ScienceLogic

IT Services

ScienceLogic Version 8.6.0

Table of Contents

Chapter

1

What Are IT Services?

Overview

An IT Service is a technical service that is provided to internal or external customers. Some examples of IT Services include Internet access, website hosting, server co-location, remote backups, and remote storage. Usually an IT Service includes an associated Service Level Agreement (SLA) that specifies the terms of the service.

An IT Service policy allows you to define an IT Service, specify the devices that are included in the IT Service, and monitor the state, availability, and risk of the IT Service. The ScienceLogic platform evaluates the current state, availability, and risk of an IT Service based on user-defined metrics that trigger user-defined events about the IT Service. You can define how often the platform evaluates the state, availability, and risk of each IT Service. When the platform evaluates the state of an IT Service, the platform generates a default event that specifies the state of the IT Service.

You can define metrics based on any performance data collected by the ScienceLogic platform, including device availability, device latency, CPU usage, memory usage, swap usage, interface utilization, data collected by a Dynamic Application, and data about network interfaces, TCP/IP ports, system processes, Windows services, email round-trip time, web-content, SOAP/XML transactions, and DNS availability. You can specify that the platform should evaluate the metric against all devices in the IT service or against one or more subsets in the IT service.

When the ScienceLogic platform evaluates a metric, it performs an aggregation, that is, the platform evaluates the data for all devices specified in the definition of the metric, over a specified time period (the *Aggregation Frequency*). Depending on the definition of the metric, the platform can calculate the average, maximum, minimum, sum, or standard deviation.

You can also create dashboards for IT Services that display information about the state, availability, risk, events, metrics, and other information about an IT Service. IT Service dashboards are defined in the **IT Service Dashboards** page (Registry > IT Services > IT Service Dashboards).

To define an IT Service policy in the ScienceLogic platform, you must perform the following tasks:

1. ***Define a service name and basic settings***. For example, we could define an IT Service policy that monitors Email service. We could call this IT Service "Email". The basic settings for an IT Service include how often the ScienceLogic platform will evaluate the state, availability, and risk of the IT Service and the data retention settings for the metrics associated with the IT Service.

2. ***Define a list of devices (the model) for the IT Service that includes all the devices associated with the IT Service***. For example, if you want to monitor Email service, you could create a device group that includes Exchange servers, DNS servers, and devices that run Email round-trip policies. You can manually assign devices to the IT Service, or you can use membership rules, like you would for a dynamic device group.

3. ***Optionally, define device subsets***. You can manually assign devices to a subset, or you can use membership rules, like you would for a dynamic device group. For example, you could define two subsets: Exchange Servers, defined by device class, and DNS servers, defined by the ports that are open on each device.

4. ***Define metrics. A metric is based on your business processes and examines all devices or one or more subsets to evaluate the state of the IT Service***. For each IT Service, the ScienceLogic platform provides a default metric called *Average Device Availability* , based on the availability of all devices in the IT Service. You can define additional metrics, based on any performance data collected by the platform, including availability, latency, CPU usage, memory usage, swap usage, interface utilization, data collected by a Dynamic Application, and data about network interfaces, TCP/IP ports, system processes, Windows services, Email round-trip time, web content, SOAP/XML transactions, and DNS availability. You can specify that the platform should evaluate the metric against all devices in the IT Service or against one or more subsets in the IT Service.

5. ***Define Key Metrics.*** Key Metrics are the standard method for describing the status of an IT Service. Key metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data.The Key Metrics are Health, Availability, and Risk. When you define a Key Metric, you are specifying how the value for a metric you created in step 4 translates to one of the standard Key Metric values. By default, all three Key Metrics are based on the default *Average Device Availability* metric.

6. ***Define alerts and associated events***. Each alert and its associated event is triggered by a metric. Although not all metrics must trigger an alert, all alerts and events for an IT Service are triggered by a metric.

Who Should Read this Manual?

This manual is intended for users who define policies and users who monitor IT Services.

This manual is will explain how to define a policy to monitor an IT Service. This manual will also explain how to view information about an IT Service in the ScienceLogic platform and create IT Service Dashboards.

Chapter

2

Creating, Editing, and Deleting IT Services

Overview

To define an IT Service policy in the ScienceLogic platform, you must perform the following tasks:

1. **Define a service name and basic settings**. For example, we could define an IT Service policy that monitors Email service. We could call this IT Service "Email". The basic settings for an IT Service include how often the ScienceLogic platform will evaluate the state, availability, and risk of the IT Service and the data retention settings for the metrics associated with the IT Service.

2. **Define a list of devices (the model) for the IT Service that includes all the devices associated with the IT Service**. For example, if you want to monitor Email service, you could create a device group that includes Exchange servers, DNS servers, and devices that run Email round-trip policies. You can manually assign devices to the IT Service, or you can use membership rules, like you would for a dynamic device group.

3. **Optionally, define device subsets**. You can manually assign devices to a subset, or you can use membership rules, like you would for a dynamic device group. For example, you could define two subsets: Exchange Servers, defined by device class, and DNS servers, defined by the ports that are open on each device.

4. **Define metrics. A metric is based on your business processes and examines all devices or one or more subsets to evaluate the state of the IT Service**. For each IT Service, the ScienceLogic platform provides a default metric called *Average Device Availability* , based on the availability of all devices in the IT Service. You can define additional metrics, based on any performance data collected by the platform, including availability, latency, CPU usage, memory usage, swap usage, interface utilization, data collected by a Dynamic Application, and data about network interfaces, TCP/IP ports, system processes, Windows services, Email round-trip time, web content, SOAP/XML transactions, and DNS availability. You can specify that the platform should evaluate the metric against all devices in the IT Service or against one or more subsets in the IT Service.

5. **Define Key Metrics.** Key Metrics are the standard method for describing the status of an IT Service. Key metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data.The Key Metrics are Health, Availability, and Risk. When you define a Key Metric, you are specifying how the value for a metric you created in step 4 translates to one of the standard Key Metric values. By default, all three Key Metrics are based on the default *Average Device Availability* metric.

6. **Define alerts and associated events**. Each alert and its associated event is triggered by a metric. Although not all metrics must trigger an alert, all alerts and events for an IT Service are triggered by a metric.

This chapter will describe:

- All the steps for creating an IT Service.
- Scheduling Downtime for an IT Service.
- Editing an IT Service Policy.
- Deleting an IT Service Policy.

Prerequisites

To create an IT Service policy, you should first determine:

- the devices that affect the service.
- the conditions that you want to monitor. These will be based upon your business processes. For example, if you provide Email service, then a failure of your primary SMTP server and backup SMTP server would constitute a critical state.

Basic Mode and Advanced Mode

When you define the *basic properties for an IT Service policy*, you can specify **Configuration Mode**. Choices are:
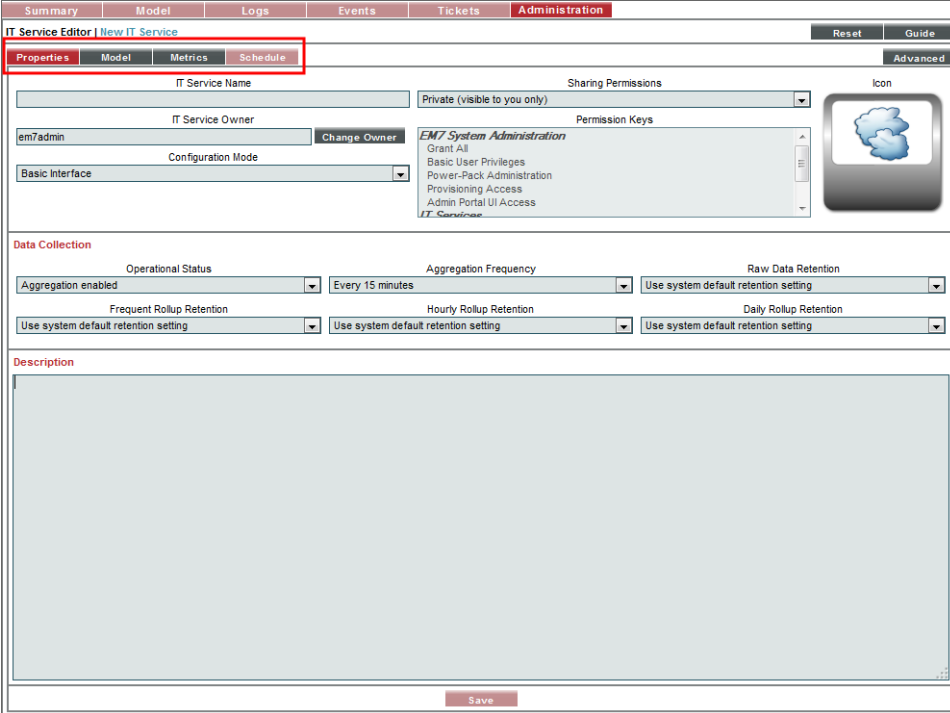
- **Basic Interface**. The *Basic Interface* allows you to quickly set up an IT Service policy.
- **Advanced Interface**. The *Advanced Interface* displays additional tabs for more granular control when defining metrics, alerts, and device groups.

If you are unsure, you can select **Basic Interface** and use the **[Advanced]** button to change to *Advanced Interface* if necessary.You can use the **[Advanced]** button from any sub-tab in the **IT Service Editor** page to toggle between **Basic Interface** and **Advanced Interface**.

Creating, Editing, and Deleting IT Services

Sub-tabs

To create or edit an IT Service policy, you must access the **IT Service Editor** page (Registry > IT Services > IT Service Manager > Create).

The **IT Service Editor** page includes the following sub-tabs, in the upper-left, under the page title:

NOTE: The *editing mode* (either *Basic Interface* or *Advanced Interface*) affects the sub-tabs that appear.

- *Properties*. This sub-tab appears in both the *Basic Interface* mode and the *Advanced Interface* mode. This sub-tab allows you to define the basic parameters (name, access permissions, collection frequency, data retention) for an IT Service policy. For details, see the section on *Defining the Basic Properties of an IT Service Policy*.

- *Model*. This sub-tab appears in both the *Basic Interface* mode and the *Advanced Interface* mode. This tab allows you to define the list of devices (device group) to include in an IT Service policy and also allows you to define Device Subsets (smaller groups of devices within the device group), and Dependencies (relationship between one or more IT Service policies). For details, see the section on *Defining Device Groups, Subsets, and Dependencies*.

- *Collection*. This sub-tab appears only in the *Advanced Interface* mode. This sub-tab allows you to fine-tune the collection and aggregation for each metric. For details, see the section on *Editing Collection and Aggregation for a Metric in Advanced Mode*.

- *Metrics*. This sub-tab appears in both the *Basic Interface* mode and the *Advanced Interface* mode. This sub-tab allows you to define the performance parameters you want to monitor for a specific IT Service and also define alerts for that parameter. For example, you might want to monitor the availability of the DNS service on each DNS server and send an alert if the DNS service is not available. For details, see the section on *Defining Metrics*.

- *Alerting*. This sub-tab appears only in the *Advanced Interface* mode. This sub-tab allows you to add details to an existing alert. For details, see the section on *Editing Alerts and Events in Advanced Mode*.

- *Schedule*. This sub-tab appears in both the *Basic Interface* mode and the *Advanced Interface* mode. This sub-tab allows you to put the IT Service policy into "maintenance mode". The ScienceLogic platform will continue to collect information from the devices in the IT Service but will not collect and aggregate information specific to the IT Service policy. During maintenance mode, the ScienceLogic platform will not evaluate or generate events about the IT Service policy. For details, see the section on *Scheduling Downtime for One or More Devices in an IT Service*.
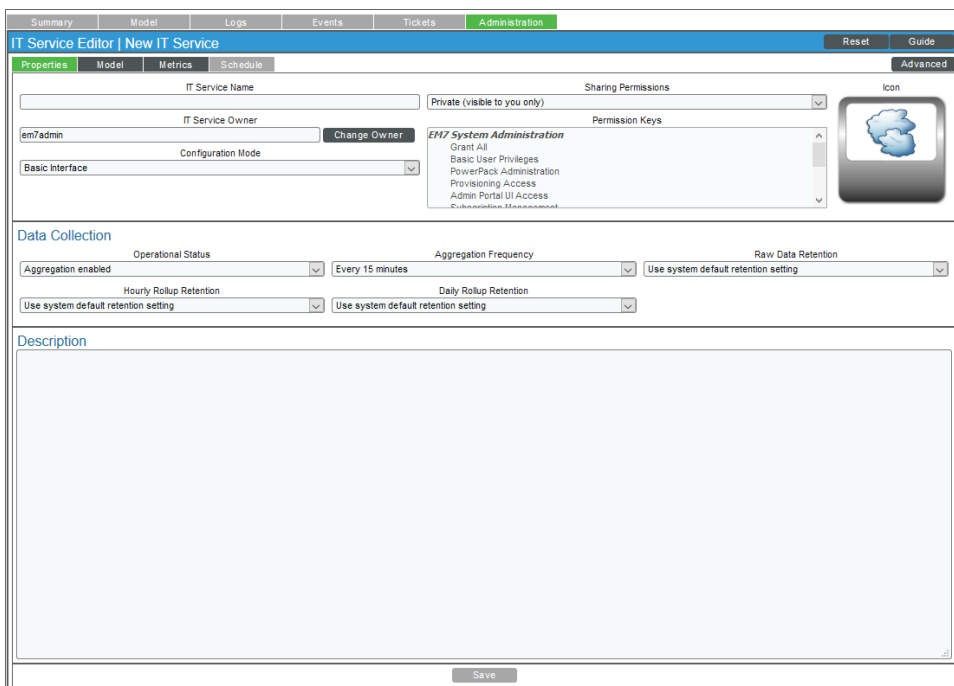
The following *additional fields* appear when you select *Advanced Interface*:

- In the *Model* sub-tab, the following additional panes appear in *Advanced Interface* mode:

  - *Service Dependencies*. A dependency is another IT Service policy.
  - *Device Subsets*. A sub-group of devices, selected from the list of all devices in the IT Service policy. You can manually assign devices to a subset, or you can use membership rules, like you would for a dynamic device group. For example, you could define a list of Exchange Servers, defined by device class, and then you could define a subset of unhealthy Exchange Servers, based on device state.

- In the *Metrics* sub-tab, when you edit an existing metric (select the wrench icon), the **Service Metric Editor** modal page appears.

  - In Advanced mode, the **Service Metric Editor** modal page allows you to define a formula for the metric, using aggregation objects and arithmetic operators (for example, devices running DNS service / devices available). In Advanced mode, you can define a metric using a page like the one for alert objects or presentation objects.
  - In Advanced mode, you must separately edit the collection object (defines data to collect and devices to collect from), the aggregation object (defines how frequently to "crunch" collected data and what to calculate from crunched data) , and the alert. You cannot edit the metric, its objects, and its alerts from a single page, like you can in Basic mode. To edit the collection object and the aggregation object, select the *Collection* sub-tab. To edit the alert, select the *Alerting* sub-tab

Defining the Basic Properties of an IT Service Policy

To define the basic properties of an IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the **[Create]** button.

3. Notice that by default the **[Administration]** tab is selected and the **[Properties]** sub-tab is selected. The **IT Service Editor** page is displayed:

4. To create a new IT Service policy, supply values in the following fields:

   - *IT Service Name*. Name of the IT Service policy.

   - *IT Service Owner*. Automatically populated with your username.

   - *Change Owner*. The **[Change Owner]** button allows the user to change the owner of an IT Service. Selecting the **[Change Owner]** button opens the **Select IT Service Owner** modal page where a new owner can be chosen by selecting the blocks icon (). The owner of an IT Service policy and the *Sharing Permissions* setting defines the users that can view and use the IT Service policy.

- **Configuration Mode**. Select *Basic Interface or Advanced Interface*.

  - The *Basic Interface* allows you to quickly setup an IT Service policy.
  - The *Advanced Interface* displays additional tabs, for more granular control when defining metrics, alerts, and device groups.
  - If you are unsure, you can select *Basic Interface* and use the **[Advanced]** button to change to *Advanced Interface* if necessary. By default, each tab in the **IT Service Editor** page will use the configuring mode you select in this field. However, you can change the configuration mode for each individual tab.

- **Sharing Permissions**. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the ScienceLogic platform where the IT Service is visible. Choices are:

  - *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the owner.
  - *Private (visible only to you)*. The IT Service policy can be viewed only by the owner and Administrator users.

- **Permission Keys**. If you selected *Shared with users in your organization* in the **Sharing Permissions** field, you can specify the required Permission Keys that an user must have to view the IT Service policy.

- **Operational Status**. Specifies whether aggregation is enabled or disabled.

- **Aggregation Frequency**. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value.

- **Raw Data Retention**. Specifies how long the ScienceLogic platform should store the raw data for the IT Service policy. You can accept the default, system-wide setting from the **Data Retention Settings** page (System > Settings > Data Retention) or you can specify a custom value that applies only to this IT Service policy. The custom value will override the system-wide value.

- **Hourly Rollup Retention**. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. You can accept the default, system-wide setting from the **Data Retention Settings** page (System > Settings > Data Retention) or you can specify a custom value that applies only to this IT Service policy. The custom value will override the system-wide value.

- **Daily Rollup Retention**. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. You can accept the default, system-wide setting from the **Data Retention Settings** page (System > Settings > Data Retention) or you can specify a custom value that applies only to this IT Service policy. The custom value will override the system-wide value.

- **Description**. Description of the IT Service policy.

5. Select the **[Save]** button to save the properties for the new IT Service policy.

Creating a New IT Service Based on an Existing IT Service

To define a new IT Service based on an existing IT Service, perform the following steps:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. In the **IT Service Manager** page, find the IT Service you want to use as a template to create a new IT Service. Select its wrench icon (🔧).

3. In the **IT Service Editor** page, supply a new value in the *IT Service Name* field.

> **CAUTION:** If you do not enter a new value in the *IT Service Name* field, the ScienceLogic platform will save the new IT Service under the same name as the existing IT Service. In some cases, this could make management of IT Services difficult. Best practice is to supply a new name for the new IT Service.

4. Edit one or more fields, if desired.

5. Select the **[Save As]** button.

6. The new IT Service will appear in the **IT Service Manager** page.

Defining the List of Devices for an IT Service

After *Defining the Basic Properties of an IT Service Policy*, you must next determine the devices to include in your IT Service policy. You do this in the **[Model]** sub-tab. When you define the list of devices to include in your IT Service policy, that list of devices appears as a device group throughout the ScienceLogic platform.

For example, if you want to monitor Email service, you could create a list of devices that includes Exchange servers, DNS servers, and devices that run Email round-trip policies.

You can manually assign devices and device groups to the IT Service device group, or you can use membership rules, like you would for a dynamic device group.

You can define optional device subsets. A device subset is a sub-group of devices from the list of all devices in the IT Service policy. You define device subsets in the **[Model]** sub-tab. A Device subset is helpful if you want to examine only some devices for a particular metric, for example, all database servers in a group of servers.

You can manually assign devices to a subset, or you can use membership rules, like you would for a dynamic device group. For example, you could define a list of Exchange Servers, defined by device class, and also define a device subset of unhealthy Exchange servers, based on device state.

Defining the List of Devices

There are three ways to add one or more devices to the list of devices for the IT Service policy:

- Add a device group to the list of devices for the IT Service policy.
- Add a static list of one or more devices to the list of devices for the IT Service policy.
- Add a dynamic list of one or more devices to the list of devices for the IT Service policy.

To create the list of devices for the IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Either create a new IT Service policy (**[Create]** button) or edit an existing policy (wrench icon 🔧).
2. After *Defining the Basic Properties of an IT Service Policy*, select the **[Model]** sub-tab.
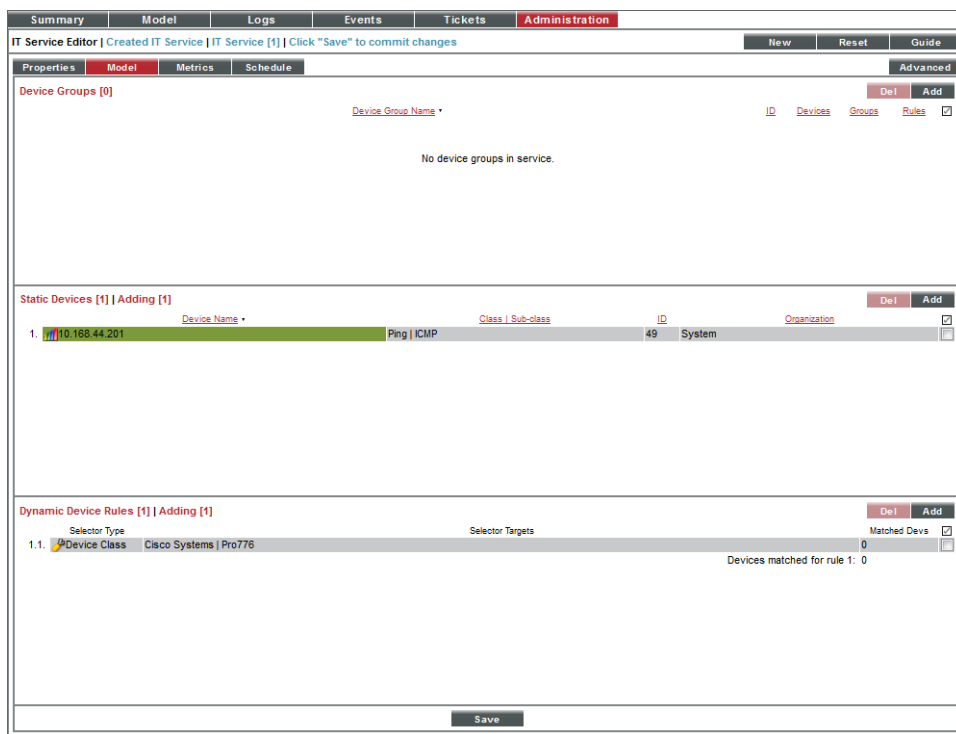
3. The following sections explain how to add device groups, a static list of devices, or a dynamic list of devices to an IT Service policy.

Adding Device Groups to the IT Service Policy

1. To add a device group to the list of devices for the IT Service policy, go to the **Device Groups** pane.
2. Select the **[Add]** button.

3. The **Device Group Alignment** modal page appears and displays a list of all devices in the ScienceLogic platform.

4. In the **Device Group Alignment** modal page, select the checkbox of each device group you want to include in the IT Service policy.

5. Select the **[Add/Remove]** button in the lower right.

6. The selected device groups will appear in the **Device Groups** pane.

7. To remove a device group from the list of devices for the IT Service policy, return to the **IT Service Editor** page, select the **[Model]** tab, select the checkbox for the device group, and then select the **[Del]** button.

Adding a Static List of Devices to the IT Service Policy

1. To add a static list of one or more devices to the list of devices for the IT Service policy, go to the **Static Devices** pane.

2. Select the **[Add]** button.

3. The **Device Alignment** modal page appears and displays a list of all devices in the ScienceLogic platform:

4. In the **Device Alignment** modal page, select the checkbox of each device group you want to include in the IT Service policy. Select the **[Add/Remove]** button in the lower right.

5. The selected devices will appear in the **Static Devices** pane.

6. To remove a device from the list of devices for the IT Service policy, go to the **IT Service Editor** page, select the **[Model]** tab, select the checkbox for the device group, and then select the **[Del]** button.

Adding a Dynamic List of Devices to the IT Service Policy
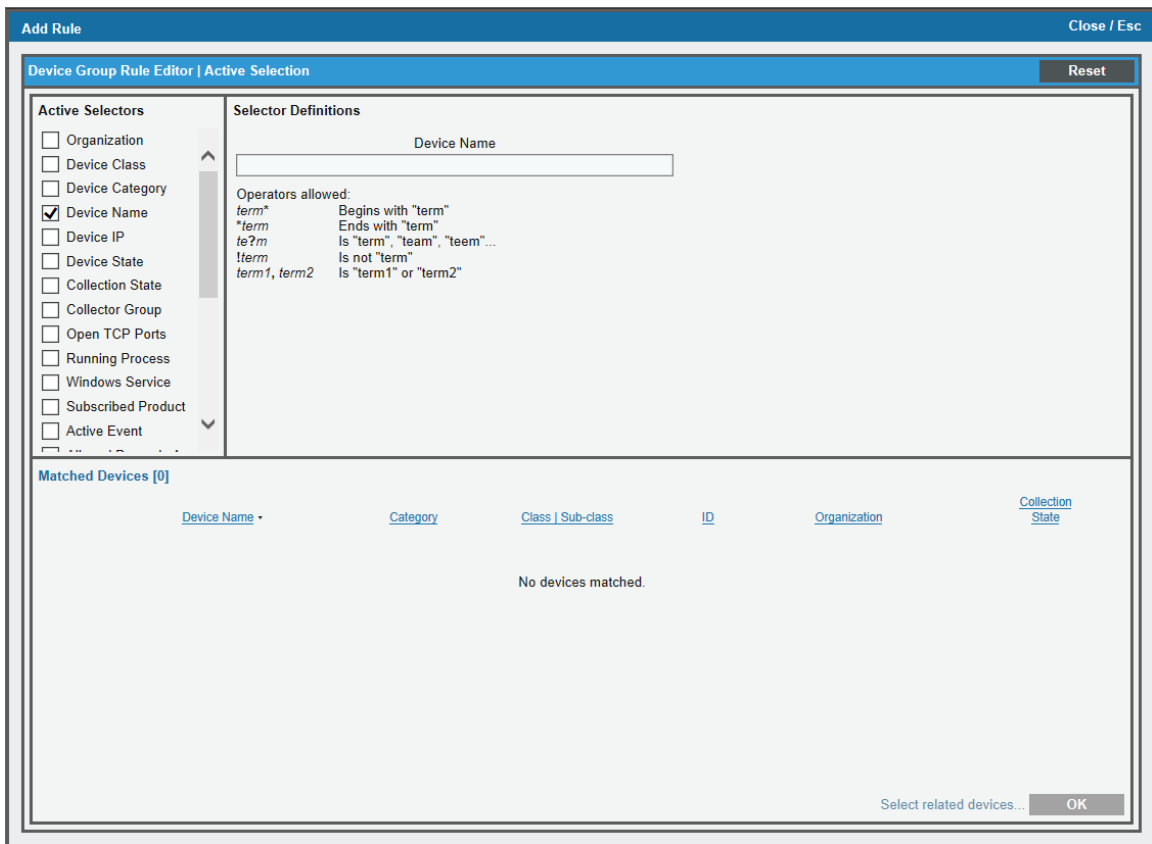
1. To add a dynamic list of one or more devices to the list of devices for the IT Service policy, go to the **Dynamic Device Rules** pane.

2. Select the **[Add]** button. The **Device Group Rule Editor** modal page appears.

3. In the **Device Group Rule Editor** modal page, define one or more criteria to include in the rule.

4. Select criteria from the **Active Selectors** field.

5. The selection appears in the **Selector Definitions** pane.

6. If you unselect a criteria in the **Active Selectors** field, the selection no longer appears in the **Selector Definitions** pane.

7. In the **Selector Definitions** pane, each criteria can include a list of selections and/or operators and wildcards:

    - *Organization*. Displays a list of all organizations in the platform.

        ○ To filter the list, you can enter a string in the field under the title.

        ○ If you select one or more organizations, the platform will search for devices that are members of at least one the selected organization(s) and include those devices in the device group. For example, if you select two organizations, all devices from each organization will be included in the device group.

        ○ If you select the *Invert* checkbox, the platform will search for devices that are *not* members of the selected organization(s) and include those devices in the device group.

- **Device Class**. Displays a list of all device classes in the platform.

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more device classes, the platform will search for devices that are members of the selected device class(es) and include those devices in the device group. For example, if you select two device classes, all devices from each device class will be included in the device group.

  - If you select the **Invert** checkbox, the platform will search for devices that are **not** members of the selected device class(es) and include those devices in the device group.

- **Device Category**. Displays a list of all device categories in the platform.

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more device categories, the platform will search for devices that are members of the selected device categories and include those devices in the device group. For example, if you select two device categories, all devices from each device category will be included in the device group.

  - If you select the **Invert** checkbox, the platform will search for devices that are **not** members of the selected device categories and include those devices in the device group.

- **Device Name**. Displays a field in which you can enter a string. The platform will use the string to search for devices with matching device names. If you do not use wildcard characters, the platform will return only devices with a device name that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any device name that begins with "term".

  - **term*. The platform searches for any device name that ends with "term".

  - *te?m*. The platform searches for any device name that contains the text "te[any single character]m".

  - *!term*. The platform searches for any device name that does not include the text "term".

  - *term1, term2*. The platform searches for any device name that contains either "term1" or "term2".

- **Device IP**. Displays a field in which you can enter a string. The platform will use the string to search for devices with matching IP addresses. If you do not use wildcard characters, the platform will return only devices with a device IP that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any device IP that begins with "term".

  - **term*. The platform searches for any device IP that ends with "term".

  - *te?m*. The platform searches for any device IP that contains the text "te[any single character]m".

  - *!term*. The platform searches for any device IP that does not include the text "term".

  - *term1, term2*. The platform searches for any device IP that contains either "term1" or "term2".

  - If you select the **Invert** checkbox, the platform will search for devices that do **not** have a matching IP address and include those devices in the device group.

- **Device State**. Displays a list of all device states in the platform (Notice, Healthy, Minor, Major, Critical). Each device's state is the same as the highest severity event associated with the device..

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more device states, the platform will search for devices that are members of the selected device states and include those devices in the device group. For example, if you select two device states, all devices from each device state will be included in the device group.

  - If you select the *Invert* checkbox, the platform will search for devices that do *not* have the selected device state and include those devices in the device group.

- **Collection State**. Displays a list of all collection states in the platform (Active, NOT Active, User-Disabled, NOT User-Disabled, Unavailable, NOT Unavailable, Maintenance, NOT Maintenance, System-Disabled, NOT System-Disabled)..

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more collection states, the platform will search for devices that are members of the selected device states and include those devices in the device group. For example, if you select two collection states, all devices with the first collection state and all devices with the second collection state will be included in the device group.

- **Collector Group**. Displays a list of all Collector Groups in the platform.

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more Collector Groups, the platform will search for devices that are members of the selected Collector Groups and include those devices in the device group. For example, if you select two Collector Groups, all devices from each Collector Group will be included in the device group.

  - If you select the *Invert* checkbox, the platform will search for devices that are *not* members of the select Collector Group(s) and include those devices in the device group.

- **Open TCP Ports**. Displays a list of all open TCP ports that the platform has discovered on at least one device.

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more ports, the platform will search for devices have that have those ports open and include those devices in the device group. For example, if you select two collection ports, all devices where the first port is open and all devices where the second port is open will be included in the device group.

  - If you select the *Invert* checkbox, the platform will search for devices that do *not* have the selected port(s) open and include those devices in the device group.

- **Running Process**. Displays a field in which you can enter a string. The platform will use the string to search for devices that are running a matching system process. If you do not use wildcard characters, the platform will return only devices running a process that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any process name that begins with "term".
  - **term*. The platform searches for any process name that ends with "term".
  - *te?m*. The platform searches for any process name that contains the text "te[any single character]m".
  - *!term*. The platform searches for any process name that does not include the text "term".
  - *term1, term2*. The platform searches for any process name that contains either "term1" or "term2".
  - If you select the **Invert** checkbox, the platform will search for devices that are **not** running the selected processes open and include those devices in the device group.

- **Windows Service**. Displays a field in which you can enter a string. The platform will use the string to search for devices that are running a matching Windows service. If you do not use wildcard characters, the platform will return only devices running a Windows services that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any Windows Service name that begins with "term".
  - **term*. The platform searches for any Windows Service name that ends with "term".
  - *te?m*. The platform searches for any Windows Service name that contains the text "te[any single character]m".
  - *!term*. The platform searches for any Windows Service name that does not include the text "term".
  - *term1, term2*. The platform searches for any Windows Service name that contains either "term1" or "term2".
  - If you select the **Invert** checkbox, the platform will search for devices that are **not** running a matching Windows service and include those devices in the device group.

- **Subscribed Product**. Displays a list of all product SKUs in the platform.

  - To filter the list, you can enter a string in the field under the title.
  - If you select one or more SKUs, the platform will search for devices that subscribe to at least one the selected SKU(s) and include those devices in the device group. For example, if you select two SKUs, all devices that subscribe to one of the two SKUs will be included in the device group.
  - If you select the **Invert** checkbox, the platform will search for devices that do **not** subscribe to the selected SKU(s) and include those devices in the device group.

- *Active Event*. Displays a list of all active events in the platform.

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more active events, the platform will search for devices for which that event is currently active and include those devices in the device group. For example, if you select two events, all devices for which one of the two events is active will be included in the device group.

  - If you select the *Invert* checkbox, the platform will search for devices for which the event is *not* currently active and include those devices in the device group.

