



Monitoring Business Services

SL1 version 8.14.0

Table of Contents

Introduction to Business Services	3
What is a Business Service?	4
Navigating Business Services in SL1	5
Example: Retail Banking	7
Using the Service Investigator	8
The Tabs on the Service Investigator Page	8
The Info Drop-down on the Service Investigator Page	9
The Constituent Services on the Service Investigator Page	10
Creating Services and Service Policies	12
Understanding Health, Availability, and Risk	13
Creating a Service	14
Selecting a Service Policy	17
Creating a Service Policy	20
Creating a Service Template	25
Creating a Service From a Template	28
Exporting a Service Template	31
Installing a Template from a PowerPack	34
Default Service Policy Settings	35
Device Service Default Policy	35
IT Service Default Policy	35
Business Service Default Policy	35
Managing Service Thresholds	36
Assigning an Icon to a Service	37
Exporting Service Data with the ScienceLogic API	38
Troubleshooting Business Services	42
Using the Root Cause Analysis Feature	43
Configuring Limits for Device Services and Constituents	44
Some services are not generating Health, Availability, or Risk values	45
All services are not generating Health, Availability, and Risk values	50
503 errors, or Health, Availability, and Risk values that are all the same or are inaccurate	51

Chapter

1


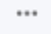
Introduction to Business Services

Overview

This manual describes how to use SL1 to create and manage business services for your company. Business services let you gauge the availability, health, and risk of your services and the devices that provide those services.

NOTE: Business Services and IT Services created in the Classic user interface are *not* included in the new Business Services, and "classic" Business Services and IT Services are not related in any way to the new business services, IT services, and device services. For more information about the Classic versions, see the **Service Provider Utilities (formerly Business Services)** and **IT Services (Classic)** manuals.

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

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

This chapter includes the following topics:

What is a Business Service?	4
Using the Service Investigator	8

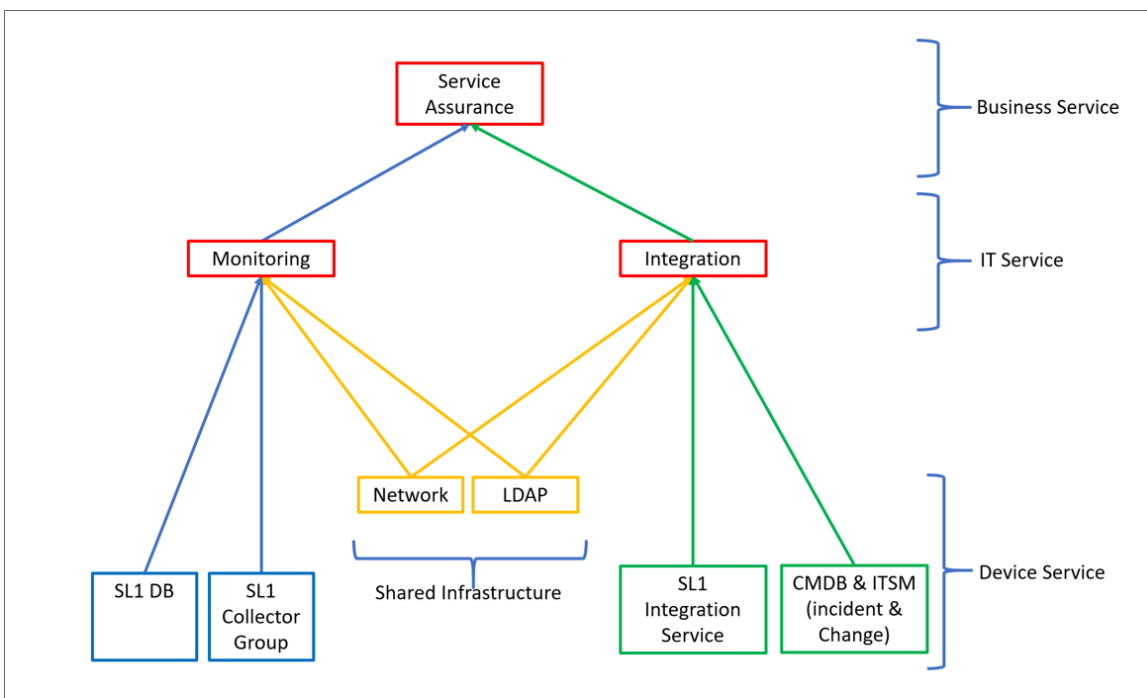
What is a Business Service?

A **business service** includes one or more technical services that provide value to internal or external customers. Some examples of business services include verifying Internet access or website hosting, online banking, remote backups, and remote storage. Usually a business service includes an associated Service Level Agreement (SLA) that specifies the terms of the service.


Create the following types of services on the **Business Services** page, in the following order:

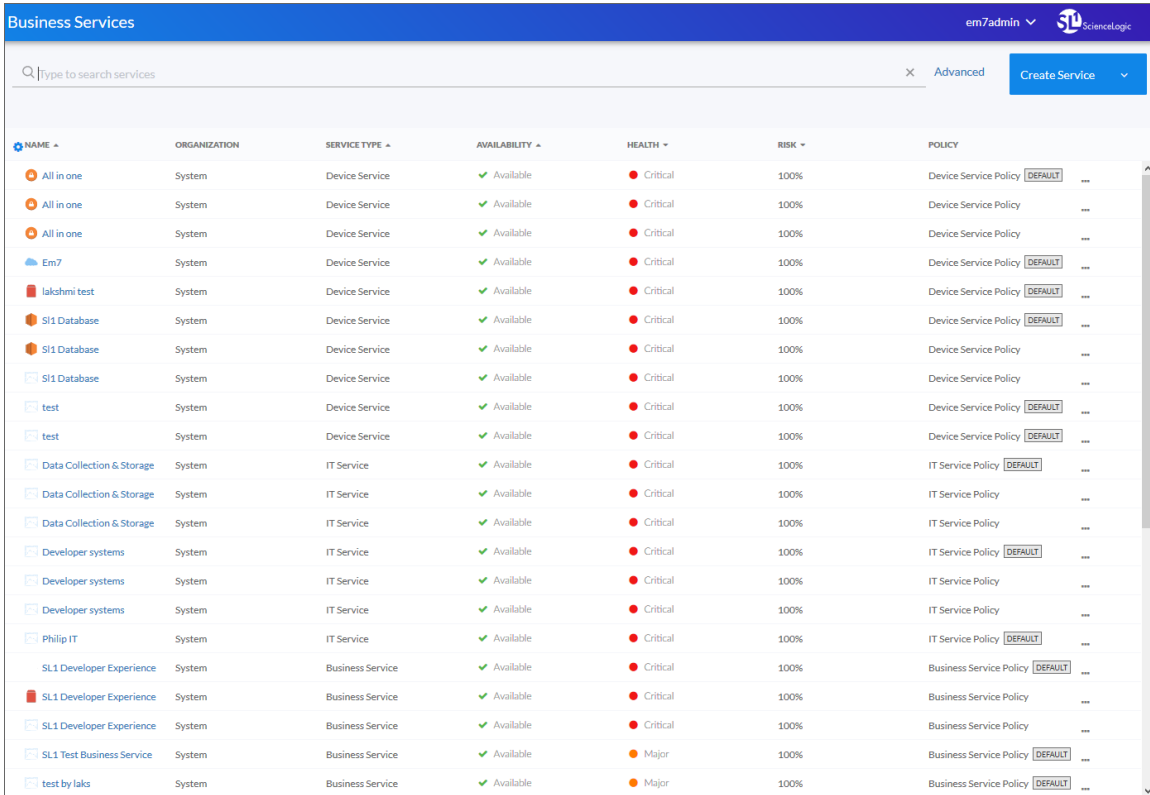
1. **Device Service**. Monitors a set of related IT infrastructure components (devices) that deliver a discrete function, such as a DNS or Collector Group, or all devices in a specific region.
2. **IT Service**. Monitors a service that IT provides to your organization. An IT Service provides a way to define how a set of related Device Services work together to power a given IT service, such as a DNS plus Collector Group plus a database.
3. **Business Service**. Monitors a service your organization provides to your customers. A business service consists of one or more IT services.

The following figure shows an example of how your business services may be organized.



Navigating Business Services in SL1


To navigate to the **Business Services** page, click the **Business Services** icon ():



NAME	ORGANIZATION	SERVICE TYPE	AVAILABILITY	HEALTH	RISK	POLICY
All in one	System	Device Service	Available	Critical	100%	Device Service Policy [DEFAULT]
All in one	System	Device Service	Available	Critical	100%	Device Service Policy
All in one	System	Device Service	Available	Critical	100%	Device Service Policy
Em7	System	Device Service	Available	Critical	100%	Device Service Policy [DEFAULT]
lakshmi test	System	Device Service	Available	Critical	100%	Device Service Policy [DEFAULT]
SI1 Database	System	Device Service	Available	Critical	100%	Device Service Policy [DEFAULT]
SI1 Database	System	Device Service	Available	Critical	100%	Device Service Policy
SI1 Database	System	Device Service	Available	Critical	100%	Device Service Policy
test	System	Device Service	Available	Critical	100%	Device Service Policy [DEFAULT]
test	System	Device Service	Available	Critical	100%	Device Service Policy [DEFAULT]
Data Collection & Storage	System	IT Service	Available	Critical	100%	IT Service Policy [DEFAULT]
Data Collection & Storage	System	IT Service	Available	Critical	100%	IT Service Policy
Data Collection & Storage	System	IT Service	Available	Critical	100%	IT Service Policy
Developer systems	System	IT Service	Available	Critical	100%	IT Service Policy [DEFAULT]
Developer systems	System	IT Service	Available	Critical	100%	IT Service Policy
Developer systems	System	IT Service	Available	Critical	100%	IT Service Policy
Philip IT	System	IT Service	Available	Critical	100%	IT Service Policy [DEFAULT]
SL1 Developer Experience	System	Business Service	Available	Critical	100%	Business Service Policy [DEFAULT]
SL1 Developer Experience	System	Business Service	Available	Critical	100%	Business Service Policy
SL1 Developer Experience	System	Business Service	Available	Critical	100%	Business Service Policy
SL1 Test Business Service	System	Business Service	Available	Major	100%	Business Service Policy [DEFAULT]
test by laks	System	Business Service	Available	Major	100%	Business Service Policy [DEFAULT]

These business services let you gauge the health, availability and risk of your services or the devices that provide those services. On the **Business Services** page, these values display in the following format and order:

1. **Availability:** The availability of a Device Service is derived from the availability rules. This may or may not be linked to device availability. A service or device is considered unavailable if SL1 is not able to collect data from the device or service, or if device is usable or not usable. A value of 0 means a device or service is unavailable, and a value of 1 means a device is available. Availability uses the following icons:



2. **Health:** Indicates the current status of a Device Service, for example, the rate of processing or throughput for the devices in the Device Service. In the case of SL1 CDB devices, the Rows Behind presentation objects can provide a good measure of how effectively the CDB is processing Collector data. Health is represented by a color-coded "severity" icon that corresponds to a numerical value between 0 and 100. For example, the Health value could indicate when a device is intermittently unavailable because of a power problem, thereby falling below the required level of performance. Health uses the following icons by default:



3. **Risk:** Displays a percentage value between 0 and 100 that indicates how close a service is to being in an undesirable state. Use risk for data that is known to cause issues if left unchecked, such as critical events, swap usage, or low database logging space. The safest possible risk value is 0%, while the worst risk value is 100%.

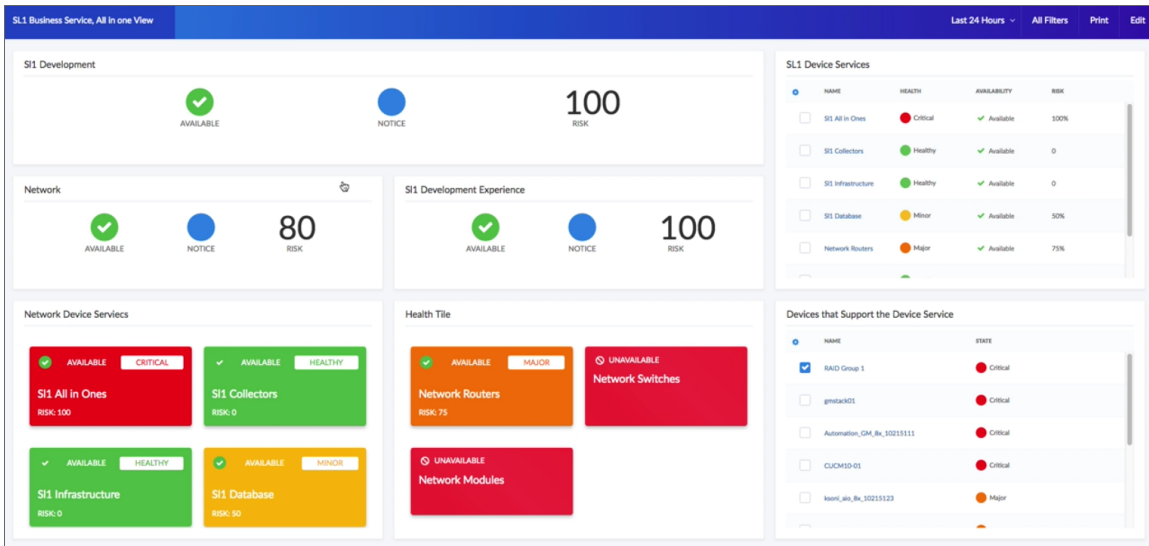
These values are computed in this order because SL1 uses *Availability* values to compute *Health* while SL1 uses both *Availability* and *Health* values to compute *Risk*.

