Machine Learning-based Anomaly Detection

Beta Version

SL1 version 10.1.0

Table of Contents

Introduction 3
What is Anomaly Detection? 4
Anomaly Detection Terminology 4
How is Anomaly Detection Different from Standard Deviation? 4
How Does SL1 Detect Anomalies? 5
What Can Anomaly Detection Do? 5
Enabling Machine Learning-based Anomaly Detection 7
Viewing the List of Devices that Have Anomaly Detection Enabled 8
Enabling Machine Learning on a Device 9
Enabling Machine Learning from the Machine Learning Page 9
Enabling Machine Learning in the Device Investigator 9
Enabling Machine Learning in the Service Investigator 10
Viewing Device Anomalies 11
Viewing Service Anomalies 12
Using Anomaly Detection to Trigger Events and Automations 15
Creating an Event Policy for Anomalies 15
Using Anomaly-related Events to Trigger Automated Run Book Actions 17

Chapter 1

Introduction

Overview

This chapter provides an overview of machine learning-based anomaly detection in SL1, including how it works and what it can do.

NOTE: Machine learning-based anomaly detection is available only in SL1 Premium solutions. To upgrade, contact ScienceLogic Customer Support.

NOTE: To use machine learning-based anomaly detection, you must be running SL1 on the SL1 Extended Architecture.

NOTE: Machine learning-based anomaly detection is a Beta feature in SL1 version 10.1.0.

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 includes the following topics:

What is Anomaly Detection? 4

Anomaly Detection Terminology 4

How is Anomaly Detection Different from Standard Deviation? 4

How Does SL1 Detect Anomalies? 5

What Can Anomaly Detection Do? *5*

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 constantly refines its original model as more data is collected about the single metric on the single device.

> **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 CPU 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.

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.

After SL1 begins performing anomaly detection for a device, you can do the following:

- You can view graphs and data about each anomaly. Graphs for anomalies appear on the following pages in SL1:
    - The **[Machine Learning]** tab in the **Device Investigator**
    - The **Anomalies** widget in the **Service Investigator** for a business, IT, or device service

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

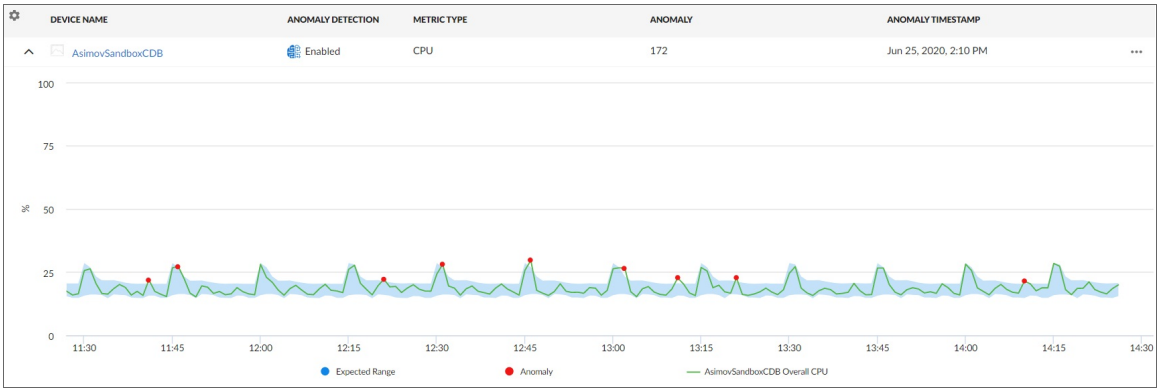

- You can use anomaly detection to trigger an event or to add extra criteria to an event policy. For example, you could specify that if an anomaly occurs five times within 10 minutes, SL1 should trigger an event.
- You can also use events based on anomaly detection to trigger Run Book Automation actions that perform further diagnostics or send notifications.

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.

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 services.