- *Aligned Dynamic App*. Displays a list of all Dynamic Applications that are currently aligned with one or more devices.

  - To filter the list, you can enter a string in the field under the title.

  - If you select one or more aligned Dynamic Applications, the platform will search for devices that are aligned with those Dynamic Applications and include those devices in the device group. For example, if you select two Dynamic Applications, all devices that are aligned with each of the Dynamic Applications will be included in the device group.

  - If you select the *Invert* checkbox, the platform will search for devices that are *not* aligned with the selected Dynamic Application(s).

- *Asset Make*. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the *Make* field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have a *Make* field that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any asset make that begins with "term".

  - **term*. The platform searches for any asset make that ends with "term".

  - *te?m*. The platform searches for any asset make that contains the text "te[any single character]m".

  - *!term*. The platform searches for any asset make that does not include the text "term".

  - *term1, term2*. The platform searches for any asset make that contains either "term1" or "term2".

- *Asset Model*. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the *Model* field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have a *Model* field that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any asset model that begins with "term".

  - **term*. The platform searches for any asset model that ends with "term".

  - *te?m*. The platform searches for any asset model that contains the text "te[any single character]m".

  - *!term*. The platform searches for any asset model that does not include the text "term".

  - *term1, term2*. The platform searches for any asset model that contains either "term1" or "term2".

- **Asset Function**. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the **Function** field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have a **Function** field that exactly matches the string. You can use the following syntax in the field:

    ○ *term**. The platform searches for any asset function that begins with "term".

    ○ **term*. The platform searches for any asset function that ends with "term".

    ○ *te?m*. The platform searches for any asset function that contains the text "te[any single character]m".

    ○ *!term*. The platform searches for any asset function that does not include the text "term".

    ○ *term1, term2*. The platform searches for any asset function that contains either "term1" or "term2".

- **Asset Owner**. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the **Management Type** field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have a **Management Type** field that exactly matches the string. You can use the following syntax in the field:

    ○ *term**. The platform searches for any asset owner that begins with "term".

    ○ **term*. The platform searches for any asset owner that ends with "term".

    ○ *te?m*. The platform searches for any asset owner that contains the text "te[any single character]m".

    ○ *!term*. The platform searches for any asset owner that does not include the text "term".

    ○ *term1, term2*. The platform searches for any asset owner that contains either "term1" or "term2".

- **Asset Location**. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the **Facility/Data Center** field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have a **Facility/Data Center** field that exactly matches the string. You can use the following syntax in the field:

    ○ *term**. The platform searches for any asset location that begins with "term".

    ○ **term*. The platform searches for any asset location that ends with "term".

    ○ *te?m*. The platform searches for any asset location that contains the text "te[any single character]m".

    ○ *!term*. The platform searches for any asset location that does not include the text "term".

    ○ *term1, term2*. The platform searches for any asset location that contains either "term1" or "term2".

- ***Asset Serial***. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the ***Serial*** field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have a ***Serial*** field that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any asset serial that begins with "term".
  - **term*. The platform searches for any asset serial that ends with "term".
  - *te?m*. The platform searches for any asset serial that contains the text "te[any single character]m".
  - *!term*. The platform searches for any asset serial that does not include the text "term".
  - *term1, term2*. The platform searches for any asset serial that contains either "term1" or "term2".

- ***Asset Tag***. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value in the ***Asset Tag*** field in their asset records (the platform automatically creates an asset record for each device during nightly auto-discovery). If you do not use wildcard characters, the platform will return only devices that have an ***Asset Tag*** field that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any asset tag that begins with "term".
  - **term*. The platform searches for any asset tag that ends with "term".
  - *te?m*. The platform searches for any asset tag that contains the text "te[any single character]m".
  - *!term*. The platform searches for any asset tag that does not include the text "term".
  - *term1, term2*. The platform searches for any asset tag that contains either "term1" or "term2".

- ***Software Title***. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching title in the list of software defined in the **Licenses** tab in the asset record for the device. If you do not use wildcard characters, the platform will return only devices that have a ***Software Title*** field that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any software title that begins with "term".
  - **term*. The platform searches for any software title that ends with "term".
  - *te?m*. The platform searches for any software title that contains the text "te[any single character]m".
  - *!term*. The platform searches for any software title that does not include the text "term".
  - *term1, term2*. The platform searches for any software title that contains either "term1" or "term2".
  - If you select the ***Invert*** checkbox, the platform will search for devices that do ***not*** have a matching software title in the list of software and include those devices in the device group.

- **Software Version**. Displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching version number in the list of software defined in the **Licenses** tab in the asset record for the device. If you do not use wildcard characters, the platform will return only devices that have a **Software Version** field that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any software version that begins with "term".
  - **term*. The platform searches for any software version that ends with "term".
  - *te?m*. The platform searches for any software version that contains the text "te[any single character]m".
  - *!term*. The platform searches for any software version that does not include the text "term".
  - *term1, term2*. The platform searches for any software version that contains either "term1" or "term2".
  - If you select the **Invert** checkbox, the platform will search for devices that do **not** have a matching version number in the list of software and include those devices in the device group.

- **Custom Attribute**. The Active Selectors field includes an entry for each custom attribute defined in your ScienceLogic system. When you select a custom attribute, the Selector Definitions pane displays a field in which you can enter a string. The platform will use the string to search for devices that have a matching value for this custom attribute. If you do not use wildcard characters, the platform will return only devices with a custom attribute that exactly matches the string. You can use the following syntax in the field:

  - *term**. The platform searches for any attribute value that begins with "term".
  - **term*. The platform searches for any attribute value that ends with "term".
  - *te?m*. The platform searches for any attribute value that contains the text "te[any single character]m".
  - *!term*. The platform searches for any attribute value that does not include the text "term".
  - *term1, term2*. The platform searches for any attribute value that contains either "term1" or "term2".

8. After you have selected an Active Selector and the Selector Definitions, you can specify that you want to *include children devices, all descendent devices, parent devices, or all ancestor devices*. To do this, **do not select the [OK] button**. Instead, select the *Select related devices* link next to the **[OK]** button.

9. If you selected the *Select related devices* link, the *Relationship Selection* modal page appears. In this modal page you can select devices by relationship. The **Matched Devices** pane displays all the devices that match all the criteria in the rule. The list of devices changes as you add and remove criteria.

10. The **[OK]** button saves the new dynamic rule or saves changes to an existing dynamic rule.

11. The new rule appears in the **Dynamic Rules** pane.

NOTE: If a single dynamic rule includes multiple criteria, a device must match **all** the criteria to be included in the device group (like the SQL AND operator). If an IT Service policy includes multiple dynamic rules, a device must match only a single rule to be included in the device group (like the SQL OR operator). To view a list of devices that are currently included in the dynamic rules, select the wrench icon (🔧) for a dynamic rule.

12. To remove a dynamic rule from the list of devices for the IT Service policy, go to the **IT Service Editor** page, select the **[Model]** tab, select the checkbox for the rule, and then select the **[Del]** button.

The Relationship Selection Page

If you click the *Select related device link* in the **Device Group Rule Editor** page, the **Relationship Selection** page appears:

The **Relationship Selection** page includes the following panes:

- *Relationship Selector*. You can choose to include children devices, all descendent devices (children, grandchildren, great grandchildren, etc), parent devices, or all ancestor devices (parents, grandparents, great grandparents, etc). You can also choose to further filter by including only the children, descendents,

parents, and ancestors devices that are related to the "Seed Devices" through one or more selected Dynamic Applications.

> **NOTE**: For details on building relationships with Dynamic Applications, see the *Dynamic Application Development* manual.

- **Seed Devices**. This is the list of devices you defined in the **Device Group Rule Editor** page. You can include these devices in the device group or include only the children, descendants, parents, ancestors.
- **Matched Devices**. This is the list of devices that will be included in the device group when you click the **[OK]** button.

> **NOTE**: As the ScienceLogic platform discovers devices and component devices that meet the criteria for the dynamic device group, the platform will automatically add those devices and component devices to the device group.

Relationship Selectors

In this pane, you can specify devices to include in the device group based on their relationships to the Seed Devices:

- **Children of**. If you select this checkbox, all children devices of the Seed Devices are included in the device group.
- **Descendents of**. This checkbox is enabled only if you select the **Children of** checkbox. If you select this checkbox, all children devices, grandchildren devices, great grandchildren devices, etc of the Seed Devices are included in the device group.
- **Parents of**. If you select this checkbox, all parent devices of the Seed Devices are included in the device group.
- **Ancestors of**. This checkbox is enabled only if you select the **Parents of** checkbox. If you select this checkbox, all parent devices, grandparent devices, great grandparent devices, etc of the Seed Devices are included in the device group.
- **Related by**. This checkbox is enabled only if you select the **Children of** checkbox or the **Parents of** checkbox. You can further filter the devices in the device group by including only the children, descendents, parents, and ancestors devices that are related to the Seed Devices through one or more selected Dynamic Applications. If you select this checkbox, you can select one or more Dynamic Applications from the list of Dynamic Applications that can create relationships. Only those devices that meet all the criteria will be included in the device group.
- **Include seed devices**. This checkbox is enabled only if you select the **Children of** checkbox or the **Parents of** checkbox. If you select this checkbox, the related Seed Devices are included in the device group. Seed Devices with no relationships are not included in the group.

Seed Devices

This is the list of devices you defined in the **Device Group Rule Editor** page. You can include these devices in the device group or include only the children, descendants, parents, and ancestors of these devices.

For each Seed Device, the **Relationship Selection** page displays.

> **TIP:** To sort the list of devices, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again.

- *Device Name*. Name of the device. For devices running SNMP or with DNS entries, the name is discovered automatically. For devices without SNMP or DNS entries, the device's IP address will appear in this field.
- *Category*. The ScienceLogic category assigned to the device. Categories include servers, routers, switches, firewalls, printers, etc. The category is automatically assigned during discovery, at the same time as the Device-Class/Sub-Class.
- *Class / Sub-class*. The manufacturer (device class) and type of device (sub-class). The Device-Class/Sub-Class is automatically assigned during discovery, at the same time as the Category.
- *ID*. Device ID. This is a unique number automatically assigned to the device by the ScienceLogic platform.
- *Organization*. The organization to which the device is assigned.
- *Collection State*. Collection state can be one of the following:

  - *Active*. The platform is currently collecting data from the device.
  - *Disabled*. The platform is not currently collecting data from the device.
  - *Unavailable*. The device is currently unavailable, so the platform cannot collect data from the device at this time.
  - *Component Vanished*. The component device has vanished, i.e. is not currently being reported by its root device. The platform cannot collect data from the device at this time.

- *Tools*. Displays icons for managing devices. The choices are:

  - *Device Management* (📊). Leads to the **Device Summary** page, where you can see reports and logs related to the device. From the **Device Summary** page, you can also access the other pages in the Device Management tools.

Matched Devices

This is the list of devices that match all the criteria in the **Relationship Selection** page.

For each Matched Device, the **Relationship Selection** page displays:

> **TIP:** To sort the list of devices, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again.

- **Device Name**. Name of the device. For devices running SNMP or with DNS entries, the name is discovered automatically. For devices without SNMP or DNS entries, the device's IP address will appear in this field.

- **Category**. The ScienceLogic category assigned to the device. Categories include servers, routers, switches, firewalls, printers, etc. The category is automatically assigned during discovery, at the same time as the Device-Class/Sub-Class.

- **Class / Sub-class**. The manufacturer (device class) and type of device (sub-class). The Device-Class/Sub-Class is automatically assigned during discovery, at the same time as the Category.

- **ID**. Device ID. This is a unique number automatically assigned to the device by the ScienceLogic platform.

- **Organization**. The organization to which the device is assigned.

- **Collection State**. Collection state can be one of the following:

    - *Active*. The platform is currently collecting data from the device.

    - *Disabled*. The platform is not currently collecting data from the device.

    - *Unavailable*. The device is currently unavailable, so the platform cannot collect data from the device at this time.

    - *Component Vanished*. The component device has vanished, i.e. is not currently being reported by its root device. The platform cannot collect data from the device at this time.

- **Tools**. Displays icons for managing devices. The choices are:

    - *Device Management* (📊). Leads to the **Device Summary** page, where you can see reports and logs related to the device. From the **Device Summary** page, you can also access the other pages in the Device Management tools.

Click the **[OK]** button to accept all changes and exit the Relationship Selection page.

Defining a Device Subset

If you select the **[Advanced]** button or are already in Advanced mode you can define device subsets for your IT Service policy. A device subset is a smaller group of devices from the list of all devices in the IT Service policy.

You can apply metrics to these device subsets, instead of applying a metric to all devices in the IT Service Policy. For example, you might want to examine only the status of the database servers in a specific IT Service. In this case, you could create a device subset that contains only the database servers. You could later apply a metric only to the database servers (for example, monitoring latency of the database servers).

To define a device subset:

1. Go to **IT Service Manager** page (Registry > IT Services > IT Service Manager). Either create a new IT Service policy (**[Create]** button) or edit an existing policy (wrench icon 🔧).

Creating, Editing, and Deleting IT Services

2. After *Defining the Basic Properties of an IT Service Policy*, select the **[Model]** sub-tab. If you are not already in Advanced mode, select the **[Advanced]** button:
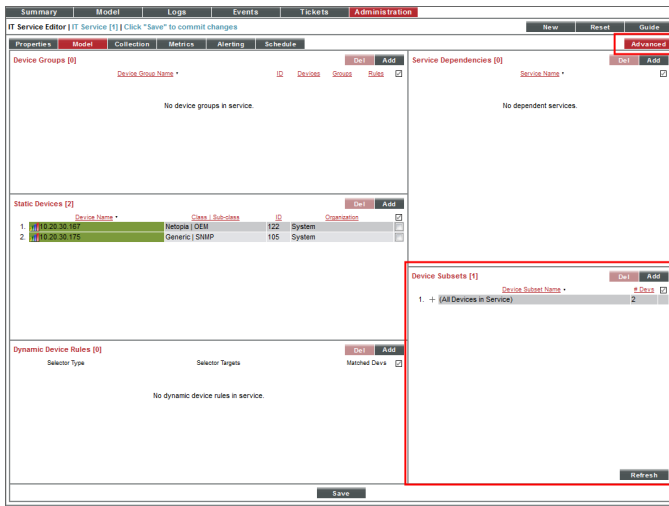

3. To add a device subset to the IT Service policy, go to the **Device Subsets** pane.
4. Select the **[Add]** button. The **Device Subset Editor** modal page appears.

5. To add a static list of one or more devices to the device subset, go to the **Static Devices** pane:

- Select the **[Add]** button.

- The **Device Alignment** modal page appears and displays a list of all devices in the ScienceLogic platform.

- In the **Device Alignment** modal page, select the checkbox of each device group you want to include in the IT Service policy. Select the **[Add/Remove]** button in the lower right.

- The selected devices will appear in the **Static Devices** pane.

- To remove a device from the list of devices in the device subset, select its checkbox and then select the **[Del]** button.

6. To add a dynamic list of one or more devices to the list of devices for the device subset, go to the **Dynamic Rules** pane:

- Select the **[Add]** button. The **Device Subset Rule Editor** modal page appears.

- In the **Device Subset Rule Editor** modal page, define one or more criteria to include in the rule. For details, see the section on *Adding a Dynamic List of Devices to the IT Service Policy*

- The new rule appears in the **Dynamic Rules** pane.

- To remove a dynamic rule from the list of devices for the IT Service policy, go to the **IT Service Editor** page, select the **[Model]**sub-tab, select the checkbox for the rule, and then select the **[Del]** button.

NOTE: If a single dynamic rule includes multiple criteria, a device must match *all* the criteria to be included in the device group (like the SQL AND operator). If an IT Service policy includes multiple dynamic rules, a device must match only a single rule to be included in the device group (like the SQL OR operator). To view a list of devices that are currently included in the dynamic rules, select the wrench icon (🔧) for a dynamic rule.


2

Defining a Service Dependency

A service dependency allows you to use a metric and list of devices from an external IT Service Policy when calculating the metrics in the current IT Service policy. In the *example in this manual*, the Acme company uses the device availability metric from an IT Service policy for Acme East Coast and the device availability metric from an IT Service policy for Acme West Coast to determine the device availability for all devices in all locations.

To define a Service Dependency:

1. Go to **IT Service Manager** page (Registry > IT Services > IT Service Manager). Either create a new IT Service policy (**[Create]** button) or edit an existing policy (wrench icon 🔧).

2. Select the **[Model]** sub-tab. If you are not already in Advanced mode, select the **[Advanced]** button.

3. To add a service dependency, go to the **Service Dependencies** pane.

4. Select the **[Add]** button. The **Service Dependency Alignment Editor** modal page appears.

Creating, Editing, and Deleting IT Services

- In the **Service Dependency Alignment Editor** modal page, select the checkbox of each external IT Service policy you want to reference in the current IT Service policy. Select the **[Add/Remove]** button in the lower right.

- The selected IT Service policies will appear in the **Service Dependencies** pane.

- To remove an IT Service policy from the list of service dependencies, select its checkbox and then select the **[Del]** button.

You can now use the metrics in each service dependency when defining the metrics for your IT Service policy. For an example, see the *example chapter* in this manual.

Defining Metrics

A metric is a performance measurement associated with an IT Service. One or more metrics are used to define the Service Health, Service Availability, and Service Risk for an IT Service.

For each IT Service, the ScienceLogic platform provides a default metric called *Average Device Availability*. This metric tells the ScienceLogic platform to collect availability data from all the devices in the IT Service and calculate the average availability.
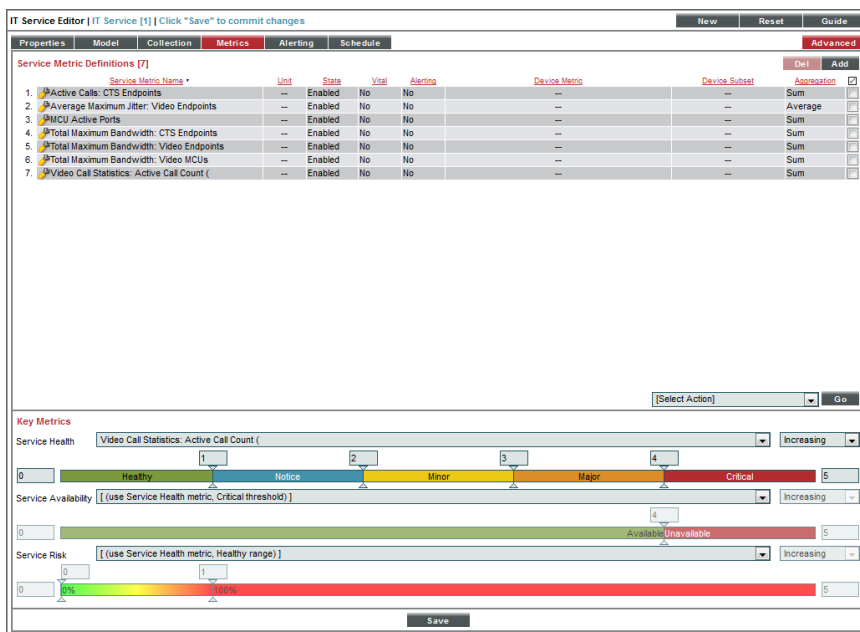
You can define additional metrics, based on any performance data collected by the platform, including:

- Device Availability
- Device Latency
- Overall CPU Usage
- Physical Memory Usage
- Swap Usage
- Device State (Condition of the device, based upon the most severe event generated by the device.)
- Device Count
- Presentation Objects from Dynamic Applications
- Network Interface Data
- TCP/IP Port Monitors
- System Process Monitors
- Windows Service Monitors
- Email Round-Trip Monitors
- Web Content Monitors
- SOAP/XML Transaction Monitors
- Domain Name Monitors
- An Aligned Service Dependency

> **NOTE**: When the ScienceLogic platform evaluates a metric, it performs an aggregation, that is, the platform evaluates the data for all devices specified in the definition of the metric, over a specified time period (the *Aggregation Frequency*). Depending on the definition of the metric, the platform can calculate the average, maximum, minimum, sum, standard deviation, count value or percentile for all devices specified in the definition.

To create a new metric:

1. Go to **IT Service Manager** page (Registry > IT Services > IT Service Manager). Either create a new IT Service policy (**[Create]** button) or edit an existing policy (wrench icon 🔧).

2. After *Defining the Basic Properties of an IT Service Policy* and *Defining the Devices for the IT Service Policy*, select the **[Metrics]** sub-tab.

3. Select the **[Add]** button. The **Service Metric Editor** modal page appears.

4. Supply a value in each field of the **Service Metric Editor** modal page.

- *Service Metric Name*. Enter a name for the metric.

- *Metric Classification*. Specifies whether the metric will be displayed in the **IT Service Summary** page in widgets that display vital metrics. Choices are:

  - *Service Vital Metric*. This metric will appear in widgets that display vital metrics.

  - *Standard Metric*. This metric will not appear in widgets that display vital metrics.

- *Active State*. Specifies whether the ScienceLogic platform should currently collect data for the metric and evaluate alerts for the metric. Choices are:

  - *Enabled*. The platform will collect data for the metric and evaluate alerts for the metric.

  - *Disabled*. The platform will not collect data for the metric and evaluate alerts for the metric.

- *Metric Type*. Specifies the type of performance data you want to use for the metric. Choices are

  - *Internal*. The metric will use data that the ScienceLogic platform automatically collects for each device (availability, latency, CPU usage, memory usage, swap usage, device state, and device count).

  - *Dynamic App*. The metric will use data collected by a specified Dynamic Application.

  - *Network Interface* . The metric will use data collected from network interfaces. For an example that uses the Network Interface metrics, see the *example chapter*.

  - *TCP/IP Port Monitor*. The metric will use data collected using a TCP/IP port monitoring policy.

- *System Process Monitor*. The metric will use data collected using a system process monitoring policy.

- *Windows Service Monitor*. The metric will use data collected using a Windows service monitoring policy.

- *Email Round-Trip Monitor*. The metric will use data collected using an Email round-trip monitoring policy.

- *Web Content Monitor*. The metric will use data collected from a using a Web Content monitoring policy.

- *SOAP/XML Transaction Monitor*. The metric will use data collected using a SOAP/XML Transaction monitoring policy.

- *Domain Name Monitor*. The metric will use data collected using a Domain Name monitoring policy.

- *Aligned Service Dependency*. The metric points to a metric and set of devices in an external IT Service policy. The external IT Service policy must first be defined as a *Service Dependency* in the current IT Service. To see an example of a metric that uses an Aligned Service Dependency, see the *example chapter*.

- *Dependency Key Metric*. The metric will inherit the values from a Key Metric (Health, Availability, Risk) defined in another IT service policy. The other IT service policy has been aligned as a dependency to this IT service.

- **Device Subset**. If you have defined one or more *device subsets*, you can select one in this field. The metric will use data from only devices in the selected subset. This field is not applicable if you selected *Aligned Service Dependency* in the **Metric Type** field.

- **Aggregation**. Specifies how the ScienceLogic platform will aggregate ("crunch") the data collected from all the devices in the IT Service or in the specified **Device Subset** into a single value. Choices are:

  - *Average*

  - *Maximum*

  - *Minimum*

  - *Sum*

  - *Std Dev*

  - *Count*

  - *Percentile*. Aggregates data by percentile. Enter the percentile value you want to monitor in the field to the right of the **Aggregation** field. For example, if you select *Percentile* and enter 65 in the field to the right of the **Aggregation** field, this metric will contain the value that is at the 65th percentile for each collection of the metric.

- **Show only metrics available for this IT Service**. Filters the succeeding fields so that they display already-defined policies aligned with one or more of the devices in the IT Service or in the specified **Device Subset**. For example, if you selected *Dynamic App* in the **Metric Type** field, and then selected this checkbox, the **Dynamic Application** field would display only Dynamic Applications that are already aligned with one or more of the devices in the IT Service or in the specified **Device Subset**.

This field is not applicable if you selected *Aligned Service Dependency* or *Dependency Key Metric* in the **Metric Type** field.

- **Device Metric**. Appears if you selected *Internal* in the **Metric Type** field. Choices are:

  - *Device Availability*. Specifies that the metric should be calculated using the availability statistics from all the devices in the IT Service or the devices specified in the **Device Subset**. To calculate a value for the metric, the ScienceLogic platform aggregates the availability value from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, Device Count, or Percentile).

  - *Device Latency*. Specifies that the metric should be calculated using the latency statistics from all the devices in the IT Service or the devices specified in the **Device Subset**. To calculate a value for the metric, the ScienceLogic platform aggregates the latency value from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, or Device Count).

  - *Overall CPU*. Specifies that the metric should be calculated using the CPU usage statistics from all the devices in the IT Service or the devices specified in the **Device Subset**. To calculate a value for the metric, the ScienceLogic platform aggregates the CPU usage value from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, Device Count, or Percentile).

  - *Physical Memory Utilization*. Specifies that the metric should be calculated using the physical memory utilization statistics from all the devices in the IT Service or the devices specified in the **Device Subset**. To calculate a value for the metric, the ScienceLogic platform aggregates the physical memory utilization value from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, or Device Count).

  - *Swap Utilization*. Specifies that the metric should be calculated using the swap utilization statistics from all the devices in the IT Service or the devices specified in the **Device Subset**. To calculate a value for the metric, the ScienceLogic platform aggregates the swap utilization value from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, Device Count, or Percentile).

  - *Device State*. Specifies that the metric should be calculated using the device state of all the devices in the IT Service or the devices specified in the **Device Subset**.To calculate a value for the metric, the ScienceLogic platform aggregates the device state value from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, or Device Count).

NOTE: For each device, device state is determined by the most severe event that is currently active on the device. Each device severity has a numeric equivalent. 0=healthy, 1=notice, 2=minor, 3=major, 4=critical.

  - *Device Count*. Device Count for an IT Service includes the number of devices in the IT Service or in the specified **Device Subset**.

- *Dynamic Application*. Appears if you selected *Dynamic App* in the **Metric Type** field. Select the Dynamic Application that includes the presentation object you want to use to aggregate data for this metric. If you selected the checkbox for **Show only metrics available for this IT Service**, this field displays a list of Dynamic Applications that are already aligned with one or more of the devices in the IT Service or in the specified **Device Subset**. If you did not select the checkbox, this field displays a list of all Dynamic Applications in the ScienceLogic platform.

- *Presentation*. Appears if you selected *Dynamic App* in the **Metric Type** field. Displays a list of all presentation objects in the Dynamic Application that you specified in the **Dynamic Application** field. Select the presentation object that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates the value for this presentation object from each device using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, Device Count, or Percentile).

- *Index/Label*. For presentation objects that include multiple data series, specifies whether you want to select all the data for aggregation, select by index, or select by object label. The choices are:

  - *All*. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate values for all indexes and labels for the presentation object you selected in the **Presentation** field for all the devices in the IT service or the devices specified in the **Device Subset**. When the data for this metric is aggregated, each index is treated as a separate data-point. For example, if you are creating a metric that will aggregate data from three devices and each device has five indexes for the selected presentation object, the aggregation will be based on the 15 separate data-points. If you select this option, the **SNMP Index** field is disabled.

  - *Index*. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate one data-point from each device. If you select this option, you must specify the single numeric index in the **SNMP Index** field.

  - *Label*. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate one data-point from each device. If you select this option, you must specify the single label in the **SNMP Index** field.

- *SNMP Index*. Appears if you selected *Dynamic App* in the **Metric Type** field. A single presentation object can include multiple data series for a given device. For example, a presentation object that displays CPU data might display a data series for each CPU on a device. In the ScienceLogic platform, a data series is called an **index**. An index can be referenced by either a numeric ID or by a label string:

  - If you select *All* in the **Index/Label** field, the **SNMP Index** field is disabled. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate values for all indexes for the presentation object you selected in the **Presentation** field for all the devices in the IT service or the devices specified in the **Device Subset**. When the data for this metric is aggregated, each index is treated as a separate data-point. For example, if you are creating a metric that will aggregate data from three devices and each device has five indexes for the selected presentation object, the aggregation will be based on the 15 separate data-points.

  - If you select *Index* in the **Index/Label** field, you must specify a single numeric index in the **SNMP Index** field. To calculate a value for this metric, the ScienceLogic platform will aggregate one data-point from each device.

Creating, Editing, and Deleting IT Services

- If you select *Label* in the **Index/Label** field, you must specify a single label string in the **SNMP Index** field. To calculate a value for this metric, the ScienceLogic platform will aggregate one data-point from each device.

- **Interface Selection**. Appears if you selected *Network Interface* in the **Metric Type** field. Select the network interfaces to include in the calculation for this metric. Choices are:

  - *All Interfaces*. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate interface utilization statistics from all interfaces on all the devices in the IT Service or the devices specified in the **Device Subset**.

  - *Management Interface*. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate interface utilization statistics from the management interface on all the devices in the IT Service or the devices specified in the **Device Subset**. The management interface is the network interface associated with the IP address that the ScienceLogic platform uses to communicate with a device.

  - *Tagged Interfaces*. Specifies that to calculate a value for this metric, the ScienceLogic platform should aggregate interface utilization statistics from the interfaces that are associated with a specific tag on all the devices in the IT Service or the devices specified in the **Device Subset**. You can manually align tags to an interface in the **Network Interfaces** page (Registry > Networks > Interfaces).

- **Interface Tag**. Appears if you selected *Network Interface* in the **Metric Type** field. Select the interface tag that must be associated with an interface for that interface to be included in the calculation for this metric. If you selected the checkbox for **Show only metrics available for this IT Service**, this field displays a list of Interface tags that are already aligned with one or more of the interfaces on the devices in the IT Service or in the specified **Device Subset**. If you did not select the checkbox, this field displays a list of all Interface tags that are associated with interfaces in the ScienceLogic platform.

- **Interface Metric**. Appears if you selected *Network Interface* in the **Metric Type** field. Select the interface measurement that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates the value for this interface measurement from all interfaces that you included in this metric using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, or Device Count). Choices are:

  - *Inbound Traffic*
  - *Outbound Traffic*
  - *Inbound Errors*
  - *Outbound Errors*
  - *Inbound Discards*
  - *Outbound Discards*

- **TCP Port**. Appears if you selected *TCP/IP Port Monitor* in the **Metric Type** field. Enter or select the TCP port that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates the availability value collected using the TCP port policy that monitors this port on each device. The availability values are aggregated using

the method specified in the ***Aggregation*** field (*Average, Minimum, Maximum, Sum, Standard Deviation, Device Count* or *Percentile*). If you selected the checkbox for ***Show only metrics available for this IT Service***, this field displays a list of port policies that are already aligned with one or more of the devices in the IT Service or in the specified ***Device Subset***. If you did not select the checkbox, enter a TCP port number in this field.

- ***Metric***. Appears if you selected *TCP/IP Port Monitor* in the ***Metric Type*** field. Specify which measurement determines the availability of a TCP port for this metric. Choice is *Port Open*.

- ***Process Name***. Appears if you selected *System Process Monitor* in the ***Metric Type*** field. Enter or select the process name that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates the availability value collected using the system process policy that monitors this process on each device. The availability values are aggregated using the method specified in the ***Aggregation*** field (*Average, Minimum, Maximum, Sum, Standard Deviation,* or *Device Count*). If you selected the checkbox for ***Show only metrics available for this IT Service***, this field displays a list of system process policies that are already aligned with one or more of the devices in the IT Service or in the specified ***Device Subset***. If you did not select the checkbox, enter a process name in this field.

- ***Metric***. Appears if you selected *System Process Monitor* in the ***Metric Type*** field. Specify which measurement determines the availability of a process for this metric. Choices are:

  - *Process Exists*. The process is considered available if it exists on a device.

  - *Process Running*. The process is considered available if it is running on a device.

- ***Service Name***. Appears if you selected *Windows Service Monitor* in the ***Metric Type*** field. Enter or select the service name that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates the availability value collected using the Windows service policy that monitors this service on each device. The availability values are aggregated using the method specified in the ***Aggregation*** field (*Average, Minimum, Maximum, Sum, Standard Deviation, Device Count,* or *Percentile*). If you selected the checkbox for ***Show only metrics available for this IT Service***, this field displays a list of Windows service policies that are already aligned with one or more of the devices in the IT Service or in the specified ***Device Subset***. If you did not select the checkbox, enter a Windows service name in this field.

- ***Metric***. Appears if you selected *Windows Service Monitor* in the ***Metric Type*** field. Specify which measurement determines the availability of a Windows service for this metric. Choice is *Service Running*.

- ***Policy Name***. Appears if you selected *Email Round-Trip Monitor* in the ***Metric Type*** field. Enter or select the Email Round-Trip policy that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates statistics collected using this Email Round-Trip policy on each device. The values for the statistic are aggregated using the method specified in the ***Aggregation*** field (*Average, Minimum, Maximum, Sum, Standard Deviation,* or *Device Count*). If you selected the checkbox for ***Show only metrics available for this IT Service***, this field displays a list of Email Round-Trip policies that are already aligned with one or more of the devices in the IT Service or in the specified ***Device Subset***. If you did not select the checkbox, enter the name of an Email Round-Trip policy in this field.