You can define metrics for *device services* based on:

- availability
- latency
- event count
- event severity
- device state
- Dynamic Application performance data collected by SL1

NOTE: IT services created in the classic user interface are *not* included in the new user interface, and "classic" IT services are not related in any way to the new business services, IT services, and device services.

You can also create dashboards for business services that display information about the state, availability, risk, events, metrics, and other information about a business service. For more information, see the **Dashboards** manual.



Example: Retail Banking

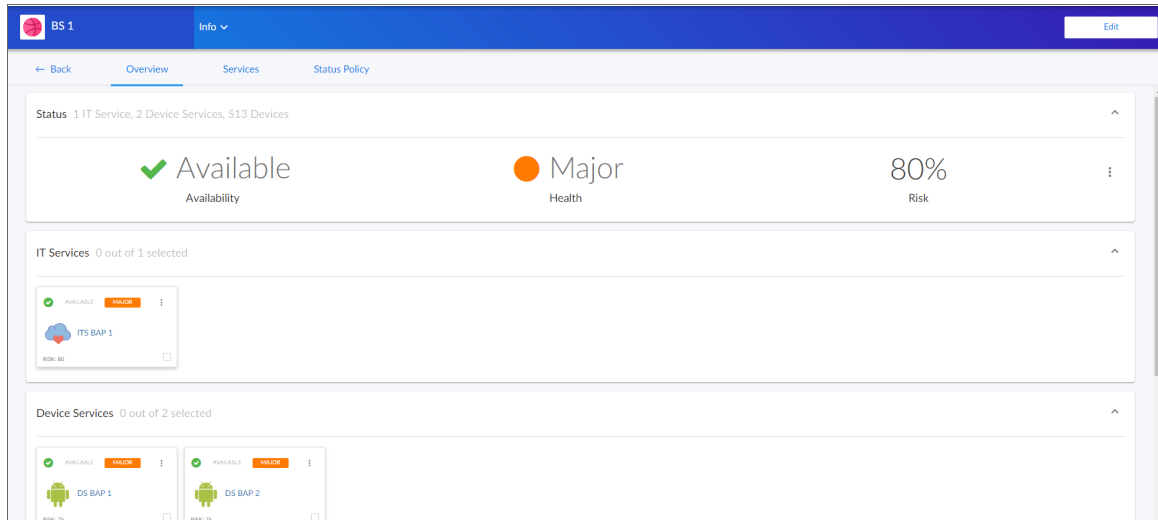
Using SL1 to monitor a business service lets you quickly see whether the service is available and working as expected for a customer or end user. For example, a banking company wants to ensure that their retail banking service is available around the world. They would use the following workflow to set up their services in SL1 :

1. Because the company has offices around the world, they create multiple **device services** that organize devices based on location or region. The company adds all of its devices to the relevant device services.
2. The company then creates multiple **IT services** to monitor the device services (from step 1), including separate IT services for online banking, teller systems, and ATM networks.
3. Next, the company creates a **business service** for its retail banking business, and this business service includes all of the IT services (from step 2) that deal with retail banking.

NOTE: As needed, the banking company repeats steps 1-3 to create additional business services (made up of IT services and device services) to monitor their commercial banking and investment banking devices and services.

Using the Service Investigator

When you select a service from the list of services on the **Business Services** page, the **Service Investigator** page appears:



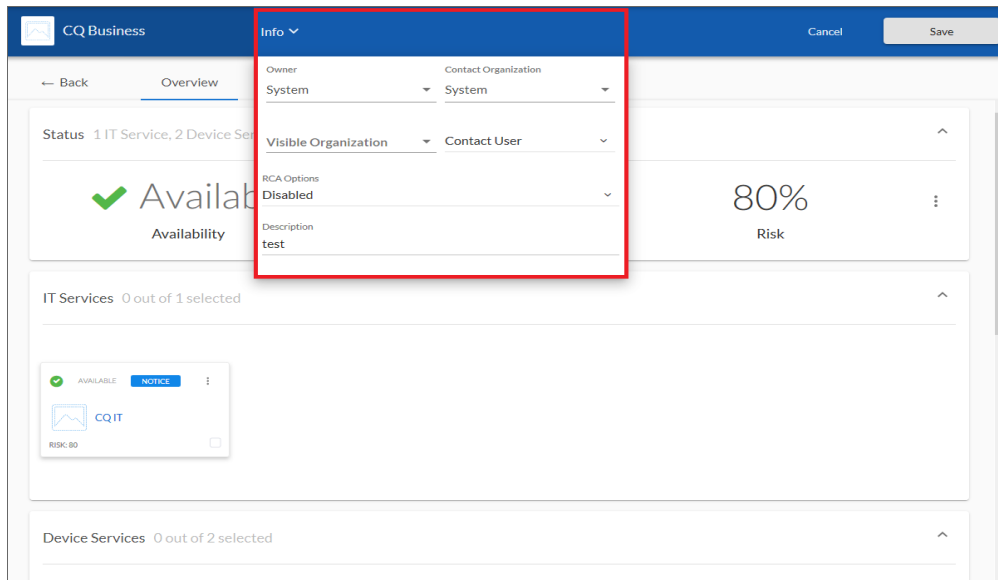
The Tabs on the Service Investigator Page

The **Service Investigator** page contains three tabs:

- **[Overview]**. Displays a "big-number" dashboard version of the most recent Availability, Health, and Risk values for the service. Below that, the tab displays a summary tile view similar to a dashboard widget for any "constituent" IT Services and Device Services that might belong to the top-level Business Service. For more information, see [Constituent Services on the Service Investigator Page](#).
- **[Services]** or **[Devices]**. Displays the services currently used in a business service or IT service, or the devices included in a device service. You can edit the search query at the top for the services or devices in the **Search** field at the top of the tab.
- **[Status Policy]**. Displays a list of all policies of that service type currently in the system and can be chosen to associate with the service being viewed. Depending on the thresholds you configured on the **Business Services Thresholds** page (Business Services > Thresholds), SL1 generates an alert message if a threshold is crossed. On this tab, you can change the policy used by a service, and you can also create a new service policy. A **Default** label appears next to the default policies.

The Info Drop-down on the Service Investigator Page

The **Info** drop-down at the top of the **Service Investigator** page displays the following:

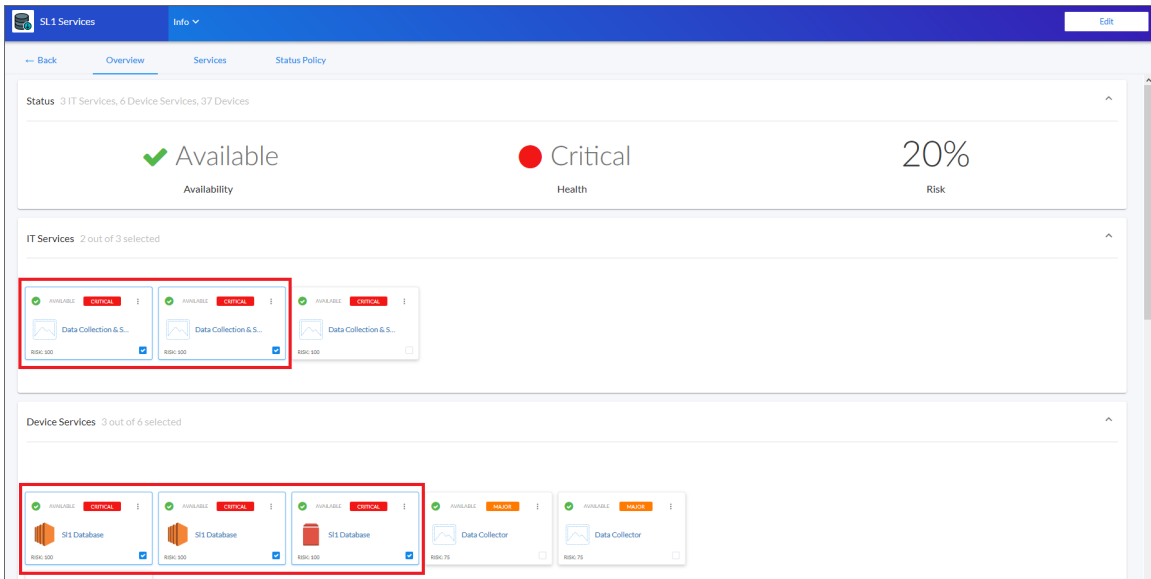


- **Owner.** The organization that owns the service.
- **Contact Organization.** A contact organization for the service.
- **Visible Organizations.** A list of organizations from which you can select devices to use in Device Services or IT Services. For example, if you selected Acme for this field, then any service that is aligned with Acme can access devices in the Acme organization. This implies the devices can be included in IT Services. There are two uses for Visible Organizations:
 1. *Device Services.* Allow the inclusion of devices from the owning organization, as well as the visible organizations.
 2. *IT Services.* Allow the inclusion of Device Services from the owning organization, as well as the visible organizations.
- **Contact User.** The contact user for the service.
- **Description.** A description of the service. You can use this field as a metadata tagging field that can be exploited in the search by a parent service. For example, if a collection of Device Services all have a description of "Shared Infrastructure", then an IT Service can search to include every Device Service in the same organization that has a description of "Shared Infrastructure". As you add more "Shared Infrastructure" device services, the IT Service will automatically expand to include them. This makes building service trees quick and self-maintaining, without resorting to rigid service names.
- **Include devices from visible organizations.** This option lets you enable or disable the ability to include devices from other organizations in a Device Service. This option appears only on **Service Investigator** pages for devices.

NOTE: Click the **[Edit]** button to edit the content on all three tabs and to edit the fields on the **Info** drop-down. You can also edit the service name and the icon associated with the service. Click **[Save]** to save your changes.

The Constituent Services on the Service Investigator Page

The **[Overview]** tab on the **Service Investigator** page for a Business Service or IT Service contains a set of tiles that represent the "constituent" services that are part of the top-level service:



The screenshot displays the 'SL1 Services' interface. At the top, there's a navigation bar with 'Back', 'Overview', 'Services', and 'Status Policy' tabs. Below this, a summary bar shows 'Status: 3 IT Services, 6 Device Services, 37 Devices'. Three key metrics are displayed: 'Availability' (green checkmark), 'Health' (red circle), and 'Risk' (20%). The main content area is divided into three sections: 'IT Services' (2 out of 3 selected), 'Device Services' (3 out of 6 selected), and a 'Device Services' section. The 'IT Services' section shows three 'Data Collection & S...' tiles, each with a checkbox. The 'Device Services' section shows three 'SI Database' tiles and two 'Data Collector' tiles, each with a checkbox. Red boxes highlight the checkboxes in the 'IT Services' and 'Device Services' sections.

Selecting the checkbox of a service tile controls which constituent service you see in the tile view below it. For example, if you have a Business Service with three IT Services, and each of those IT Services has two Device Services, selecting the checkbox of one of the IT Services will filter the Device Service tile panel down to only the Device Services that belong to the selected IT Service.

