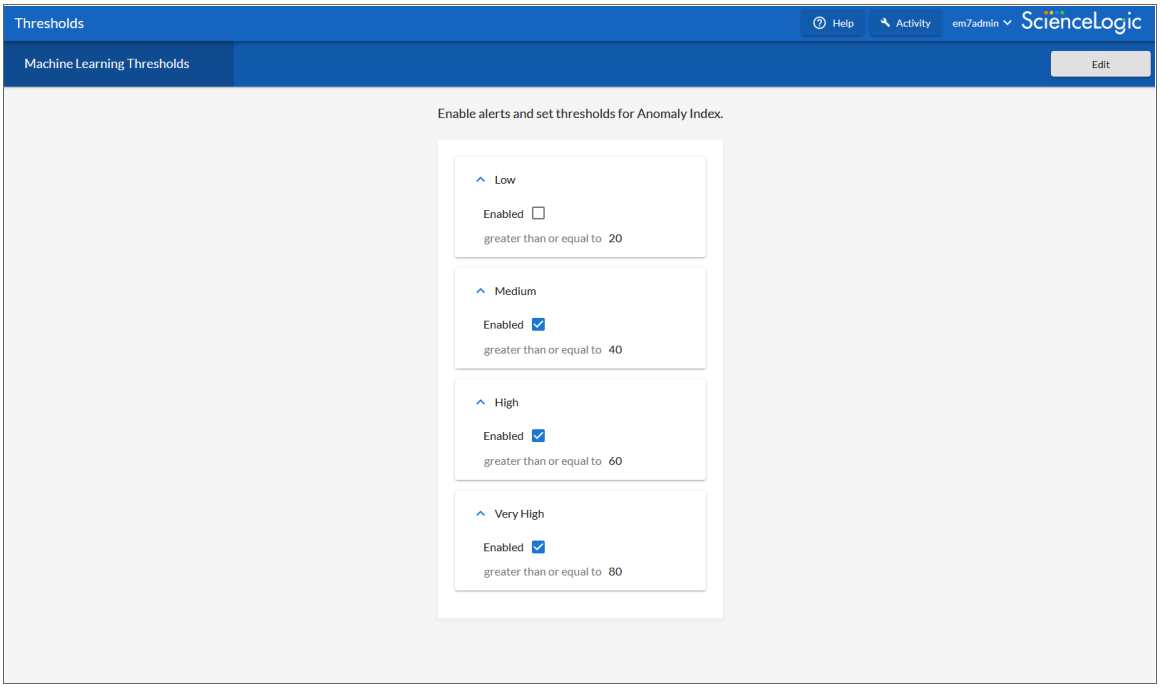

You can define which value in the Anomaly Index will trigger an alert, and the severity level of the alert. These settings are used by all devices that have enabled anomaly detection.

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

NOTE: An Anomaly Index severity level of Normal is assigned to a value in the chart that is lower than the lowest enabled alert level. In the example above, the threshold for the Low severity is enabled and set to 20 or higher, so the Anomaly Index value for that specific point in time has a severity level of Normal.

To edit the Anomaly Index thresholds:

1. On the **Machine Learning Thresholds** page (Machine Learning > Thresholds), 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 Index 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.
   For example, if you want to enable a Low level alert when the Anomaly Index value is between 25 and 39 for the SL1 system in the image above, you would go to the **Low** panel, select ***Enabled***, and update the value from "20" to "25".
4. Click **[Save]**.
5. 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*.

Chapter 3

Using Anomaly Detection to Trigger Events and Automations

Overview

This chapter describes how to use machine learning-based anomaly detection to trigger events and automations in SL1.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ( ).
- To view a page containing all of the menu options, click the Advanced menu icon (•••).

This chapter covers the following topics:

Creating an Event Policy for Anomalies

After you have enabled machine learning-based anomaly detection for devices, you can create 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":