- *Metric*. Appears if you selected *Email Round-Trip Monitor* in the **Metric Type** field. Specify the Email round-trip statistic to aggregate for this metric. Choices are:

  - *Round-Trip Completed*. Indicates whether or not the round-trip was successfully completed (1) or failed (0).

  - *Round-Trip Time*. Indicates the number of seconds for an Email message to be sent from the ScienceLogic platform to the external server and back to the platform.

- *Policy Name*. Appears if you selected *Web Content Monitor* in the **Metric Type** field. Enter or select the Web Content policy that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates statistics collected using this Web Content policy on each device. The values for the statistic are aggregated using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, Device Count, or Percentile). If you selected the checkbox for **Show only metrics available for this IT Service**, this field displays a list of Web Content policies that are already aligned with one or more of the devices in the IT Service or in the specified **Device Subset**. If you did not select the checkbox, enter a the name of a Web Content policy in this field.

- *Metric*. Appears if you selected *Web Content Monitor* in the **Metric Type** field. Specify the Web Content statistic to aggregate for this metric. Choices are:

  - *Verification Succeeded*. Indicates whether or not the content was found on the website (1) or was not found on the website (0).

  - *Connection Time*

  - *Domain Name Lookup Time*

  - *Page Size*

  - *Download Speed*

  - *Transaction Time*

- *Policy Name*. Appears if you selected *SOAP/XML Transaction Monitor* in the **Metric Type** field. Enter or select the SOAP/XML Transaction policy that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates statistics collected using this SOAP/XML Transaction policy on each device. The values for the statistic are aggregated using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, or Device Count). If you selected the checkbox for **Show only metrics available for this IT Service**, this field displays a list of SOAP/XML Transaction policies that are already aligned with one or more of the devices in the IT Service or in the specified **Device Subset**. If you did not select the checkbox, enter the name of a SOAP/XML Transaction policy in this field.

- *Metric*. Appears if you selected *SOAP/XML Transaction Monitor* in the **Metric Type** field. Specify the SOAP/XML Transaction statistic to aggregate for this metric. Choices are:

  - *Transaction Verification Succeeded*. Indicates whether the transaction was successfully verified (1) or was not successfully verified (0).

  - *Connection Time*

  - *Domain Name Lookup Time*

- ○ *Download Size*

- ○ *Download Speed*

- ○ *Transaction Time*

- **Domain Name**. Appears if you selected *Domain Name Monitor* in the **Metric Type** field. Enter or select the Domain Name policy that the ScienceLogic platform should use to calculate the value for this metric. To calculate a value for the metric, the ScienceLogic platform aggregates statistics collected using this Domain Name Transaction policy on each device. The values for the statistic are aggregated using the method specified in the **Aggregation** field (Average, Minimum, Maximum, Sum, Standard Deviation, Device Count or Percentile). If you selected the checkbox for **Show only metrics available for this IT Service**, this field displays a list of Domain Name Transaction policies that are already aligned with one or more of the devices in the IT Service or in the specified **Device Subset**. If you did not select the checkbox, enter the name of a Domain Name Transaction policy in this field.

- **Metric**. Appears if you selected *Domain Name Monitor* in the **Metric Type** field. Specify the Domain Name statistic to aggregate for this metric. Choices are:

  - ○ *Domain Lookup Succeeded*. Indicates whether the domain lookup was successful (1) or failed (0).

  - ○ *Domain Lookup Time*. Indicates the lookup time, in seconds, for a specified record.

- **Dependent Service**. Appears if you selected *Aligned Service Dependency* or *Dependency Key Metric* in the **Metric Type** field. Select the IT Service that includes the metric you want to use.

- **Service Metric**. Appears if you selected *Aligned Service Dependency* or *Dependency Key Metric* in the **Metric Type** field. Select a metric aligned with the IT Service policy you selected in the **Dependent Service** field. This metric will inherit its values from the selected metric. If you selected *Dependency Key Metric* in the **Metric Type** field, the choices are the Key Metrics:

  - ○ *Health*. The health of the dependent IT services. Health values are represented by integers 0 (healthy) through 4 (critical).

  - ○ *Availability*. The availability of the dependent IT services. Availability values are 0 (unavailable) and 1 (available).

  - ○ *Risk*. The risk of the dependent IT services. Risk values are percentages (0 - 100).

- **Guide Text**. Optionally, enter a description for the metric.

5. Select the **[Save]** button to save your new metric.

NOTE: You can define advanced metrics. Advanced metrics allow you to perform arithmetic operations on the aggregated data and combine multiple sources of aggregated data together. For example, you could specify an advanced metric that combines the download speed from the web content monitors and the SOAP/XML transaction monitors associated with the devices in the IT Service policy. For details on advanced metrics, see the section *Defining Advanced Metrics*.

---

6. You can use the remaining fields to define an optional alert and optional event associated with the metric.

7. To define an associated alert and event, supply a value in each of the following fields:

- **Metric Alerting**. Specifies whether or not you want the ScienceLogic platform to generate an alert for this metric. Choices are:

  ○ *No Alert Policy*. The platform will not trigger alerts for this metric.

  ○ *Single Threshold*. The platform will trigger an alert when the metric meets or exceeds a specified threshold. The platform will clear the alert when the metric no longer meets or exceeds the threshold.

  ○ *Trigger/Reset Thresholds*. The platform will trigger an alert when the metric falls within the "Critical" threshold. The platform will reset the alert when the metric falls within the "Healthy" threshold.

- **Alert Policy Name**. Name of the alert. When you define an alert, the ScienceLogic platform automatically creates an event policy that corresponds to this alert. This name will appear in the name of the event policy.

- **Event Severity**. When the alert is generated, the ScienceLogic platform will trigger an event with the selected event severity. Choices are: *Critical*, *Major*, *Minor*, *Notice*, or *Healthy*.

- *Decreasing/Increasing*. Toggles whether the alert is triggered when the value for the metric is above a specific threshold (Increasing) or below a specific threshold (Decreasing).

- *Alert Threshold*. Use sliders to define the threshold at which the alert should be generated and trigger an event and the threshold at which the alert should be reset and no longer trigger an event.

- *Event Policy Description*. Optionally enter cause and resolution text for the event. When the ScienceLogic platform automatically creates an event policy for this alert, the text you supply in this field will be used to populate the *Policy Description* field in the **Event Policy Editor** for the event. If this event is triggered, the text you supply in this field will be displayed in the **Event Information** modal page for the event.

8. Select the **[Save]** button to save your new alert.

Editing Metrics

You can edit one or more metrics from the **Service Metric Editor** modal page. To do so:

1. Go to **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the IT Service policy you want to edit. Select its wrench icon ( ).

2. Select the **[Metrics]** sub-tab.

3. In the top pane, find the metric you want to edit, and select its wrench icon( ).

4. If you are editing in *Basic mode*, the **Service Metric Editor** modal page appears:

5. If you are editing in *Advanced mode*, the **Service Metric Editor (Advanced)** page appears:

6. You can edit the values in one or more fields. Select the **[Save]** button to save your changes to the metric.

Key Metrics

Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data. For example, you can define "health" for a remote backup service and also define "health" for an Internet bandwidth service, even though you would use different criteria to measure the health of those two services.

All IT Service policies define how the ScienceLogic platform should calculate the following Key Metrics for the IT Service:

- *Health*. The health of an IT Service can be one of the five standard severity values: Healthy, Notice, Minor, Major, or Critical.

- *Availability*. The availability of an IT Service can be either *available* or *unavailable*.

- *Risk*. The risk of an IT Service is a percentage value that indicates how close an IT Service is to being in an undesirable state.

The definition of a Key Metric specifies:

- The metric the ScienceLogic platform should examine to determine a value for the Key Metric.
- One or more threshold values. These thresholds translate values from the metric into values for the Key Metric. For example, "5 or greater" could translate to "critical".

The ScienceLogic platform collects and calculates the value for each Key Metric at the frequency specified in the *Aggregation Frequency* field in the **[Properties]** sub-tab in the **IT Service Editor** page.

The ScienceLogic platform generates an event if the *Service Health* Key Metric has a value of Notice, Minor, Major, or Critical.

To view the value for each Key Metric for an IT Service, see the corresponding column values in the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

Each IT Service policy includes default definitions for each Key Metric. The default definition for each Key Metric uses the default metric *Average Device Availability* (automatically included with each new IT Service policy). You can edit the default definitions to suit your business needs.

To view and edit the definitions of each Key Metric for an IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Either create a new IT Service policy (**[Create]** button) or edit an existing policy (wrench icon 🔧).

2. After *Defining the Basic Properties of an IT Service Policy* and *Defining the List of Devices for an IT Service Policy*, select the **[Metrics]** sub-tab.

3. In the top pane, you will see the default metric, *Average Device Availability*. If you have already defined additional custom metrics, they will also appear in the top pane.

Creating, Editing, and Deleting IT Services

4. In the bottom pane, you will see the three Key Metrics:

- *Service Health*. Appears in the *Health* column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Healthy*, *Notice*, *Minor*, *Major*, and *Critical*.

  - By default, the *Service Health* metric is aligned with the *Average Device Availability* metric.

  - By default, the *Service Health* metric has a range of 0 - 100 and has thresholds set at 25, 50, 75, 90, and 100.

  - To align the *Service Health* metric with another *custom metric*, select that custom metric from the drop-down list that appears above the *Service Health* Key Metric.

  - To change the minimum range, enter a new value in the field the to the far left of the threshold slider. To change the maximum range, enter a new value in the field to the far right of the threshold slider.

  - To customize the thresholds, use the sliders or manually enter values to align with *Healthy*, *Notice*, *Minor*, *Major*, and *Critical*.

- *Service Availability*. Appears in the *Availability* column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Available* and *Unavailable*.

  - By default, this Key Metric is aligned with the same metric as *Service Health*, converting *Critical Service Health* to *Unavailable* and all other *Service Health* values to *Available*.

  - By default, the *Service Availability* metric has a range of 0 - 100 and has a single threshold set at 25.

  - To align the *Service Availability* metric with another *custom metric*, select that custom metric from the drop-down list that appears above the *Service Availability* Key Metric.

  - To customize the thresholds, use the sliders or manually enter values to align with *Unavailable* and *Available*.

  - To change the minimum range, enter a new value in the field the to the far left of the threshold slider. To change the maximum range, enter a new value in the field to the far right of the threshold slider.

- *Service Risk*. Appears as a percentage in the *Risk* column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are 0% - 100%.

  - By default, this Key Metric is aligned with the same metric as *Service Health*, converting the threshold between *Healthy* and *Notice* *Service Health* to *100%* and the healthiest possible value to *0%*.

  - By default, the *Service Risk* metric has a range of 0 - 100 and has a single threshold set at 90.

  - To align the *Service Risk* metric with another *custom metric*, select that custom metric from the drop-down list that appears above the *Service Risk* Key Metric.

  - To customize the thresholds, use the sliders or manually enter values to align with *Critical Risk* (red colored) and *Acceptable Risk* (all other colors).

  - You can also customize the thresholds for how the ScienceLogic platform translates the value for the selected metric into a percentage value.

  - To change the minimum range, enter a new value in the field to the far left of the threshold slider. To change the maximum range, enter a new value in the field to the far right of the threshold slider.

Defining Advanced Metrics

After defining a metric, you can edit the metric in Advanced mode. In Advanced mode, you can perform arithmetic operations on the aggregated data for a metric and combine multiple sources of aggregated data together. The formula for an advanced metric is similar to the formulas you would define for an alert object or a presentation object in a Dynamic Application.

- The **Service Metric Editor (Advanced)** page *does not appear when you create a new metric*. When you create a new metric, the **Service Metric Editor** page appears.

- The **Service Metric Editor (Advanced)** page appears only when you edit an existing metric in *Advanced* mode. To toggle on and off *Advanced* mode, select the **[Advanced]** button.

- After you save a metric with advanced features, you can no longer edit that metric in **[Basic]** mode.

To define an advanced metric:

1. Go to **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the IT Service policy you want to edit. Select its wrench icon (🔧).

2. In the **IT Service Editor** page, select the **[Metrics]** sub-tab.

3. If you have not yet created a metric, perform the steps to *create a metric*.

4. In the **Service Metric Definitions** pane, select the **[Advanced]** button to enable *Advanced* mode.

5. In the **Service Metric Definitions** pane, in the top pane, find the metric you want to edit, and select its wrench icon (🔧).

6. The **Service Metric Editor (Advanced)** page appears:

7. You can edit the values in one or more fields:

   - *Service Metric Name*. Enter a name for the metric.

   - *Metric Classification*. Specifies whether the metric will be displayed in the **IT Service Summary** page in widgets that display vital metrics. Choices are:

     ○ *Service Vital Metric*. This metric will appear in widgets that display vital metrics.

- *Standard Metric*. This metric will not appear in widgets that display vital metrics.

- **Active State**. Specifies whether the ScienceLogic platform should currently collect data for the metric and evaluate alerts for the metric. Choices are:

  - *Enabled*. The platform will collect data for the metric and evaluate alerts for the metric.

  - *Disabled*. The platform will not collect data for the metric and evaluate alerts for the metric.

- **Metric is Percentage Value**. If selected, the next two fields are populated automatically. If not selected, you can supply custom values in the next two fields.

- **Abbreviation/Suffix**. Abbreviation for the unit of measure used in the metric.

- **Data Unit Description**. Description of the unit of measure used in the metric.

- **Metric Formula**. Area where you can perform arithmetic operations on one or more aggregation objects. For details, see the section on the **Metric Formula** pane.

- **Aggregation Objects**. After you define a metric, the ScienceLogic platform creates a collection object and an aggregation object for the metric. The aggregation object tells the ScienceLogic platform the type of performance data to aggregate, the devices for which data will be aggregated, and the method of aggregation (average, maximum, minimum, sum, standard deviation, count) to perform on the data.

- **Guide Text**. Optionally, enter a description for the metric.

8. Select the **[Save]** button to save your changes to the metric. The metric will now appear as an Advanced Metric. You can no longer edit this metric in **Basic** mode.

Metric Formula pane

The **Metric Formula** pane allows you to define which aggregation object(s) the ScienceLogic platform will use for each metric and also allows you to perform manipulations on those aggregation object(s).

- The scrolling list below the **Metric Formula** pane contains a list of all aggregation objects in the current IT Service policy. To include an object in the **Metric Formula** pane, double-click on it.

- To use the calculated value of an aggregation object, you can enter only the aggregation object ID in the **Metric Formula** pane.

- The ScienceLogic platform can perform additional processing to calculate the values for the metric. To specify additional calculations, you can use any combination of arithmetic operations, numeric values, and aggregation object IDs. Parentheses are used to group and set precedence for operators.

- You can use the PHP ternary operator when applying calculations to a metric.

For example, suppose you want a metric that shows the percentage of database servers that are up and running.

- Suppose object o_4095 specifies the sum of all mysqld processes running on all devices in the Database subset.

- Suppose object o_4096 specifies the device count for all devices in the Database subset.

- You could enter the following in the **Metric Formula** pane:

  ```
  ((o_4095/o_4096) * 100)
  ```

- This formula allows you to use aggregation objects to calculate a ratio and convert the ratio to a percentage.

Suppose you want assign a secondary value to a metric, based on the current value of the aggregation object, o_10187. For example, suppose you want to assign values like this:

- If o_10187 > 95% then value = 1
- If o_10187 > 75 % but < 95% then value = 0.8
- If o_10187 > 50 % but < 75% then value = 0.5
- Else value = 0

You could use ternary operators to do this:

```
(o_10187 > 95) ? 1 : ((o_10187 > 75) ? 0.8 : ((o_10187 > 50) ? 0.5 : 0))
```

Editing Collection and Aggregation for a Metric in Advanced Mode

When you define and save a metric, the ScienceLogic platform creates the following objects for the metric:

- *Device Metric Collection object*. A collection object for a metric specifies the type of performance data that will be aggregated for the metric and the devices in the IT Service for which data will be aggregated.

- *Aggregation object*. An aggregation object for a metric specifies the collection object and the method of aggregation (average, maximum, minimum, sum, standard deviation, count, and percentile) to perform on the collected values.

- *Metric definition*. The metric definition (which can be edited in the *Service Metric Editor (Advanced)* page) specifies the name of the metric, the active state for the metric, the unit of measure for the metric, the aggregation object to include in the metric, and any advanced calculations for the metric.

In *Advanced* mode, you can edit the collection object and aggregation object for an existing metric. In *Advanced* mode, you can add and edit collection objects, aggregation objects, and the metric formula independently of each other. You can edit the device metric collection object and/or aggregation object for a metric even if you don't want to define an advanced metric formula.

Editing a Collection Object for a Metric

To add or edit a device metric collection object for an existing metric:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the IT Service policy you want to edit. Select its wrench icon (🔧).

2. Select the **[Metrics]** tab.

3. In the **Service Metrics Definitions** pane, find the metric for which you want to edit the collection object. Note the value in the *Device Metric* column.

4. Select the **[Advanced]** button to enable *Advanced* mode.

5. Select the **[Collection]** tab.

6. In the **Device Metric Collections** pane, find the collection object you want to edit. It will have the same name as the value you previously noted from the *Device Metric* column. Select its wrench icon (🔧). If you want to create a new device metric collection object, select the **[Add]** button.

7. The **Device Metric Collection Editor** modal page appears. In this page, you can define or edit one or more of the following fields:

   - *Metric Type*. Specifies the type of performance data you want to use for the metric. For a description of each choice, see the section on *Defining Metrics.*

   - *Device Subset*. If you have defined one or more *device subsets*, you can select one in this field. The metric will use data from only devices in the selected subset.

   - *Show only metrics available for this IT Service*. Filters the succeeding fields so that they display already-defined policies aligned with one or more of the devices in the IT Service or in the specified *Device Subset*. For example, if you selected *Dynamic App* in the *Metric Type* field, and then selected this checkbox, the *Dynamic Application* field would display only Dynamic Applications that are already aligned with one or more of the devices in the IT Service or in the specified *Device Subset*.

   - The remaining fields provide further parameters for the *Metric Type*. For a detailed description of each possible field and field option, see the section on *Defining Metrics.* If you want to change the parameters of a metric, you can do so in these fields.

8. Select the **[OK]** button to save your changes to the collection object.

9. Select the **[Reset]** button to clear your changes from the device metric collection object and return to the previous values.

10. Select the **[Save]** button in the **IT Service Editor** page to save your changes.

Editing an Aggregation Object for a Metric

To add or edit an aggregation object for an existing metric:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the IT Service policy you want to edit. Select its wrench icon (🔧).

2. Select the **[Metrics]** tab.

3. In the **Service Metrics Definitions** pane, find the metric for which you want to edit the aggregation object. Note the value in the *Device Metric* column.

4. Select the **[Advanced]** button to enable *Advanced* mode.

5. Select the **[Collection]** tab.

6. In the **Aggregation Objects** pane, find the aggregation object you want to edit. The name for the aggregation object will begin with the value you previously noted from the *Device Metric* column. Select its wrench icon (🔧). If you want to create a new aggregation object, select the **[Add]** button.

Creating, Editing, and Deleting IT Services

7. The **Device Metric Aggregation Editor** modal page appears. In this page, you can edit one or more of the following fields:

- *Metric Collection*. Displays the name of the device metric collection object and the associated device subset in parentheses. You can select from a list of all device metric collection objects in the current IT Service policy.

- *Aggregation*. Specifies how the ScienceLogic platform will aggregate ("crunch") the data specified in the device metric collection object into a single value. Choices are:

  - Average
  - Maximum
  - Minimum
  - Sum
  - Std Dev
  - Count
  - *Percentile*. Aggregates data by percentile. Enter the percentile value you want to monitor in the field to the right of the *Aggregation* field. For example, if you select *Percentile* and enter *65* in the field to the right of the *Aggregation* field, this metric will contain the value that is at the 65th percentile for each collection of the metric.

- *Aggregation Name*. The name of the aggregation object. You can edit this value. By default, the value is the name of the collection object and the value from the Aggregation field in parentheses.

8. Select the **[OK]** button to save your changes to the aggregation object.

9. Select the **[Reset]** button to clear your changes from the aggregation object and return to the previous values.

10. Select the **[Save]** button in the **IT Service Editor** page to save your changes.

Editing Alerts and Events in Advanced Mode

In *Advanced* mode, you can edit or define an alert separately from its associated metric.

- If you created an alert for a metric (when you created the metric), you can **edit the alert** in *Advanced* mode.

- If you did not create an alert for a metric when you created the metric, you can **add an alert** in *Advanced* mode.

To edit an existing alert or create an alert for an IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the IT Service policy you want to edit. Select its wrench icon ( ).

2. Select the **[Advanced]** button to enable *Advanced* mode.

3. Select the **[Alerting]** tab.

4. To **edit an existing alert**, select its wrench icon ( ).

5. To *create a new alert*, select the **[Add]** button. The **Service Alert Policy Editor** modal page appears:

6. To edit an alert, edit the value in one or more of the following fields. To create an alert, supply a value in each of the following fields.

- **Alert Policy Name**. Name of the alert. When you define an alert, the ScienceLogic platform automatically creates an event policy that corresponds to this alert. The value from this field will appear in the name of the event policy.

- **Service Metric Name**. Name of the metric associated with the alert. You can select from a list of all metrics in the current IT Service policy.

- **Policy State**. Specifies whether the ScienceLogic platform will evaluate this alert every time data is aggregated for the metric. Choices are:

  ◦ *Enabled*. The platform will evaluate this alert.

  ◦ *Disabled*. The platform will not evaluate this alert.

- ***Threshold Type***. Specifies how the ScienceLogic platform will trigger an alert and reset the alert.

    ○ *Single Threshold*. The platform will trigger an alert when the metric meets or exceeds a specified threshold. The platform will clear the alert when the metric no longer meets or exceeds the threshold.

    ○ *Trigger/Reset Thresholds*. The platform will trigger an alert when the metric falls within the "Critical" threshold. The platform will reset the alert when the metric falls within the "Healthy" threshold.

- ***Event Severity***. When the alert is generated, the ScienceLogic platform will trigger an event with the selected event severity. Choices are: *Critical*, *Major*, *Minor*, *Notice*, or *Healthy*.

- ***Decreasing/Increasing***. Toggles whether the alert is triggered when the value for the metric is above a specific threshold (Increasing) or below a specific threshold (Decreasing).

- ***Alert Threshold***. Use sliders to define two thresholds: the threshold at which the alert should be generated and trigger an event and the threshold at which the alert should be reset and no longer trigger an event. You can edit the maximum and minimum values for the threshold by editing the fields that appear at each end of the threshold slider.

- ***Event Policy Description***. Optionally enter cause and resolution text for the event. When the ScienceLogic platform automatically creates an event policy for this alert, the text you supply in this field will be used to populate the **Policy Description** field in the **Event Policy Editor** for the event. If this event is triggered, the text you supply in this field will be displayed in the **Event Information** modal page for the event.

7. Select the **[OK]** button to save your changes to the alert.

8. Select the **[Reset]** button to clear your changes from the alert and return to the previous values.

9. Select the **[Save]** button in the **Alert Policies** page to save your changes.
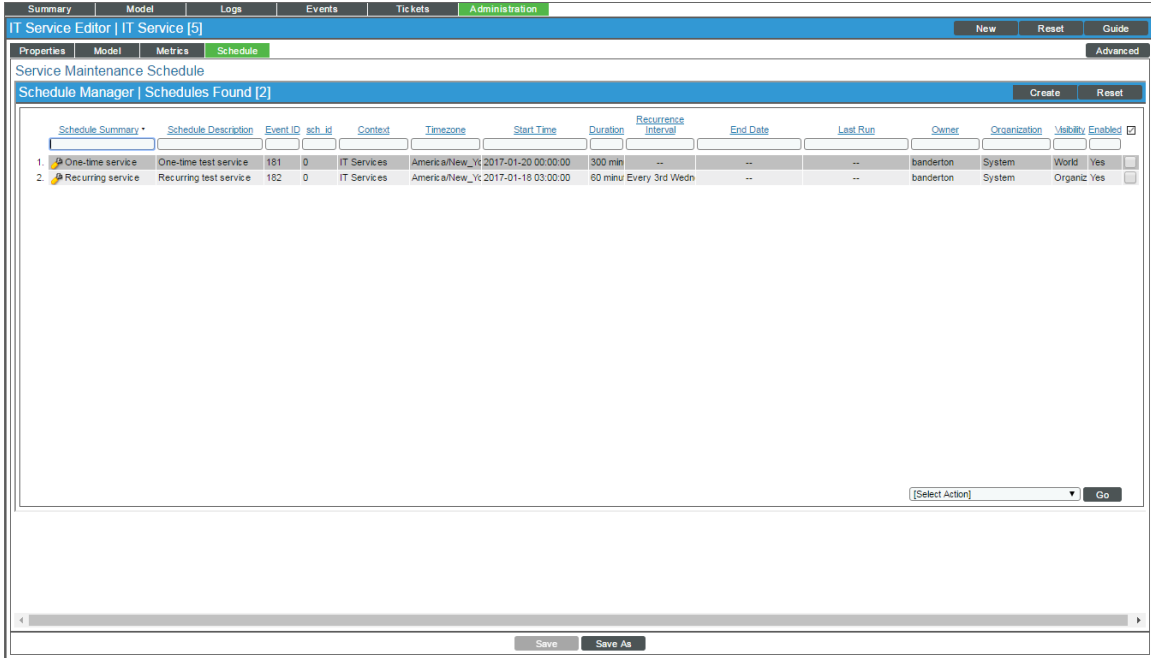
NOTE: To view the event policy associated with an alert, go to the **Event Policy Manager** page (Registry > Events > Event Manager). In the **Event Policy Name** column, search for the value from the **Alert Policy Name** field for the current alert.

Scheduling Downtime for an IT Service

During maintenance mode, the ScienceLogic platform will not aggregate data for the IT Service or generate alerts and events about the IT Service. The ScienceLogic platform will continue to collect information from the devices in the IT Service but will not collect and aggregate information specific to the IT Service policy.

Viewing the Schedule Manager

The **Schedule Manager** page (Registry > IT Services > IT Services Manager > wrench icon > Schedule) displays the following information about each scheduled or recurring maintenance period for an IT Service period:

- **Schedule Summary**. Displays the name assigned to the scheduled process.

- **Schedule Description**. Displays a description of the scheduled process.

- **Event ID**. Displays a unique, numeric ID for the scheduled process. The ScienceLogic platform automatically created this ID for each scheduled process.

- **sch id**. Displays a unique, numeric ID for the schedule. The ScienceLogic platform automatically created this ID for each schedule.

- **Context**. Displays the area of the ScienceLogic platform upon which the schedule works.

- **Timezone**. Displays the time zone associated with the scheduled process.

- **Start Time**. Displays the date and time at which the scheduled process will begin.

- **Duration**. Displays the duration, in minutes, which the scheduled process occurs.

- **Recurrence Interval**. If applicable, displays the interval at which the scheduled process recurs.

- **End Date**. If applicable, displays the date and time on which the scheduled process will recur.

- **Last Run**. If applicable, displays the date and time the scheduled process most recently ran.

- **Owner**. Displays the username of the owner of the scheduled process.

- **Organization**. Displays the organization to which the scheduled process is assigned.

- *Visibility*. Displays the visibility level for the scheduled process. Possible values are "Private", "Organization", or "World".

- *Enabled*. Specifies if the scheduled process is enabled. Possible values are "Yes" or "No".

To edit a scheduled or recurring maintenance period for an IT Service, click its wrench icon (🔧) and update the information as needed on the **Schedule Editor** modal page. (For more information, see the section *Defining a Scheduled or Recurring IT Service Maintenance Period*.)

Defining a Scheduled or Recurring IT Service Maintenance Period

You can schedule an maintenance period for an IT Service on the **Schedule Manager** page. The ScienceLogic platform will automatically set the service to maintenance mode at the scheduled time.

To define a scheduled or recurring IT Service maintenance period:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Click the wrench icon (🔧) for the IT Service for which you want to schedule maintenance.

3. Click the **[Schedule]** sub-tab. The **Schedule Manager** modal page appears.

4. Click **[Create]**. The **Schedule Editor** modal page appears.

5. On the **Schedule Editor** modal page, enter values in the following fields:

Basic Settings

- *Schedule Name*. Type a name for the scheduled process.

- *Schedule Type*. Indicates the scheduled process type (such as Tickets, Reports, or Devices).

- *Visibility*. Select the visibility for the scheduled process. You can select one of the following:

  - *Private*. The scheduled process is visible only to the owner selected in the **Owner** field.

- *Organization*. The scheduled process is visible only to the organization selected in the **Organization** field.

- *World*. The scheduled process is visible to all users.

- **Organization**. Select the organization to which you want to assign the scheduled process.

- **Owner**. Select the owner of the scheduled process. The default value is the username of the user who created the scheduled process.

- **Description**. Type a description of the scheduled process.

Time Settings

- **Start Time**. Type the date and time at which you want the scheduled process to start. The date and time must be in *YYYY-MM-DD HH:MM:SS* format.

- **End Time**. Type the date and time at which you want the scheduled process to end. The date and time must be in *YYYY-MM-DD HH:MM:SS* format.

- **Time Zone**. Select the region or time zone for the scheduled start time.

NOTE: If you want the ScienceLogic platform to automatically adjust for daylight savings time (if applicable), then you must select a named region (such as *America/New York*) in the **Time Zone** field. If you select a specific time zone (such as *EST*) or a specific time offset (such as *GMT-5*), then the ScienceLogic platform will not automatically adjust for daylight savings time.

- **All Day**. Select this checkbox if the scheduled process occurs all day rather than during a specific period of time. If you do so, the **End Time** field becomes disabled.

- **Recurrence**. Select whether you want the scheduled process to occur once or on a recurring basis. You can select one of the following:

  - *None*. The scheduled process occurs only once.
  - *By Interval*. The scheduled process recurs at a specific interval.

If you select *By Interval*, the following additional fields appear:

- **Interval**. In the first field, enter a number representing the frequency of the scheduled process, then select the time interval in the second field. Choices are *Hours*, *Days*, *Weeks*, or *Months*. For example:

  - If you specify "6 Hours", then the scheduled process recurs every six hours from the time listed in the **Start Date** field.

  - If you specify "10 Days", then the scheduled process recurs every 10 days from the date listed in the **Start Date** field.

  - If you specify "2 Weeks", then the scheduled process recurs every two weeks, on the same day of the week as the **Start Date**.

  - If you specify "3 Months" the ticket recurs every three months, on the same day of the month as the **Start Date**.

- **Recur Until**. Specifies when the scheduled process stops recurring. You can select one of the following:

  ○ *No Limit*. The scheduled process recurs indefinitely until it is disabled.

  ○ *Specified Date*. The scheduled process recurs until a specific date and time. If you select *Specified Date*, you must enter a date and time in the **Last Recurrence** field.

- **Last Recurrence**. Specify the date and time at which you want the scheduled process to stop recurring. The date and time must be in *YYYY-MM-DD HH:MM:SS* format.

6. Click **[Save]**.

Enabling or Disabling One or More Scheduled IT Service Maintenance Periods

You can enable or disable one or more scheduled or recurring IT Service maintenance periods from the **Schedule Manager** page (Registry > IT Services > IT Services Manager > wrench icon > Schedule). To do this:

1. Go to the **Schedule Manager** page (Registry > IT Services > IT Services Manager > wrench icon > Schedule).
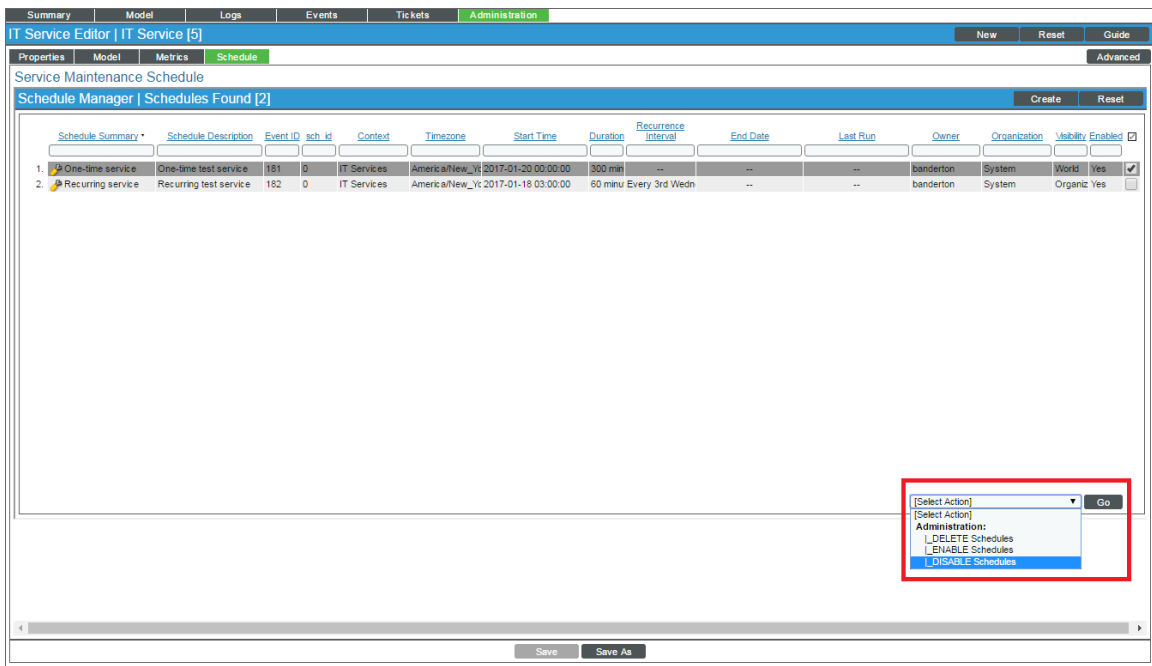
2. Select the checkbox icon for each scheduled process you want to enable or disable.