At the bottom of the **[Overview]** tab, you can review a list of devices that belong to the services you selected on the various tiles. You can then select one or more devices to see events for those devices:

Devices 3 out of 26 selected

Q Type to search devices Advanced

<input type="checkbox"/>	NAME	STATE	IP ADDRESS	CATEGORY	CLASS	SUB-CLASS	ORGANIZATION	ID
<input type="checkbox"/>	DB1	Major	192.168.33.211	System.EM7	ScienceLogic, Inc.	EM7 Database	System	23
<input checked="" type="checkbox"/>	DB2	Major	192.168.33.222	System.EM7	ScienceLogic, Inc.	EM7 Database	System	41
<input checked="" type="checkbox"/>	em7-hadr-db1	Minor	192.168.33.141	System.EM7	ScienceLogic, Inc.	EM7 Database	System	84
<input checked="" type="checkbox"/>	em7-hadr-db2	Healthy	192.168.33.146	System.EM7	ScienceLogic, Inc.	EM7 Database	System	85
<input type="checkbox"/>	gzyrb890	Major	192.168.33.129	System.EM7	ScienceLogic, Inc.	EM7 Database	System	81
<input type="checkbox"/>	gzi-bb	Major	192.168.33.185	System.EM7	ScienceLogic, Inc.	EM7 Database	System	17
<input type="checkbox"/>	ij83oq2-134	Major	192.168.33.134	System.EM7	ScienceLogic, Inc.	EM7 Database	System	35
<input type="checkbox"/>	kk-hadr01-236	Minor	192.168.33.236	System.EM7	ScienceLogic, Inc.	EM7 Database	System	31
<input type="checkbox"/>	kk-hadr02-237	Major	192.168.33.160	System.EM7	ScienceLogic, Inc.	EM7 Database	System	7
<input type="checkbox"/>	klw85-d8Q2-224	Major	192.168.33.224	System.EM7	ScienceLogic, Inc.	EM7 Database	System	38

Events 3 events

Q Type to search events Advanced

<input type="checkbox"/>	ID	NAME	SEVERITY	MESSAGE	ORGANIZATION	LAST DETECTED	AGE
<input type="checkbox"/>	486705	DB2	Major	Device Failed Availability Check: UDP-524...	System	Apr 1st 2019, 4:47:10	4 days 6 hours
<input type="checkbox"/>	466841	DB2	Minor	Network latency exceeded threshold: No Re...	System	Apr 1st 2019, 4:47:10	4 days 5 hours
<input type="checkbox"/>	496905	em7-hadr-db1	Minor	Physical Memory has exceeded threshold: 8...	System	Apr 1st 2019, 4:46:12	18 hours 33 minutes


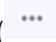
NOTE: The **Service Investigator** page for an IT Service displays Device Services, devices, and events, while the **Service Investigator** page for a Device Service displays devices and events.

Creating Services and Service Policies

Overview

This chapter describes how to create the three types of services you can monitor with SL1: business services, IT services, and devices services. This chapter also describes how to create and use policies for each service to assist with monitoring those services.

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

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

This chapter includes the following topics:

<i>Understanding Health, Availability, and Risk</i>	13
<i>Creating a Service</i>	14
<i>Selecting a Service Policy</i>	17
<i>Creating a Service Policy</i>	20
<i>Creating a Service Template</i>	25
<i>Creating a Service From a Template</i>	28
<i>Exporting a Service Template</i>	31
<i>Default Service Policy Settings</i>	35
<i>Managing Service Thresholds</i>	36
<i>Assigning an Icon to a Service</i>	37
<i>Exporting Service Data with the ScienceLogic API</i>	38

Understanding Health, Availability, and Risk

NOTE: None of the metrics described in these examples actually pinpoint the exact cause of the unavailability, degradation in health, or increase in risk, but they do bring it to your attention quickly and with a minimal level of administration. When you use key performance indicators (KPI) for responsiveness or availability, you may find it much easier than trying to model every way a service can break

Understanding Availability

Availability assesses whether something is reachable or is performing at a level to be useful. Here are a few examples to help you understand availability:

- **Website.** The URL for a website must be responsive, that is, it must respond either with the expected page or with an error page indicating that the site is unreachable (up/down). The web site's response also needs to be fast enough that users will not leave the page due to a slow response time. This should be considered when defining availability.
- **Cluster of database servers.** Assume one database server can process 1,000 transactions per second with good response times. To maintain those response times with 3,000 transactions per second, four equivalently configured database servers are put into a cluster. This method allows for any one database server to be down without losing acceptable throughput and responsiveness. If three servers in the database cluster become unavailable, the one remaining database server will not be able to maintain throughput or responsiveness, so the cluster is effectively unavailable.
- **Processes.** Consider that Process A passes work to Process B by way of a queue. If the queue depth sits at zero, it indicates that Process A is not passing any new work and is considered to be unavailable. If the queue grows to a specified threshold, it indicates that Process B is not pulling work from the queue and is considered to be unavailable.

Understanding Health

A decline in health for a given service or device means that one or more key performance indicators (KPI) are degrading. Left unchecked, this can be expected to degrade throughput or responsiveness. Here are a few examples of issues that impact health:

- **Database Server.** On an SL1 Database Server, a key database function is to retrieve and store events and Dynamic Applications data. You can create Device Service policies that degrade health as the volume of high frequency (HF) rows climbs, as this indicates the CDB is becoming overloaded or slow to process incoming data. This could lead to delays in events from Collectors being presented to automation actions or the Events page, and can impact overall system performance.
- **Windows server.** In some cases, the CPU Queue depth on a Windows server starts to increase, indicating the CPU has insufficient bandwidth to process its workload. When this happens, all processes or applications running on the Windows server will run slowly, impacting either responsiveness or throughput. You can build a policy that lets you know if this is happening on any Windows servers.

- **Website.** A website that is the face of an application has increasing web URL response times, indicating stress in the delivery of the URL. If it is known that the URL becomes functionally unavailable at 5 seconds, meaning that your customer may give up and goes to another vendor, then setting health to degrade for 1 to 4 seconds will give notice that the service health is degrading and investigations and resolution can be performed before the URL reaches an unavailable state.

Understanding Risk

In considering risk, think of the consequences of a KPI degrading. If a selected KPI is known to indicate situations that, if left unaddressed, will impact Health or Availability, you will want to create a policy for that. Some examples:

- On an SL1 Database Server, if the InnoDB table runs out of space, MariaDB will stop, which leads the Database Server to become unavailable. A shrinking level of available InnoDB space will not degrade the responsiveness and throughput of MariaDB, and therefore the Database Server, but it can indicate that your Database Server availability is at risk.
- Another way to measure risk for devices in a service is by monitoring the level of severity for events. This provides a reasonable baseline for risk. For example, many critical events for a device either indicates a false positive that should be suppressed or that monitoring has found a condition that is deemed to be unacceptable.

Creating a Service

You can create a new business service to monitor a specific set of IT services and devices for Availability, Health, and Risk values. A good design principle is to begin with the end in mind. To create a new business service, you should first determine the following:

- *Stakeholders.* Who are you creating the business service for?
- *Purpose.* What problem are you trying to solve for your stakeholders?
- *Visibility.* Who needs to see which services?
- *Workflow.* How are your stakeholders currently performing fault isolation?
- *Right-sizing.* What is the right number of device, IT, and business services? Consider the following:
 - The devices that impact the business service
 - The IT services that impact the business service
 - The specific conditions that you want to monitor, based on your business processes

If you follow the design flow described above, you will have an outline of what you need to build. For example, if you provide email service, then a failure of your primary SMTP server and backup SMTP server would constitute a Critical status.

The next consideration is to determine which devices share a common description of health, availability, and risk rules. If two devices need different rules, you will need to create two Device Services.

TIP: You can copy an existing service on the **Business Services** page by clicking the **[Actions]** button (☰) for that service and selecting *Duplicate*.

To create a Business, IT, or Device Service:

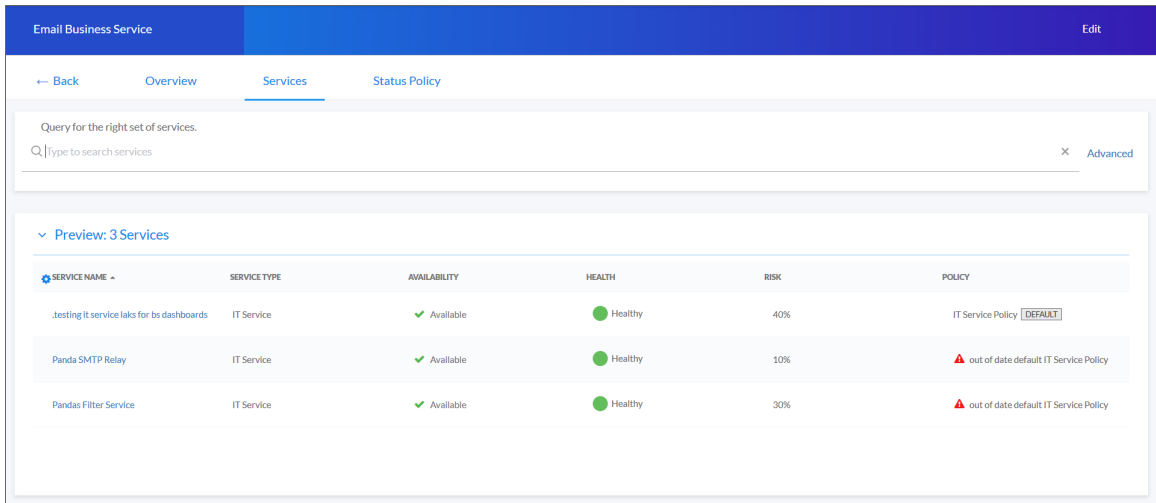
1. On the **Business Services** page, click the **[Create Service]** button. The **New Service** page appears:

The screenshot shows the 'New Service' page with the following elements:

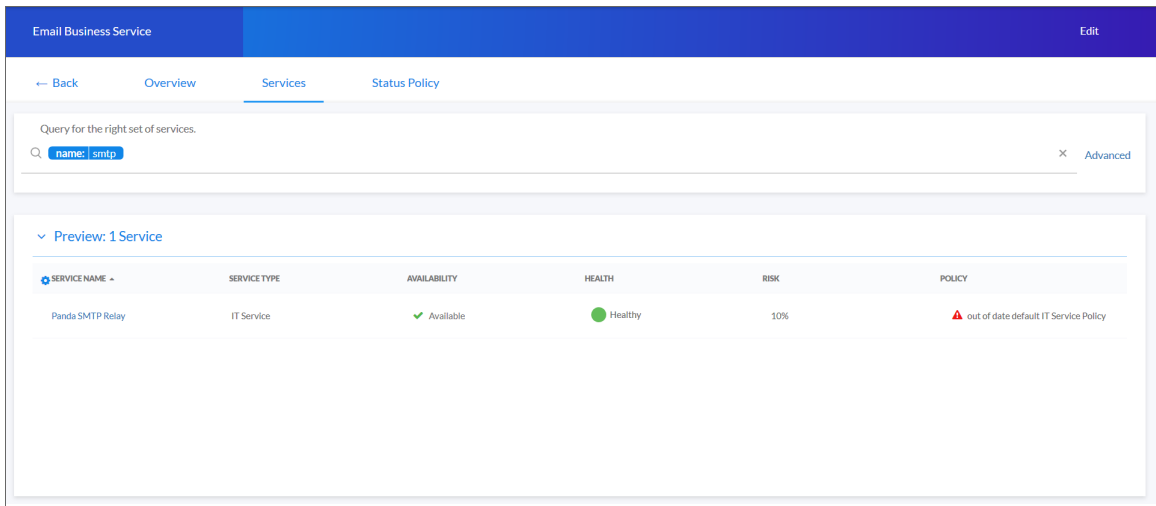
- Title:** New Service (with X and ESC icons in the top right corner)
- Section:** Select a service type
- Options:** Three service type cards: Business Service (See how your company provides business value to your customers), IT Service (Show how IT delivers value to the business), and Device Service (Aggregate status of similar devices). Each card has a diagram showing the relationship between BS, IT, and DS nodes.
- Form Fields:**
 - Service Name:** Retail Banking
 - What organization manages this service?:** System (dropdown menu)
 - Service Description:** Monitors retail banking IT Services, including online banking, the teller systems, and the ATM network. Ensures devices are responsive and performing as expected.
- Buttons:** A blue Save button at the bottom right.

2. Select a service type. You should start by creating your device services, then your IT services, and then finally your business service. Your options include:
 - **Device Service.** Monitors a set of related devices.
 - **IT Service.** Monitors a service that IT provides to your to your organization. An IT service includes one or more device services.
 - **Business Service.** Monitors a service your organization provides to your customers. A business service includes one or more IT services.
3. Complete the remaining fields:
 - **Service Name.** Type a unique name for this service.
 - **What organization manages this service?.** Select the name of the organization that owns this service.
 - **Service Description.** (Optional) Type a short description of this service and its purpose. You can use the text in this description to search for this service on the **Business Services** page. For example, if a collection of Device Services all have a description of "Shared Infrastructure", then an IT Service can search to include every Device Service in the same organization that has a description of "Shared Infrastructure". As you add more "Shared Infrastructure" device services, the IT Service will automatically expand to include them. This makes building service trees quick and self-maintaining, without resorting to rigid service names.


- Click the **[Save]** button. If you selected *Device Service* in step 2, the **[Devices]** tab appears, with a list of available devices in the **Preview** section. If you selected *Business Service* or *IT Service* in step 2, the **[Services]** tab appears, with a list of available services in the **Preview** section.



- In the **Search** field, type search criteria for the services or devices you want to monitor. A list of services or devices that match your search criteria appears in the **Preview** section:



NOTE: If your search for devices to be included in a Device Service exceeds 100 or the total number of defined services exceeds 100, then you must increase the Business Service default limits. For more information see, [Configuring Limits for Device Services and Constituents](#).

TIP: : If you are looking for a very specific set of services or devices, click the gear icon () to the right of the **Search** field and select *Advanced*. In this mode, you can create an advanced search using AND or OR for multiple search criteria.