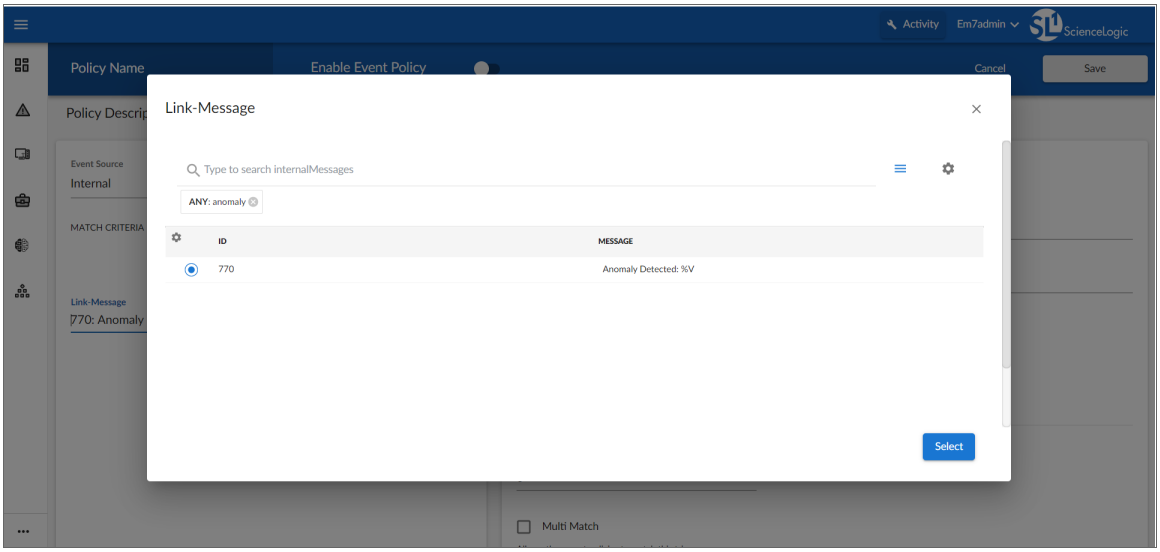

8. Click the radio button for the message "Anomaly Detected: %V", and then click **[Select]**.
9. Complete the remaining fields and tabs in the **Event Policy Editor** based on the specific parameters that you want to establish for the event. For more information about the fields and tabs in the **Event Policy Editor**, see the chapter on "Defining and Editing Event Policies" in the ***Events*** manual.
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:

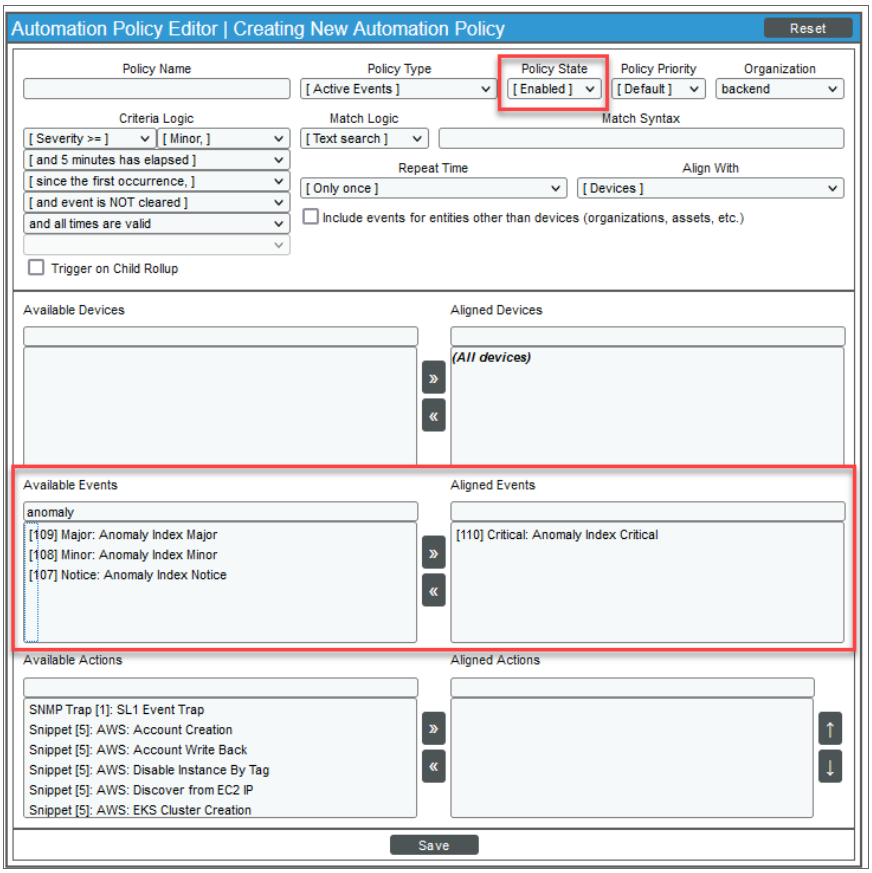

3. In the ***Policy State*** field, select *Enabled*.
4. In the ***Available Events*** field, search for and select an anomaly-related event policy, and then click the right-arrow icon to move it to the ***Aligned Events*** field. For more information about anomaly-related events, see *Creating an Event 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]**.

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.

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

International: +1-703-354-1010