3. Click the **Select Action** menu and choose *Enable Schedules* or *Disable Schedules*.

4. Click the **[Go]** button.

Deleting One or More Scheduled IT Service Maintenance Periods

You can delete one or more scheduled or recurring IT Service maintenance periods from the **Schedule Manager** page (Registry > IT Services > IT Services Manager > wrench icon > Schedule). To do this:

1. Go to the **Schedule Manager** page (Registry > IT Services > IT Services Manager > wrench icon > Schedule).

2. Select the checkbox icon for each scheduled process you want to delete.

3. Click the **Select Action** menu and choose *Delete Schedules*.

4. Click the **[Go]** button.

Creating, Editing, and Deleting IT Services

Editing an IT Service Policy

To edit the properties of an IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Find the IT Service policy you want to edit. Select its wrench icon ( ).

3. The **IT Service Editor** page is displayed:

4. Notice that by default the **[Administration]** tab is selected and the **[Properties]** sub-tab is selected.

5. You can edit one or more fields in the **IT Service Editor** page, its sub-tabs, or any of the main tabs for the IT Service.

Deleting an IT Service Policy

To delete one or more IT Service policies:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Find the IT Service policy you want to delete. Select its checkbox (☑).

3. Select the checkbox for each IT Service policy you want to delete.

4. Go to the **Select Action** field in the lower right of the page. Select *Delete IT Services*. Select the **[Go]** button.

5. Each selected IT Service policy is removed from the ScienceLogic platform.

System Settings that Affect IT Services

Some of the parameters in the **Data Retention Settings** page affect IT services in the platform.

To define or edit the settings that affect discovery in the **Data Retention Settings** page:

1. Go to the **Data Retention Settings** page (System > Settings > Data Retention).

2. In the **Data Retention Settings** page, edit the values in one or more of the following fields:

   - *Raw ITSM Data*. Before the value for a metric in an IT service policy is calculated, a copy of all the device data that will be aggregated is saved. This setting is the number of days to retain the un-aggregated copies of device data associated with each IT service. The default value is 14 days.

Creating, Editing, and Deleting IT Services

- *ITSM Service Metrics Data*. Number of days to retain values for metrics in IT service policies. The default value is 30 days.

- *Frequent Rollup ITSM Service Metrics Data*. Deprecated field no longer used by the ScienceLogic platform.

- *Hourly Rollup ITSM Service Metrics Data*. Number of days to retain hourly normalized values for metrics in IT service policies. The default value is 120 days.

- *Daily Rollup ITSM Service Metrics Data*. Number of days to retain daily normalized values for metrics in IT service policies. The default value is 365 days.

- *ITSM Key Metrics Data*. Number of days to retain values for key metrics in IT service policies (Health, Availability, and Risk). The default value is 120 days.

- *Frequent Rollup ITSM Key Metrics Data*. Deprecated field no longer used by the ScienceLogic platform.

- *Hourly Rollup ITSM Key Metrics Data*. Number of days to retain hourly normalized values for key metrics in IT service policies (Health, Availability, and Risk). The default value is 365 days.

- *Daily Rollup ITSM Key Metrics Data*. Number of days to retain daily normalized values for key metrics in IT service policies (Health, Availability, and Risk). The default value is 720 days.

3. Select the **[Save]** button to save changes in this page.

IT Service Policies in PowerPacks

IT Service policies can be included in PowerPacks for export and import between ScienceLogic systems.

When an IT Service policy is included in a PowerPack, the following properties of the IT Service policy are included:

- The settings defined in the **[Properties]** sub-tab in the **IT Service Editor**. However, the following system-specific values are substituted when the PowerPack is installed:

  - The *IT Service Owner* field is set to the user that installed the PowerPack.

  - The values that were selected in the *Permission Keys* field are removed.

- The dynamic device rules defined in the **[Model]** sub-tab in the **IT Service Editor**. Statically aligned devices, device groups, and service dependencies are not included.

- The device subsets defined in the **[Model]** sub-tab in the **IT Service Editor**. However, only dynamic device rules are included in those subsets. Statically aligned devices are not included.

- All settings defined in the **[Metrics]**, **[Collection]**, and **[Alerting]** sub-tabs in the **IT Service Editor**.

When a PowerPack that includes an IT Service policy is installed on a ScienceLogic system, the following items are also installed by the PowerPack:

- All event policies associated with the IT Service policy.

- All interface tags that are specified in a metric in the IT Service policy.

Additionally, if a metric uses a Dynamic Application that is not included in the same PowerPack, the ScienceLogic platform will automatically include the Dynamic Application in the PowerPack. When the PowerPack is installed on a ScienceLogic system, the ScienceLogic platform will install the Dynamic Application (with their aligned PowerPack GUIDs). However, if the platform discovers one or more of the Dynamic Applications in the PowerPack already exists on the ScienceLogic system, the platform will not overwrite the existing Dynamic Applications.

Chapter

3

Viewing IT Services
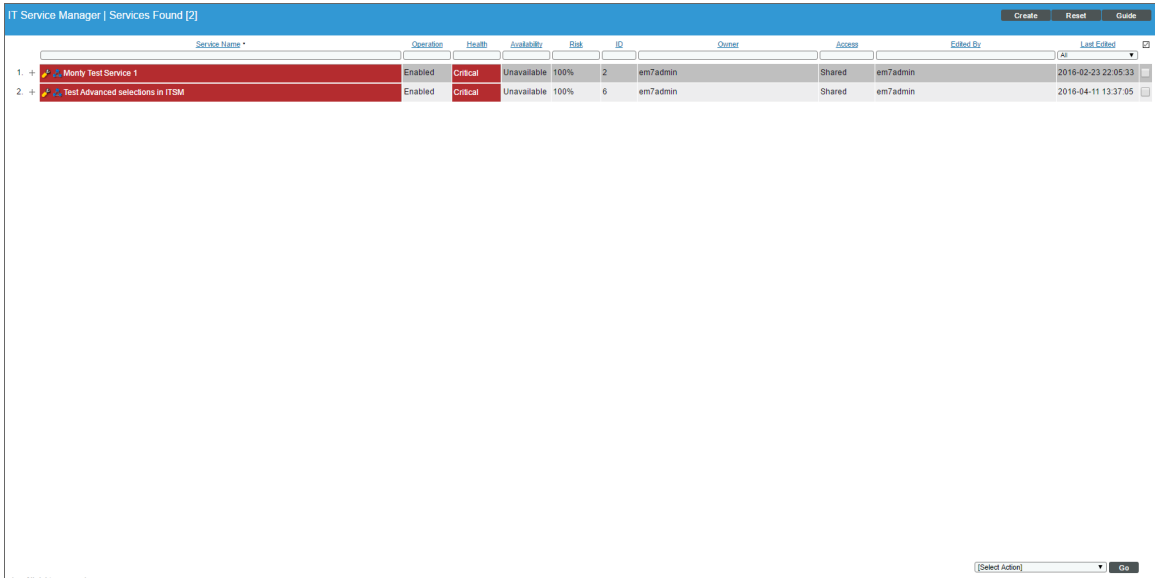
Overview

This chapter describes the pages in which you can view information about an IT Service.

Viewing the List of IT Service Policies

The **IT Service Manager** page (Registry > IT Services > IT Service Manager) displays the following about each IT Service policy:

> **TIP:** To sort the list of policies, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The *Last Edited* column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

- *Service Name*. Name of the policy.
- *Health*. This is a default Key Metric for each IT Service policy. This metric specifies the overall health of the IT Service. Possible values are: critical, major, minor, notice, and healthy.
- *Availability*. This is a Key Metric for each IT Service policy. This metric specifies the overall availability of the IT Service. Possible values are: available or unavailable.
- *Risk*. This is a Key Metric for each IT Service policy. This metric specifies the overall risk to the IT Service. Possible values are 0% - 100%, in integer values.
- *ID*. Unique numeric identifier, automatically assigned to the policy by the ScienceLogic platform.
- *Owner*. User who created the policy.
- *Access*. Specifies whether other users can view and use the policy. Shared policies can be viewed by other users who belong to the same organization as the creator. A private policy can be viewed only by the creator of the policy and administrators.
- *Edited By*. User who created or last edited the policy.
- *Last Edited*. Date and time the policy was created or last edited.

Filtering the List of IT Services

The Filter-While-You-Type fields appear as a row of blank fields at the top of the list. These fields allow you to filter the items that appear in the list.

The list is dynamically updated as you select each filter. For each filter, you must make a selection from a drop-down menu or type text to match against. The ScienceLogic platform will search for entries that match the text, including partial matches. Text matches are not case-sensitive, and you can use special characters in each text field.

By default, the cursor is placed in the first Filter-While-You-Type field. You can use the <Tab> key or your mouse to move your cursor through the fields.

You can filter by one or more of the following parameters. Only items that meet all of the filter criteria are displayed on the page.

The following describes each filer on the **IT Service Manager** page:

- *Service Name*. Filter by the name of the IT service policy. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching service name.

- **Operation**. Specifies whether the IT service policy is enabled or disabled. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching operational state (Enabled or Disabled).

- **Health**. This is a default Key Metric for each IT service policy and specifies the overall health of the IT service. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching health status.

- **Availability**. his is a Key Metric for each IT service policy and specifies the overall availability of the IT service. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching availability status.

- **Risk**. This is a Key Metric for each IT service policy and specifies the overall risk to the IT service, in percent. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching risk status.

- **ID**. Unique numeric identifier, automatically assigned to the policy by EM7. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching policy ID.

- **Owner**. User who created the policy. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching value for owner.

- **Access**. Specifies whether other EM7 users can view and use the policy. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching value for access.

- **Edited By**. EM7 user who created or last edited the policy. You can enter text to match, including special characters, and the IT Service Manager page will display only IT service policies that have a matching value for edited by.

- **Last Edited**. Date and time the IT service policy. You can select from a list of time periods. The IT Service Manager page will display only IT service policies that have been edited within that time period.

SpecialCharacters

You can include the following special characters to filter by each column except those that display date and time:

> **NOTE**: When searching for a string, the ScienceLogic platform will match substrings by default, even if you do not include any special characters. For example, searching for "hel" will match both "hello" and "helicopter". When searching for a numeric value, the ScienceLogic platform will not match a substring unless you use a special character.

String and Numeric

- , (comma). Specifies an "OR" operation. Works for string and numeric values. For example:

  "dell, micro" matches all values that contain the string "dell" OR the string "micro".

- & (ampersand). Specifies an "AND " operation. Works for string and numeric values. For example:

  "dell & micro" matches all values that contain both the string "dell" AND the string "micro", in any order.

- ! (exclamation point). Specifies a "not" operation. Works for string and numeric values. For example:

  "!dell" matches all values that do not contain the string "dell".

  "!^micro" would match all values that do not start with "micro".

  "!fer$" would match all values that do not end with "fer".

  "!^$" would match all values that are not null.

  "!^" would match null values.

  "!$" would match null values.

  "happy, !dell" would match values that contain "happy" OR values that do not contain "dell".

NOTE: You can also use the "!" character in combination with the arithmetic special characters (min-max, >, <, >=, <=, =) described below.

- * (asterisk). Specifies a "match zero or more" operation. Works for string and numeric values. For a string, matches any string that matches the text before and after the asterisk. For a number, matches any number that contains the text. For example:

  "hel*er" would match "helpers" and "helicopter" but not "hello".

  "325*" would match "325", "32561", and "325000".

  "*000" would match "1000", "25000", and "10500000".

- ? (question mark). Specifies "match any one character". Works for string and numeric values. For example:

  "l?ver" would match the strings "oliver", "levers", and "lover", but not "believer".

  "135?" would match the numbers "1350", "1354", and "1359", but not "135" or "13502"

String

- ^ (caret). For strings only. Specifies "match the beginning". Matches any string that begins with the specified string. For example:

  "^sci" would match "scientific" and "sciencelogic", but not "conscious".

  "^happy$" would match only the string "happy", with no characters before or after.

  "!^micro" would match all values that do not start with "micro".

  "!^$" would match all values that are not null.

  "!^" would match null values.

- $ (dollar sign). For strings only. Specifies "match the ending". Matches any string that ends with the specified string. For example:

  "ter$" would match the string "renter" but not the string "terrific".

  "^happy$" would match only the string "happy", with no characters before or after.

  "!fer$" would match all values that do not end with "fer".

  "!^$" would match all values that are not null.

  "!$" would match null values.

NOTE: You can use both ^ and $ if you want to match an entire string and only that string. For example, "^tern$" would match the strings "tern" or "Tern" or "TERN"; it would not match the strings "terne" or "cistern".

Numeric

- min-max. Matches numeric values only. Specifies any value between the minimum value and the maximum value, including the minimum and the maximum. For example:

  "1-5 "would match 1, 2, 3, 4, and 5.

- - (dash). Matches numeric values only. A "half open" range. Specifies values including the minimum and greater or including the maximum and lesser. For example:

  "1-" matches 1 and greater. So would match 1, 2, 6, 345, etc.

  "-5" matches 5 and less. So would match 5, 3, 1, 0, etc.

- > (greater than). Matches numeric values only. Specifies any value "greater than". For example:

  ">7" would match all values greater than 7.

- < (less than). Matches numeric values only. Specifies any value "less than". For example:

  "<12" would match all values less than 12.

- >= (greater than or equal to). Matches numeric values only. Specifies any value "greater than or equal to". For example:

  "=>7" would match all values 7 and greater.

- <= (less than or equal to). Matches numeric values only. Specifies any value "less than or equal to". For example:

  "=<12" would match all values 12 and less.

- = (equal). Matches numeric values only. For numeric values, allows you to match a negative value. For example:

  "=-5 " would match "-5" instead of being evaluated as the "half open range" as described above.

Additional Examples

- "aio$". Matches only text that ends with "aio".
- "^shu". Matches only text that begins with "shu".
- "^silo$". Matches only the text "silo", with no characters before or after.
- "!silo". Matches only text that does not contains the characters "silo".
- "!^silo". Matches only text that does not start with "silo".
- "!0$". Matches only text that does not end with "0".
- "!^silo$". Matches only text that is not the exact text "silo", with no characters before or after.
- "!^". Matches null values, typically represented as "--" in most pages.
- "!$". Matches null values, typically represented as "--" in most pages.
- "!^$". Matches all text that is not null.
- silo, !aggr". Matches text that contains the characters "silo" and also text that does not contain "aggr".
- "silo, 02, !aggr". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr".
- "silo, 02, !aggr, !01". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr" and also text that does not contain "01".
- "^s*i*l*o$". Matches text that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXlXo" would match.
- "!^s*i*l*o$". Matches all text that does not that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXlXo" would not match.
- "!vol&!silo". Matches text that does not contain "vol" AND also does not contain "silo". For example, "volume" would match, because it contains "vol" but not "silo".
- "!vol&02". Matches text that does not contain "vol" AND also contains "02". For example, "happy02" would match, because it does not contain "vol' and it does contain "02".
- "aggr,!vol&02". Matches text that contains "aggr" OR text that does not contain "vol" AND also contains "02".
- "aggr,!vol&!infra". Matches text that contains "aggr" OR text that does not contain "vol" AND does not contain "infra".
- "*". Matches all text.
- "!*". Matches null values, typically represented as "--" in most pages.
- "silo". Matches text that contains "silo".
- " !silo ". Matches text that does not contain "silo".
- " !^silo$ ". Matches all text except the text "silo", with no characters before or after.
- "-3,7-8,11,24,50-". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50.

- "-3,7-8,11,24,50-,a". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50, and text that includes "a".

- "?n". Matches text that contains any single character and the character "n". For example, this string would match "an", "bn", "cn", "1n", and "2n".

- "n*SAN". Matches text the contains "n", zero or any number of any characters and then "SAN". For example, the string would match "nSAN", and "nhamburgerSAN".

- "^?n*SAN$". Matches text that begins with any single character, is following by "n", and then zero or any number of any characters, and ends in "SAN".

Viewing an IT Service Dashboard

The **IT Service Summary** page allows you to view the IT Service Dashboards that have been configured for the selected IT Service. For more information about creating and customizing IT Service Dashboards, see the *Creating and Editing IT Service Dashboards* chapter.

To view an IT Service dashboard in the ScienceLogic platform:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the View IT Service Icon (👤) for the IT Service you want to view. The **IT Service Summary** page is displayed with the default dashboard for this IT Service selected:
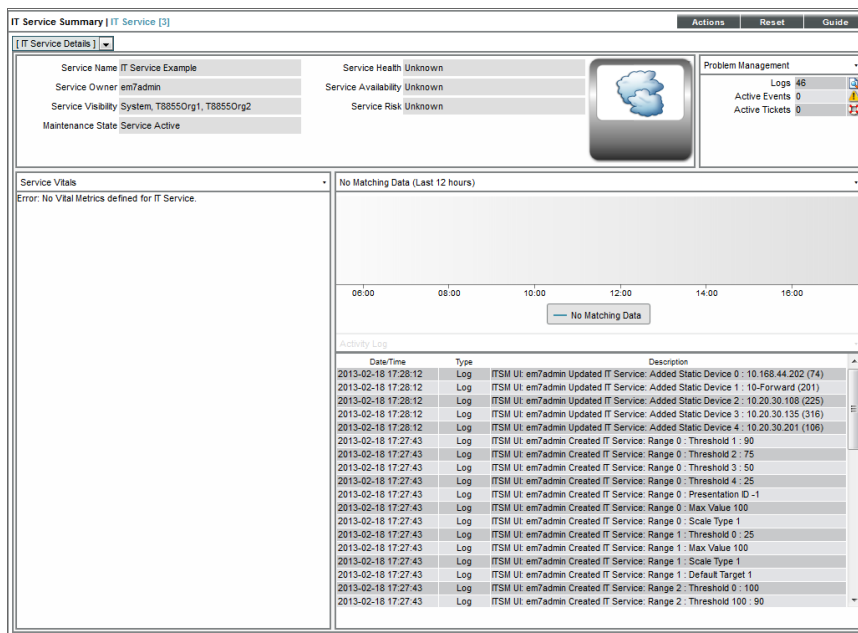
3. The main pane of the dashboard will display one or more graphs, charts, and tables, called *widgets*. The widgets in an IT Service dashboard are configured to automatically display information about the IT Service you selected.

The ScienceLogic platform includes the following widgets that are specifically designed for IT Service dashboards:

- **IT Service Details**. Displays the following information about an IT Service:

    - *Service Name*

    - *Service Owner*

    - *Service Visibility*

    - *Maintenance State*

    - *Service Health*

    - *Service Availability*

    - *Service Risk*

- **IT Service Vitals**. Displays the current value for each Key Metric defined for an IT Service.

- **IT Service Problem Management**. Displays the number of logs, active events, and active tickets associated with an IT Service.

- **IT Service Health Last 12 Hours**. Displays a graph of the Availability metric. The y-axis displays percent availability. The x-axis displays time in one-hour increments.

- **IT Service Activity Log**. Displays a list of all current and past alerts and events associated with an IT Service.

The [**Actions**] menu allows you to perform many dashboard-related tasks, directly from the current page. The [**Actions**] menu looks like a button and is located in the upper right of the page.

The [**Actions**] menu in the **IT Service Summary** page contains the following entries:

- *Configure dashboard...*. Displays the **Dashboard Settings** modal page, where you can edit the layout properties and Access Keys for the currently displayed dashboard. For a full description of the **Dashboard Settings** modal page, see the *Creating and Editing IT Service Dashboards* chapter.

- *Set as default*. Sets the current dashboard as the default dashboard for this IT Service.

- *Use system default*. If the default dashboard for this IT Service is not the system default, this option appears. This option sets the system default IT Service dashboard as the default dashboard for this IT Service.

- *Create Dashboard*. Creates a new IT Service dashboard. The new dashboard is configured to be viewable only for the current IT Service by default.

- *Copy Dashboard (Save As)*. Saves a copy of the currently selected dashboard.
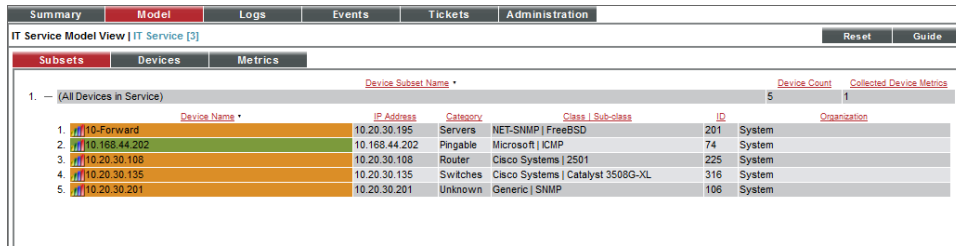
Other Views for IT Services

The **IT Service Model View** page allows you to view:

- A list of all devices that are currently members of an IT Service and a list of devices in each subset of the IT Service.

- A simplified, sortable view of the of all devices that are currently members of an IT Service.

- A relational map of the subsets and metrics that have been defined for an IT Service.
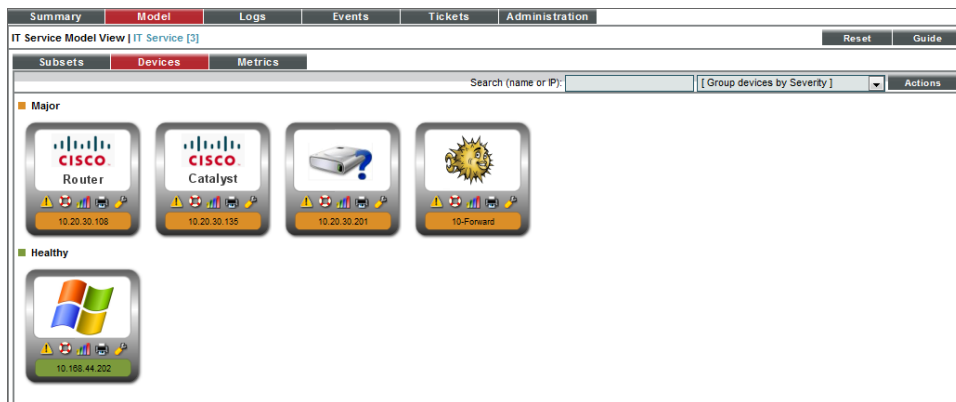
To view the **IT Service Model View** page for an IT Service:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the View IT Service icon (🔺) for the IT Service you want to view.

3. Select the **[Model]** tab.

4. Select one of the following sub-tabs:

   - **[Subsets]**. Displays a list of all devices that are currently members of an IT Service and a list of devices in each subset of the IT Service:

   - **[Devices]**. Displays a simplified, sortable view of all devices that are currently members of an IT Service:

- **[Metrics]**. Displays a relational map of the subsets and metrics that have been defined for an IT Service:

Viewing Events for an IT Service

To view the events associated with an IT Service:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the View IT Service icon ( ) for the IT Service you want to view.

3. Select the **[Events]** tab.

4. The **IT Service Events** page displays the following about each event:

> **TIP:** To sort the list of events, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The *Last Detected* column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

- *Event Message|Severity*. Message generated for the event. The Message is color-coded fro severity.
- *Acknowledged*. If the event has been acknowledged, this column displays a red checkmark character and specifies the user who acknowledged the event. If the event has not been acknowledged, this field displays a gray checkmark character. To acknowledge an event, click in this column.
- *Ticket*. If a ticket has been created for the event, this column displays the ticket ID of that ticket.
- *External Ticket*. The numeric ID associated with a ticket from an external ticketing system (that is, a ticket that was not created in the ScienceLogic platform). If this field displays a value, you can click on that value to spawn a new window and view the external ticket.

> **NOTE:** To link an external ticket to an event, you must create a custom Run Book Automation policy and a custom Run Book Action or use the ScienceLogic APIs. For help with these tasks, contact ScienceLogic Customer Care.

- *Aged/Elapsed*. Number of days, hours, and minutes since the last occurrence of the event.
- *Last Detected*. Date and time the event last occurred on this entity.
- *EID*. Unique ID for the event, generated by the ScienceLogic platform.
- *Source*. System or application that generated this event. Choices are:
  - *Syslog*. Event was generated from system log generated by a device.
  - *Email*. Event was generated by an Email from an external agent, for example, Microsoft Operations Manager (MOM).
  - *Internal*. Event was generated by the ScienceLogic platform.
  - *Trap*. Event was generated by an SNMP trap.
  - *Dynamic*. Event was generated by a Dynamic Application collecting data from the device.
  - *API*. Event was generated by another application.
  - *Count*. Number of times this event has occurred.

You can perform the following actions from the **IT Service Events** page:

- To acknowledge an event, select the checkmark icon (☑) in the *Acknowledged* column.

- To add or edit a note to an event, select the wrench icon (🔧) in the **Notes** column. Select the disk icon (💾) to save your changes to the note.

- To view a summary of an event, select the information icon (ℹ️).

- To view a summary of the automation actions that were triggered by an event, select the notification icon (📨).

- To create a new ticket about an event, select the life-ring icon (🛟).

Filtering the List of Events

The **IT Service Events** page includes a filter for each column, except **Age/Elapsed**. You can specify one or more parameters to filter the display of events. Only events that meet all the filter criteria will be displayed in the **IT Service Events** page.

The list of events is dynamically updated as you select each filter.

> **TIP**: To return to the default list of events, select the **[Reset]** button.

- For each filter (except **Last Detected**), you must enter text to match against. The ScienceLogic platform will search for events that match the text, including partial matches. Text matches are not case-sensitive. You can use the following special characters in each filter:

  - , (comma). Specifies an "or" operation. For example:

    "dell, micro" matches all values that contain the string "dell" OR the string "micro".

  - & (ampersand). Specifies an "and" operation. For example:

    "dell & micro" matches all values that contain the string "dell" AND the string "micro".

  - ! (exclamation mark). Specifies a "not" operation. For example:

    "!dell" matches all values that do not contain the string "dell".

- **Event Message**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have a matching event message.

- **Acknowledged**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have been acknowledged by a matching user account.

- **Ticket**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have a matching ticket ID.

- **External Ticket**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have a matching external ticket name or ID.

- **Last Detected**. Only those events that match the specified detection date will be displayed. The choices are:

  - *All*. Display all events that match the other filters.
  - *Last Minute*. Display only events that have been detected within the last minute.
  - *Last Hour*. Display only events that have been detected within the last hour.
  - *Last Day*. Display only events that have been detected within the last day.
  - *Last Week*. Display only events that have been detected within the last week.
  - *Last Month*. Display only events that have been detected within the last month.
  - *Last Year*. Display only events that have been detected within the last year.

- **EID**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have a matching event ID.

- **Source**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have a matching source.

- **Count**. You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **IT Service Events** page will display only events that have a matching count number.

You can also apply advanced filters to the list of events. To apply advanced filters:

1. Click on the funnel icon ( ).
2. In the advanced filter fields, supply values.
3. Select the **[Apply]** button to apply the filters.

Viewing the Logs for an IT Service

To view the logs that the ScienceLogic platform has generated about an IT Service:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).
2. Select the View IT Service icon ( ) for the IT Service you want to view.

3. Select the **[Logs]** tab.

4. The **IT Service Logs** page displays the following about each log entry:

- *Log Timestamp*. The date and time the entry was made in the log.

- *Source*. The process that generated the log entry. This column will contain the value *ITSM Event*.

- *Message*. The text of the log entry.

Viewing Tickets for an IT Service

To view the tickets associated with an IT Service:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the View IT Service icon (🎋) for the IT Service you want to view.

3. Select the **[Tickets]** tab.

Summary	Model	Logs	Events	Tickets	Administration				

IT Service Tickets | IT Service [3]   Actions   Reset   Guide

	Description / Severity ▼	TID	Queue	Status	
1.	ℹ️ TICKET FOR IT SERVICE: IT Service Example	ID: 3	2	Asset Management	Open	⚒️
2.	ℹ️ TICKET FOR IT SERVICE: IT Service Example	ID: 3	3	Asset Management	Open	⚒️

4. The **IT Service Tickets** page displays the following about each ticket:

> **TIP:** To sort the list of tickets, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again.

- *Description/Severity*. Description of the problem or request. This field is color coded based on the ticket's severity.
- *TID*. Unique ID assigned to the ticket by the ScienceLogic platform.

- *Queue*. Ticket Queue to which the ticket is assigned.

- *Status*. Status of the ticket.

5. You can perform the following actions from the **IT Service Tickets** page:

    - To view a summary of a ticket, select the information icon ( ).

    - To open the **Ticket Editor** for a ticket, select the life-ring icon ( ).

    - To create a new ticket about the IT Service, select the **[Actions]** menu, and then select *Create Ticket*.

Chapter

4

Creating and Editing IT Service Dashboards

Overview

The **IT Service Summary** page allows you to view one or more dashboards that have been configured for the selected IT Service. A dashboard is a page that displays one or more graphical reports, called widgets. Each widget is displayed in its own page.

The ScienceLogic platform includes a default IT Service dashboard that displays general information that is available for all IT Services. You can create additional IT Service dashboards that can be made available in the **IT Service Summary** for one, multiple, or all IT Services.

Viewing the List of IT Service Dashboards

The **IT Service Dashboards** page (Registry > IT Services > IT Service Dashboards) displays the following about each IT Service dashboard in your ScienceLogic system:

> **TIP:** To sort the list of dashboards, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The *Last Edited* column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

- *Dashboard Name*. Name of the dashboard.
- *Service Association*. Specifies the IT Services associated with this dashboard. If a dashboard is associated with an IT Service, the dashboard can be viewed in the **IT Service Summary** page for that IT Service. The possible values in this field are:
  - *System Default*. The dashboard is associated with all IT Services and is displayed when the **IT Service Summary** page loads.
  - *All Services*. The dashboard is associated with all IT Services.
  - *Selected Services*. The dashboard is associated with multiple, but not all, IT Services.
  - *The name of a specific IT Service*. The dashboard is associated with only the IT Service specified in this field.
- *Edited By*. User who created or last edited the dashboard.
- *Last Edited*. Date and time the dashboard was created or last edited.

Creating an IT Service Dashboard

To create an IT Service dashboard:

1. Go to the **IT Service Dashboards** page (Registry > IT Services > IT Service Dashboards).

2. Select the **[Create]** button.

3. The **IT Service Dashboard Editor** page is displayed:

4. The following sections describe how to configure an IT Service dashboard.

Configuring Dashboard Settings

If you want to edit the name, layout, or IT Service association for an IT Service dashboard, you can edit the general settings for that dashboard. To do this from the **IT Service Dashboard Editor** page:

1. Select the **[Actions]** menu and choose *Configure dashboard…*.

2. In the **Dashboard Settings** modal page, supply a value in the following fields:

   - *Dashboard Title*. Enter a name for the dashboard. This name is displayed in the *Select Dashboard* field in the top-left of the **IT Service Summary** page.

   - *Screen Width*. Enter the minimum screen width required to view the whole dashboard at once. If a user views the dashboard at a resolution that has a smaller screen width than the value in this field, a horizontal scroll bar will be displayed. For convenience, the current size of the dashboard as viewed in your monitor is displayed in parentheses above the *Screen Width* and *Screen Height* fields.

   - *Screen Height*. Enter the minimum screen height required to view the whole dashboard at once. If a user views the dashboard at a resolution that has a smaller screen height than the value in this field, a vertical scroll bar will be displayed. For convenience, the current size of the dashboard as viewed in your monitor is displayed in parentheses above the *Screen Width* and *Screen Height* fields.