For example, to search for devices with a Device Class of "network.router", use: `deviceClass has (deviceCategory has (name contains 'network.router'))` For more information, see the "Advanced Search" chapter in the *Introduction to SL1* manual.'

TIP: : If you want to search for devices that have specific custom attributes, use Advanced Search. Use the following format:

```
attribute has (id == custom attribute and value == value)
```

Note that search cannot process colons (:) in strings. The presence of a colon in service inclusion searches will stop the HAR engine that calculates HAR for that service. For more information, see the "Advanced Search" chapter in the *Introduction to SL1* manual.'


6. When you have the right combination of services or devices, click the **[Save]** button. The default policy for the type of service you selected is automatically added to the new service.
7. If you want to use a different business policy with the new service, see [Selecting a Business Service Policy](#).
8. If you want to create a new business policy to use with the new service, see [Creating a Business Service Policy](#).
9. Repeat this process until you have the right combination of device services and IT services in your business service (or business services, if needed).

Selecting a Service Policy

Each service type (device service, IT service, and business service) requires a **policy** that determines what it monitors. A business service policy contains a set of rules and conditions that define the Availability, Health, and Risk values for the service, depending on your business needs. Each service requires that one policy be associated with a service at a time.

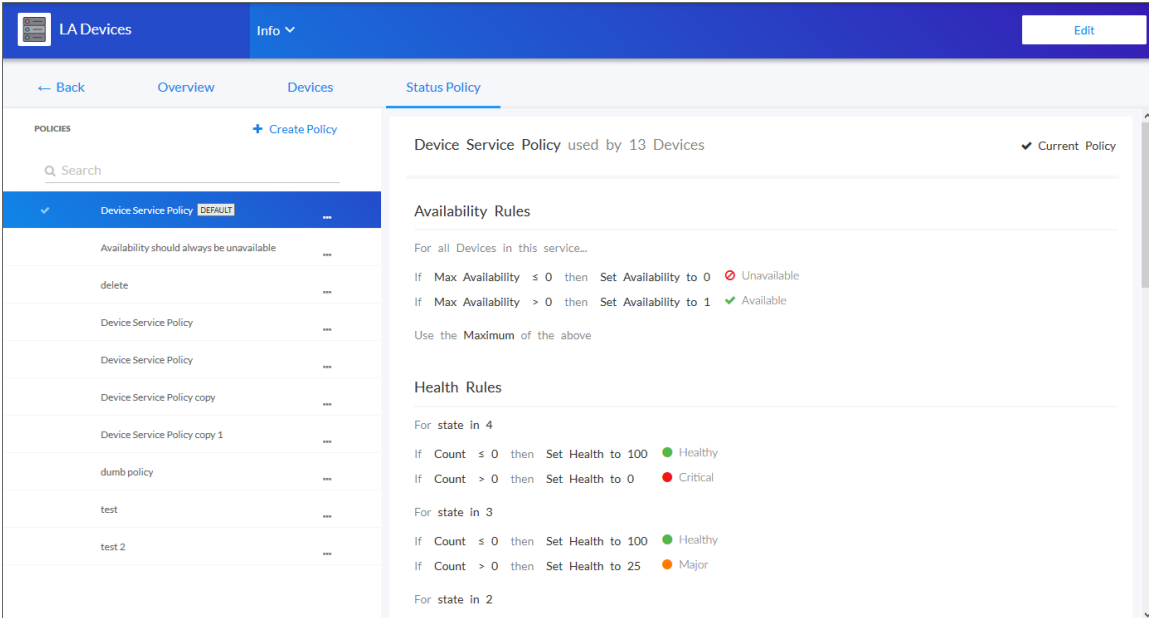
NOTE: The *Business Services PowerPack* contains a set of new business service policies you can use for your services.

When you create a business service of any type, SL1 automatically uses the *default* policy for that particular type of business service. You can remove the default policy after you create a new policy. The default policies cannot be edited.

TIP: If a policy contains errors, an error icon () appears next to the policy name. To view details about what makes the policy invalid, select the policy and hover over the error icon next to the policy name in the right-hand section. A pop-up window lists the problems with the policy. Note that most Status Policies will display the icon during the time between a save and the next HAR aggregation cycle. For best results, wait for the next HAR cycle before investigating whether there is a true error.

To select an existing business service policy:


1. On the **Business Services** page, select the service that needs a policy. The **[Overview]** tab for the service appears.
2. Click the **[Status Policy]** tab:




The screenshot displays the 'LA Devices' interface. The top navigation bar includes 'LA Devices', 'Info', and 'Edit'. Below the navigation, there are tabs for 'Back', 'Overview', 'Devices', and 'Status Policy'. The 'Status Policy' tab is active, showing a list of policies on the left and configuration rules on the right. The selected policy is 'Device Service Policy [DEFAULT]'. The configuration rules include 'Availability Rules' and 'Health Rules'.

Availability Rules

For all Devices in this service...


If Max Availability \leq 0 then Set Availability to 0  Unavailable


If Max Availability $>$ 0 then Set Availability to 1  Available

Use the Maximum of the above

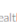
Health Rules

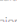
For state in 4

If Count \leq 0 then Set Health to 100  Healthy

If Count $>$ 0 then Set Health to 0  Critical

For state in 3

If Count \leq 0 then Set Health to 100  Healthy


If Count $>$ 0 then Set Health to 25  Major

For state in 2

3. In the **Policies** section on the left, select the policy you want to use.

TIP: You can type basic search criteria in the **Search** field to locate a specific policy in the list.

- To view the details of a selected policy, click the **[Actions]** button (☰) for that policy and select *Edit* (or *View* for the default policy). The **Policy Editor** page appears:



- Click the **[Cancel]** button when you are done viewing the details for that policy.

TIP: You can copy an existing service policy on the **Business Services** page by clicking the **[Actions]** button (☰) for that policy and selecting *Duplicate*.

- To add a policy to the service, select the policy in the **Policies** section and click the **[Use Policy]** button in the right-hand section. A check mark icon (☑) appears next to that policy in the **Policies** section, and the words "Current Policy" replace the **[Use Policy]** button in the right-hand section.
- To make a copy of a policy, click the **[Actions]** button (☰) for that policy and select *Duplicate*.
- To *delete* a policy you no longer want to use, click the **[Actions]** button (☰) for that policy, select *Delete*, and then click **[Delete Policy]**. If that policy is used by any other services, those services are assigned the default policy type. You cannot delete a default policy.

Creating a Service Policy

When you create a business service of any type, SL1 automatically uses the *default* policy for that particular type of business service. You can create a new policy to replace the default policy. When you create a new policy, the new policy uses the values from the default policy for that type of service as a starting point.

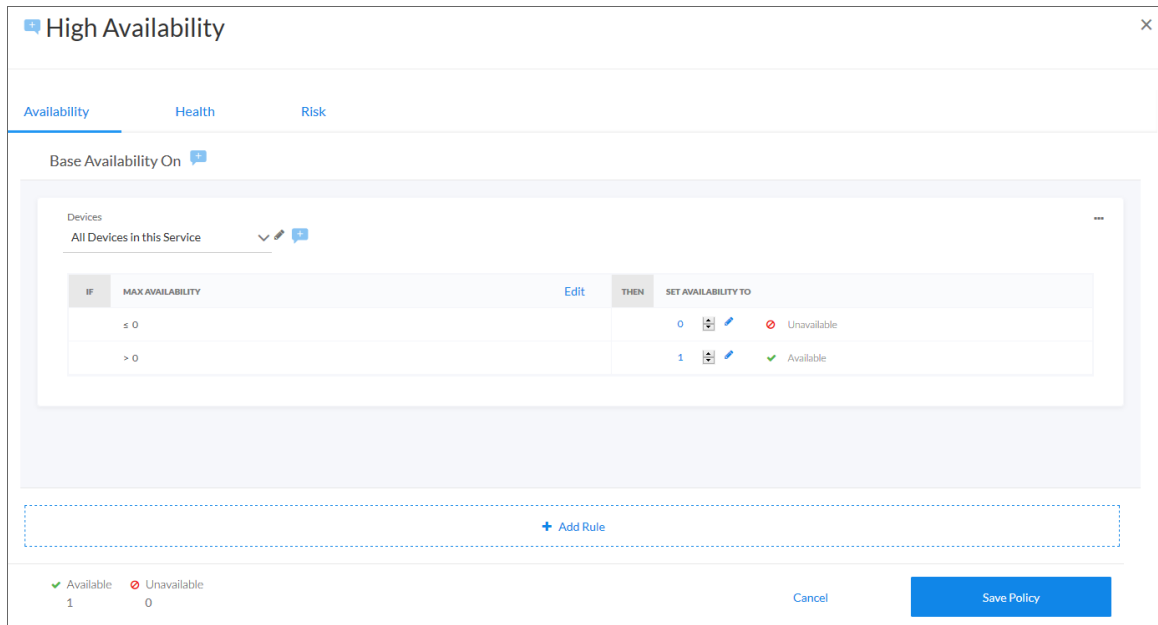
A policy includes a set of **rules**, and each rule can include one to three **conditions**. If you have multiple rules and conditions, *all* rules and conditions on a tab must be met to generate the Availability, Health, or Risk value. In other words, if a rule had three conditions, you would set up the conditions for that rule as an IF, AND, AND, THEN statement.

NOTE: Before you configure your service policy, it is important to understand why each severity is set as a range. For example, Critical for Risk is 81 - 100. The range allows one rule to be more causal or important than another. For example, suppose a Device Service for Linux servers has two risk rules: one for memory utilization and one for swap utilization. A server that has exhausted memory but still has free swap space to expand into will stay running but will slow down. A server that has exhausted swap space is likely to fail. Therefore, we understand that both statuses can be bad, but the lack of free swap space is worse than having low memory. When building Risk rules, we could set 95% memory utilization as Critical with a score of 85, but set Swap at 95% utilization to Critical with a score of 95. This will indicate that swap space is more causal than memory, and that as soon as you fix the swap space issue, you will need to check into the problems with memory.

To create a policy:


1. On the **Business Services** page, select the service for which you want to create a policy. The **Service Investigator** page appears.
2. Click the **[Status Policy]** tab, and then click **Create Policy** in the **Policies** section. A **Create Policy** window appears.
3. Type a policy name and click **[Create Policy]**. The new policy is added to the **Policies** section on the **[Status Policy]** tab.

- Click the **[Actions]** button (☰) for the new policy and select *Edit*, or click the **[Edit Policy]** button. The **Service Policy Editor** page appears, with a default rule already configured on each tab for Availability, Health, and Risk:









- On the **[Availability]**, **[Health]**, and **[Risk]** tabs, edit the rules and conditions for each of the three values that make up this policy. Each tab uses the same layout.

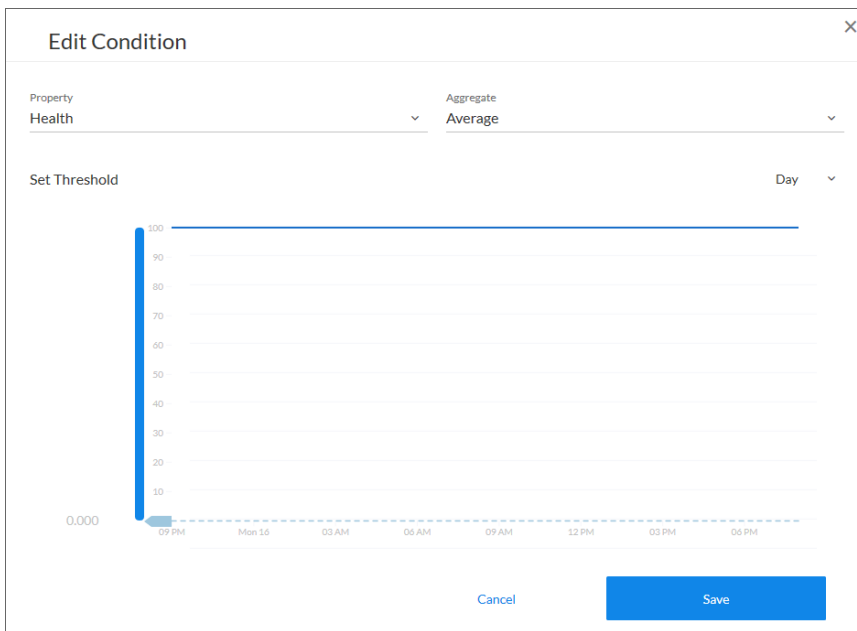
NOTE: Availability is not populated for component devices. Therefore, Availability will have a null value for any Device Service that includes component devices. The null value is displayed as a hyphen. A solution is to change the rule from Availability to Count, because Count only considers devices that are shown to be available from a collection perspective.