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 includes the following topics:

Viewing the List of Devices that Have Anomaly Detection Enabled 8
Enabling Machine Learning on a Device 9
Enabling Machine Learning from the Machine Learning Page 9
Enabling Machine Learning in the Device Investigator 9
Enabling Machine Learning in the Service Investigator 10
Viewing Device Anomalies 11
Viewing Service Anomalies 12

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 ( ):

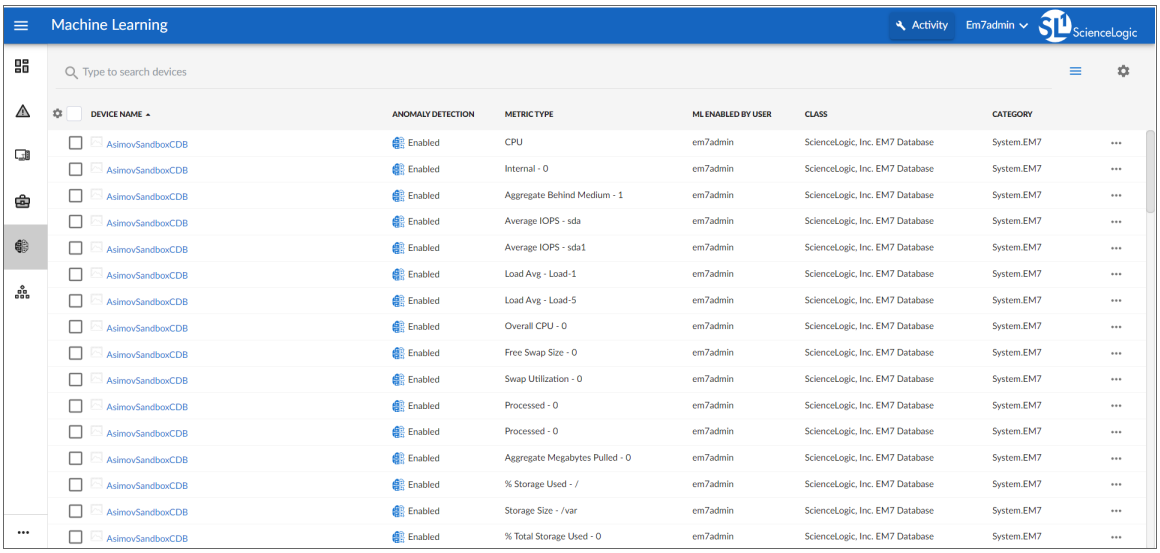

For each device in the list, the **Machine Learning** page displays the following information:

- ***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 whether anomaly detection is enabled or disabled for the device and metric.
- ***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's Device Class.
- ***Category.*** Displays the device's Device Category.

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 (and metric, if applicable) on which you want to enable machine learning.
3. Click the **[Actions]** icon (⋯) for that device and metric and then select *Enable*. The **Select Metric to Enable Machine Learning** modal page appears.

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

4. Click on the ***Metric*** drop-down field. You can either 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 as "enabled" on the **Machine Learning** page.

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

Enabling Machine Learning in the Device Investigator

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

1. Click the Devices icon (). The **Devices** page displays.
2. On the **Devices** page, click the ***Device Name*** for the device on which you want to enable anomaly detection. The **Device Investigator** page displays.
3. 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.

4. On the **[Machine Learning]** tab of the **Device Investigator** page, 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.
5. Click on the ***Metric*** drop-down field. You can either 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.
6. Click the name of the metric on which you want to enable machine learning for the device.
7. 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.
8. Click **[Enable Machine Learning]**. 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. Click on the ***Metric*** drop-down field. You can either 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** page, you can view a list of machine learning metrics that are enabled for the device:

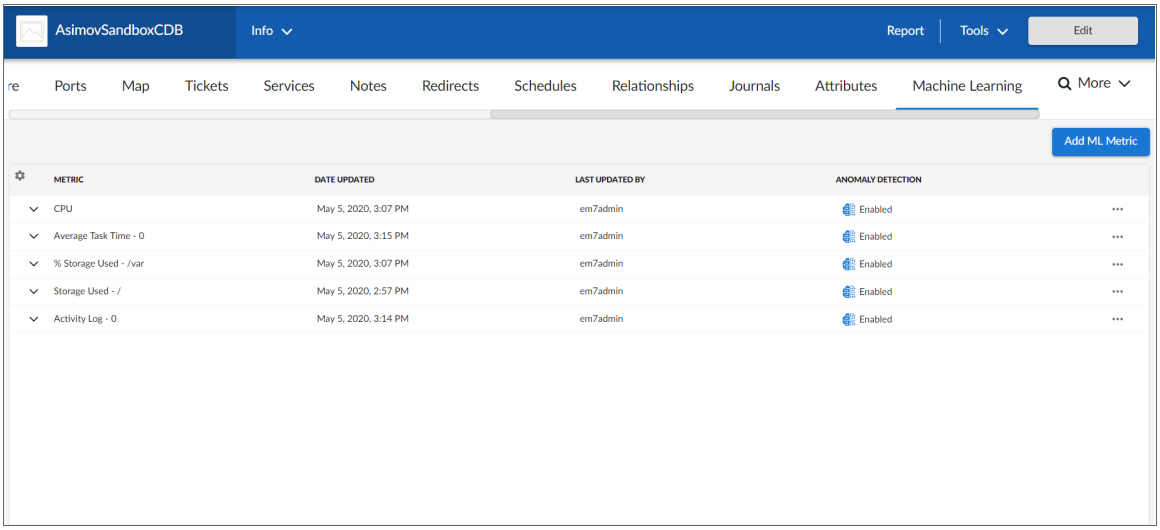

For each machine learning metric, the **[Machine Learning]** tab displays the following information:

- ***Metric***. The metric label.
- ***Date Updated***. The date and time at which machine learning was enabled for the metric or the metric's machine learning settings were last updated.
- ***Last Updated By***. The username of the user who enabled machine learning for the metric or last updated the metric's machine learning settings.
- ***Anomaly Detection***. Indicates whether anomaly detection is enabled or disabled for the metric.

Click the expand icon (⌄) next to a metric label to open a graph showing the anomalies detected for the device metric during the previous three hours. This graph displays:

- 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.