   - *Lock Dashboard Layout*. Select this checkbox to prevent the dashboard from being edited. When this checkbox is selected, a user viewing the dashboard cannot add, move, or modify widgets while viewing the dashboard. To edit a dashboard that has the *Lock Dashboard Layout* checkbox selected, a user must access the **Dashboard Settings** page and uncheck the *Lock Dashboard Layout* checkbox. This option is most useful for preventing users from accidentally modifying a dashboard.

- *Merge Adjacent Borders*. If the **Lock Dashboard Layout** checkbox is selected, this option is available. If you select this checkbox, widgets that appear next to each other will share a border. If this checkbox is not selected, a gap appears between each widget. If this checkbox is selected, the borders of adjacent widgets are merged.

- *IT Service*. Select how this dashboard will be associated with IT Services. If a dashboard is associated with an IT Service, the dashboard can be viewed in the **IT Service Summary** page for that IT Service. Choices are:
    - *All Services*. Select this option to associate the dashboard with all IT Services. You can optionally select one or more IT Services that will be excluded from this association in the **Excluded Services** field.
    - *Selected Services*. Select this option to associate the dashboard with multiple, but not all, IT Services. If you select this option, select one or more IT Services to associate with this dashboard in the **Included Services** field.
    - *The name of a specific IT Service*. Select a single IT Service to associate with the dashboard.

- *Category*. Select one or more categories to associate with the dashboard. To select multiple categories, hold down the **<Ctrl>** key (or **Command** on Apple computers) when you select the categories.

- *Keywords (comma separated)*. Enter a comma-delimited list of keywords to associate with the dashboard.

3. Select the **[Save]** button to save your changes to the dashboard settings.

Context in IT Service Dashboards

The *context* of a dashboard is a set of values that control what is displayed in one or more widgets on the dashboard. Widgets can be configured to:

- Define the context, by allowing a user to select one or more values.

- Read the context, by updating to include only information about the selected values.

- Both define and read the context.

When a user views an IT Service dashboard in the **IT Service Summary** page, the selected IT Service is automatically defined in the context, i.e. all widgets in the dashboard that are configured to read the context will automatically include only information about the selected IT Service.

For example, suppose an instance of the **Leaderboard/Top-N** widget is included on an IT Service dashboard. Suppose that in the **Widget Configuration** page for the widget, the *Use Device-related Context* checkbox is selected. If you view the IT Service dashboard in the **IT Service Summary** page for an IT Service, the **Leaderboard/Top-N** widget will include only devices in that IT Service. If the IT Service dashboard is configured to display for multiple IT Services, the **Leaderboard/Top-N** widget will display different data for each IT Service.

Adding Widgets to a Dashboard

To add a widget to an IT Service dashboard from the **IT Service Dashboard Editor** page left-click and drag with your mouse to draw a rectangle in the main dashboard pane. This shape will determine the initial size and position of the widget in your dashboard.

1. In the **Dashboards** page (Dashboards tab), in the selection field in the upper left of the page, select the dashboard to which you want to add a widget.

2. Click the **[Actions]** button, and then select *Add Widget*.

Or:

1. In the **Dashboards** page (**[Dashboards]** tab), in the selection field in the upper left of the page, select the dashboard to which you want to add a widget.

2. In the **Dashboards** page, left-click and drag with your mouse to draw a rectangle. This shape will determine the initial size and position of the widget in your dashboard.

3. The **New Widget Configuration** modal page is displayed:
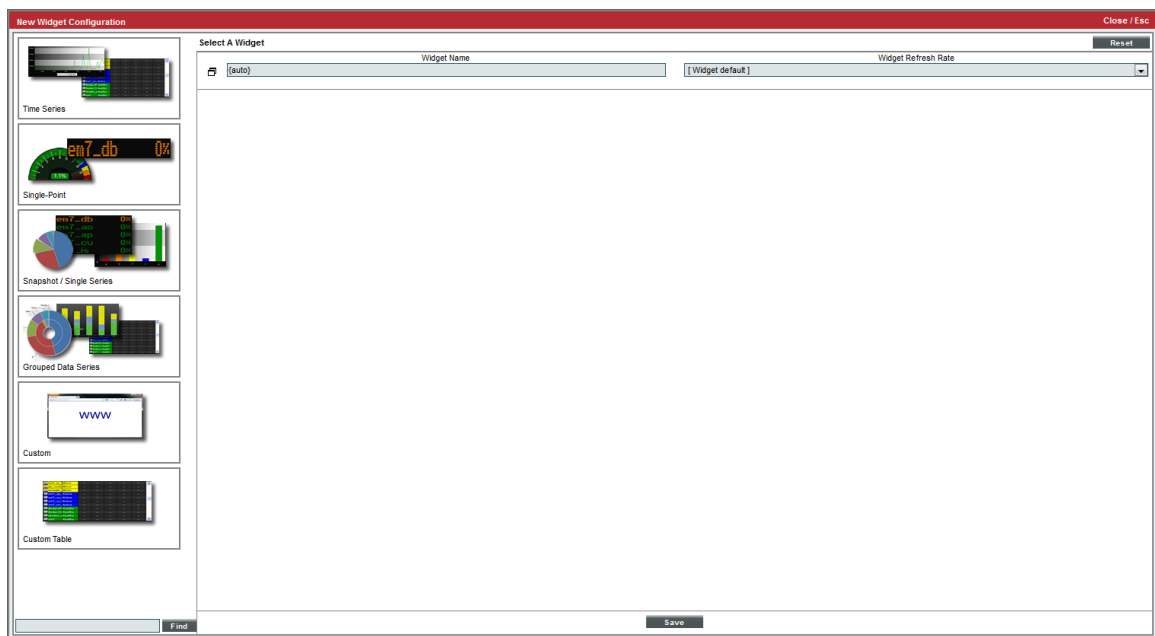
4. To configure a new widget, use the left NavBar to navigate to the widget you want to include in the dashboard.

> **NOTE**: If you are editing an *existing widget*, the **Widget Configuration** page displays the configuration panel for the widget with the *left NavBar and data type selection buttons automatically hidden*. If you want to select a new widget, you can show the left NavBar and data type selection buttons by clicking the window icon to the left of the *Widget Name* field.

- The left NavBar includes an icon for each type of data that can be displayed in a widget:

  ○ **[Time Series]**. Expands a list that includes widgets that display one or more metrics over time.

  ○ **[Single-Point]**. Expands a list that includes widgets that display a single metric.

  ○ **[Snapshot/Single Series]**. Expands a list that includes widgets that display single instances of a metric for multiple entities.

  ○ **[Grouped Data Series]**. Expands a list that includes widgets that display a single metric for multiple entities, with each metric sub-divided between multiple entities of another types. For example, a grouped data series could display a bar graph of the number of tickets in each state, with each bar in the graph divided by organization.

  ○ **[Custom]**. Expands a list that includes widgets that display custom HTML content.

  ○ **[Custom Table]**. Expands a list that includes widgets that display tabular data.

- When you expand the list of widgets that display a type of data, the NavBar displays a list of categories that are associated with the widget definitions in your system. You can expand a category to view the list of all widgets associated with that category. If a widget is associated with multiple categories, the widget will appear under each category it is associated with.

- You can search the list of widgets by entering a search term in the field that appears in the bottom-left of the page. When you click the **[Find]** button, the widgets that have a name or associated keywords that match your search term are highlighted in the left NavBar. The categories in the left NavBar will be automatically expanded and collapsed so that only categories that include a matching widget are expanded.

5. When you select a widget in the left NavBar, the right pane of the **Widget Configuration** page displays the configuration fields for the selected widget. Each widget definition has a different list of configuration fields.

> **TIP**: After you select a widget, you can hide the left NavBar and data type selection buttons by selecting the window icon to the left of the *Widget Name* field.

6. The following fields appear in *all* widget configuration panes:

- *Widget Name*. Enter a title for the widget. This title is displayed in the header that appears at the top of the widget. If you leave the default value of "{auto}" in this field, the ScienceLogic platform will automatically generate a title for the widget based on what is currently being displayed in the widget.

- **Widget Refresh Rate**. Specify how frequently the widget will be automatically updated with new data. The choices are:

    - *Widget Default*. The widget will refresh at its default refresh rate, as defined by the widget developer. You can view and edit the default refresh rate in the **Dashboard Widgets** page (System > Customize > Dashboard Widgets) by clicking the wrench icon (🔧) for a widget.

    - *Auto-refresh disabled*. The widget will not automatically refresh.

    - *1 minute*. The widget will automatically refresh every minute.

    - *5 minutes*. The widget will automatically refresh every 5 minutes.

    - *10 minutes*. The widget will automatically refresh every 10 minutes.

    - *15 minutes*. The widget will automatically refresh every 15 minutes.

    - *30 minutes*. The widget will automatically refresh every 30 minutes.

    - *45 minutes*. The widget will automatically refresh every 45 minutes.

    - *1 hour*. The widget will automatically refresh once an hour.

- **Create Template**. When selected, this checkbox allows you to save the current configuration as a **Quick-Add** option.

7. Each widget contains additional fields.

> **NOTE**: In widgets that allow you to filter the list of devices by the device class or device category, merged devices include special behavior. For merged devices, you can select either the device class or device category of the physical device or the device class or device category of the component device. If both device classes or device categories are selected, a merged device will appear twice in a single widget.

The ScienceLogic platform includes several widgets that are specifically designed for IT Service dashboards. The configuration panels for these widgets include only the default fields in the **Widget Configuration** page. These widgets always read the IT Service value from the dashboard context. The widgets for IT Services are:

- **IT Service Details**. Displays the following information about an IT Service:
    - *Service Name*
    - *Service Owner*
    - *Service Visibility*
    - *Maintenance State*
    - *Service Health*
    - *Service Availability*
    - *Service Risk*

- **IT Service Vitals**. Displays the current value for each Key Metric defined for an IT Service.

- **IT Service Problem Management**. Displays the number of logs, active events, and active tickets associated with an IT Service.

- **IT Service Health Last 12 Hours**. Displays a graph of the Availability metric. The y-axis displays percent availability. The x-axis displays time in one-hour increments.

- **IT Service Activity Log**. Displays a list of all current and past alerts and events associated with an IT Service.

You can also use any other widgets that are installed in your ScienceLogic system in an IT Service dashboard. For a description of the other default widgets supplied by ScienceLogic, see the **Dashboards** manual. To ensure that a user viewing the dashboard can view only information that they have access to, you should configure all the widgets on an IT Service dashboard to read the IT Service value from the context.

For example, suppose that on an IT Service dashboard that will be used for multiple IT Services, you want to include a line graph that displays the Health, Availability, and Risk values for an IT Service over time. To do this, you would add an instance of the **Multi-series Performance** widget to the IT Service dashboard. To ensure that the widget displays only information about the IT Service for which a user is viewing the dashboard, you would make the following selections for the **Series Selections** option in the configuration panel for the **Multi-series Performance** widget:

- **Type**. Select *IT Service*.

- **Element**. Select *Contextual Service 1*. Selecting this option tells the widget to read the IT Service value that is set in the context.

- **Collection**. Select *Service Availability*. The widget will display the availability values for the IT Service that is set in the context.

To add *Service Health* and *Service Risk* values to the graph, you can select the **Add another series** button, select the same **Type** and **Element** values, and then select the appropriate value in the **Collection** field.

Editing the Widgets in a Dashboard

You can edit the size and position of the widgets in a dashboard by manipulating the pane in which that widget appears:

- To move a widget, click and hold on the header of the widget, then drag the widget to a new position on the dashboard.

- To resize a widget, click and hold on the border of the widget, then drag the border to adjust the size.

Each widget has an options menu that appears in the top-right corner of the widget. The widget options menu includes the following options for editing a widget:

- **Configure**. Leads to the **Widget Configuration** modal page, where you can edit the parameters of the widget.

- **Copy To…**. Leads to the **Copy Widget** modal page, where you can copy the widget to another dashboard. To copy the widget to another dashboard, select the dashboard you want to copy the widget to in the **Copy Widget To** drop-down list and select the **[Copy]** button.

- *Lower*. If widgets are stacked on top of each other in the dashboard, this option sends the widget to the bottom of the stack.

- *Duplicate*. Creates a copy of the widget in the current dashboard, using the same widget definition and same parameters as the original widget.

- *Remove*. Deletes the widget from the dashboard.

Setting a Dashboard as the System Default Dashboard

In the ScienceLogic platform, there is always a default dashboard for IT Services. The default dashboard is initially displayed when a user views the **IT Service Summary** page. The default IT Service dashboard must be visible for all IT Services. If the dashboard you want to set as the default is not currently visible for all IT Services, perform the following steps from the **IT Service Dashboard Editor** page for that dashboard:

1. Select the **[Actions]** menu and choose *Configure dashboard…*.

2. In the **Dashboard Settings** modal page, select *All Services* in the **IT Service** field.

3. Select the **[Save]** button.

After you have configured the dashboard to be visible for all IT Services, select the **[Actions]** menu in the **IT Service Dashboard Editor** page and choose *Set as system default*.

Editing an IT Service Dashboard

To edit an IT Service dashboard:

1. Go to the **IT Service Dashboards** page (Registry > IT Services > IT Service Dashboards).

2. Find the IT Service dashboard you want to edit. Select its wrench icon (🔧).

3. The **IT Service Dashboard Editor** page is displayed. For a description of the **IT Service Dashboard Editor** page, see the *Creating an IT Service Dashboard* section.

Deleting an IT Service Dashboard

To delete one or more IT Service dashboards:

1. Go to the **IT Service Dashboards** page (Registry > IT Services > IT Service Dashboards).
2. Find the IT Service dashboard you want to delete. Select its checkbox.
3. Select the checkbox for each IT Service dashboard you want to delete.
4. Go to the **Select Action** field in the lower right of the page. Select *Delete Dashboards*. Select the **[Go]** button.
5. Each selected IT Service dashboard is removed from the ScienceLogic platform.

Chapter

5

SLA Definitions, Reports, and Widgets

What is an SLA?

For service providers, an SLA is a Service-Level Agreement. An SLA is a contract between a service provider and a customer. The contract specifies the service that the provider will deliver to the customer. Some SLAs include penalties for non-compliance.

For example, an Internet service provider might guarantee 99.99% uptime for all customers.

What is an SLA Definition?
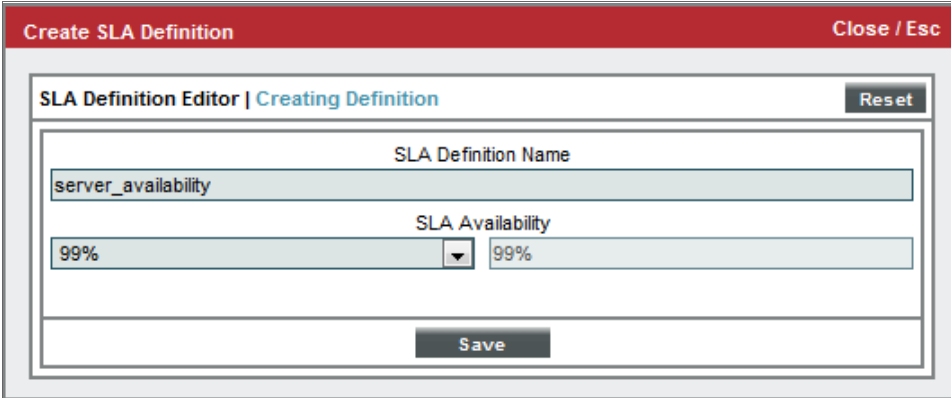
In the ScienceLogic platform, you can create an SLA Definition. The SLA Definition is a threshold. The threshold is applied to the Availability Key Metric of an IT Service policy. For example:

- Suppose we define an IT Service policy to monitor web servers.
- Suppose we define the Availability Key Metric for that IT Service policy to map to the metric that aggregates the average availability of all web servers.
- Suppose that the metric that aggregates the average availability of all web servers is set to "Critical = anything less than average availability of 99%".
- Suppose we create an SLA definition that says that the web servers must be available 99% of the time. 99% uptime allows for 432 minutes of downtime per month.
- If the IT Service policy has a polling frequency of 5 minutes, the web servers can be unavailable (average availability is less than 99%) no more than 86 polling periods per month (and still meet 99% uptime).

Creating an SLA Definition

To create an SLA Definition:

1. Go to the **Service Level Agreement Definitions** page (Registry > IT Services > SLA Definitions).
2. In the **Service Level Agreement Definitions** page, select the **[Create]** button.
3. The **SLA Definition Editor** page is displayed.

4. In the **SLA Definition Editor** page, supply values in the following fields:

   - *SLA Definition Name*. The name of the SLA Definition. Can be any combination of numbers, letters, and symbols.
   - *SLA Availability*. The threshold that will be evaluated using this SLA definition. You can select from six predefined percentage values or specify a custom value. If you select *Custom* in the drop-down list, enter a percentage value in the text field. You can evaluate an IT Service policy using this threshold; you can do this in a dashboard widget or in an SLA report.

5. Select the **[Save]** button to save your new SLA Definition.

Viewing the List of SLA Definitions

The **Service Level Agreement Definitions** page displays the list of existing SLA definitions. To view the list of existing SLA definitions:

1. Go to the **Service Level Agreement Definitions** page (Registry > IT Services > SLA Definitions).

2. The **Service Level Agreement Definitions** page displays the following about each SLA Definition:

> **TIP:** To sort the list of SLA Definitions, click on a column heading. The list will be sorted by the column value, in ascending order. To sort by descending order, click the column heading again. The **Last Edited** column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

- **SLA Definition Name**. The name of the SLA Definition.
- **Availability**. The availability, in percent, specified in the service level agreement.
- **ID**. Unique numeric ID for the SLA Definition, generated by the ScienceLogic platform.
- **Edited By**. Name of the user who created or last edited the SLA Definition.
- **Last Edited**. Date and time the SLA Definition was created or last edited.

Using an SLA Widget in a Dashboard

You can use a Dashboard Widget to evaluate an existing IT Service policy using an existing SLA Definition. The Dashboard Widget will then display the results. To do this, you must perform two steps:

- Create the Dashboard
- Configure the SLA Widget

The following sections describe how to perform these two tasks.

Create the Dashboard

To use the SLA widget in a dashboard:

1. Select the **[Dashboards]** tab.
2. Select the **[New]** button. The system will create a blank dashboard with a default name.
3. Select the **[Actions]** menu, then select *Configure Dashboard...* The **Dashboard Settings** modal page is displayed, where you can configure the basic settings for the dashboard:

4. You can enter values in one or more of the following fields, or you can accept the default values:

- *Dashboard Title*. Enter a name for the dashboard. This name is displayed in the **Select Dashboard** field in the top-left of the **Dashboards** page.

- *Screen Width*. Enter the minimum screen width required to view the whole dashboard at once. If a user views the dashboard at a resolution that has a smaller screen width than the value in this field, a horizontal scroll bar will be displayed. For convenience, the current size of the dashboard as viewed in your monitor is displayed in parentheses above the **Screen Width** and **Screen Height** fields.

- *Screen Height*. Enter the minimum screen height required to view the whole dashboard at once. If a user views the dashboard at a resolution that has a smaller screen height than the value in this field, a vertical scroll bar will be displayed. For convenience, the current size of the dashboard as viewed in your monitor is displayed in parentheses above the **Screen Width** and **Screen Height** fields.

- *Lock Dashboard Layout*. Select this checkbox to prevent the dashboard from being edited. When this checkbox is selected, a user viewing the dashboard cannot add, move, or modify widgets while viewing the dashboard. To edit a dashboard that has the **Lock Dashboard Layout** checkbox selected, a user must access the **Dashboard Settings** page and uncheck the **Lock Dashboard Layout** checkbox. This option is most useful for preventing users from accidentally modifying a dashboard.

- *Merge Adjacent Borders*. If the **Lock Dashboard Layout** checkbox is selected, this option is available. If you select this checkbox, widgets that appear next to each other will share a border.

  - If this checkbox is not selected, a gap appears between each widget:

  - If this checkbox is selected, the widget borders are merged:

- **Access Control**. Specifies whether the dashboard is viewable only by the creator or if the dashboard is viewable by other users. Choices are:

  - *Share with organizations*. Allows other members of the creator's organizations to view the dashboard.
  - *Private dashboard*. Only the dashboard's creator and administrators can view the dashboard.

- **Access Keys**. This field is applicable only if the dashboard is Shared. Specifies the Access Keys required to view the dashboard. If you don't select any Access Keys, no specific keys are required to view the dashboard.

  If you select an Access Key in the **Required Keys** field, each user must meet the following criteria to use the dashboard:

  - The user must have at least one of the Access Keys selected in the **Required Keys** field for the dashboard.
  - The user must be granted an Access Key that includes the "Dash:View" and "Dash:View Shared" Access Hooks.
  - The user and the creator of the dashboard must be members of the same organization.

  If you do not select any Access Keys in the **Required Keys** field, any user meeting the following two requirements may access the dashboard:

  - The user must be granted an Access Key that includes the "Dash:View" and "Dash:View Shared" Access Hooks.
  - The user and the creator of the dashboard must be members of the same organization.

CAUTION: If a user meets the above requirements and also has been granted an Access Key that includes the "Dash:Edit Shared" Access Hook, that user will be able to edit the shared dashboard. If a user has been granted an Access Key that includes the "Dash:Add/Rem Shared" Access Hook, that user may delete shared dashboards.

- **Category**. Select one or more categories to associate with the dashboard. To select multiple categories, hold down the **[Ctrl]** key (or **[Command]** on Apple computers) when you select the categories. Categories are used to arrange the dashboard selection drop-down list in the **Dashboards** page. If you do not select a category in this field, the dashboard will appear under the "Other" category in the drop-down list.
- **Keywords (comma separated)**. Enter a comma-delimited list of keywords to associate with the dashboard.

5. Select the **[Save]** button to save your changes to the dashboard settings.

Configure the SLA Widget

To add the SLA widget to your dashboard:

1. In the **Dashboards** page (**[Dashboards]** tab), in the selection field in the upper left of the page, select the dashboard to which you want to add a widget.
2. Select the **[Actions]** button, and then select *Add Widget*.

Or:

1. In the **Dashboards** page (**[Dashboards]** tab), in the selection field in the upper left of the page, select the dashboard to which you want to add a widget.
2. In the **Dashboards** page, left-click and drag with your mouse to draw a rectangle. This shape will determine the initial size and position of the widget in your dashboard.
3. The **New Widget Configuration** modal page is displayed:

4. In the NavBar, expand the Single Point icon. Expand the SLA category. Select the **(base) SLA Gauge**.
5. Enter the values in the following fields:

   - *Widget Name*. Enter a title for the widget. This title is displayed in the header that appears at the top of the widget. If you leave the default value of "{auto}" in this field, the ScienceLogic platform will automatically generate a title for the widget based on what is currently being displayed in the widget.

- *Widget Refresh Rate*. Specify how frequently the widget will be automatically updated with new data. The choices are:

  - *Widget Default*. The widget will refresh at its default refresh rate, as defined by the widget developer. You can view and edit the default refresh rate in the **Dashboard Widgets** page (System > Customize > Dashboard Widgets) by selecting the wrench icon (🔧) for a widget.

  - *Auto-refresh disabled*. The widget will not automatically refresh.

  - *1 minute*. The widget will automatically refresh every minute.

  - *5 minutes*. The widget will automatically refresh every 5 minutes.

  - *10 minutes*. The widget will automatically refresh every 10 minutes.

  - *15 minutes*. The widget will automatically refresh every 15 minutes.

  - *30 minutes*. The widget will automatically refresh every 30 minutes.

  - *45 minutes*. The widget will automatically refresh every 45 minutes.

  - *1 hour*. The widget will automatically refresh once an hour.

- *SLA Definition*. Select the SLA you want to use as a threshold and monitor with this widget.

- *Service*. Select the IT Service you want to monitor with this widget.

- *Compliance Period*. Specify the time period you want to monitor with this widget. Choices are:

  - *Current*.

  - *Last (most recently ended)*.

  - *Last (ending in context range)*.

- *Display Type*. Select how the metric will be displayed in the widget:

  - *Gauge*

  - *Column*

  - *Horizontal Bar*

  - *Scoreboard*

  - *Waterline Gauge*
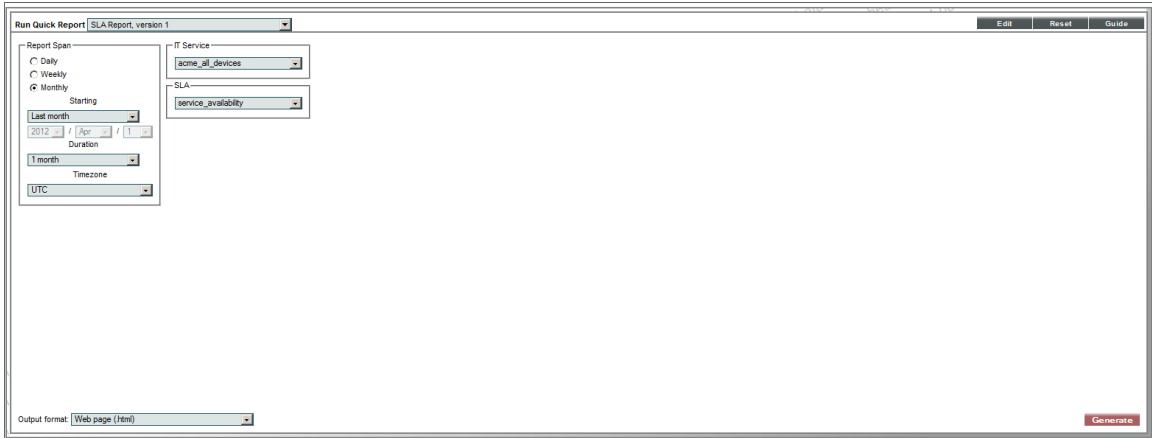
6. You should see a report like the following:

Generating an SLA Report

You can use a Quick Report to evaluate an existing IT Service policy using an existing SLA Definition. The Quick Report will then display the results. To do this:

1. Go the **Run Quick Report** page (Reports > Create Report > Quick Report).

2. In the **_Run Quick Report_** drop-down list, select the report **_SLA Report_**.

3. Enter a value in each of the following fields:

    - **_Report Span_**. Specify a Daily, Weekly, or Monthly span to include in the report.
    - **_Starting_**. This field allows you to choose a start date. Selecting a different **_Report Span_** will change the options in this drop-down list.
    - **_Duration_**. This field allows you to specify the duration for the report. Selecting a different **_Report Span_** will change the options in this drop-down list.
    - **_Timezone_**. Specify the time zone to display in the report.
    - **_IT Service_**. Select the IT Service you want to monitor with this report.
    - **_SLA._** Select the SLA you want to use as a threshold and monitor with this report.
    - **_Output Format_**. Specify an output format for the report.

4. Select the **[Generate]** button to generate the report.

5. The generated SLA report displays the target threshold and the percentage of polls that were successful. The report displays the days violations occurred and the number of minutes each violation lasted.

Sheet 1: *SLA Report*

.∎l ScienceLogic™

SLA Report

Summary			
Agreement Name	**IT Service Name**	**Target %**	**Compliance %**
website_response_sla	web_hosting_home	99.99000%	99.80159%

Violation Periods			
Start Date/Time	**End Date/Time**	**Duration**	**Cumulative**
2012/05/02 19:00:00	2012/05/02 19:20:00	00:20	00:20
Total Violation Time:			00:20

Violations By Day

Excluded Periods			
Start Date/Time	**End Date/Time**	**Duration**	**Type**
No Excluded Periods			
Total Downtime:		00:00	

Effective Times		
Day	**Start**	**End**
Sunday	00:00	23:59
Monday	00:00	23:59
Tuesday	00:00	23:59
Wednesday	00:00	23:59
Thursday	00:00	23:59
Friday	00:00	23:59
Saturday	00:00	23:59

Generated On: 05/03/2012

Sheet 2: *Control*

.∎l ScienceLogic™

Control – SLA Report

Description:	SLA compliance metrics
Report Version:	1
Generated On:	05/03/2012
SLA:	website_response_sla
IT Service:	web_hosting_home
Start Date:	Apr 28, 2012
Duration:	1 week

Generated On: 05/03/2012

Chapter

6

Events and Run Book Policies for IT Services

Events

The ScienceLogic platform includes event definitions that can be triggered when the state of an IT Service is examined. These events are generated in addition to any custom events that you create for an IT Service.

If the state of an IT Service becomes non-healthy (notice, minor, major, or critical), the following events will be generated and displayed in the **Event Console** (**[Events]** tab):

- IT Service State Critical: *Name of IT Service*
- IT Service State Major: *Name of IT Service*
- IT Service State Minor: *Name of IT Service*
- IT Service State Notice: *Name of IT Service*

In the **Event Console**, you will see something like this:
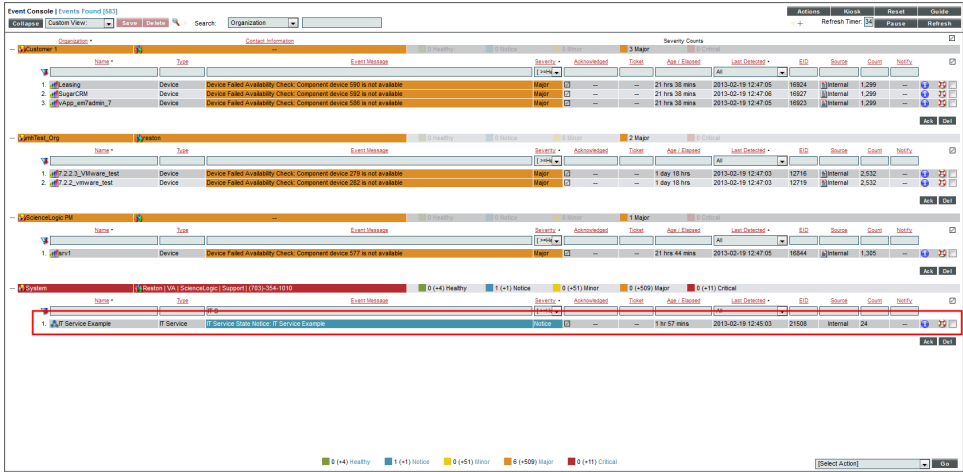
- ![icon]. Leads to the **IT Service Summary** page for the IT Service.
- *Name*. Name of IT Service.
- *Type*. Type of IT Service.
- *Event Message*. IT Service State [*severity*]: [*name of IT Service*].

If the state of an IT Service becomes healthy after being in a non-healthy state, the event "IT Service State Healthy: *Name of IT Service*" will be generated and will clear the event for the non-healthy state.

> **NOTE**: Events are not generated when the availability and risk values for an IT Service change.

If you have defined alerts for metrics, additional events for IT services will appear in the **Event Console**.

Run Book Automation

In an Automation Policy (Automation Policy Manager), you can select an IT Service and trigger an automated action when a specified event occurs on that IT Service.

For details on defining Run Book Automation policies and Run Book Actions, see the manual **Run Book Automation**.

Example

1

1

Example Using Device Availability, Device Latency, and Process Availability

Overview

This chapter will provide a working example of an IT Service policy.

This example will monitor two hosted, mission-critical, MS SQL database servers for availability and latency. The MS SQL database servers are leased by the company Acme and hosted by a service provider.

Creating an IT Service Policy

To define an IT Service policy, you must:

1. *Define a service name and basic properties*. This example monitors two MS SQL database servers. The name of the IT Service policy will be "Acme: MS SQL Database Servers".

2. *Define a list of devices (model) for the IT Service that includes all the devices associated with the IT Service*. This example includes two MS SQL servers in the IT Service.

3. *Optionally, define service sets. A service set is a sub-group of devices*. This example does not use service sets.

4. *Define metrics. A metric is based on your business processes and examines all devices or one or more service sets to evaluate the state of the IT Service*. For each IT Service, the ScienceLogic platform provides a default metric called *Average Device Availability*, based on the availability of all devices in the IT Service. You can define additional metrics, based on default data collected by the ScienceLogic platform (availability, latency, CPU usage, memory usage, swap usage, device state, and device count), data collected by a Dynamic Application, and data about network interfaces, TCP/IP ports, system processes, Windows services, Email round-trip time, web content, SOAP/XML transactions, and DNS availability.

> **NOTE**: When the ScienceLogic platform evaluates a metric, it performs an aggregation, that is, the platform evaluates the data for all devices specified in the definition of the metric, over a specified time period (the **Aggregation Frequency**). Depending on the definition of the metric, the platform calculates the average, maximum, minimum, sum, standard deviation, or count value for all devices specified in the definition of the metric.

5. **Define Key Metrics.** Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data. The Key Metrics are Health, Availability, and Risk. When you define a Key Metric, you are specifying how the value for a metric you created in step 4 translates to one of the standard Key Metric values. By default, all three Key Metrics are based on the default *Average Device Availability* metric.