- In the **Services** or **Devices** drop-down list, select one of the following options to filter the services for this policy, as needed:
 - All Services in this Service* or *All Devices in this Service*. This default setting uses all services or devices that are included in the service.
 - Queried Services* or *Queried Devices*. This setting uses only the devices or services you specify in the **Search** field that appears when you select this option. This setting lets you filter the list of devices or services for this policy.
 - Edit* (pencil icon). Click the Edit icon to specify a query to find specific devices. To filter health, availability, or risk based on a specific message text mask, click the  icon to allow for an advanced search. Search using the following format:
`event has (message contains 'text mask')`

7. To update an Availability, Health, or Risk value for a rule, edit the value in the **SET <VALUE> TO** column:

Edit	THEN	SET HEALTH TO
		100    Healthy
		25    Major




8. To edit the default conditions for an existing rule, click the **[Edit]** button for that rule. The **Edit Condition** window appears:




9. Complete the following fields:

- **Property.** Select the metric you want to monitor for this condition:
 - If this is a business service or an IT service, your options include *Availability*, *Health*, and *Risk* for the services you want to monitor.
 - If this is a device service, select a device metric, such as Vitals like *Availability* and *Latency*, performance metrics, metrics collected by the SL1 Agent, or Dynamic Application metrics.
- **Aggregate.** Select an aggregation method for the data for this condition. Your options include *Average*, *Minimum*, *Maximum*, *Count*, and *Sum*. For example, suppose you have a web server farm consisting of three web servers. You have created a rule for web response time and are building for *Health*.
 - *Minimum* will drive health based on the fastest responding web server.
 - *Maximum* will drive health based on the slowest responding web server.
 - *Average* will drive health based on the average between slowest and fastest. This may give false positives. For example, assume that 5 seconds is the ideal target response time. If web server 1 gives a .1-second response time, web server 2 gives a 5-second response time, and web server 3 gives a 10-second response time, then the average will be 5 seconds, masking the fact that one of the response times is grossly unacceptable.
 - *Count* determines how many devices are currently being included in the Device Service. (The devices must be available as seen on the Device page). This is useful if we need at least 2 out of our 3 web servers to be active at any one time.
 - *Sum* is the result of adding up the value of the metric from all devices currently included in the Device Service. This is useful when you need to know how many devices are available across all the devices in the Device Service.
- **Day.** Select a time frame for the data in the graph in the **Set Threshold** section, below. You can use this graph to select reasonable thresholds for your condition. Your options include *Day*, *Week*, and *Month*.

10. In the **Set Threshold** section, click and drag the slider to specify a threshold for this condition. A small **Threshold** window appears, where you can specify the following threshold details:

- The upper threshold icon () lets you set the highest acceptable number for that condition, including any numbers less than that number. For example, $x \leq 80$.
- The lower threshold icon () lets you set the lowest acceptable number for that condition, including any numbers greater than that number. For example, $x \geq 60$.
- The equals icon () in conjunction with a number lets you set a specific number only for this condition. For example, $x = 75$.
- You can specify a range of values by clicking to add a second slider to the **Set Threshold** graph. For example, $40 < x < 60$.
- You can type a number in the **Threshold** window instead of using the slider.

- If needed, you can add a threshold that extends past the existing Y-axis of the table. The scale of the table automatically adjust to the new value.
- The different ranges for your conditions display in alternating shades of dark blue and light blue:



TIP: If the line below the number in the **Threshold** window is red, then your current threshold is invalid. Click the icons or adjust the slider to make sure the line is not red under the threshold value.

11. To save the conditions and threshold settings and close the **Edit Condition** window, click the **[Save]** button.
12. To add more conditions to a rule, click **Edit** on the **Service Policy Editor** page and follow the instructions in steps 8-11.

TIP: To remove a condition from a rule, click the **[Actions]** button (≡) for that condition and select *Delete*. To copy a condition, click the **[Options]** button (≡) for that condition and select *Duplicate*.

13. If you have more than one rule, select the type of aggregation you want to use in the **Use <type> of rules** field. You can choose to use the minimum, maximum, or average value for the rules.

NOTE: The Availability value calculates only the minimum and maximum values for rules.

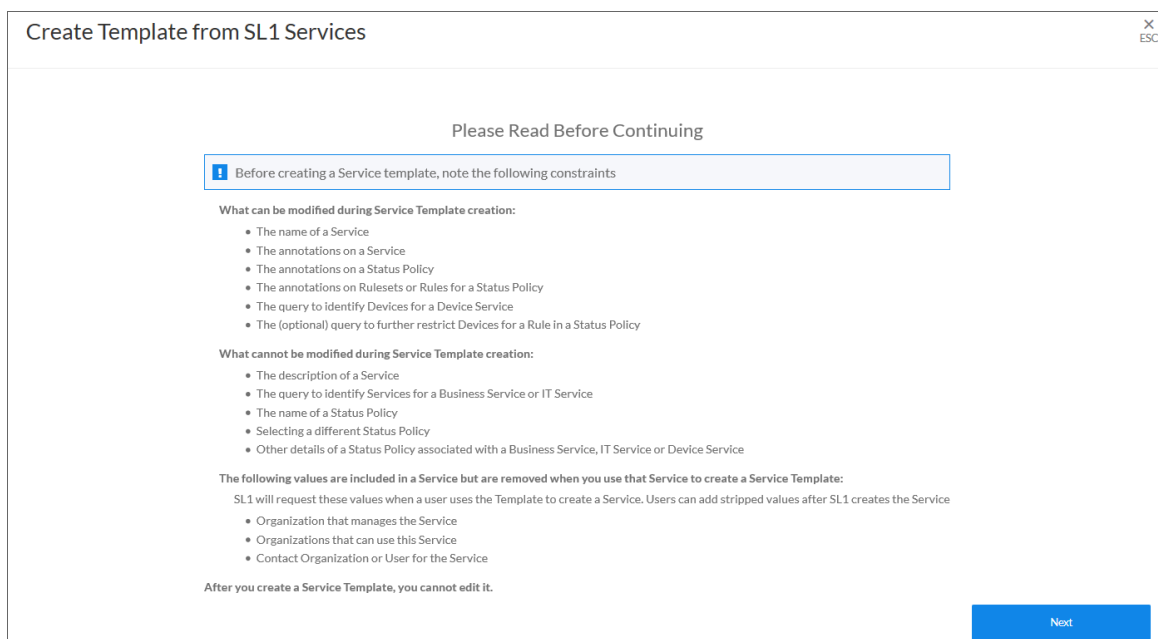
14. Edit any additional conditions or rules on the remaining tabs for this policy, and then click the **[Save Policy]** button.

Creating a Service Template

You can create a **service template** from an existing service to simplify the process of replicating an entire service or service hierarchy on another SL1 system. For example, if you want to create the same service hierarchy, but only change the owner of the service hierarchy, creating a service template from an existing service streamlines this process.

To create a service template:

1. On the **Business Services** page, click the **[Actions]** button (☰) for the service you want to use as the basis for your template and select **Create Template**. The **Create Template From Service** window appears:



- This window contains important information about what you can and cannot do with a service template. After reading this information, click **[Next]**. The next **Create Template From Service** window appears:

- Type a name for the template in the **Template Name** field, and type a description of the template in the **Description** field, if needed. Click **[Next]**. The next **Create Template From Service** window appears:

SERVICE NAME	TYPE	POLICY
DC IT Services	IT Service	IT Service Policy

- The left side of the window displays the tree for the service hierarchy that is being made into a template. You can select each service in the tree to see information related to that service on the right side of the window. For example, if you select a device service, the **Devices** tab displays the search query used for the devices included in that service. If you select a business service or an IT service, the **Services** tab displays the search query for that service. Note the following about the **Dynamic?** slider.
 - If **Dynamic?** is disabled (default), the template inherits the result of the services inclusion search. This is useful if you want to lock the service tree at the time of template creation. For example, a Managed

Service Provider (MSP) might do this to allow end customers to create services from the template but not to modify them. Another use case is if you want to use searches for tags to lock in a set of services that matched the rules at template creation time.

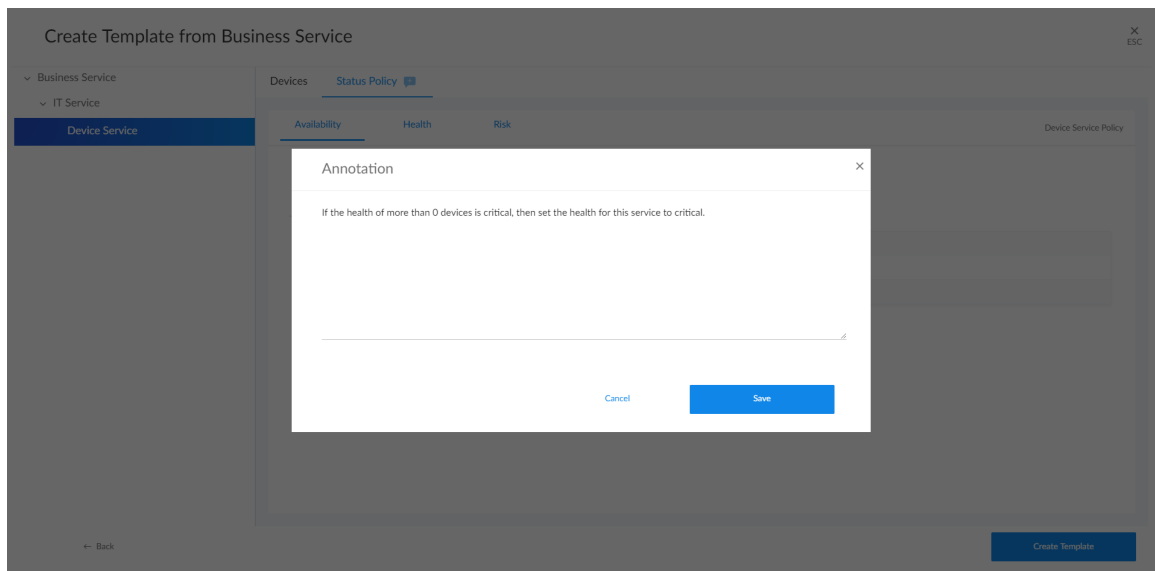
- If **Dynamic?** is enabled, the original rule is maintained in the template, so every service tree created from the template will be dynamic based on the services that match the rules.



Query for the right set of services.

Q (id contains 'dc' or name contains 'dc' or policy has (name contains 'dc') or organization has (company contains 'dc')) Dynamic ?

5. Click the **Status Policy** tab to view the status policy definition for Availability, Health and Risk for that service.
6. On the **Status Policy** tab for a device service, you can add annotations for the policies in the template. When a new user uses the template on another system, your annotations can help that user understand the purpose of this status policy.



7. To leave an annotation for a status policy or rule, click the talk bubble icon () next to the rule or tab. Type your annotation text in the **Annotation** window and click **[Save]**. The talk bubble icon now displays as solid blue, while empty talk bubble icons contain a plus sign.
8. Click **[Create Template]**. A confirmation window appears stating that you created the template. Click **[Close]**. The template appears on the **Service Templates** page (Business Services > Templates).

Creating a Service From a Template

To create a service from a template:

1. Go to the **Service Templates** page (Business Services > Templates) and click the **[Actions]** button (☰) for the template you want to use and select *Create Service*. The **Create Service from Template** window appears:

NAME	TYPE	DESCRIPTION
<input checked="" type="checkbox"/> BS template	businessService	BS temp
<input type="checkbox"/> DS Template	deviceService	DS temp
<input type="checkbox"/> ITS template	ITService	ITS Temp
<input type="checkbox"/> Templating a service that has been created from template	deviceService	

TIP: You can also go to the **Business Services** page, click the down arrow on the **[Create Service]** button, and select *Create Service from Template*.

2. Select a template and click **[Next]**. The next **Create Service from Template** window appears:

Create Service from Template

Template Name
Template Example

Description (Optional)

What organization manages this service?

Select Organization

System

Next

3. Select an organization from the **What organization manages this service?** drop-down list and click **[Next]**. The next **Create Service from Template** window appears:

Create Service from Template

test.bs by laks

ITS by laks

Test Device Service by laks

Services Status Policy

Query for the right set of services.

id in c7f6cc00498534e21e7b343e


Preview: 1 Service

SERVICE NAME	SERVICE TYPE	POLICY
ITS by laks	IT Service	IT Service Policy copy

← Back


Create Service from Template

- To edit the names of the services in the hierarchy at the left, click the service name and update the name. Updating the service names is recommended if you are creating the new service on the same system from which the template was created.
- Any annotations for a device service that were added when the template was created will be present, and you can edit them and add new annotations.
- You can edit the rules for Availability, Health, and Risk for a device service in the template.