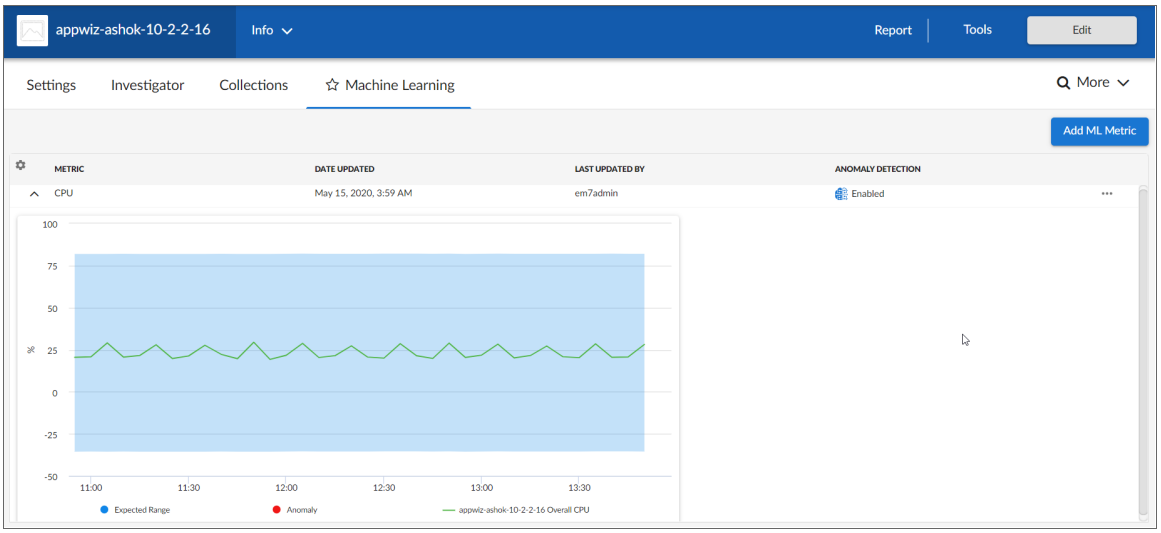

TIP: You can hover over any part of the graph with your mouse to see the exact historical values for each polling cycle. Also, you can zoom in on a shorter time frame by clicking and dragging your mouse over the part of the chart representing that time frame, and you can return to the original time span by clicking the **[Reset zoom]** button.

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

Viewing 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.

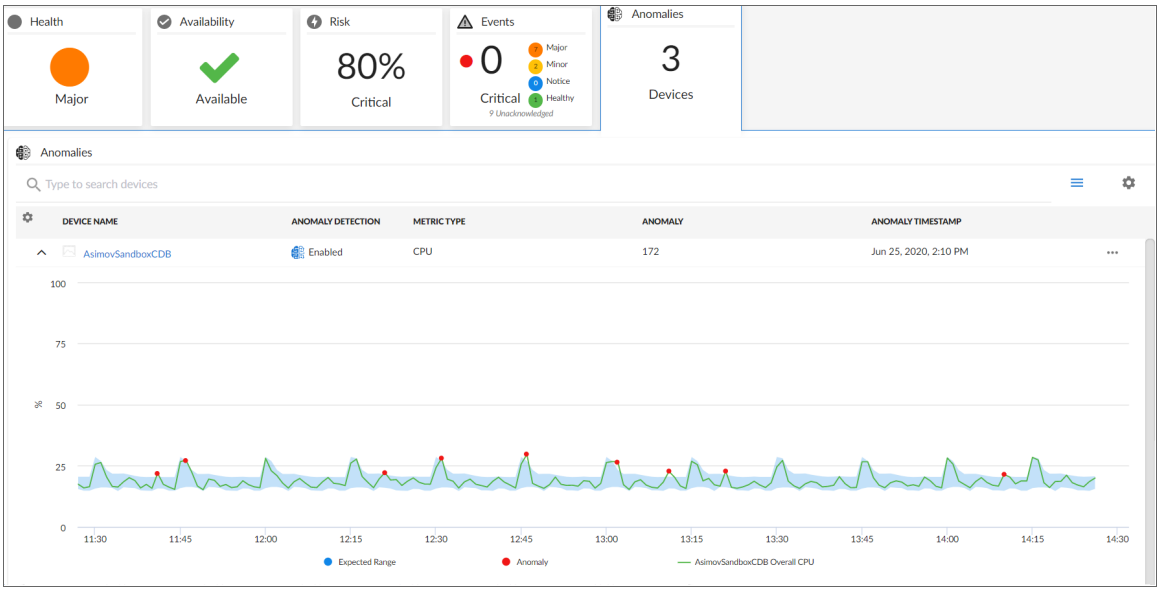

To view the **Service Investigator** page, select a service from the list of business, IT, and device services on the **Business Services** page ( ). The **[Overview]** tab opens by default. This tab provides a single-page view of the selected service, including key metrics, events, and anomalies that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

The **Anomalies** widget displays the following information about each device and metric in the list:

- ***Device Name***. Displays the name of the device. Click the hyperlink to go to the **[Machine Learning]** tab of the **Device Investigator** for the device.
- ***Anomaly Detection***. Indicates whether anomaly detection is enabled or disabled for the device and metric.
- ***Metric Type***. Indicates the metric that SL1 is evaluating for anomalies on the device.
- ***Anomaly***. The most recent anomaly value for the metric listed in the *Metric Type* column.
- ***Anomaly Timestamp***. The time at which the most recent anomaly occurred.