6. **Define alerts and associated events**. Each alert and its associated event is triggered by a metric. In our example, we will define alerts for each metric.

Defining the Name of the IT Service Policy and its Basic Properties

To define the basic parameters of our example IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the **[Create]** button. The **IT Service Editor** page appears, with the **[Administration]** tab and **[Properties]** sub-tab selected:

Example Using Device Availability, Device Latency, and Process Availability

3. Supply the following values in the following fields:

- *IT Service Name*. Name of the IT Service policy. We entered "Acme: MS SQL database_servers".

- *IT Service Owner*. Automatically populated with your username.

- *Configuration Mode*. We selected *Basic Interface*. The *Basic Interface* allows you to quickly setup an IT Service policy.

- *Sharing Permissions*. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the ScienceLogic platform where the IT Service is visible. We selected *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the creator.

- *Permission Keys*. We did not select any permission keys.

- *Operational Status*. We selected *Aggregation enabled*.

- *Aggregation Frequency*. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value. We specified *Every 3 minutes*.

- *Raw Data Retention*. Specifies how long the ScienceLogic platform should store the raw data for the IT Service Policy. We accepted the default value.

- *Frequent Rollup Retention*. Deprecated field no longer used by the ScienceLogic platform.

- *Hourly Rollup Retention*. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. We accepted the default value.

- *Daily Rollup Retention*. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. We accepted the default value.

- *Description*. We did not enter a description.

4. Select the **[Save]** button to save the values in the **[Properties]** tab.

Defining a List of Devices for the IT Service Policy

After defining the name and basic properties of an IT Service Policy, you must next determine the devices to include in your IT Service policy. You do this in the **[Model]** sub-tab.

For example, if you want to monitor Email service, you could create a list of devices that includes Exchange servers, DNS servers, and devices that run Email round-trip policies.

You can manually assign devices and device groups to the IT Service device group, or you can use membership rules, like you would for a dynamic device group.

When you define the list of devices to include in your IT Service policy, that list of devices appears as a device group throughout the ScienceLogic platform.

There are three ways to add a device to the list of devices for the IT Service policy.

- Add a device group to the list of devices for the IT Service policy.

- Add a static list of one or more devices to the list of devices for the IT Service policy.
- Add a dynamic list of one or more devices to the list of devices for the IT Service policy.

In our example, we will add a *static list of devices* that includes two MS SQL servers to the IT Service policy.

To create the list of devices for the IT Service policy.

1. After performing the tasks in the *previous section*, select the **[Model]** sub-tab.
2. To add a static list of one or more devices to the list of devices for the IT Service policy, go to the **Static Devices** pane.
3. Select the **[Add]** button. The **Device Alignment** modal page appears and displays a list of all devices in the ScienceLogic platform.
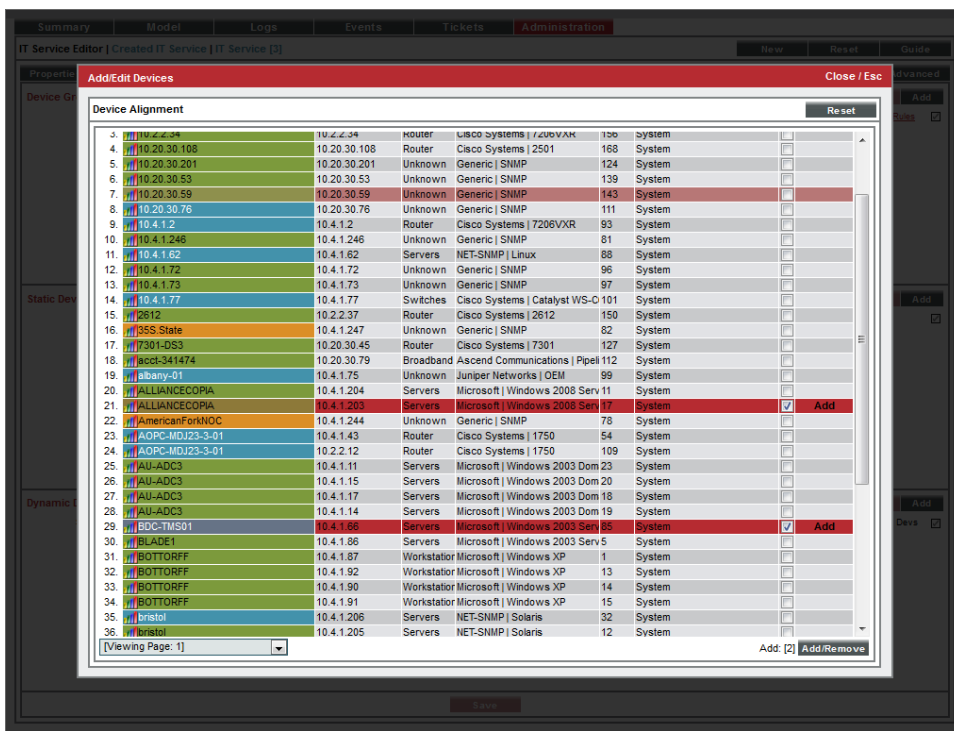
4. In the **Device Alignment** modal page, we selected the checkbox for devices "ALLIANCECOPIA" and "BDC-TMS01". Each is a device running the Windows operating system and MS SQL database server.
5. Select the **[Add/Remove]** button in the lower right.

6. The selected devices appear in the **Static Devices** pane:

7. Select the **[Save]** button to save the list of devices.

Defining Metrics for the IT Service Policy

A metric is a measurement that helps determine the status of an IT Service.

The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The Average Device Availability metric aggregates the current availability value (0 or 1) of all devices in the IT Service and calculates the average value. The aggregation is performed at the frequency specified in the *Aggregation Frequency* setting in the basic properties for the IT Service policy. The availability of a device is determined every 5 minutes.

Before you can define a metric, you must determine what parameters you want to monitor for the IT Service policy. You can use data from the following sources to monitor the IT Service:

- Device Availability
- Device Latency
- Overall CPU Usage
- Physical Memory Usage
- Swap Usage
- Device State (Condition of the device, based upon the most severe event generated by the device.)

- Device Count

- Presentation Objects from Dynamic Applications

- Network Interface

- TCP/IP Port Monitor

- System Process Monitor

- Windows Service Monitor

- Email Round Trip Monitor

- Web Content Monitor

- SOAP/XML Transaction Monitor

- Domain Name Monitor

Our example includes three metrics:

- Device Availability

- Device Latency

- System Process availability

We will create our metrics in **Basic mode**. We will edit the default metric and create two additional metrics. We will also define an alert/event for each metric.

Device Availability Metric

1. After performing the tasks in the *previous section*, select the [**Metrics**] sub-tab.

2. Ensure that you are in **Basic mode**. If you see the [**Alerting**] sub-tab, you are not in **Basic mode**. Click on the [**Advanced**] button to toggle to **Basic mode**.

Example Using Device Availability, Device Latency, and Process Availability

3. First, we will edit the default metric. In the **Service Metric Definitions** pane, find the metric *Average Device Availability* and select its wrench icon (🔧).
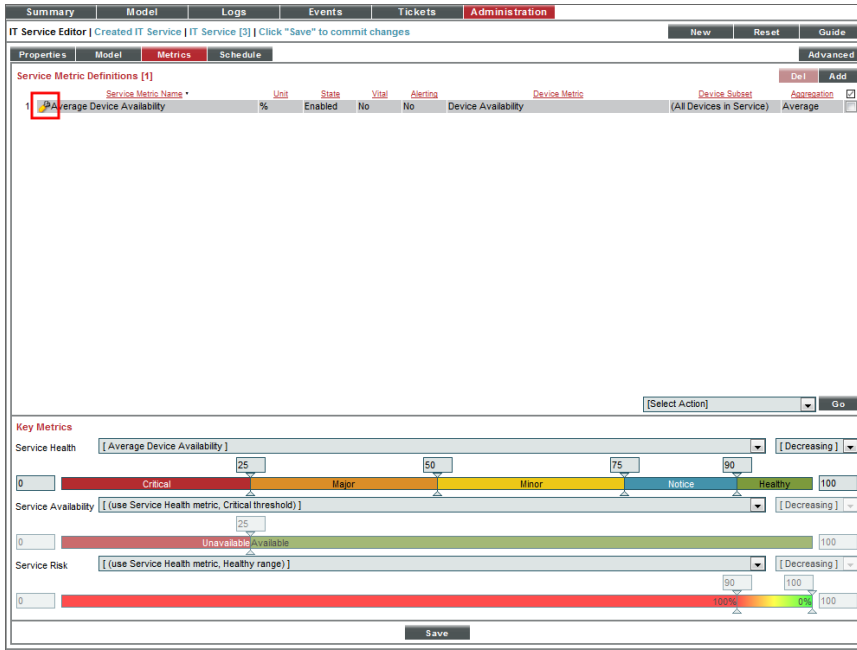
4. The **Service Metric Editor** modal page appears:

5. In the **Service Metric Editor** page, edit the following field:

- *Service Metric Name*. Enter *Average DB Device Availability*. This lets us know that the metric is measuring the device's availability, not the database server's availability.

6. In the lower pane, we'll define an alert for the metric. This alert specifies that when the availability of the two database servers falls below 75%, trigger an event with a severity of *Critical*. To define this alert and event, supply values in the following fields:

- *Metric Alerting*. Select *Single Threshold*.
- *Alert Policy Name*. Enter "DB_availability_too_low".
- *Event Severity*. Select *Critical*.
- *Increasing/Decreasing*. Select *Decreasing*.
- *Threshold*. Drag the slider to *75*.
- Select the **[OK]** button to save the metric.

Device Latency Metric

Next, we will define a new metric that examines the **latency of the devices** where the MS SQL servers reside. To do this:

1. Go to **Service Metric Definitions** pane and select the **[Add]** button.

2. The **Service Metric Editor** modal page appears. In this page, we'll define a metric that measures the latency of the two devices in our IT Service policy. We'll also define an alert that will trigger an event if the average latency of the two devices is greater than 30 milliseconds.

3. To create the new metric, supply the following values in the fields:

    - **Service Metric Name**. Enter "Average DB device latency".

    - **Metric Type**. Select *Internal*.

    - **Device Metric**. Select *Device Latency*.

    - For all other fields in the top pane, you can accept the default values.

4. In the lower pane, we'll define an alert for the metric. This alert specifies that when the average latency of the two devices where the database servers reside is greater than 30 milliseconds, it will trigger an event with a severity of *Critical*. To define this alert and event, supply values in the following fields:

    - **Metric Alerting**. Select *Single Threshold*.

    - **Alert Policy Name**. Enter "db_too_slow".

    - **Event Severity**. Select *Critical*.

    - **Increasing/Decreasing**. Select *Increasing*.

    - **Threshold**. Drag the slider to *30*.

5. Select the **[OK]** button to save the metric.

System Process Metric

Next, we will define a metric that makes sure that the process **sqlservr.exe** is running on both devices where the MS SQL databases reside.

1. Before we can define this metric in the **IT Service Editor**, we must tell the ScienceLogic platform to monitor the sqlservr.exe process, outside of the IT Service policy. To do this:

2. Go to the **Device Processes** page (Registry > Devices > Processes).

3. Use the *Device Name* column to search for the device *ALLIANCECOPIA*.

4. Use the *Process* column to search for the process *sqlservr.exe*.

    - Select the checkbox for **sqlservr.exe** running on **ALLIANCECOPIA**.

    - Select the *Select Actions* menu. Select *Enable (Create Policy)*.

    - Select the **[Go]** button.

5. Use the *Device Name* column to search for the device *BDC-TMS01*.

6. Use the *Process* column to search for the process *sqlservr.exe*.

    - Select the checkbox for **sqlservr.exe** running on **BDC-TMS01**.

    - Select the *Select Actions* menu. Select *Enable (Create Policy)*.

    - Select the **[Go]** button.

7. Go back to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the IT Service policy **Acme: MS SQL database server**. Select its wrench icon (🔧).

8. Select the **[Metrics]** tab. Select the **[Add]** button.

9. The **Service Metric Editor** modal page appears. In this page, we'll define a metric ensures that the process *sqlservr.exe* is running on both devices where the MS SQL database resides. We'll also define an alert that will trigger an event if the availability of the process *sqlservr.exe* is less than 99%. In other words, with the exceptions of very, very brief outages, the process *sqlservr.exe* should be running.

10. To create the new metric, supply the following values in the fields:

    - *Service Metric Name*. Enter *SQL Server process*.

    - *Metric Type*. Select *System Process Monitor*.

    - *Process Name*. Enter *sqlservr.exe*.

    - *Metric*. Select *Process Running*.

    - For all other fields in the top pane, you can accept the default values.

11. In the lower pane, we'll define an alert for the metric. This alert specifies that if the process *sqlservr.exe* is not running on both devices in the IT Service policy, trigger an event with a severity of *Critical*. To define this alert and event, supply values in the following fields:

    - *Metric Alerting*. Select *Single Threshold*.

    - *Alert Policy Name*. Enter "SQL Server process is not running".

    - *Event Severity*. Select *Critical*.

    - *Increasing/Decreasing*. Select *Increasing*.

    - *Threshold*. Drag the slider to 99.

12. Select the **[OK]** button to save the metric.

Defining Key Metrics for the IT Service Policy

Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data. For example, you can define "health" for a remote backup service and also define "health" for an Internet bandwidth service, even though you would use different criteria to measure the health of those two services.

All IT Service policies define how the ScienceLogic platform should calculate the following Key Metrics for the IT Service:

NOTE: The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The *Average Device Availability* metric specifies that the ScienceLogic platform should aggregate the availability data for all the devices in the policy and calculate the average availability.

- *Service Health*. The health of an IT Service can be one of the five standard severity values: Healthy, Notice, Minor, Major, or Critical. By default, the *Service Health* metric is aligned with the *Average Device Availability* metric.

- *Service Availability*. The availability of an IT Service can be either *available* or *unavailable*. By default, the *Service Availability* metric is aligned with the same metric as *Service Health*, converting *Critical* *Service Health* to *Unavailable* and all other *Service Health* values to *Available*.

- *Service Risk*. The risk of an IT Service is a percentage value that indicates how close an IT Service is to being in an undesirable state. By default, the *Service Risk* metric is aligned with the same metric as *Service Health*, converting the threshold between *Healthy* and *Notice* *Service Health* to *100%* and the healthiest possible value to *0%*.

The ScienceLogic platform generates an event if the *Service Health* Key Metric has a value of Notice, Minor, Major, or Critical, and/or if the *System Availability* key metric has a value of *unavailable*.

For more details on Key Metrics, see the chapter on *Creating, Editing, and Deleting IT Services*.

Using the three metrics we created in the *previous section*, we'll define the Key Metrics for our IT Service Policy:

1. Select the **[Metrics]** sub-tab.

2. In the top pane, you will see the default metric, *Average Device Availability*. If you have already defined additional custom metrics, they will also appear in the top pane.

3. In the bottom pane, you will see the three Key Metrics:
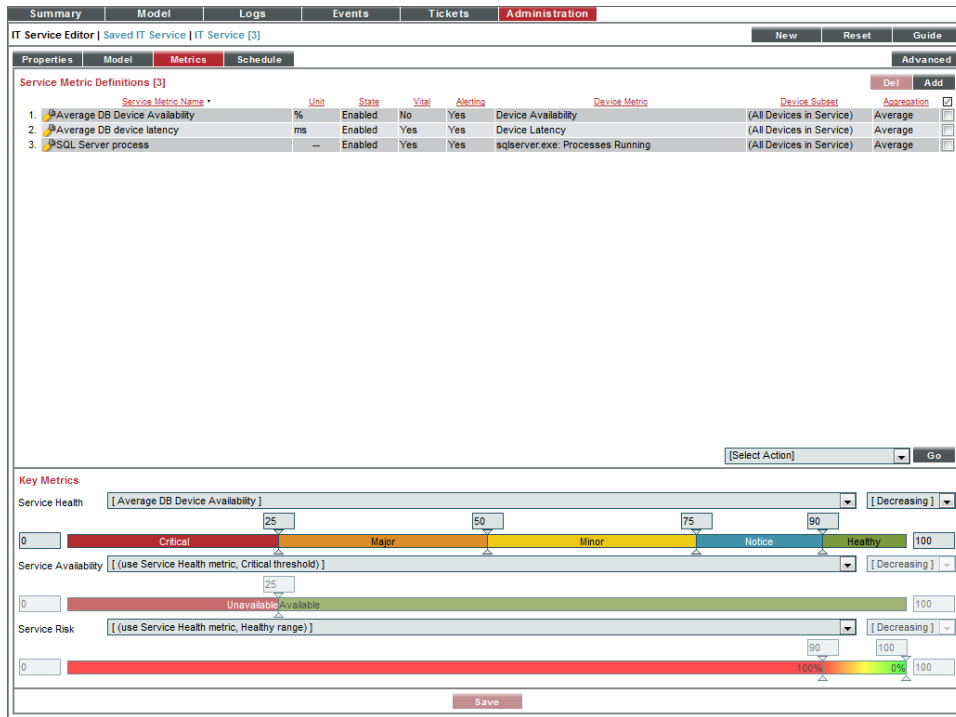
4. To edit each metric, supply the following values:

- *Service Health*. Appears in the **Health** column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Healthy*, *Notice*, *Minor*, *Major*, and *Critical*. By default, the Service Health Key Metric is based on the metric for Average Device Availability, with values set at 0-24 is Critical, 25-49 is Major, 50-74 is Minor, 75-89 is Notice, and 90 and above is Healthy.

  ○ In the drop-down list that appears above the **Service Health** Key Metric, select *Average DB device latency*.

  ○ Select *Increasing*.

  ○ Enter the following thresholds: *30, 40, 50, 60*.

  ○ If average latency is below 30 milliseconds, the IT Service policy will have a Health value of *Healthy*.

  ○ If average latency is between 30 - 40 milliseconds, the IT Service policy will have a Health value of *Notice*.

  ○ If average latency is between 40 - 50 milliseconds, the IT Service policy will have a Health value of *Minor*.

  ○ If average latency is between 50 - 60 milliseconds, the IT Service policy will have a Health value of *Major*.

Example Using Device Availability, Device Latency, and Process Availability

- If average latency is greater than 60 milliseconds, the IT Service policy will have a Health value of *Critical*.

- **Service Availability**. Appears in the ***Availability*** column in **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Available* and *Unavailable*. By default, the Service Availability Key Metric is based on the same metric as is used for the Service Health Key Metric. By default, 0 - 24 is *Unavailable* and 25 - 100 is *Available*.

  - In the drop-down list that appears above the **Service Availability** Key Metric, select *Average DB Device Availability*.

  - Select *Decreasing*.

  - Enter the threshold *75*.

  - If the average availability of the two devices in the IT Service policy falls below 75%, the IT Service policy will have an Availability value of *Unavailable*.

- **Service Risk**. Appears as a percentage in the ***Risk*** column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are 0% - 100%. By default, the Service Risk Key Metric is based on the same metric as is used for the Service Health Key Metric. By default, 0 - 89 is *Critical* and 90 - 100 is *Healthy*.

  - In the drop-down list that appears above the **Service Risk** Key Metric, select *SQL Server process*.

  - Select *Decreasing*.

  - Enter the threshold *99*.

  - If the process sqlservr.exe is not running an average of 99% of the time, the IT Service policy will have a Risk value of *100%*.

5. Select the **[Save]** button to save your changes to the Key Metrics.

Events for the IT Service Policy

To see the definitions for the events associated with each metric, go to the **Event Policy Manager** page (Registry > Events > Event Manager).

To find the event definitions, filter the **Event Policy Name** field by the name of the IT Service policy.

When an event for an IT Service is triggered, it displays the following message in the **Event Console**:
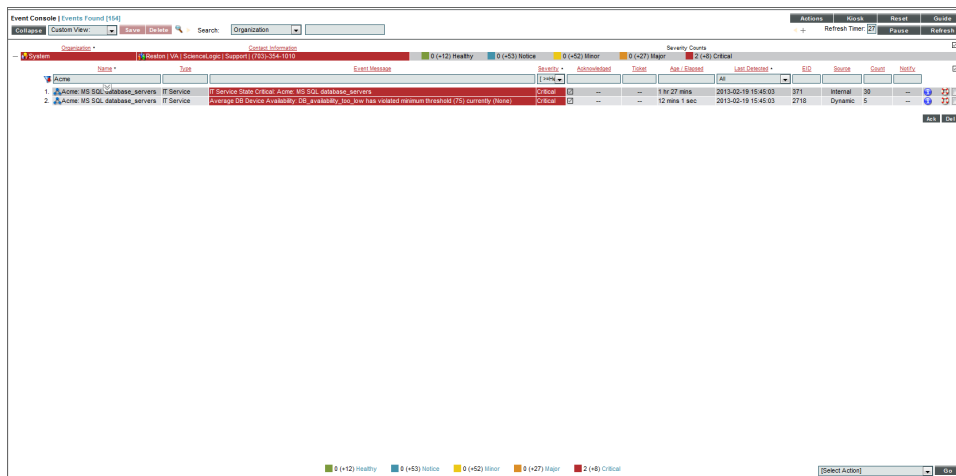
[name of metric] has violated threshold (%T) currently (%V).

where %T is the threshold you defined for the alert and %V is the current value for the metric.

The ScienceLogic platform generates an event if the **Service Health** key metric has a value of *Notice, Minor, Major, or Critical*. In the event above the metric that we associated with the Service Health Key Metric exceeded the threshold for the metric.

IT Service Dashboard

If you select the **[Summary]** tab for our example IT Service policy, you'll see something like the following:
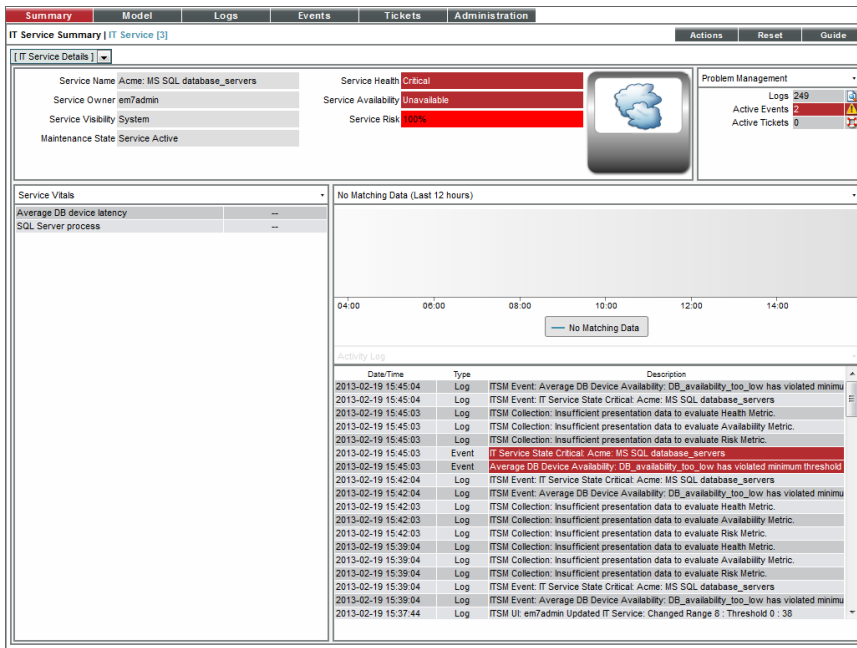
The **IT Service Summary** page displays the following:

- **IT Service Details**. Displays the following information about an IT Service:

  - *Service Name*

  - *Service Owner*

  - *Service Visibility*

  - *Maintenance State*

  - *Service Health*

  - *Service Availability*

  - *Service Risk*

- **IT Service Vitals**. Displays the current value for each Key Metric defined for an IT Service.

- **IT Service Problem Management**. Displays the number of logs, active events, and active tickets associated with an IT Service.

- **IT Service Health Last 12 Hours**. Displays a graph of the Availability metric. The y-axis displays percent availability. The x-axis displays time in one-hour increments.

- **IT Service Activity Log**. Displays a list of all current and past alerts and events associated with an IT Service.

Example

2

Example Using Device Availability and Interface Monitoring

Overview

This chapter will describe an example of an IT Service policy.

This example will monitor a leased WAN circuit. The example IT Service policy includes two routers that are at the two ends of the circuit.

Creating an IT Service Policy

To define an IT Service policy, you must:

1. *Define a service name and basic properties*. In this example, we will monitor the routers at both ends of a leased WAN circuit. The name of the IT Service policy will be "wan_circuit_1".

2. *Define a list of devices (model) for the IT Service that includes all the devices associated with the IT Service*. For example, if you want to monitor a WAN circuit, you could create a device group that includes the routers at each end of the circuit. You could create another device group that includes the switches that are connected to those routers. In our example, we'll select two routers to monitor.

3. *Optionally, define service sets. A service set is a sub-group of devices*. You can manually assign devices to a service set, or you can use membership rules, like you would for a dynamic device group. For example, you could define two service sets: Exchange Servers, defined by device class, and DNS servers, defined by the DNS Server running on each device. We don't use service sets in this example.

4. *Define Interface Tags for Interface Metrics*. If your IT Service policy will include interface metrics, you can use interface tags to create groups of interfaces. You can then apply a metric to a group of interfaces. Each interface can belong to multiple interface tags.

5. ***Define metrics. A metric is based on your business processes and examines all devices or one or more service sets to evaluate the state of the IT Service***. For each IT Service, the ScienceLogic platform provides a default metric called *Average Device Availability*, based on the availability of all devices in the IT Service. You can define additional metrics, based on default data collected by the ScienceLogic platform (availability, latency, CPU usage, memory usage, swap usage, device state, and device count), data collected by a Dynamic Application, and data about network interfaces, TCP/IP ports, system processes, Windows services, Email round-trip time, web content, SOAP/XML transactions, and DNS availability. For our example, we'll examine the traffic and errors on the interfaces on the routers at each end of the WAN circuit.

> **NOTE**: When the ScienceLogic platform evaluates a metric, it performs an aggregation, that is, the platform evaluates the data for all devices specified in the definition of the metric, over a specified time period (the **Aggregation Frequency**). Depending on the definition of the metric, the platform calculates the average, maximum, minimum, sum, standard deviation, or count value for all devices specified in the definition of the metric.

6. ***Define Key Metrics.*** Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data.The Key Metrics are Health, Availability, and Risk. When you define a Key Metric, you are specifying how the value for a metric you created in step 4 translates to one of the standard Key Metric values. By default, all three Key Metrics are based on the default *Average Device Availability* metric.

7. ***Define alerts and associated events***. Each alert and its associated event is triggered by a metric. In our example, we will define alerts for each metric.

Defining the Name of the IT Service Policy and its Basic Properties

To define the basic parameters of our example IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the **[Create]** button.

3. The **IT Service Editor** page appears, with the **[Administration]** tab selected.

4. Select the **[Properties]** sub-tab. Supply the following values in the following fields:

- **IT Service Name**. Name of the IT Service policy. We entered "wan_circuit_1"

- **IT Service Owner**. Automatically populated with your username.

- **Configuration Mode**. We selected *Basic Interface*. The *Basic Interface* allows you to quickly setup an IT Service policy.

- **Sharing Permissions**. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the platform where the IT Service is visible. We selected *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the creator.

- **Permission Keys**. We did not select any permission keys.

- **Operational Status**. We selected *Aggregation enabled*.

- **Aggregation Frequency**. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value. We specified *Every 2 minutes*.

- **Raw Data Retention**. Specifies how long the ScienceLogic platform should store the raw data for the IT Service Policy. We accepted the default value.

- **Frequent Rollup Retention**. Deprecated field no longer used by the ScienceLogic platform.

- *Hourly Rollup Retention*. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. We accepted the default value.

- *Daily Rollup Retention*. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. We accepted the default value.

- *Description*. We did not enter a description.

5. Select the **[Save]** button to save the values in the **[Properties]** tab.

Defining a List of Devices for the IT Service Policy

After defining the name and basic properties of an IT Service Policy, you must next determine the devices to include in your IT Service policy. You do this in the **[Model]** sub-tab.

For example, if you want to monitor Email service, you could create a list of devices that includes Exchange servers, DNS servers, and devices that run Email round-trip policies.

You can manually assign devices and device groups to the IT Service device group, or you can use membership rules, like you would for a dynamic device group.

When you define the list of devices to include in your IT Service policy, that list of devices appears as a device group throughout the ScienceLogic platform.

There are three ways to add a device to the list of devices for the IT Service policy.

- Add a device group to the list of devices for the IT Service policy.
- Add a static list of one or more devices to the list of devices for the IT Service policy.
- Add a dynamic list of one or more devices to the list of devices for the IT Service policy.

In our example, we will create a *static list of devices*.

To create the list of devices for the IT Service policy:

1. After performing the tasks in the *previous section*, select the **[Model]** sub-tab.

2. We will statically add two routers to our policy. To add a static list of one or more device to the list of devices for the IT Service policy, go to the **Static Devices** pane:

3. Select the **[Add]** button. The **Device Alignment** modal page displays a list of all devices in the ScienceLogic platform.
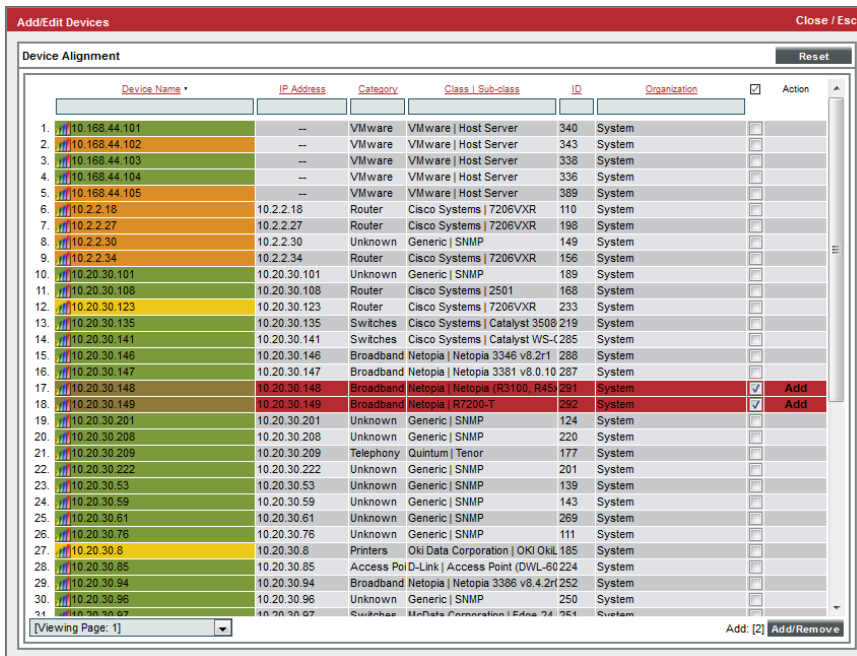
4. In the **Device Alignment** modal page, we selected the checkbox for devices "10.20.30.148" and "10.20.30.149". Each is a Netopia router with one WAN (ATM) interface and one Ethernet interface.

5. Select the **[Add/Remove]** button in the lower right. The selected devices will appear in the **Static Devices** pane.

6. Select the **[Save]** button to save the list of devices.

Defining Interface Tags for Interface Metrics

You can define interface metrics that monitor the following:

- Inbound Traffic
- Outbound Traffic
- Inbound Errors
- Outbound Errors
- Inbound Discards
- Outbound Discards

You can apply these metrics to:
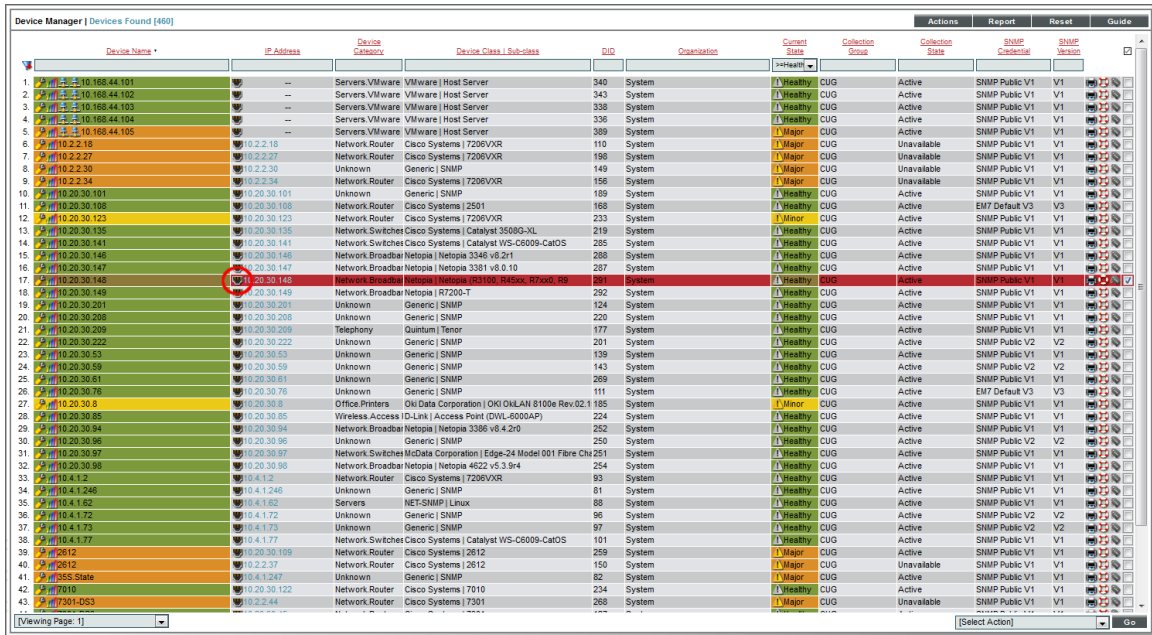
- All Interfaces
- Management Interface
- Tagged Interfaces

Interface Tags allow you to create one or more groups of interfaces. You can then apply an interface metric to that group. If *All Interfaces* or *Management Interface* doesn't suit your needs, you can define and apply interface tags.