- To edit a rule, click the gray pencil icon () next to the rule, and an edit window appears where you can update the rule:

NAME	STATE	IP ADDRESS	CATEGORY	CLASS	SUB-CLASS	ORGANIZATION	ID
SAC-IS03-DB-9-56-60093	Minor	10.140.234.220	Network.Router	Cisco Systems	12410 GSR	System	10
SAC-IS03-DB-9-56-60094	Minor	10.140.234.221	Network.Router	Cisco Systems	12410 GSR	System	11
SAC-IS03-DB-9-56-60098	Minor	10.140.234.225	Network.Router	Cisco Systems	12410 GSR	System	13
SAC-IS03-DB-9-56-60097	Minor	10.140.234.224	Network.Router	Cisco Systems	12410 GSR	System	14
SAC-IS03-DB-9-56-60100	Minor	10.140.234.227	Network.Router	Cisco Systems	12410 GSR	System	15

8. Click the **[Save]** button to close the edit window.
9. Click the **[Create Service from Template]** button to save your service. A confirmation window appears:



10. Click the **[Close]** button. The new services appear on the **Business Services** page.


Exporting a Service Template

If you want to use a business service template on another SL1 system, you can package that template into a PowerPack and export it to the other system.

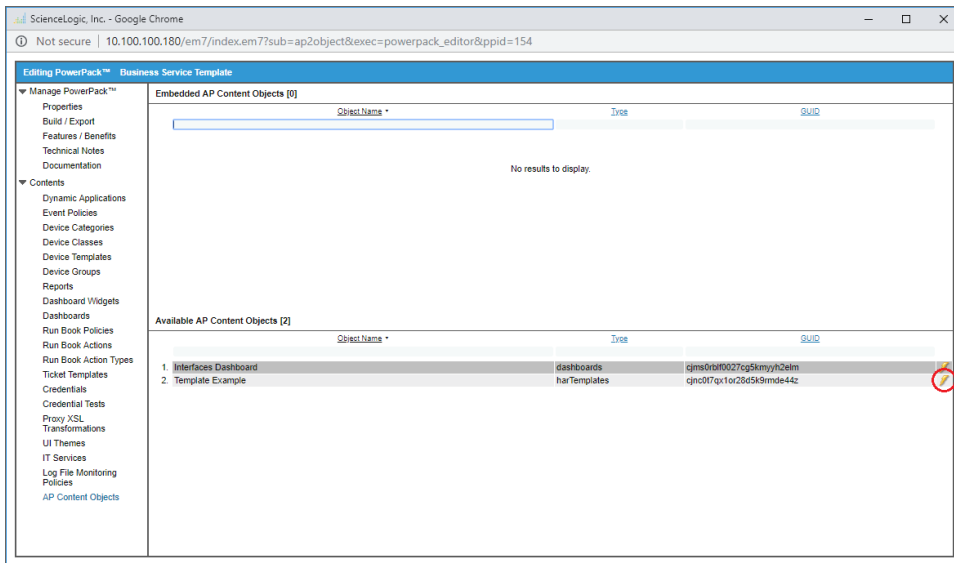
To package and export a service template:

1. Go to **The PowerPack Manager** page (System > Manage > PowerPacks).
2. Click the **[Actions]** button and select *Create a New PowerPack*.


- On the **PowerPack Properties** page, type a name for the PowerPack in the **Name** field and click **[Save]**.




- Select **AP Content Objects** from the left-nav on the **PowerPack Properties** page. Your template appears in the **Available AP Content Objects** pane:




5. Click the lightning bolt icon (⚡) next to the template to add it to the PowerPack. The template moves up to the **Embedded AP Content Objects** pane:




6. Select *Build/Export* from the left-nav to open the **Compiled PowerPacks** window, and then click the *Create a new build* link:




7. In the **Configure New Export File** window, select *Administrative (including export & license)* from the **Embedded license key** drop-down list. Click **[Build]**.



8. When the PowerPack finishes building, you can download the build with the download icon () and use that file to upload the template to a new SL1 system.

Installing a Template from a PowerPack

1. On the SL1 system where you want to install the template, import the PowerPack on the **PowerPack Manager** page (System > Manage > PowerPacks).
2. After you have imported the PowerPack, click the **[Actions]** button and select *Install PowerPack*.
3. Locate the PowerPack you created in the **Imported PowerPacks** window and click its lightning bolt icon ().
4. When the **Install PowerPack** window appears, click the **[Install]** button.
5. After you install the PowerPack, you can access the template on the Service Templates (Business Services > Templates).

Default Service Policy Settings

The following sections describe how the three default service policies calculate Availability, Health, and Risk:

Device Service Default Policy

Availability: Maximum available: if one device is available, then all are available

Health: Based upon the worst device severity, then uses the following settings:

- Critical = 0-20
- Major = 21-40
- Minor = 41-60
- Notice = 61-80
- Healthy = 81-100

Risk: Based upon the worst device severity, then uses the following percentages:

- Healthy= 0-20%
- Notice = 21-40%
- Minor = 41-60%
- Major = 61-80%
- Critical = 81-100%

IT Service Default Policy

Availability: Maximum available: if one service is available, then all are available

Health: Average Health value of all services

Risk: Maximum Risk value of any service

Business Service Default Policy


Availability: Maximum available : if one service is available, then all are available

Health: Average Health value of all services

Risk: Maximum Risk value of any service

Managing Service Thresholds

When SL1 evaluates the state of a service, it reviews the Health, Availability, and Risk values produced by your business services, IT services, and devices services. SL1 then compares those values against the alert thresholds that are defined on the **Business Service Thresholds** page (Business Services > Thresholds):



If any of the thresholds on the **Business Service Thresholds** page are crossed, SL1 generates an alert message. For an event to be produced, you need to create or install an event policy that watches for that alert message and produces an event when it sees that alert message.

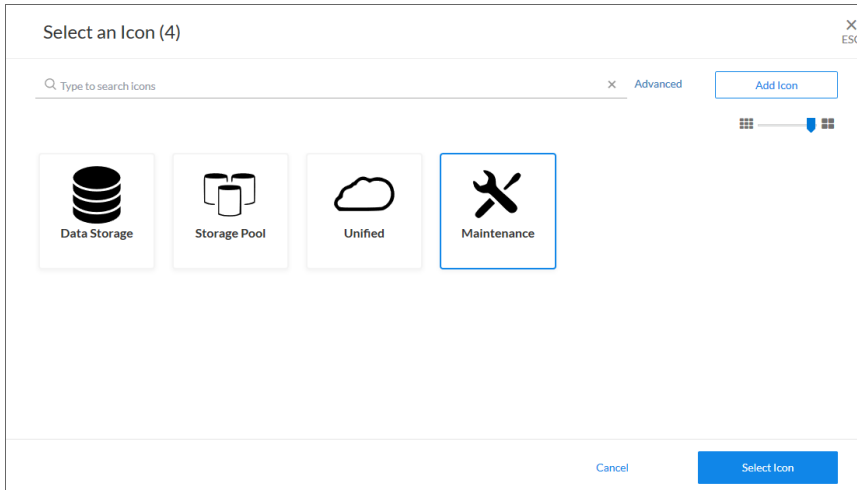
TIP: To update the thresholds on this tab, click the **[Edit]** button, select which thresholds should generate an alert message, and then click **[Save]**.

By monitoring the events tied to your business services, you can act quickly if one of your services is unavailable, unhealthy, or potentially at risk.

Assigning an Icon to a Service

To assign an icon to a service:

1. On the **Business Services** page, locate the service to which you want to add an icon.
2. Click the **[Actions]** button (☰) for that service and select *Assign Icon*. The **Select an Icon** window appears:



3. To use an existing icon, select that icon from the list of icons and click the **[Select Icon]** button.

TIP: If an icon includes a tag, you can search for that icon by typing some or all of the tag text in the **Search** field.

4. To upload an icon from your local drive, make sure that the image file meets the following criteria:
 - The image file should be in .SVG format.
 - The file should not be larger than 40 KB.
 - The file should not be animated.
 - The file should not contain bitmaps

5. To start the upload process, click the **[Add Icon]** button. The **Add an Icon** window appears:

The screenshot shows a modal window titled "Add an Icon" with a close button (X ESC) in the top right corner. The window contains the following elements:

- An "Icon name" input field.
- An "ADD TAGS" section with a "# New tag" input field.
- A "REUSE TAGS" section.
- A large dashed box containing the text "Browse or Drop".
- A list of requirements for icons: "Icons must: Be SVG format, Be no more than 40kb, Not be animated, Not contain bitmaps".
- Two buttons at the bottom: "Cancel" and "Add Icon".

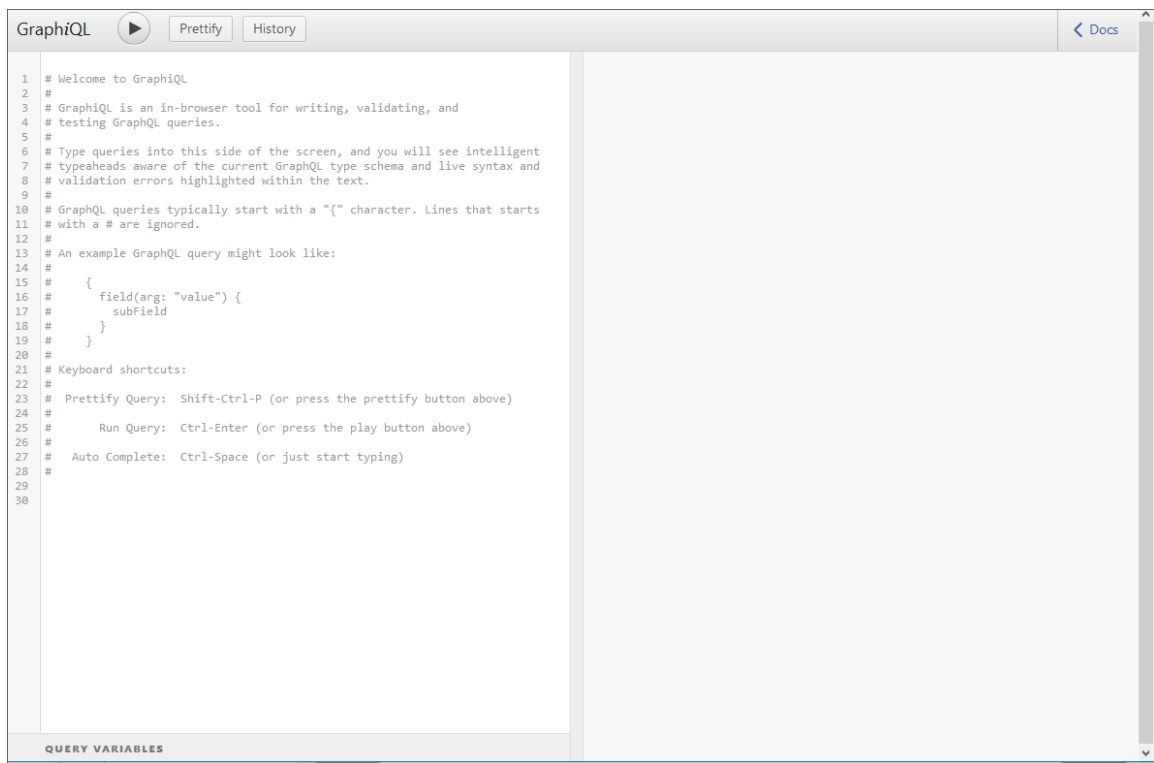
6. In the **Icon name** field, type a name for the icon you want to upload.
7. In the **Add Tags** field, type a short descriptor for the icon, without spaces. You can use this tag for searching later.
8. You can click the **Browse or Drop** area to browse for and select the icon, or you can drag and drop the icon file onto the **Add an Icon** window.
9. Click the **[Add Icon]** button. The icon is added to the **Select an Icon** window.
10. Click the **[Select Icon]** button to add the icon to the service.

Exporting Service Data with the ScienceLogic API

By navigating to the GraphQL interface, you can export business service data with the ScienceLogic API. GraphQL is a user interface for interactively exploring the capabilities of, and executing queries against, a GraphQL API.

To access the GraphiQL interface:

1. In a browser, type the URL or IP address for SL1.
2. Type `/gql` at the end of the URL or IP address. For example, you could type <https://sl1.sciencelogic.com/gql>. The GraphiQL interface appears:



3. In SL1, make a note of the URL that displays for the service you want to export. For example, if you have a service named "East Coast Tech," and its URL in SL1 is <http://sl1.sciencelogic.com/inventory/services/cjunt2se20p3izg6lmiqool5b/overview>. Make a note of the unique value between `/services` and `/overview`. In this example, the value you need is **`cjunt2se20p3izg6lmiqool5b`**.
4. In the GraphiQL interface, create a `harProvider` query for the service you want to export, using the following format:

```
query {harProvider (id:"<Service_URI>") { name} }
```

- where `Service_URI` is the value found in the URL for the Service you want to export.