From the **Anomalies** widget, you can do the following:

- Use the search field to search for specific devices.
- Click the expand icon ( ) next to the device name to open a graph showing the anomalies detected for the device during the previous four hours. This graph displays:
    - 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 hover over any part of the graph with your mouse to see the exact historical values for each polling cycle. Also, 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.

- Enable or disable anomaly detection on the devices listed by clicking the **[Actions]** icon (⋯) and selecting *Enable* or *Disable*.

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

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 includes the following topics:

Creating an Event Policy for Anomalies *15*
Using Anomaly-related Events to Trigger Automated Run Book Actions *17*

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 **Event Policies** page (Events > Event Policies).
2. In 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":

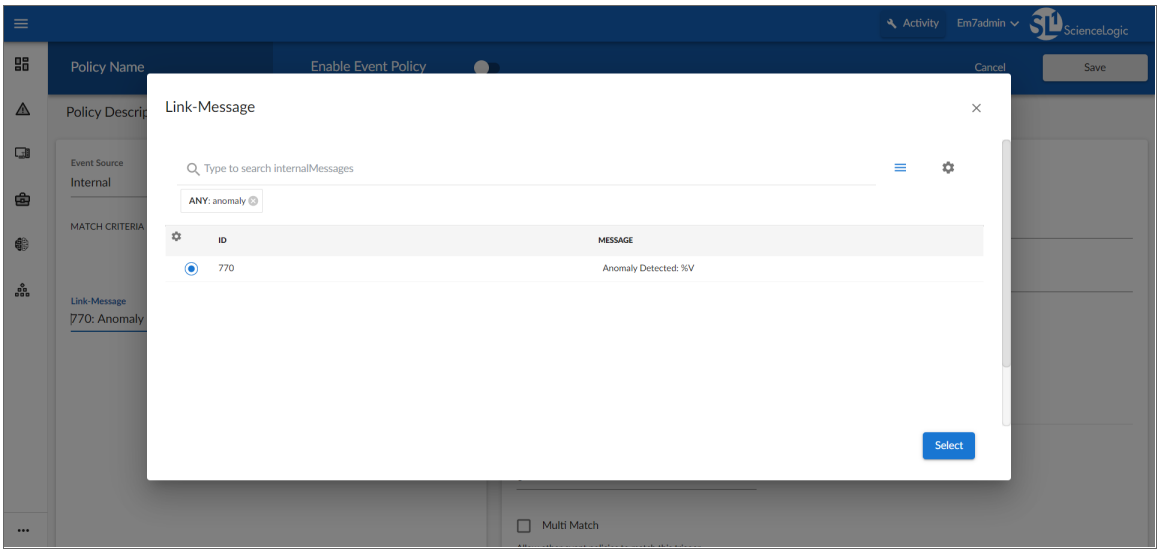

8. Click the radio button for the message "Anomaly Detected: %V", and then click **[Select]**.
9. Complete the remaining fields and tabs in the **Event Policy Editor** based on the specific parameters that you want to establish for the event. For more information about the fields and tabs in the **Event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