In our example, we will create an interface tag for the interfaces at each end of the WAN link. When we define a metric, we will specify that interface tag to include only data from the interfaces at each end of the WAN link. To create the interface tags:
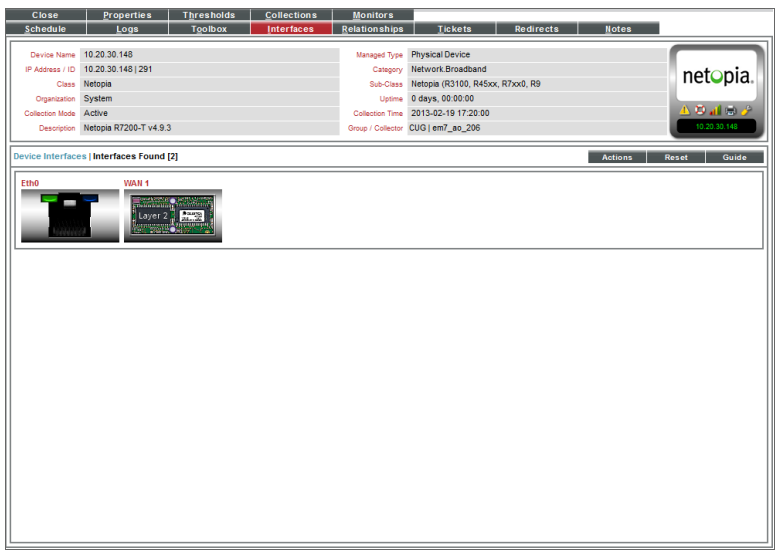
1. Go to the **Device Manager** page (Registry > Devices > Device Manager). Find the device where the interface resides. In our example, we'll search for *10.20.30.148* and *10.20.30.149*.

2. In the **Device Manager** page, find the first device (*10.20.30.148*). In the *IP Address* column, click on the interface icon ( ). The **Device Interfaces** page appears.

3. In our example, we will add the tag *wan_link* to the WAN interface on *10.20.30.148*. Click on the WAN interface. The **Interface Properties** page appears.

4. In the **Interface Properties** page, find the *Interface Tags* field. Select the wrench icon (🔧) to the right of the field.

5. The **Edit Network Interface Tags** modal page appears. In this page, enter the following:

- *Tags (comma separated)*. Enter *wan_link*.
- Select the **[Save]** button.
- Select the **[Close]** button.

Example Using Device Availability and Interface Monitoring

6. In the **Interface Properties** page, notice that the **Interface Tags** field contains the entry *wan_link*.

7. Repeat steps 1-6 for the WAN interface on the second device (**10.20.30.149**).

Defining Metrics for the IT Service Policy

A metric is a measurement that helps determine the status of an IT Service.

The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The *Average Device Availability* metric examines the availability of all devices in the IT Service. By default, the *Average Device Availability* metric is collected from every device every minute and "crunched" and averaged every 15 minutes.

Before you can define additional metrics for an IT Service policy, you must determine what parameters you want to monitor for the IT Service policy. You can use data from the following sources to monitor the IT Service:

- Device Availability
- Device Latency
- Overall CPU Usage
- Physical Memory Usage
- Swap Usage
- Device State (Condition of the device, based upon the most severe event generated by the device.)
- Device Count
- Presentation Objects from Dynamic Applications
- Network Interface
- TCP/IP Port Monitor
- System Process Monitor
- Windows Service Monitor
- Email Round Trip Monitor
- Web Content Monitor
- SOAP/XML Transaction Monitor
- Domain Name Monitor

Our example uses data from *Network Interface* monitoring. We will create our metrics in *Basic mode*. We will create a metric called *wan_inbound_errors* that will examine specified (tagged) interfaces for errors. We will also define an alert that tells the ScienceLogic platform to generate an event when the interfaces exceed the threshold for acceptable errors.

1. After performing the tasks in the previous sections, select the **[Metrics]** sub-tab.
2. Ensure that you are in *Basic mode*. If you see the **[Alerting]** sub-tab, you are not in *Basic mode*. Click on the **[Advanced]** button to toggle to *Basic mode*.
3. Select the **[Add]** button.

4. The **Service Metric Editor** modal page appears.

5. We will create the metric *wan_inbound_errors*. This metric will measure the number of inbound errors on the interfaces at each end of the WAN link. To create this metric, enter the following values in the **Service Metric Editor** modal page:

-   *Service Metric Name*. Enter "wan_inbound_errors".

-   *Metric Classification*. Specifies whether the metric will be displayed in the **IT Service Summary** page in widgets that display vital metrics. Select *Service Vital Metric*. The metric will appear in widgets that display vital metrics.

-   *Active State*. Specifies whether the ScienceLogic platform should currently collect data for the metric and evaluate alerts for the metric. Select *Enabled*.

-   *Metric Type*. Specifies the type of performance data you want to use for the metric. Select *Network Interface*. Our metric will examine data from network interfaces.

-   *Device Subset*. We have not define any *device subsets*. Select *All Devices in Service*.

-   *Aggregation*. Specifies how the ScienceLogic platform will aggregate ("crunch") the data collected from all the devices in the IT Service into a single value. Select *Average*.

- *Show only metrics available for this IT Service*. Leave this checkbox unselected. This checkbox filters the succeeding fields so that they display already-defined policies aligned with one or more of the devices in the IT Service or in the specified *Device Subset*. For example, if you selected *Dynamic App* in the *Metric Type* field, and then selected this checkbox, the *Dynamic Application* field would display only Dynamic Applications that are already aligned with one or more of the devices in the IT Service or in the specified *Device Subset*.

- *Interface Selection*. Select the network interfaces to include in the calculation for this metric. Select *Tagged Interfaces*. To calculate a value for this metric, the ScienceLogic platform should aggregate interface utilization statistics from the interfaces that are associated with a specific *tag* on all the devices in the IT Service.

- *Interface Tag*. Appears if you selected *Network Interface* in the *Metric Type* field. Select *wan_link*. This is the interface tag that we assigned to the interfaces at each end of the WAN circuit.

- *Interface Metric*. Select the interface measurement that the ScienceLogic platform should use to calculate the value for this metric. Select *Inbound Errors*. To calculate a value for the metric, the ScienceLogic platform aggregates the value for this interface measurement from all interfaces that you included in this metric using the method specified in the *Aggregation* field (Average, Minimum, Maximum, Sum, Standard Deviation, or Device Count).

6. Select the **[Save]** button to save your new metric.

Defining Alerts for the IT Service Policy

For each metric in an IT Service policy, you can define an associated alert and event. In our example, we will create an alert for the metric we created in the *previous section*. The alert will trigger an event when the inbound errors on the development router exceed the threshold of acceptable errors.

1. If the **IT Service Editor** page is not still open, go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the policy *WAN_circuit_1*. Select its wrench icon (🔧).

2. In the **IT Service Editor** page, select the **[Metrics]** sub-tab.

3. Ensure that you are in *Basic mode*. If you see the **[Alerting]** sub-tab, you are not in *Basic mode*. Click on the **[Advanced]** button to toggle between *Basic mode* and *Advanced mode*.

4. In the **Service Metrics Definitions** pane, find the metric *wan_inbound_errors*. Select its wrench icon (🔧).

5. In the **Service Metric Editor** modal page, go to the bottom pane. We will use the fields in the bottom pane to define an optional alert and optional event associated with the metric.

6. Enter values in the following fields:

    - *Alert Policy Name*. Enter "dev_router_too_many_inbound_errors". The ScienceLogic platform will automatically create an event policy that corresponds to this alert. This name will appear in the name of the event policy.

    - *Event Severity*. When the alert is generated, the ScienceLogic platform will trigger an event with the selected event severity. Select *Critical*.

    - *Decreasing/Increasing*. Toggles whether the alert is triggered when the value for the metric is above a specific threshold (Increasing) or below a specific threshold (Decreasing). Select *Increasing*.

    - *Alert Threshold*. Use sliders to define the threshold at which the alert should be generated and trigger an event and the threshold at which the alert should be reset and no longer trigger an event. Select *25*.

    - *Alert Range*. Accept the default values.

- *Event Policy Description*. Optionally enter cause and resolution text for the event. The text you supply in this field will be used to populate the **Policy Description** field in the **Event Policy Manager** for the event. If this event is triggered, the text you supply in this field will be displayed in the **Event Information** modal page for the event.

7. Select the **[Save]** button to save your new alert.

Defining Key Metrics for the IT Service Policy

Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services include metrics that aggregate very different performance data. For example, you can define "health" for a remote backup service and also define "health" for an Internet bandwidth service, even though you would use different criteria to measure the health of those two services.

All IT Service policies define how the ScienceLogic platform should calculate the following Key Metrics for the IT Service:

NOTE: The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The *Average Device Availability* metric specifies that the ScienceLogic platform should aggregate the availability data for all the devices in the policy and calculate the average availability.

- *Service Health*. The health of an IT Service can be one of the five standard severity values: Healthy, Notice, Minor, Major, or Critical. By default, the **Service Health** metric is aligned with the *Average Device Availability* metric.
- *Service Availability*. The availability of an IT Service can be either *available* or *unavailable*. By default, the **Service Availability** metric is aligned with the same metric as **Service Health**, converting *Critical* **Service Health** to *Unavailable* and all other **Service Health** values to *Available*.
- *Service Risk*. The risk of an IT Service is a percentage value that indicates how close an IT Service is to being in an undesirable state. By default, the **Service Risk** metric is aligned with the same metric as **Service Health**, converting the threshold between *Healthy*and *Notice* **Service Health** to *100%* and the healthiest possible value to *0%*.

For more details on Key Metrics, see the chapter on *Creating, Editing, and Deleting IT Services*.

To edit the definitions of each Key Metric for our example IT Service policy:

1. If the **IT Service Editor** page is not still open, go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the policy *WAN_circuit_1*. Select its wrench icon (🔧).
2. In the **IT Service Editor** page, select the **[Metrics]** sub-tab.

3. In the bottom pane, you will see the three Key Metrics:

4. To edit the Key Metrics for our example IT Service policy:

- **Service Health**. *This example uses the default values* for this Key Metric. This Key Metric appears in the **Health** column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Healthy*, *Notice*, *Minor*, *Major*, and *Critical*.

- **Service Availability**. This Key Metric appears in the **Availability** column in **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Available* and *Unavailable*.

  ○ From the drop-down list that appears above the **Service Availability** Key Metric, select **wan_inbound_errors**. The **Service Availability** Key Metric will now examine the metric **wan_inbound_errors** to determine the availability of the IT Service.

  ○ From the drop-down list that appears to the right of the **Service Availability** Key Metric, select **Increasing**.

  ○ Move the slider to **25**. If there are more than 25 errors, the service will be considered unavailable.

  ○ Accept the default minimum range and maximum range.

- **Service Risk**. This Key Metric appears as a percentage in the **Risk** column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are 0% - 100%.

  ○ From the drop-down list that appears above the **Service Risk** Key Metric, select **wan_inbound_errors**. The **Service Risk** Key Metric will now examine the metric **wan_inbound_errors** to determine the risk to the IT Service.

  ○ From the drop-down list that appears to the right of the **Service Risk** Key Metric, select **Increasing**.

  ○ Move the 0% slider to **0**. Move the 100% slider to **25**. The **Service Risk** metric will now show how at risk the service is, with 0% risk being completely healthy (no errors) and 100% risk being unavailable (25 errors).

  ○ Accept the default minimum range and maximum range.
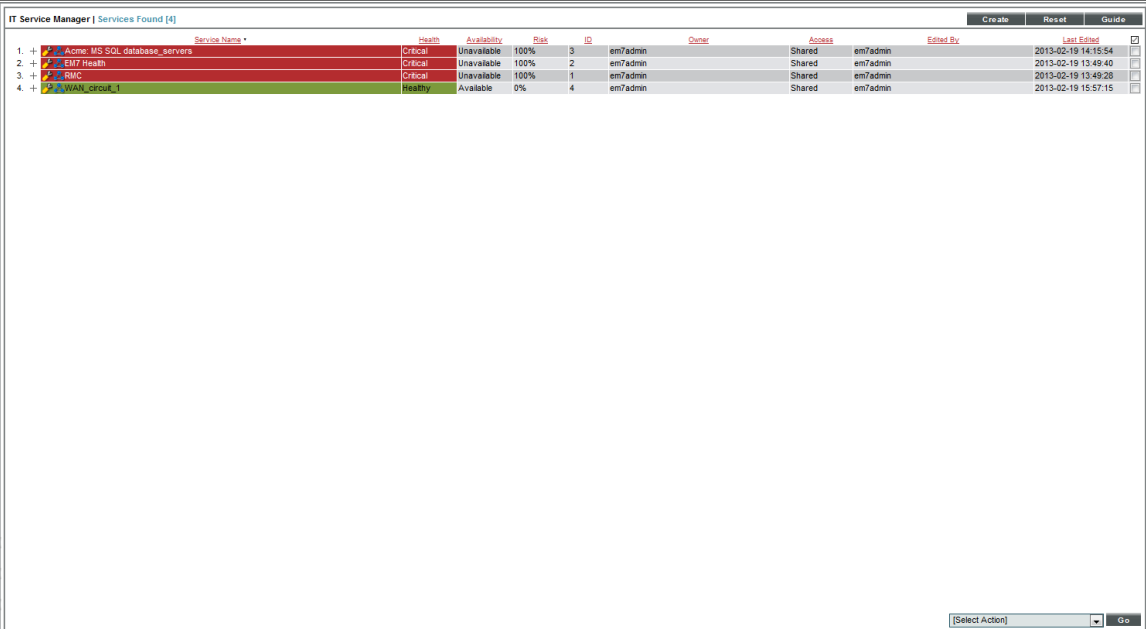
5. Select the **[Save]** button to save the changes to the Key Metrics.

Viewing Information about the IT Service Policy

IT Service Manager

The **IT Service Manager** page displays overview information each IT Service policy. To view the **IT Service Manager** page:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).
2. Find the policy *WAN_circuit_1*.

3. The **IT Service Manager** displays the following about the IT Service policy.

   - *Service Name*. Name of the policy. The color indicates the severity of the most severe event associated with the IT Service policy.
   - *Health*. This is a default Key Metric for each IT Service policy. This metric specifies the overall health of the IT Service. Possible values are: *Critical*, *Major*, *Minor*, *Notice*, and *Healthy*.
   - *Availability*. This is a Key Metric for each IT Service policy. This metric specifies the overall availability of the IT Service. Possible values are: *Available* or *Unavailable*.
   - *Risk*. This is a Key Metric for each IT Service policy. This metric specifies the overall risk to the IT Service. Possible values are 0% - 100%, in integer values

Example Using Device Availability and Interface Monitoring

IT Service Summary

The **IT Service Summary** page allows you to view the IT Service Dashboards that have been configured for the selected IT Service. By default, each IT Service Policy includes the *IT Service Details* dashboard. To access the **IT Service Summary** page:

To view the **IT Service Summary** page:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Find the policy *WAN_circuit_1*. Select its map icon ( ).

3. The **IT Service Summary** page for our example IT Service policy displays the default *IT Service Details* dashboard for the IT Service:

Viewing Additional Information

For instructions on how to view *information about the devices* in an IT Service policy, view the *events* associated with an IT Service policy, view the *tickets* associated with an IT Service, and view the *log messages* associated with an IT Service policy, see the chapter on *Viewing IT Services*.

Example

3

Example Using Service Dependencies

Overview

This chapter will provide a working example of an IT Service policy.

This example will include a **Service Dependency**. The example will create a metric for availability for all devices in the Acme company. The example will use metrics from other IT Service policies to calculate device availability for all devices in the company. Specifically, the example policy **acme_all_devices**, will use the availability metrics and list of devices from the IT Service policies **acme_east_coast_devices** and **acme_west_coast_devices** to determine the availability of all devices in the Acme company.

This example is abbreviated and specifically highlights using a Service Dependency. For detailed examples of other aspects of IT Service policy, see the chapter on *Creating, Editing, and Deleting IT Services* or *Example 1* or *Example 2*.

Creating an IT Service Policy

To define the example IT Service policy in this example, you must:

1. Define two simple IT Service policies to use as Service Dependencies.

2. **Define a service name and basic properties**. In this example, we will monitor the availability of devices at two corporate offices. The name of the IT Service policy will be *acme_all_devices*.

3. **Define a list of Service Dependencies (model) for the IT Service that includes all the external IT Service policies you want to reference**. In our example, we'll use the data from two different IT Service policies to create a metric in our example IT Service policy. The metric in our example IT Service policy will include all the devices and availability data from our two Service Dependencies.

4. ***Define metrics. A metric is based on your business processes and examines all devices or one or more service sets to evaluate the state of the IT Service***. For our example, we'll examine the availability of multiple groups of devices and create a metric for availability of the sum of all devices.

> **NOTE**: When the ScienceLogic platform evaluates a metric, it performs an aggregation, that is, the platform evaluates the data for all devices specified in the definition of the metric, over a specified time period (the ***Aggregation Frequency***). Depending on the definition of the metric, the platform calculates the average, maximum, minimum, sum, standard deviation, or count value for all devices specified in the definition of the metric.

5. ***Define alerts and associated events***. Each alert and its associated event is triggered by a metric. In our example, we will not define alerts and events. For examples of defining alerts and events, see the chapter on *Creating, Editing, and Deleting IT Services* or *Example 1* or *Example 2*.

Defining the Two External IT Service Policies

First, we must define the two IT Service policies that we will use as Service Dependencies. Each IT Service policy will monitor the availability of three devices.

To define the basic parameters of our two external IT Service policies, follow the steps in the following two sub-sections.

Defining "acme_east_coast_devices"

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).
2. Select the **[Create]** button. This new IT Service policy will monitor the availability of devices in a branch office of the Acme company.

Example Using Service Dependencies

3. The **IT Service Editor** page appears, with the **[Administration]** tab selected.

4. Select the **[Properties]** sub-tab. Supply the following values in the following fields:

- *IT Service Name*. Name of the IT Service policy. We entered "acme_east_coast_devices".

- *IT Service Owner*. Automatically populated with your username.

- *Configuration Mode*. We selected *Basic Interface*. The *Basic Interface* allows you to quickly setup an IT Service policy.

- *Sharing Permissions*. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the ScienceLogic platform where the IT Service is visible. We selected *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the creator.

- *Permission Keys*. We did not select any permission keys.

- *Operational Status*. We selected *Aggregation enabled*.

- *Aggregation Frequency*. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value. We specified *Every 5 minutes*.

- *Raw Data Retention*. Specifies how long the ScienceLogic platform should store the raw data for the IT Service policy. We accepted the default value.

- *Frequent Rollup Retention*. Deprecated field no longer used by the ScienceLogic platform.

- *Hourly Rollup Retention*. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. We accepted the default value.

- *Daily Rollup Retention*. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. We accepted the default value.

- *Description*. We did not enter a description.

5. Select the **[Save]** button to save the values in the **[Properties]** tab.

6. Select the **[Model]** tab.

7. We will statically add three servers to our policy. To add a static list of one or more devices to the list of devices for the IT Service policy, go to the **Static Devices** pane.

8. Select the **[Add]** button. The **Device Alignment** modal page displays a list of all devices in the ScienceLogic platform.

9. In the **Device Alignment** modal page, we selected the checkbox for devices *ALLIANCECOPIA*, *AU-ADC3*, and *BLADE1*. Each is a Windows server.

10. Select the **[Add/Remove]** button in the lower right. The selected devices will appear in the **Static Devices** pane.

11. Select the **[Save]** button to save the list of devices

12. We will use the default metric automatically created by the ScienceLogic platform, *Average Device Availability*, so we don't have to create a custom metric.

Defining "acme_west_coast_devices"

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the **[Create]** button. This new IT Service policy will monitor the availability of devices in a branch office of the Acme company.

3. The **IT Service Editor** page appears, with the **[Administration]** tab selected.

4. Select the **[Properties]** sub-tab. Supply the following values in the following fields:

- *IT Service Name*. Name of the IT Service policy. We entered "acme_west_coast_devices".

- *IT Service Owner*. Automatically populated with your username.

- *Configuration Mode*. We selected *Basic Interface*. The *Basic Interface* allows you to quickly setup an IT Service policy.

- *Sharing Permissions*. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the ScienceLogic platform where the IT Service is visible. We selected *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the creator.

- *Permission Keys*. We did not select any permission keys.

- *Operational Status*. We selected *Aggregation enabled*.

- *Aggregation Frequency*. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value. We specified *Every 5 minutes*.

- *Raw Data Retention*. Specifies how long the ScienceLogic platform should store the raw data for the IT Service policy. We accepted the default value.

- *Frequent Rollup Retention*. Deprecated field no longer used by the ScienceLogic platform.

- *Hourly Rollup Retention*. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. We accepted the default value.

- *Daily Rollup Retention*. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. We accepted the default value.

- *Description*. We did not enter a description.

5. Select the **[Save]** button to save the values in the **[Properties]** tab.

6. Select the **[Model]** tab.

7. We will statically add three devices to our policy. To add a static list of one or more devices to the list of devices for the IT Service policy, go to the **Static Devices** pane.

8. Select the **[Add]** button. The **Device Alignment** modal page displays a list of all devices in the ScienceLogic platform.


9. In the **Device Alignment** modal page, we selected the checkbox for devices "10.2.2.13", "10.2.2.37", and "10.2.2.9". Each is a Cisco network device.

10. Select the **[Add/Remove]** button in the lower right. The selected devices will appear in the **Static Devices** pane.

11. Select the **[Save]** button to save the list of devices

12. We will use the default metric automatically created by the ScienceLogic platform, *Average Device Availability*, so we don't have to create a custom metric.

Defining the Name of the IT Service Policy and its Basic Properties

Our example IT Service policy, *acme_all_devices*, will examine the devices and metrics from the *two external IT Service policies* to calculate availability for all devices in the Acme company. To define the basic parameters of our example IT Service policy:

1. Go to the **IT Service Manager** page Registry > IT Services > IT Service Manager).

2. Select the **[Create]** button.

3. The **IT Service Editor** page appears, with the **[Administration]** tab selected.

Example Using Service Dependencies

4. Select the **[Properties]** sub-tab. Supply the following values in the following fields:
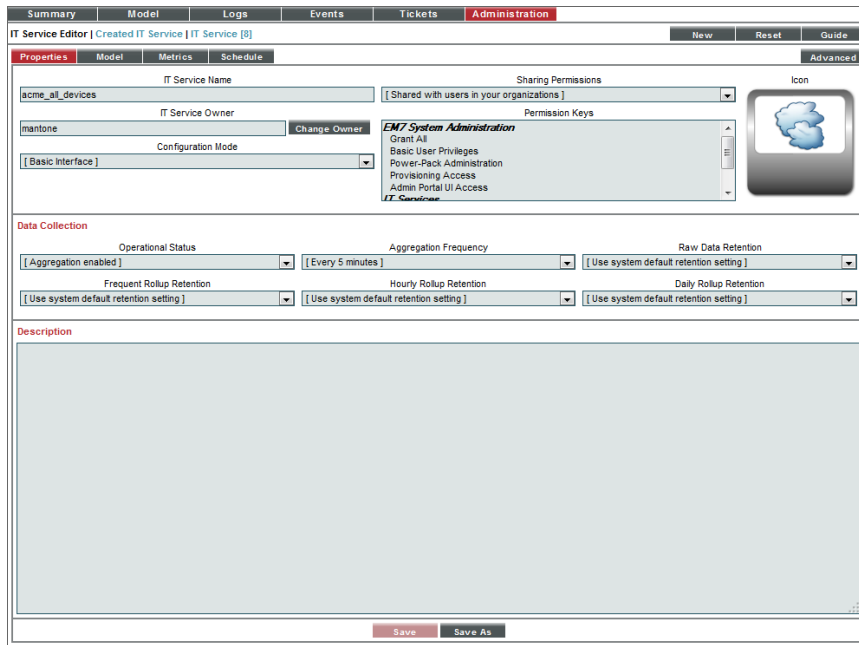
- *IT Service Name*. Name of the IT Service policy. We entered "acme_all_devices".

- *IT Service Owner*. Automatically populated with your username.

- *Configuration Mode*. We selected *Basic Interface*. The *Basic Interface* allows you to quickly setup an IT Service policy.

- *Sharing Permissions*. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the ScienceLogic platform where the IT Service is visible. We selected *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the creator.

- *Permission Keys*. We did not select any permission keys.

- *Operational Status*. We selected *Aggregation enabled*.

- *Aggregation Frequency*. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value. We specified *Every 5 minutes*.

- *Raw Data Retention*. Deprecated field no longer used by the ScienceLogic platform.

- *Frequent Rollup Retention*. Specifies how long the ScienceLogic platform should store the "frequent" normalized data for the IT Service policy. We accepted the default value.

- *Hourly Rollup Retention*. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. We accepted the default value.

- *Daily Rollup Retention*. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. We accepted the default value.

- *Description*. We did not enter a description.

5. Select the **[Save]** button to save the values in the **[Properties]** tab.

# Defining a List of Service Dependencies for the IT Service Policy

After defining the name and basic properties of the example IT Service policy, we next define the Service Dependencies. You do this in the **[Model]** sub-tab.

To create the Service Dependencies for the IT Service policy:

1. After performing the tasks in the previous sections, go to the **IT Service Editor** page for the example IT Service policy and select the **[Model]** sub-tab.

2. Select the **[Advanced]** button.

3. Go to the **Service Dependencies** pane. Select the **[Add]** button.
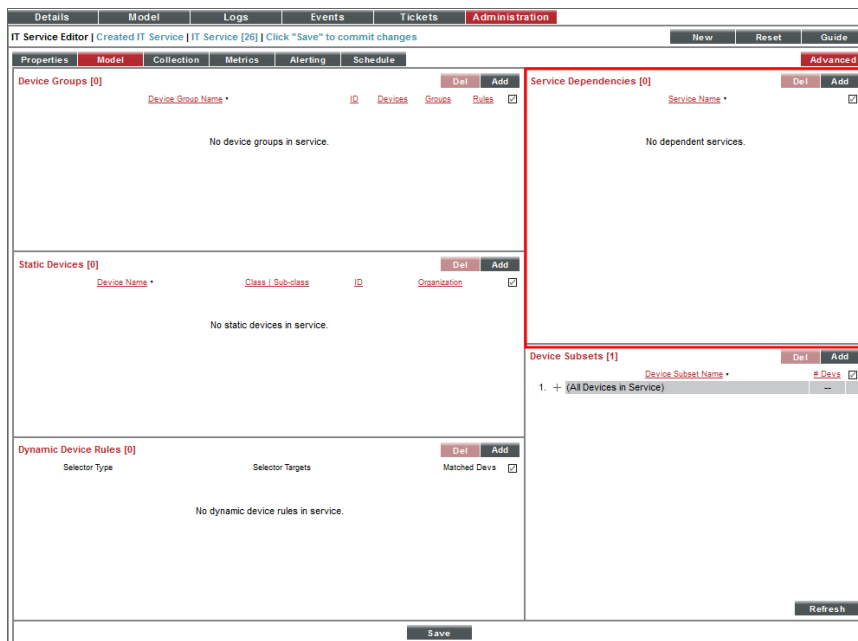
4. The **Service Dependency Alignment Editor** modal page appears. The **Service Dependency Alignment Editor** modal page displays a list of all IT Services in the ScienceLogic platform that you are allowed to view.

5. To align the two external IT Services in this example, we selected the checkbox for "acme_east_coast_devices " and "acme_west_coast_devices"

6. Select the **[Add/Remove]** button in the lower right. The selected IT Service policies will appear in the **Service Dependencies** pane.

7. Select the **[Save]** button to save the list of service dependencies.

Defining Metrics for the IT Service Policy

A metric is a measurement that helps determine the status of an IT Service.

The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The *Average Device Availability* metric examines the availability of all devices in the IT Service. By default, the *Average Device Availability* metric is collected from every device every minute and "crunched" and averaged every 15 minutes.

Our example uses data from two external IT Service policies (acme_east_coast_devices and acme_west_coast_devices) to determine the availability of all devices in the Acme network.

- We will define a metric for *east_coast_device_availability* that is a link to the *acme_east_coast_devices* IT Service policy.

- We will define a metric for *west_coast_device_availability* that is a link to the *acme_west_coast_devices* IT Service policy.

- We will define a metric for **all_devices_availability** that will examine **east_coast_device_availability** and **west_coast_device_availability** and calculate the average availability for the sum of all devices on the east coast plus all the devices on the west coast.

Defining east_coast_device_availability

1. Go to the **IT Service Editor** page and select the **[Metrics]** sub-tab.
2. Ensure that you are in *Basic mode*. If you see the **[Alerting]** sub-tab, you are not in *Basic mode*. Click on the **[Advanced]** button to toggle to *Basic mode*.
3. Select the **[Add]** button.
4. The **Service Metric Editor** modal page appears.

5. We will create the metric **east_coast_device_availability**. This metric will point to the service dependency for **acme_east_coast_devices** and calculate the average device availability of the devices in that external IT Service policy. Because we defined **acme_east_coast_devices** as a service dependency, the ScienceLogic platform will use the availability metric defined in the **acme_east_coast_devices** IT Service policy and will use the list of devices in the **acme_east_coast_devices** IT Service policy.

6. To create this metric, enter the following values in the **Service Metric Editor** modal page:

   - *Service Metric Name*. Enter "east_coast_device_availability".

   - *Metric Classification*. Specifies whether the metric will be displayed in the **IT Service Summary** page in widgets that display vital metrics. Select *Service Vital Metric*. The metric will appear in widgets that display vital metrics.

Example Using Service Dependencies

- **Active State**. Specifies whether the ScienceLogic platform should currently collect data for the metric and evaluate alerts for the metric. Select *Enabled*.

- **Metric Type**. Specifies the type of performance data you want to use for the metric. Select *Aligned Service Dependency*. Our metric will examine devices in a service dependency. Because the metric is based on a metric that is defined in another IT Service policy, the **Device Subset**, **Aggregation**, and **Show only metrics available for this IT Service** fields are not applicable and are grayed out.

- **Service Dependency Metric**. This field displays a list of service dependencies for the current IT Service policy. Select *acme_east_coast_devices*.

- **Presentation**. This field displays a list of all metrics in the **acme_east_coast_devices** IT Service policy. Select *Average Device Availability*.

7. Select the **[Save]** button to save your new metric.

Defining west_coast_device_availability

1. Go to the **IT Service Editor** page and select the **[Metrics]** sub-tab.

2. Ensure that you are in **Basic mode**. If you see the **[Alerting]** sub-tab, you are not in **Basic mode**. Click on the **[Advanced]** button to toggle to *Basic mode*.