- Click the **[Execute Query]** (Play) button to tell GraphQL to send the query to the GraphQL server and get the results. Using the example service from step 3, the query and its data appear in the following format:

The screenshot shows the GraphQL interface with a query on the left and its JSON response on the right. The query is:

```

1 {
2   harProvider(id:"cjumt2se20p3izg6lmiqool5b") {
3     name
4   }
5 }

```

The JSON response is:

```

{
  "data": {
    "harProvider": {
      "name": "East Coast Tech"
    }
  }
}

```

- To export additional data, use the filter-while-you-type capabilities of the GraphQL interface to gather other information, such as the collection timestamp, health, availability, and risk:

The screenshot shows the GraphQL interface with a query on the left and its JSON response on the right. The query is:

```

1 {
2   harProvider(id:"cjumt2se20p3izg6lmiqool5b") {
3     name
4     collectionTime
5     health
6     availability
7     risk
8   }
9 }

```

A dropdown menu is open under the 'risk' field, showing a list of options: 'risk', 'origin', 'description', 'organization', 'referenceFilter', 'contactOrganization', 'additionalOrganizations', and 'id'. The 'risk' option is selected and highlighted in blue. Below the dropdown, there is a tooltip that reads: 'Float The most recently computed risk for this provider'.

The JSON response is:

```

{
  "data": {
    "harProvider": {
      "name": "East Coast Tech"
    }
  }
}

```

- After you finish updating your query, click the **[Execute Query]** button.

The screenshot shows the GraphQL interface with the updated query on the left and its JSON response on the right. The query is:

```

1 {
2   harProvider(id:"cjumt2se20p3izg6lmiqool5b") {
3     name
4     collectionTime
5     health
6     availability
7     risk
8   }
9 }

```

The JSON response is:

```

{
  "data": {
    "harProvider": {
      "name": "East Coast Tech",
      "collectionTime": 1558531800,
      "health": 50,
      "availability": 1,
      "risk": 80
    }
  }
}

```

- To return to the SL1 user interface, replace the "gql" and any text after it in the URL with "ap2", such as <https://sl1.sciencelogic.com/ap2>.


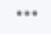
TIP: For more information about GraphQL and the GraphiQL user interface, see the **ScienceLogic GraphQL API Quick Start Guide**.

Troubleshooting Business Services

Overview

This chapter covers some of the issues you might encounter while working with services and policies on the **Business Services** page, and how to resolve those issues.

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

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

This chapter includes the following topics:

<i>Using the Root Cause Analysis Feature</i>	43
<i>Configuring Limits for Device Services and Constituents</i>	44
<i>Some services are not generating Health, Availability, or Risk values</i>	45
<i>All services are not generating Health, Availability, and Risk values</i>	50
<i>503 errors, or Health, Availability, and Risk values that are all the same or are inaccurate</i>	51

Using the Root Cause Analysis Feature

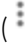
Advanced SL1 users can use the **Root Cause Analysis** feature to determine what is causing a service to be unhealthy, troubleshoot that service, and refine your policies.

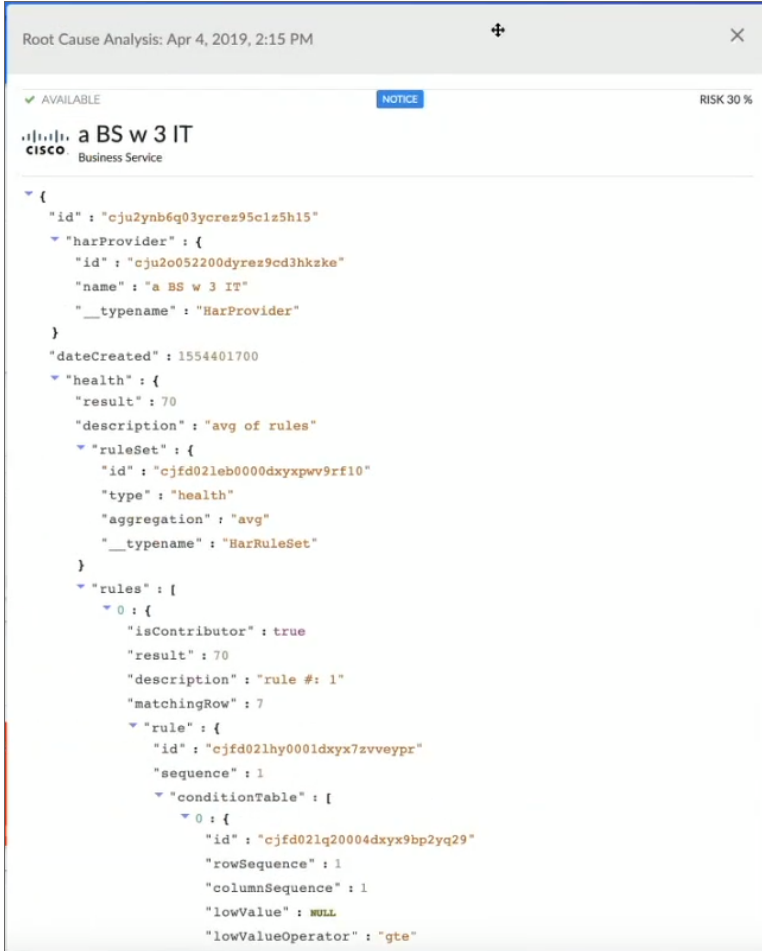
To enable Root Cause Analysis:

1. Open the **Service Investigator** page for a service and click **[Edit]**.
2. Click the **Info** drop-down and select one of the following from the **RCA Options** drop-down:
 - *Disabled*
 - *Enabled (contributors only)*
 - *Enabled (next run only)*
 - *Enabled*

NOTE: You might experience performance slowdown if Root Cause Analysis is continuously enabled.

3. Click **[Save]**.

You can view a **Root Cause Analysis** window by clicking the Actions icon () next to the service's **Status** panel on the **Service Investigator** page, or in a service file. A **Root Cause Analysis** window appears with more data as a JSON object:



```

Root Cause Analysis: Apr 4, 2019, 2:15 PM
✓ AVAILABLE NOTICE RISK 30 %
a BS w 3 IT
Business Service
{
  "id": "cju2ynb6q03ycrez95clz5h15"
  "harProvider": {
    "id": "cju2o052200dyrez9cd3hkske"
    "name": "a BS w 3 IT"
    "__typename": "HarProvider"
  }
  "dateCreated": 1554401700
  "health": {
    "result": 70
    "description": "avg of rules"
    "ruleSet": {
      "id": "cjfd021eb000dxyxpwv9rf10"
      "type": "health"
      "aggregation": "avg"
      "__typename": "HarRuleSet"
    }
    "rules": [
      0: {
        "isContributor": true
        "result": 70
        "description": "rule #: 1"
        "matchingRow": 7
        "rule": {
          "id": "cjfd021hy0001dxyx7zvveypr"
          "sequence": 1
          "conditionTable": [
            0: {
              "id": "cjfd021q20004dxyx9bp2yq29"
              "rowSequence": 1
              "columnSequence": 1
              "lowValue": null
              "lowValueOperator": "gte"
            }
          ]
        }
      }
    ]
  }
}

```

The data is a JSON object of the processing information that was used to compute the Health, Availability and Risk for that service. Within the Root Cause Analysis JSON object, the arrays named "contributors" describe which child services contributed in the calculation for the resulting Health, Availability or Risk for the current service.

Configuring Limits for Device Services and Constituents

You can configure the limit of device services processed and constituents in a service. The default limit for each is set at 100.

To configure the limit the number of device services and constituents in a service:


1. Open the `nextui.env` file by navigating to `/opt/em7/nextui/nextui.env`.
2. Edit the value for `BUSINESS_SERVICES_MAX_SERVICES` to set the maximum number of services processed.

3. Edit the value for `BUSINESS_SERVICES_MAX_CONSTITUENTS` to set the maximum number of constituents in any one service.
4. Save the file.

NOTE: Constituents can be devices in a device service or other services in a business or IT service.

Some services are not generating Health, Availability, or Risk values

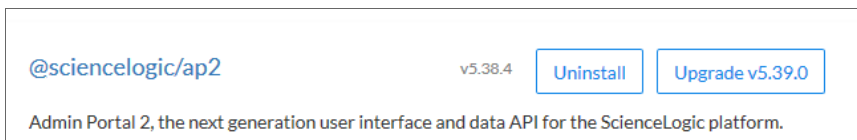
In this situation, some services in SL1 do not generate any values for Health, Availability, or Risk. For example, a dash might appear instead of a value in the **Status** table on the **Service Investigator** page:



To address this issue, review the following settings and suggestions:

Step 1: Confirm you have the latest code for the new user interface:

1. Navigate to the **[Content Management]** page (Manage > Content Management).
2. Click the **[Install/Upgrade Packages]** button. The Install Packages page appears.
3. If needed, upgrade to the latest version of `@sciencelogic/ap2` to potentially resolve any issues that might have caused this issue.
4. For example, in the following image, the *installed* version of `@sciencelogic/ap2` is 5.38.4, while the *latest* version is 5.39.0:



Step 2: Turn up the log level to trace:

1. Either go to the console of the SL1 server or use SSH to access the SL1 appliance.
2. Log in as user **em7admin**.
3. Open the file `/usr/local/silo/nextui/nextui.env` with `vi` or another text editor:

```
sudo vi /usr/local/silo/nextui/nextui.env
```

4. Change the log setting to the following: **NEXT_UI_LOG_LEVEL=all:trace**

5. Restart SL1 and GraphQL with the following command:

```
sudo systemctl restart nextui
```

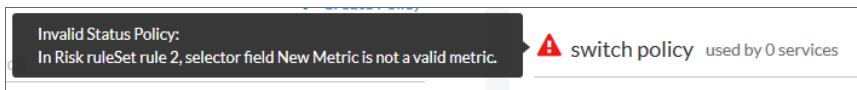
6. Tail the log with the following command:

```
sudo journalctl -u nextui -f
```

Step 3: Ensure that your service policy is valid:

1. In SL1, navigate to your service on the **Business Services** page.

2. Review the policy used by that service for any validation errors, as in the following example:

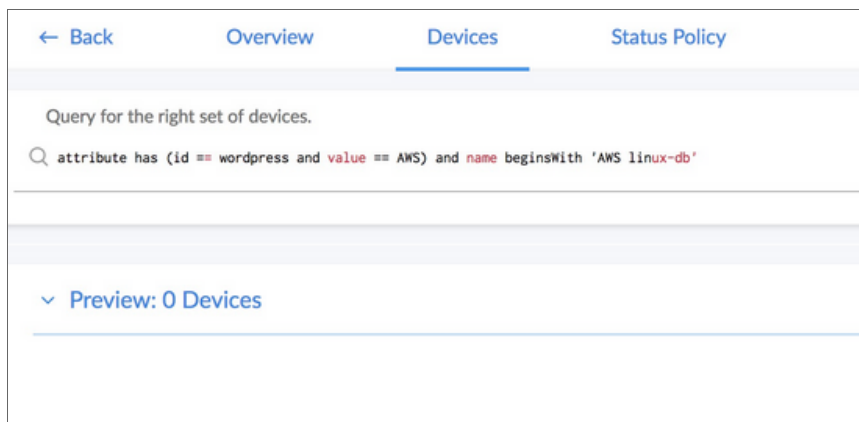


3. Address any errors in the service policy.

Step 4: Ensure that your service contains at least one service or device:

1. Navigate to the **Business Services** page.

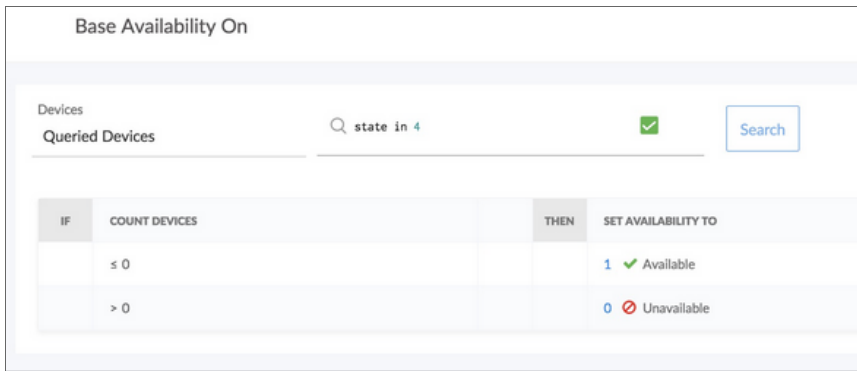
2. Navigate to the **[Devices]** or **[Services]** tab for the service or services that are not displaying values.



3. Ensure that at least one device or service appears in the **Preview** section. If not, create a new search for devices or services.