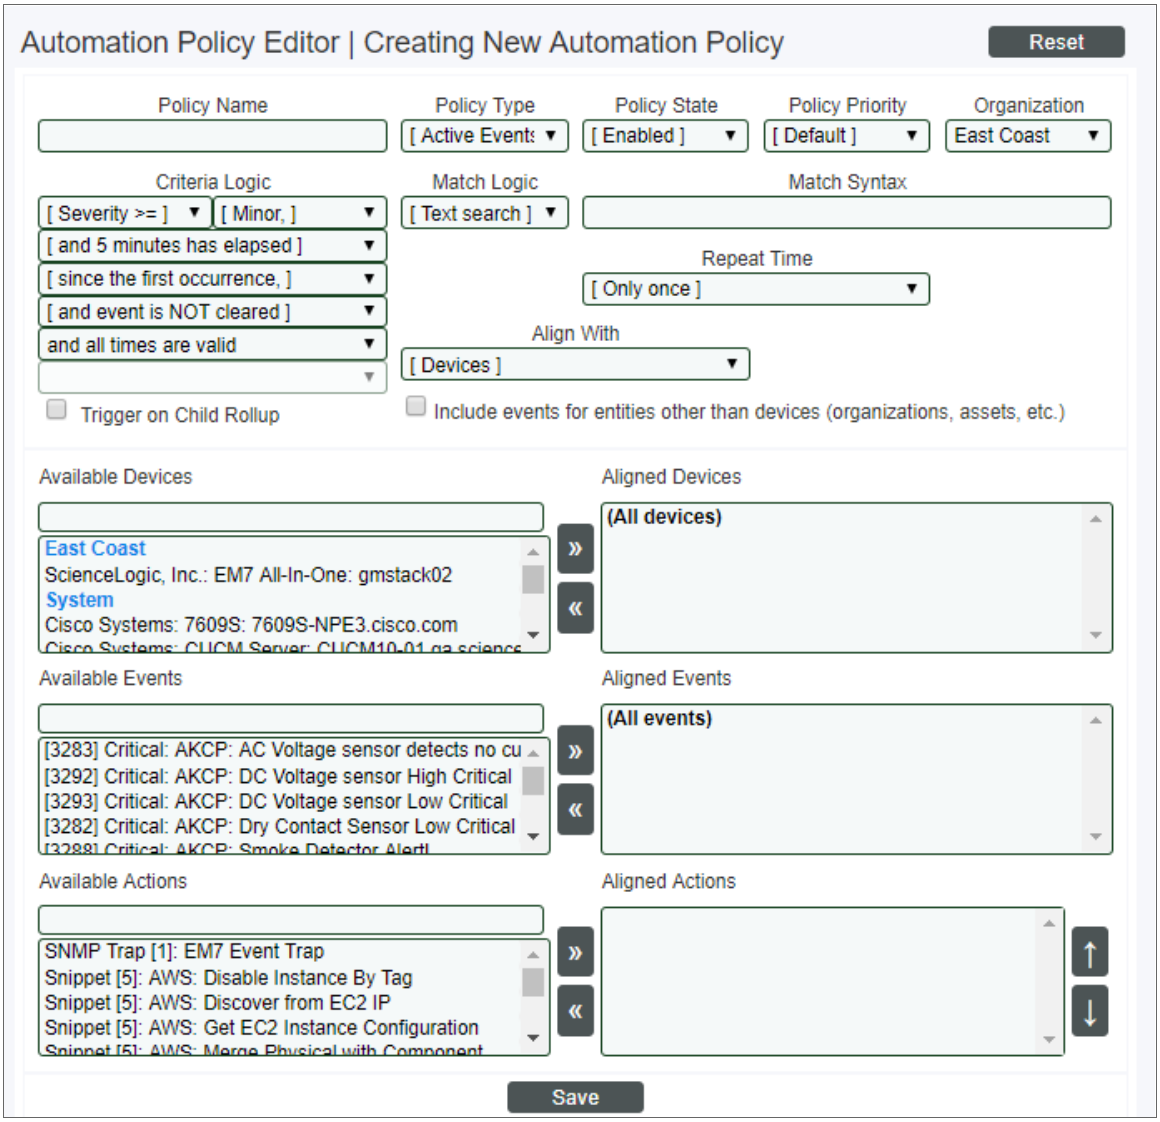

3. In the ***Available Events*** field, select an anomaly-related event policy and then click the right-arrow icon to move it to the ***Aligned Events*** field. For more information about anomaly-related events, see the section on *Creating an Event Policy for Anomalies*.
4. Complete the remaining fields 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 the chapter on "Automation Policies" in the ***Run Book Automation*** manual.
5. When you are finished, click **[Save]**.

© 2003 - 2020, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

Copyrights and Trademarks

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

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

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

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

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

Other

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

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

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

International: +1-703-354-1010