3. Select the **[Add]**button.

4. The **Service Metric Editor** modal page appears.
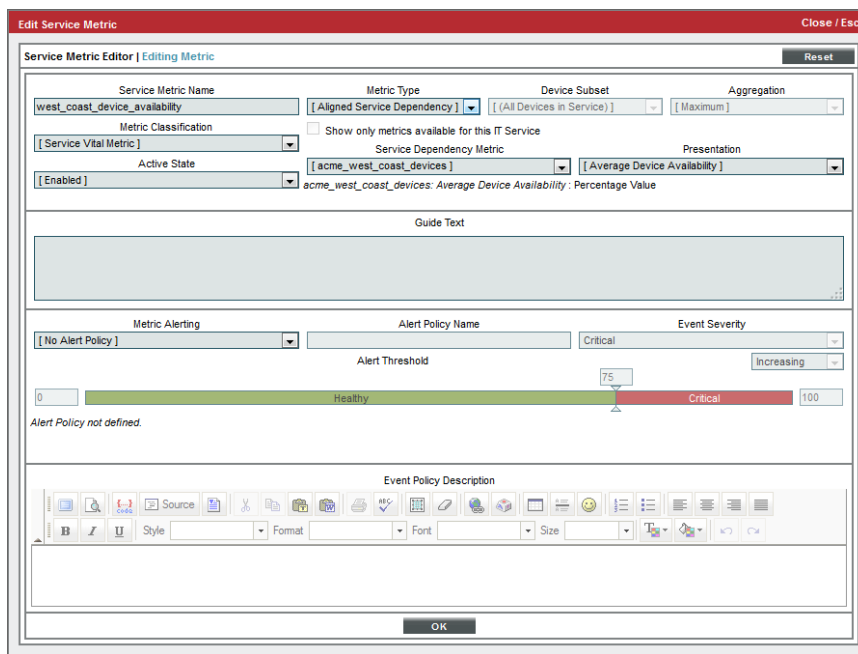
5. We will create the metric **west_coast_device_availability**. This metric will point to the service dependency for **acme_west_coast_devices** and calculate the average device availability of the devices in that external IT Service policy. Because we defined **acme_west_coast_devices** as a service dependency, the ScienceLogic

platform will use the availability metric defined in the **_acme_west_coast_devices_** IT Service policy and will use the list of devices in the **_acme_west_coast_devices_** IT Service policy

6. To create this metric, enter the following values in the **Service Metric Editor** modal page:

- *Service Metric Name*. Enter "west_coast_device_availability".

- *Metric Classification*. Specifies whether the metric will be displayed in the **IT Service Summary** page in widgets that display vital metrics. Select *Service Vital Metric*. The metric will appear in widgets that display vital metrics.

- *Active State*. Specifies whether the ScienceLogic platform should currently collect data for the metric and evaluate alerts for the metric. Select *Enabled*.

- *Metric Type*. Specifies the type of performance data you want to use for the metric. Select *Aligned Service Dependency*. Our metric will examine devices in a service dependency.

> **NOTE**: Because this metric is based on a metric that is defined in another IT Service policy, the **_Device Subset_**, **_Aggregation_**, and **_Show only metrics available for this IT Service_** fields are not applicable and are grayed out.

- *Service Dependency Metric*. This field displays a list of service dependencies for the current IT Service policy. Select *acme_west_coast_devices*.

- *Presentation*. This field displays a list of all metrics in the **_acme_west_coast_devices_** IT Service policy. Select *Average Device Availability*.

7. Select the **[Save]** button to save your new metric.

Defining all_devices_availability

1. Go to the **IT Service Editor** page and select the **[Metrics]** sub-tab.
2. Select the **[Advanced]** button. Ensure that you are in **_Advanced_** mode. If you see the **[Alerting]** sub-tab, you are in **_Advanced_** mode.
3. Select the **[Add]** button.
4. The **Service Metric Editor** modal page appears.
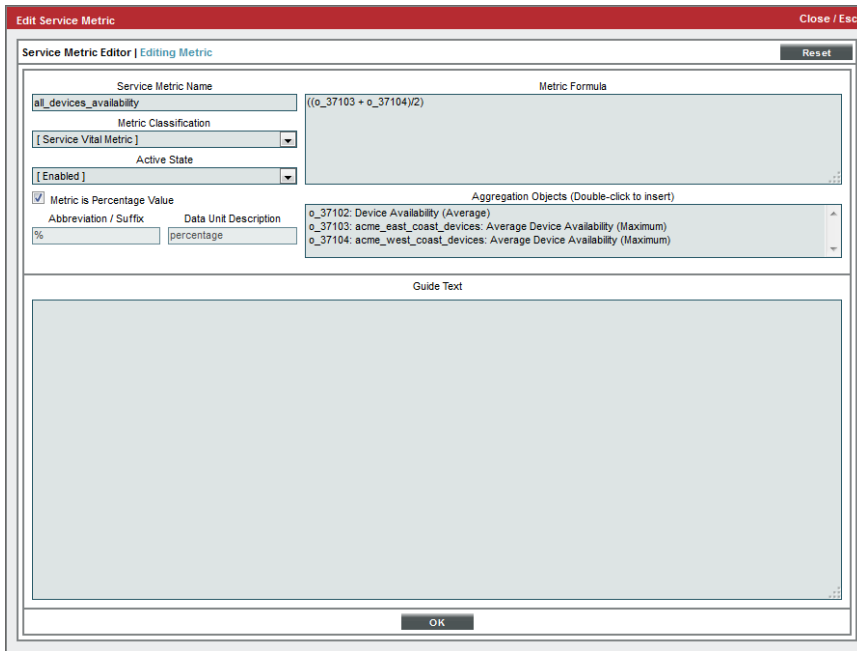
5. We will create the metric **all_devices_availability**. This metric will use the aggregation objects from the metrics **east_coast_device_availability** and **west_coast_device_availability** to calculate the average availability for the sum of all devices on the east coast plus all the devices on the west coast

6. To create this metric, enter the following values in the **Service Metric Editor** modal page:

    - *Service Metric Name*. Enter "all_devices_availability".

    - *Metric Classification*. Specifies whether the metric will be displayed in the **IT Service Summary** page in widgets that display vital metrics. Select *Service Vital Metric*. The metric will appear in widgets that display vital metrics.

    - *Active State*. Specifies whether the ScienceLogic platform should currently collect data for the metric and evaluate alerts for the metric. Select *Enabled*.

    - *Metric is Percentage Value*. If selected, the next two fields are populated automatically. If not selected, you can supply custom values in the next two fields.

    - *Abbreviation/Suffix*. Populated automatically.

    - *Data Unit Description*. Populated automatically.

    - *Metric Formula*. Enter:

      $((o\_XXXX + o\_YYYY)/2)$

      Where XXXX is the number displayed for east_coast_devices in the **Aggregation Objects** field and YYYY is the number displayed for the west_coast_devices in the **Aggregation Objects** field. This formula tells the ScienceLogic platform to add the average for all devices in the east_coast_ availability metric with the average for all devices in the west_coast_availability metric and divide the sum by two.

> **NOTE**: If you don't want to manually enter the object ID for each aggregation object, you can add the object ID for by double-clicking on the aggregation object in the **Aggregation Objects** field.

- **Guide Text**. Leave blank.

7. Select the **[Save]** button to save your new metric.

> **NOTE**: We didn't define any custom alerts for this example IT Service policy. You could optionally define alerts in the **[Alerts]** tab.

Defining Key Metrics for the IT Service Policy

Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services include metrics that aggregate very different performance data. For example, you can define "health" for a remote backup service and also define "health" for an Internet bandwidth service, even though you would use different criteria to measure the health of those two services.

All IT Service policies define how the ScienceLogic platform should calculate the following Key Metrics for the IT Service:

> **NOTE**: The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called **Average Device Availability**. The **Average Device Availability** metric specifies that the ScienceLogic platform should aggregate the availability data for all the devices in the policy and calculate the average availability.

- **Service Health**. The health of an IT Service can be one of the five standard severity values: Healthy, Notice, Minor, Major, or Critical. By default, the **Service Health** metric is aligned with the **Average Device Availability** metric.
- **Service Availability**. The availability of an IT Service can be either *available* or *unavailable*. By default, the **Service Availability** metric is aligned with the same metric as **Service Health**, converting *Critical* **Service Health** to *Unavailable* and all other **Service Health** values to *Available*.
- **Service Risk**. The risk of an IT Service is a percentage value that indicates how close an IT Service is to being in an undesirable state. By default, the **Service Risk** metric is aligned with the same metric as **Service Health**, converting the threshold between *Healthy* and *Notice* **Service Health** to *100%* and the healthiest possible value to *0%*.

For more details on Key Metrics, see the chapter on *Creating, Editing, and Deleting IT Services*.

We will edit the definition of the **Service Health** metric. We will accept the default value of the other metrics.

To edit the definition of the *Service Health* key metric:

1. If the **IT Service Editor** page is not still open, go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Find the policy *all_devices_availability*. Select its wrench icon (🔧).

2. In the **IT Service Editor** page, select the **[Metrics]** sub-tab.

3. In the bottom pane, you will see the three Key Metrics:
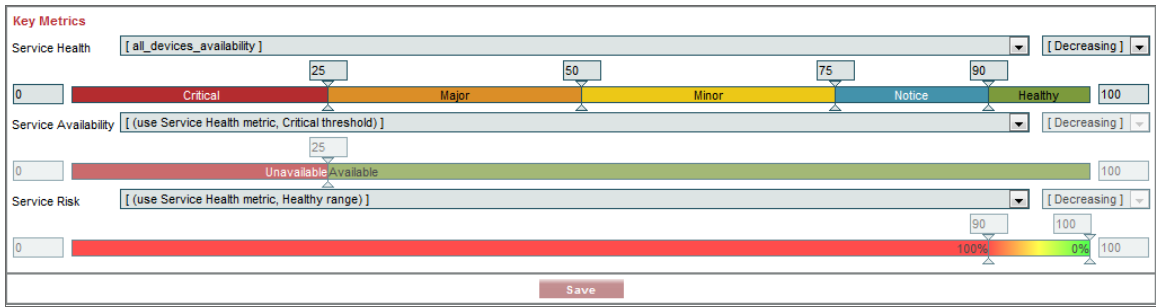
4. To edit the Key Metrics for our example IT Service policy:

   - *Service Health*. This Key Metric appears in the *Health* column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Healthy*, *Notice*, *Minor*, *Major*, and *Critical*.

     ○ From the drop-down list that appears above the *Service Health* Key Metric, select *all_devices_ availability*. The *Service Health* Key Metric will now examine the metric *all_devices_availability* to determine the availability of the IT Service.

     ○ From the drop-down list that appears to the right of the *Service Availability* Key Metric, select *Decreasing*.

   - *Service Availability*. This Key Metric appears in the *Availability* column in **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Available* and *Unavailable*. For this example, the default values are accepted.

   - *Service Risk*. This Key Metric appears as a percentage in the *Risk* column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are 0% - 100%. For this example, the default values are accepted.

Example

4

Using an SLA Definition with an IT Service Policy

Overview

In this chapter, we will define a policy to monitor web content, define an IT Service policy that uses the web-content monitor, define an SLA Definition, and then generate a dashboard report and a quick report for the IT Service policy and the SLA Definition.

In our example:

- We will create an IT Service policy that monitors a web server. We want to ensure that the web server will return a request within 8 seconds. We want to see this type of performance 99.99% of the time.

- We will first create a web-content monitoring policy. This policy will be aligned with the web server we want to monitor. We will search a web site for a string and monitor the time it takes to send the request and receive a reply.

- We will create an IT Service policy. This policy will monitor the web server that we aligned with the web content monitoring policy.

- The IT Service policy will include a metric that is based on the web content monitoring policy.

- The IT Service policy will use the metric (based on the web content monitoring policy) to define availability of the IT Service.

- We will define an SLA that says "the web server should return a request within 8 seconds. We want to see this type of performance 99.99% of the time".

- 99.99% uptime allows for 432 minutes of downtime per month.

- Because our IT Service policy has a polling frequency of 5 minutes, the web servers can be unavailable (transaction time is greater than 8 seconds) no more than 86 polling periods per month (and still meet 99.99% uptime).

- We will define a dashboard widget and generate a report that shows whether the IT Service policy is complying with the SLA.

- We will define and generate a spreadsheet that shows whether the IT Service policy is complying with the SLA.

Creating the Web-Content Monitoring Policy

The ScienceLogic platform allows users to create policies that monitor a web site for specific content. This is helpful:

- To determine if a web site is up and running.

- To determine if the connection between a webserver and a database is up and running.

- To monitor system tools that can be accessed through a browser.

- To monitor content on a web site.

If the ScienceLogic platform cannot match the expression in the content policy with the text on the web site, the platform generates an event.

The ScienceLogic platform uses cURL to send and receive data from the web site.

There are two places in the ScienceLogic platform from which you can define a policy for monitoring web content:

1. From the **Device Manager** page (Registry > Devices > Device Manager):

   - In the **Device Manager** page, find the device that you want to associate with the monitoring policy. Select the wrench icon (🔧) for the device.

   - In the **Device Administration** panel, select the **[Monitors]** tab.

   - From the **[Create]** menu in the upper right, select *Create Web Content Policy*.

Or:

2. From the **Web Content Monitoring** page (Registry > Monitors > Web Content):

   - In the **Web Content Monitoring** page, select the **[Create]** button.

3. The **Web Content Policy** modal page appears.

4. In the **Web Content Policy** modal page, supply a value in the following fields:

- **Select Device**. From this drop-down list, select a device to align with this policy. By default, the current device is selected in this field. We selected the web server "hq-w2k3-jump01".

> **NOTE**: Before you can define a content policy, you must decide which managed device you want to associate with the policy. You might want to associate the policy with the device where the referenced web server resides, but you aren't required to do so. Alternately, you might want to create a virtual device to associate with a content policy (for details on defining a virtual device, see the manual **Device Management**). Although the ScienceLogic platform will not use the device name to determine where to send the policy data, the reports that result from the policy will be aligned with the device you specify in the **Select Device** field.

- **Policy Name**. Name of the new policy. Can be any combination of letters and numbers. We entered "website_response_policy".

- **State**. Specifies whether the ScienceLogic platform should start collecting data specified in this policy from the device. We selected *Enabled*.

- **Port**. Port on web server to which the platform will send queries. We left this blank.

- **Timeout**. After a specified number of seconds, the platform should stop trying to connect to the server. We accepted the default value.

- **Error Codes**. Specify the HTTP status code you expect to receive in the response. We accepted the default value.

- **Proxy Server:Port**. For companies or organizations that use proxy servers, enter the URL and port for the proxy server in this field. We left this field blank.

- **Proxy Username:Password**. For companies or organizations that use proxy servers, enter the username and password for the proxy server in this field. we left this field blank.

- **Uniform Resource Locator (URL)**. URL or URI of the server to send the transaction to. For this example, we wanted to use a web site we knew we could always reach and that would always have content. We entered "http://www.cnn.com/US".

- **Post String**. If the URL is very long or requires data that cannot be transferred with a standard "GET" request (that is, data that cannot be included in the URL), you can enter a POST string in this field. We left this field blank.

- **Cookie Value**. For pages that require a cookie value to be set, enter the cookie value in this field. We left this field blank.

- **Browser Emulation**. Specifies how to format the query. Select the agent that is compatible with the web server. We accepted the default value.

- **HTTP Auth Username:Password**. For web sites that pop up a dialog box asking for username and password, use this field. We left this field blank.

- **SSL Encryption**. Specifies whether the platform should use SSL when communicating with the web site. If login for the web site is forms-based, enable this option. We accepted the default value.

- **Expression Check #1**. Regular expression to search for. For this example, we wanted to search for a word that would appear within a news web site. We entered "Obama"

- **Expression Check #2**. Another regular expression to search for. We left this field blank.

- **Referrer String**. URL of the web site. Some load-balanced configurations will not allow a request for a specific IP address. If you entered a specific IP address in the URL field, you can spoof a URL in this field. We left this field blank.

- **Host Resolution**. Host name of the web site. Some load-balanced configurations will not allow a request for a specific IP address. If you entered a specific IP address in the URL field, you can spoof a fully-qualified host name in this field. We left this field blank.

- **Min Page size (Kb)**. Page size means the size of the page, in Kb, specified in the URL of the policy. If the returned page is not at least the size specified in this field, the ScienceLogic platform generates an event. This threshold triggers the event "Page size below minimum threshold." We left this field blank.

- **Max Page size (Kb)**. Page size means the size of the page, in Kb, specified in the URL of the policy. If the returned page is larger than the size specified in this field, the platform generates an event. This threshold triggers the event "Page size above maximum threshold." We left this field blank.

- **Min Download speed (kb/s)**. Download speed is the speed, measured in Kb/s, at which data was downloaded from the server (specified in the policy) to the platform. If the download speed is not at least the speed specified in this field, the platform generates an event. This threshold triggers the event "Download speed below threshold."

- **Max nslookup time (msec)**. NSlookup speed is the speed at which your DNS system was able to resolve the name of the server specified in the policy. If the lookup time exceeds the value in this field, the platform generates an event. This threshold triggers the event "DNS hostname resolution time above threshold." We left this field blank.

- *Max TCP connect time (msec)*. TCP connect time is the time it takes for the platform to establish communication with the external server. In other words, the time it takes from the beginning of the HTTP request to the TCP/IP connection. If the connection time exceeds the value in this field, the platform generates an event. This threshold triggers the event "TCP connection time above threshold." We left this field blank.

- *Max Overall transaction time (msec)*. Overall transaction time is the total time it takes to make a connection to the external server, send the HTTP request, wait for the server to parse the request, receive the requested data from the server, and close the connection. If the overall transaction time exceeds the value in this field, the platform generates an event. This threshold triggers the event "Total transaction time above threshold." Because our SLA requires that the home page respond within 8 seconds, we entered "8000".

5. Select the **[Save]** button to save the new policy.

Creating the IT Service Policy

To define an IT Service policy, you must:

1. ***Define a service name and basic properties***. This example monitors a single web server . The name of the IT Service policy will be "web_hosting_home".

2. ***Define a list of devices (model) for the IT Service that includes the web server we want to monitor.***. This example includes the web server that is also associated with the web-content monitoring policy we created in the previous section.

3. ***Optionally, define service sets. A service set is a sub-group of devices***. This example does not use service sets.

4. ***Define metrics. A metric is based on your business processes and examines all devices or one or more service sets to evaluate the state of the IT Service***. For each IT Service, the ScienceLogic platform provides a default metric called *Average Device Availability*, based on the availability of all devices in the IT Service. You can define additional metrics, based on default data collected by the ScienceLogic platform (availability, latency, CPU usage, memory usage, swap usage, device state, and device count), data collected by a Dynamic Application, and data about network interfaces, TCP/IP ports, system processes, Windows services, Email round-trip time, web content, SOAP/XML transactions, and DNS availability. ***Our example will use data collected by a web-content monitoring policy***.

> **NOTE**: When the ScienceLogic platform evaluates a metric, it performs an aggregation, that is, the platform evaluates the data for all devices specified in the definition of the metric, over a specified time period (the ***Aggregation Frequency***). Depending on the definition of the metric, the platform calculates the average, maximum, minimum, sum, standard deviation, or count value for all devices specified in the definition of the metric.

5. ***Define Key Metrics.*** Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data. The Key Metrics are Health, Availability, and Risk. When you define a Key Metric, you are specifying how the value for a metric you created in step 4 translates to one of the standard Key Metric values. By default, all three Key Metrics are based on the default *Average Device Availability* metric.

6. ***Define alerts and associated events***. This example does not include alerts and events.

Defining the Name of the IT Service Policy and its Basic Properties

To define the basic parameters of our example IT Service policy:

1. Go to the **IT Service Manager** page (Registry > IT Services > IT Service Manager).

2. Select the **[Create]** button. The **IT Service Editor** page appears, with the **[Administration]** tab and **[Properties]** sub-tab selected:

3. Supply the following values in the following fields:

   - ***IT Service Name***. Name of the IT Service policy. We entered "web_hosting_home".

   - ***IT Service Owner***. Automatically populated with your username.

Using an SLA Definition with an IT Service Policy

- **Configuration Mode**. We selected *Basic Interface*. The *Basic Interface* allows you to quickly setup an IT Service policy.

- **Sharing Permissions**. Specifies whether other users can view and use the IT Service policy, in both the **IT Service Manager** page, **IT Service Editor** page, and in the pages in the ScienceLogic platform where the IT Service is visible. We selected *Shared with users in your organization*. The IT Service policy can be viewed and used by other users who belong to the same organization as the creator.

- **Permission Keys**. We did not select any permission keys.

- **Operational Status**. We selected *Aggregation enabled*.

- **Aggregation Frequency**. Frequency at which the ScienceLogic platform will collect data from all devices in the IT Service and "crunch" the data for each metric into a single value. We specified *Every 5 Minutes*.

- **Raw Data Retention**. Specifies how long the ScienceLogic platform should store the raw data for the IT Service policy. We accepted the default value.

- **Frequent Rollup Retention**. Deprecated field no longer used by the ScienceLogic platform.

- **Hourly Rollup Retention**. Specifies how long the ScienceLogic platform should store the "hourly" normalized data for the IT Service policy. We accepted the default value.

- **Daily Rollup Retention**. Specifies how long the ScienceLogic platform should store the "daily" normalized data for the IT Service policy. We accepted the default value.

- **Description**. We did not enter a description.

4. Select the **[Save]** button to save the values in the **[Properties]** tab.

Defining a List of Devices for the IT Service Policy

After defining the name and basic properties of an IT Service policy, you must next determine the devices to include in your IT Service policy. You do this in the **[Model]** sub-tab.

For example, if you want to monitor Email service, you could create a list of devices that includes Exchange servers, DNS servers, and devices that run Email round-trip policies.

You can manually assign devices and device groups to the IT Service device group, or you can use membership rules, like you would for a dynamic device group.

When you define the list of devices to include in your IT Service policy, that list of devices appears as a device group throughout the ScienceLogic platform.

There are three ways to add a device to the list of devices for the IT Service policy.

- Add a device group to the list of devices for the IT Service policy.

- Add a static list of one or more devices to the list of devices for the IT Service policy.

- Add a dynamic list of one or more devices to the list of devices for the IT Service policy.

In our example, we will add a *static list of devices* that includes a single web server to the IT Service policy.

To create the list of devices for the IT Service policy.

1. After performing the tasks in the *previous section*, select the **[Model]** sub-tab.

2. To add a static list of one or more devices to the list of devices for the IT Service policy, go to the **Static Devices** pane.

3. Select the **[Add]** button. The **Device Alignment** modal page appears and displays a list of all devices in the ScienceLogic platform.

4. In the **Device Alignment** modal page, we selected the checkbox for devices "HQ-W2K3-JUMP01".

5. Select the **[Add/Remove]** button in the lower right.

6. The selected devices appear in the **Static Devices** pane:

7. Select the **[Save]** button to save the list of devices.

Defining Metrics for the IT Service Policy

A metric is a measurement that helps determine the status of an IT Service.

The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The Average Device Availability metric aggregates the current availability value (0 or 1) of all devices in the IT Service and calculates the average value. The aggregation is performed at the frequency specified in the *Aggregation Frequency* setting in the basic properties for the IT Service policy. The availability of a device is determined every 5 minutes.

Before you can define a metric, you must determine what parameters you want to monitor for the IT Service policy. In our example, we will create a single custom metric. We will use data from the following sources to monitor the IT Service:

- Web Content Monitor

Our custom metric is:

- website_response (based on a Web Content Monitor created in a *previous section*)

We will create our metric in **Basic mode**. We will leave the default metric unchanged and create an additional metric.

1. After performing the tasks in the *previous section*, select the **[Metrics]** sub-tab.

2. Ensure that you are in **Basic mode**. If you see the **[Alerting]** sub-tab, you are not in **Basic mode**. Click on the **[Advanced]** button to toggle to **Basic mode**.

3. Next, we will define a new metric that examines the *response time of a web-content policy*. The web content policy is associated with the web server we added in the **[Model]** tab. Go to **Service Metric Definitions** pane and select the **[Add]** button.

4. The **Service Metric Editor** modal page appears. In this page, we will define a metric that measures the latency of the two devices in our IT Service policy. We will also define an alert that will trigger an event if the average latency of the two devices is greater than 30 milliseconds.

Using an SLA Definition with an IT Service Policy

5. To create the new metric, supply the following values in the fields:

- **Service Metric Name**. We entered "website_response".

- **Metric Type**. We selected *Web Content Monitor*.

- **Device Metric**. We entered "website_response_policy". This is the policy we created in a *previous section* of this chapter.

- **Metric**. We selected *Transaction Time*.

- For all other fields in the top pane, you can accept the default values.

Defining Key Metrics for the IT Service Policy

Key Metrics are the standard method for describing the status of an IT Service. Key Metrics allow you to quickly gauge the status of multiple IT Services, even if those IT Services require very different metrics that aggregate very different performance data. For example, you can define "health" for a remote backup service and also define "health" for an Internet bandwidth service, even though you would use different criteria to measure the health of those two services.

All IT Service policies define how the ScienceLogic platform should calculate the following Key Metrics for the IT Service:

> **NOTE**: The ScienceLogic platform automatically includes a default metric with each IT Service policy. The default metric is called *Average Device Availability*. The *Average Device Availability* metric specifies that the ScienceLogic platform should aggregate the availability data for all the devices in the policy and calculate the average availability.

- **Service Health**. The health of an IT Service can be one of the five standard severity values: Healthy, Notice, Minor, Major, or Critical. By default, the **Service Health** metric is aligned with the *Average Device Availability* metric.

- **Service Availability**. The availability of an IT Service can be either *Available* or *Unavailable*. By default, the **Service Availability** metric is aligned with the same metric as **Service Health**, converting *Critical* **Service Health** to *Unavailable* and all other **Service Health** values to *Available*.

- **Service Risk**. The risk of an IT Service is a percentage value that indicates how close an IT Service is to being in an undesirable state. By default, the **Service Risk** metric is aligned with the same metric as **Service Health**, converting the threshold between *Healthy* and *Notice* **Service Health** to *100%* and the healthiest possible value to *0%*.

The ScienceLogic platform generates an event if the **Service Health** Key Metric has a value of Notice, Minor, Major, or Critical, and/or if the **System Availability** Key Metric has a value of *unavailable*.

For more details on Key Metrics, see the chapter on *Creating, Editing, and Deleting IT Services*.

Using the metric we created in the *previous section*, we'll define the Key Metrics for our IT Service policy:

1. Select the **[Metrics]** sub-tab.

2. In the top pane, you will see the default metric, *Average Device Availability*. If you have already defined additional custom metrics, they will also appear in the top pane.

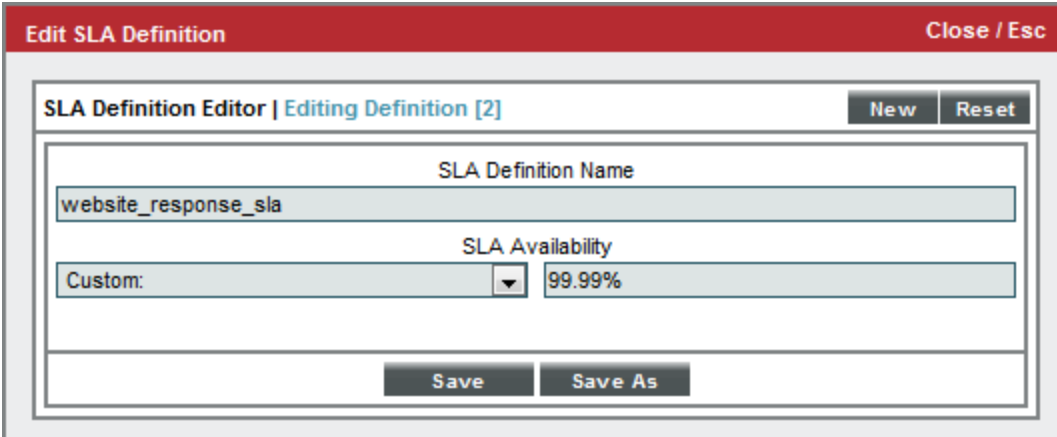3. In the bottom pane, you will see the three Key Metrics:

4. To edit each metric, supply the following values:

- *Service Health*. Appears in the **Health** column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Healthy*, *Notice*, *Minor*, *Major*, and *Critical*. Accept the default values for this Key Metric.

- *Service Availability*. Appears in the **Availability** column in **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are *Available* and *Unavailable*.

  - In the drop-down list that appears above the **Service Availability** Key Metric, select *website_response*.
  - Select *Increasing*.
  - Set the bottom of the range to "0". Set the top of the range to "9000".
  - Enter the threshold *8000*.
  - If the website_response metric has a transaction time greater than 8000 ms, the IT Service policy will have an Availability value of *Unavailable*.

- *Service Risk*. Appears as a percentage in the **Risk** column in the **IT Service Manager** page (Registry > IT Services > IT Service Manager). Possible values are 0% - 100%. Accept the default values for this Key Metric.

5. Select the **[Save]** button to save your changes to the Key Metrics.

Creating an SLA Definition

In the ScienceLogic platform, you can create an SLA Definition. The SLA Definition is a threshold. The threshold is applied to the Availability Key Metric of an IT Service policy.

To create an SLA Definition:

1. Go to the **Service Level Agreement Definitions** page (Registry > IT Services > SLA Definitions).
2. In the **Service Level Agreement Definitions** page, select the **[Create]** button.
3. The **SLA Definition Editor** page is displayed.

4. In the **SLA Definition Editor** page, supply values in the following fields:

   - *SLA Definition Name*. The name of the SLA Definition. Can be any combination of numbers, letters, and symbols. We entered "website_response_sla".

   - *SLA Availability*. The threshold that will be evaluated using this SLA Definition. You can select from six predefined percentage values or specify a custom value. If you select *Custom* in the drop-down list, enter a percentage value in the text field. You can evaluate an IT Service policy using this threshold; you can do this in a dashboard widget or in an SLA report. We selected *Custom* and then entered "99.99%".

5. Select the **[Save]** button to save your new SLA Definition.

Generating the SLA Widget

To add the SLA widget to your dashboard:

1. In the **Dashboards** page (**[Dashboards]** tab), in the selection field in the upper left of the page, select the dashboard to which you want to add a widget.

2. Select the **[Actions]** button, and then select *Add Widget*.

3. The **New Widget Configuration** modal page is displayed:

4. In the NavBar, expand the *Single Point* icon. Expand the **SLA** category. Select *(base) SLA Gauge*.

5. Enter values in the following fields:

   - *Widget Name*. Enter a title for the widget. This title is displayed in the header that appears at the top of the widget. If you leave the default value of "{auto}" in this field, the ScienceLogic platform will automatically generate a title for the widget based on what is currently being displayed in the widget. We left this field blank.

   - *Widget Refresh Rate*. Specify how frequently the widget will be automatically updated with new data. We selected *5 minutes*.

   - *SLA Definition*. Select the SLA you want to use as a threshold and monitor with this widget. We selected *website_response_sla*.

   - *Service*. Select the IT service you want to monitor with this widget. We selected *web_hosting_home*.

   - *Compliance Period*. Specify the time period you want to monitor with this widget. We selected *Last (most recently ended)*.

   - *Display Type*. Select how the metric will be displayed in the widget. We selected *Gauge*.

6. You should see a widget like the following:

Generating the SLA Report

You can use a Quick Report to evaluate an existing IT Service policy using an existing SLA definition. The Quick Report will then display the results. To do this:

1. Go the **Run Quick Report** page (Reports > Create Report > Quick Report).

2. In the **Run Quick Report** drop-down list, select the report **SLA Report**.

3. Enter a value in each of the following fields:

   - **Report Span**. Specify a Daily, Weekly, or Monthly span to include in the report. We selected *Monthly*.

   - **Starting**. This field allows you to choose a start date. Selecting a different **Report Span** will change the options in this drop-down list. We selected *Last Month*.

   - **Duration**. This field allows you to specify the duration for the report. Selecting a different **Report Span** will change the options in this drop-down list. We selected *1 month*.

   - **Timezone**. Specify the time zone to display in the report. We accepted the default value (*UTC*).

   - **IT Service**. Select the IT Service you want to monitor with this report. We selected *web_hosting_home*.

   - **SLA.** Select the SLA you want to use as a threshold and monitor with this report. We selected *website_response_sla*.

   - **Output Format**. Specify an output format for the report.

4. Select the **[Generate]** button to generate the report.

5. The generated SLA report displays the target threshold and the percentage of polls that were successful. The report displays the days violations occurred and the number of minutes each violation lasted.

Sheet 1: *SLA Report*

.ııl **ScienceLogic**™

SLA Report

Summary			
Agreement Name	IT Service Name	Target %	Compliance %
website_response_sla	web_hosting_home	99.990000%	99.801599%

Violation Periods			
Start Date/Time	End Date/Time	Duration	Cumulative
2012/05/02 19:00:00	2012/05/02 19:20:00	00:20	00:20
Total Violation Time:			00:20

Violations By Day

Excluded Periods			
Start Date/Time	End Date/Time	Duration	Type
No Excluded Periods			
Total Downtime:		00:00	

Effective Times		
Day	Start	End
Sunday	00:00	23:59
Monday	00:00	23:59
Tuesday	00:00	23:59
Wednesday	00:00	23:59
Thursday	00:00	23:59
Friday	00:00	23:59
Saturday	00:00	23:59

Generated On: 05/03/2012

Sheet 2: *Control*

.ııl **ScienceLogic**™

Control – SLA Report

Description:	SLA compliance metrics
Report Version:	1
Generated On:	05/03/2012
SLA:	website_response_sla
IT Service:	web_hosting_home
Start Date:	Apr 28, 2012
Duration:	1 week

Generated On: 05/03/2012

Using an SLA Definition with an IT Service Policy