Step 5: Ensure that your service policy *rules* contain at least one service or device:

1. Rule filters select a subset of the devices or services defined by the service filter. If a device service filter results in five devices, the rule filter selects some subset of those five devices. You might create rule filters that exclude all devices or services in the service, resulting in no metric values.
2. The following rule filter only selects the devices with a state of 4, or Critical. If no devices have a state of 4, the resulting list of devices for that filter will be empty, and you cannot get any device metric values:



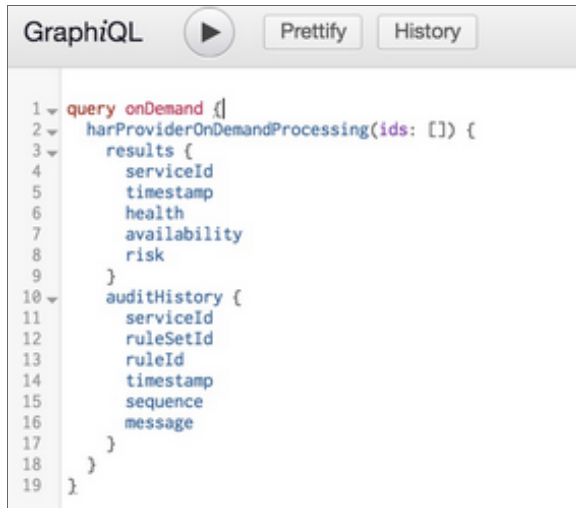
3. In this case, we are counting devices, so the count is zero and produces a value based in the condition table.
4. If the metric had been a normal device metric like latency, the result would have been "null," because getting the average latency from zero devices results in null.

Step 6: Generate audit data by running onDemandProcessing with the GraphQL interface:

1. In a browser, type the URL or IP address for the new user interface, and then type **/gql** at the end of the URL or IP address. The GraphQL interface appears.
2. On the left side of the GraphQL editor, type the following query:

```
query onDemand {
  harProviderOnDemandProcessing(ids: []) {
    results { serviceId timestamp health availability risk }
    auditHistory { serviceId ruleSetId ruleId timestamp sequence message }
  }
}
```

3. Click the **[Execute Query]** (Play) button to tell GraphQL to send the query to the GraphQL server and get the results:



The screenshot shows the GraphQL editor interface. At the top, there is a title bar with the text "GraphQL" and three buttons: a play button (labeled "Execute Query"), "Prettify", and "History". Below the buttons is a text area containing a GraphQL query:

```
1 query onDemand {  
2   harProviderOnDemandProcessing(ids: []) {  
3     results {  
4       serviceId  
5       timestamp  
6       health  
7       availability  
8       risk  
9     }  
10    auditHistory {  
11      serviceId  
12      ruleSetId  
13      ruleId  
14      timestamp  
15      sequence  
16      message  
17    }  
18  }  
19 }
```

4. Review the resulting audit information on the right side of the GraphQL editor:
5. If you know the service ID you are looking for, search for it by clicking inside the right pane and typing **cmd+f**. The GraphQL interface highlights the services that match the ID you looked for:



The screenshot shows the GraphQL editor interface with a search bar at the top. The search bar contains the text "Search: cjc9k2fcw0022r2qim00m52vq (Use /re/ syntax for regexp search)". Below the search bar, the query results are displayed in a tree view. The results are highlighted in red, indicating a search match. The results are:

```
{  
  "data": {  
    "harProviderOnDemandProcessing": {  
      "results": [  
        {  
          "serviceId": "cjc9k2fcw0022r2qim00m52vq",  
          "timestamp": 1524698040,  
          "health": 100,  
          "availability": null,  
          "risk": 0  
        }  
      ]  
    }  
  }  
}
```


6. Scroll down to see the audit information for this service (look for the highlighted information):

```

},
"auditHistory": [
  {
    "serviceId": "cjk9k2fcw0022r2qim00m52vq",
    "ruleSetId": "cjfcyh40m00a31byxi5chr1u5",
    "ruleId": "cjfcyh48300a41byxqcw5tqx4",
    "timestamp": 1524698040,
    "sequence": 1,
    "message": "Service has no constituents for rule. Service: Web DS Cloud Policy: Device Service Policy RuleSet:
availability Rule: 1"
  },
  {
    "serviceId": "cjk9k2fcw0022r2qim00m52vq",
    "ruleSetId": "cjfcyh40m00a31byxi5chr1u5",
    "ruleId": "cjfcyh48300a41byxqcw5tqx4",
    "timestamp": 1524698040,
    "sequence": 2,
    "message": "No matching row found in condition table Result: null Service: Web DS Cloud Policy: Device Service
Policy RuleSet: availability Rule #: 1 Matching Row #: none Constituents: 0 Values: {max availability: null}"
  },
  {
    "serviceId": "cjk9k2fcw0022r2qim00m52vq",
    "ruleSetId": "cjfcyh40m00a31byxi5chr1u5",
    "ruleId": null,
    "timestamp": 1524698040,
    "sequence": 3,
    "message": "RuleSet Result: null Service: Web DS Cloud Policy: Device Service Policy RuleSet: availability
Aggregation: max Values: []"
  },
  {
    "serviceId": "cjk9k2fcw0022r2qim00m52vq",
    "ruleSetId": "cjfcyglb00931byxmyu8zdm",
    "ruleId": "cjfcygos00941byxg2o5k3hu",
    "timestamp": 1524698040,
    "sequence": 4,
    "message": "Service has no constituents for rule. Service: Web DS Cloud Policy: Device Service Policy RuleSet:
health Rule: 1"
  },
  {
    "serviceId": "cjk9k2fcw0022r2qim00m52vq",
    "ruleSetId": "cjfcyglb00931byxmyu8zdm",
    "ruleId": "cjfcygos00941byxg2o5k3hu",
    "timestamp": 1524698040,
    "sequence": 5,
    "message": "Rule Result: 100 Service: Web DS Cloud Policy: Device Service Policy RuleSet: health Rule: 1
Matching Row #: 1 Matching Row: [IF (-Infinity <= count <= 0) THEN 100] Constituents: 0 Values: {count : 0}"
  },
  {
    "serviceId": "cjk9k2fcw0022r2qim00m52vq",
    "ruleSetId": "cjfcyglb00931byxmyu8zdm",
    "ruleId": "cjfcygtf00981byxam86mbiv",
    "timestamp": 1524698040,
    "sequence": 6,
    "message": "Service has no constituents for rule. Service: Web DS Cloud Policy: Device Service Policy RuleSet:
health Rule: 5"
  }
]

```

7. After running onDemandProcessing with the GraphiQL interface and updating the log settings on the server to do `all:trace`, you can now see trace-level log messages in the terminal where you ran `sudo journalctl -u nextui -f`.

- Review the log messages for errors and warnings:

```

user_id=asc
Apr 26 00:22:03 dc2-s11-db01 node[25004]: 00:22:03.169 <warn> dao.js:327 (Object.getMetricValuesForConstituents) [ { GraphQLError: Variable "$metricSearch" got invalid value {"first":{"guid":{"eq":"d check"}}}; Field "guid" is not defined by type MetricSearch at value.first; did you mean id?
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at coercionError (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/utilities/coerceValue.js:179:10)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at coerceValue (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/utilities/coerceValue.js:148:36)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at coerceValue (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/utilities/coerceValue.js:132:30)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at coerceValue (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/utilities/coerceValue.js:55:12)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at getVariableValues (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/execution/values.js:74:53)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at buildExecutionContext (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/execution/execute.js:246:63)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at executeImpl (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/execution/execute.js:148:17)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at execute (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/execution/execute.js:131:229)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at graphqlImpl (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/graphql.js:112:31)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at /var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/graphql.js:66:223
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at new Promise (<anonymous>)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at graphql (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/graphql/graphql.js:63:10)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at Object.gqlLocal [as graphql] (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/sl-em7-og1/build/middleware/gql.js:116:33)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at Object.gqlLocal [as graphql] (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/sl-em7-og1/build/middleware/gql.js:116:33)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at Object.getMetricValuesForConstituents (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/@sciencelogic/sl-em7-og1/build/lib/businessServices/dao.js:321:26)
Apr 26 00:22:03 dc2-s11-db01 node[25004]: at Object.getMetricValuesForConstituents (/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2/node_modules/@sciencelogic/sl-em7-og1/build/lib/businessServices/dao.js:321:26)

```


All services are not generating Health, Availability, and Risk values

In this situation, *all* of your services in SL1 fail to generate any values for Health, Availability, or Risk.

To address this issue, review the following settings and suggestions.

Step 1: Confirm that the Business Services process exists:

- Go to the **Process Manager** page (System > Settings > Admin Processes) and start typing "Business" in the **Process Name** filter:



Step 2: Follow the steps in [Generate audit data using the GraphQL user interface](#), above. If the process times out, then the processing has taken more than two minutes to complete, and no computed results are stored.

Step 3: Look for logs from the python process:

1. The python process calls the onDemandProcessing GraphQL query. If python is having trouble connecting to GraphQL, it could be an authentication problem or some other code-related issue.
2. Look in `/var/log/em7` for newly created logs, and `ls -lrt` to see if any new error logs were created with "business" in the file name.
3. Also check the **silob.log** for messages related to the `business_service_management` process:

```
grep service /var/log/em7/silo.log
```

503 errors, or Health, Availability, and Risk values that are all the same or are inaccurate

Step 1: Check the number of services you have configured. If you are seeing 503 errors in the nextui log or within the SL1 user interface, use the following procedure to check the number of services you have configured on your ScienceLogic SL1 system.

To determine the number of services you have:

1. Open the GraphiQL editor on your system:

```
http://<SL1_IP_address>/gql
```

2. Enter the following query:

```
query harProviders {
  harProviders {
    pageInfo {
      matchCount
    }
  }
}
```

3. Click **[Execute Query]** (Play) to see the number of services. In this example, the results shows that 10 services are configured.

```
query harProviders {
  harProviders {
    pageInfo {
      matchCount: 10
    }
  }
}
```

Step 2: (Optional) If the number of services is greater than 100, check the following feature toggle:

1. Open the GraphiQL editor on your system:

```
http://<SL1_IP_address>/gql
```

2. Enter the following query:

```
query featureToggle {
  featureToggle {id: "system.BUSINESS_SERVICES_MAX_SERVICES"} {
    value
  }
}
```

3. Click **[Execute Query]** (Play) to see the number of services. In this example, the results shows null. A value of "null" means that this value is set to the default of "100". The count must be greater than the number of services configured on your system. It is recommended that you keep the number of services at 1000 or less.

```
query featureToggle {
  featureToggle: {
    value: null
  }
}
```

4. To increase the value of this setting, modify the `nextui.env` file as described in [Configuring Limits for Device Services and Constituents](#).

Step 3: (503 Errors) Confirm that the nginx configuration has an appropriate limit set. In some cases, the `limit_conn` value might be set to 20. Increase the value to 200.

To address this issue:

1. Either go to the console of the SL1 server or use SSH to access the SL1 appliance.
2. Log in as user **em7admin**.
3. Confirm that the nginx config file has the `limit_conn perip` value set to 200 instead of 20:
`sudo vi /etc/nginx/conf.d/em7_limits.conf`
4. If needed, update the line to say:
`limit_conn perip 200;`
5. Run the following command:
`sudo systemctl restart nginx`

Step 4: (503 Errors) Check to see if the nginx server is rate-limiting you.

1. Either go to the console of the SL1 server or use SSH to access the SL1 appliance.
2. Log in as user **em7admin**.
3. Enter the following command:
`sudo grep excess /var/log/em7/nginx.log`
4. If you see any results from the above command, then the nginx proxy is rate-limiting requests to your database. In that case, you should increase the rate limit to 100 requests per second. Edit the **em7_limits.conf** file:
`sudo vi /etc/nginx/conf.d/em7_limits.conf`
5. Change the following line to **100r/s** from the default 5 r/s.
`limit_req_zone $binary_remote_addr zone=addr_req:10m rate=100r/s;`

- Restart your SL1 system.
`sudo systemctl restart nextui`

Step 5: (502 Errors) Check node memory usage.

- Either go to the console of the SL1 server or use SSH to access the SL1 appliance.
- Log in as user **em7admin**.
- Enter the following command:
`sudo journalctl -u nextui|grep "JavaScript heap out of memory"`
- If you see any results from the above command, the `node.js` process is running out of memory. In that case, you should increase the space limit allocated. Edit the **nextui.service** to increase memory to 4096 or 8192 MB, depending on how much memory you have at your disposal.
`ExecStart=/usr/bin/node --max-old-space-size=4096 /usr/local/silo/nextui/index.js`
- Restart your SL1 system.
`sudo systemctl restart nextui`

Step 6: (504 Errors) Check Nginx timeout.

- Either go to the console of the SL1 server or use SSH to access the SL1 appliance.
- Log in as user **em7admin**.
- Edit the `nextui.fragment` file:
`sudo vi /opt/em7/share/config/nginx.d/nextui.fragment`
- Change the `proxy_read_timeout` under "location /gq|" to **900** as follows:
`proxy_read_timeout 900;`
- Restart your SL1 system.
`sudo systemctl restart nextui`

© 2003 - 2020, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

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

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

Copyrights and Trademarks

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

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

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

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

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

Other

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

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



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

International: +1-703-354-1010