



Monitoring Google Cloud Platform

Google Cloud Platform *BETA* PowerPack version 101

Table of Contents

Introduction	3
What is Google Cloud Platform?	3
What Does the Google Cloud Platform *BETA* PowerPack Monitor?	4
What is the GCP Resource Hierarchy?	4
What are GCP Regions and Zones?	5
Installing the Google Cloud Platform *BETA* PowerPack	6
Configuring Google Cloud Platform for Monitoring	7
Creating a Google Cloud Platform Service Account	7
Enabling Google Cloud Platform APIs	10
Creating a SOAP/XML Credential for Google Cloud Platform	11
Creating a Google Cloud Platform Virtual Device	13
Aligning the Google Cloud Platform Dynamic Applications	13
Discovering Google Cloud Platform Component Devices	14
Viewing Google Cloud Platform Component Devices	16
Relationships Between Component Devices	18
Google Cloud Platform Dashboards	19
Device Dashboards	19
Google Backend Service	20
Google Compute Instance	21
Google Compute Instance Service	22
Google Global HTTP(S) Load Balancer	23
Google Persistent Disk	24
Google Persistent Disk Service	25
Google Project	26
Google Regional TCP/UDP Load Balancer	27
Google Storage Bucket	28
Google Storage Bucket Service	29

Introduction

Overview

This manual describes how to monitor Google Cloud Platform (GCP) resources in SL1 using the *Google Cloud Platform *BETA* PowerPack*.

The following sections provide an overview of GCP and the *Google Cloud Platform *BETA* PowerPack*:

<i>What is Google Cloud Platform?</i>	3
<i>What Does the Google Cloud Platform *BETA* PowerPack Monitor?</i>	4
<i>What is the GCP Resource Hierarchy?</i>	4
<i>What are GCP Regions and Zones?</i>	5
<i>Installing the Google Cloud Platform *BETA* PowerPack</i>	6

NOTE: ScienceLogic provides this documentation for the convenience of ScienceLogic customers. Some of the configuration information contained herein pertains to third-party vendor software that is subject to change without notice to ScienceLogic. ScienceLogic makes every attempt to maintain accurate technical information and cannot be held responsible for defects or changes in third-party vendor software. There is no written or implied guarantee that information contained herein will work for all third-party variants. See the End User License Agreement (EULA) for more information.

What is Google Cloud Platform?

Google Cloud Platform is a suite of modular, cloud-based products and services that enables users to build, test, deploy, and manage applications for web, mobile, and back-end solutions. It combines physical assets and virtual resources that users can utilize for computing, data storage, networking, and other solutions.

What Does the Google Cloud Platform *BETA* PowerPack Monitor?

To monitor GCP resources using SL1, you must install the *Google Cloud Platform *BETA* PowerPack*. This PowerPack enables you to discover, model, and collect performance and configuration data about GCP resources.

The *Google Cloud Platform *BETA* PowerPack* includes:

- Dynamic Applications to discover, model, and collect configuration data and performance metrics for GCP services and resources
- Device Classes for each type of GCP device that SL1 monitors, plus Device Classes for each region and zone the *Google Cloud Platform *BETA* PowerPack* supports
- Event Policies that are triggered when GCP resources meet certain status criteria
- A sample Credential that you can use to create SOAP/XML credentials to monitor GCP devices
- The ScienceLogic Libraries that are utilized by the PowerPack

What is the GCP Resource Hierarchy?

The GCP resource hierarchy is a tiered methodology that determines GCP resource ownership, in which child resources inherit the access control policies and configuration settings of their parent resources.

The GCP resource hierarchy consists of the following levels:

- **Organization.** The Organization resource represents an organization, such as a company, and is the top level of the hierarchy for your GCP account. An account can have only one Organization associated with it. Organizations are not required resources, but when an account includes one, every project created by users of that account will belong to that Organization resource by default.
 - **Folders.** Folders are essentially a tool for organizing resources and creating borders between Projects. For example, they could represent different departments or teams within an Organization. Folders can include one or more Projects or additional sub-folders. Like Organizations, Folders are not required resources.
 - **Projects.** Projects are the primary resources used for organizing other GCP resources. Unlike Organizations and Folders, Projects are required. From the Project level, GCP users can create, enable, and use all GCP services and resources; manage APIs and permissions; and perform other managerial tasks.

NOTE: The *Google Cloud Platform *BETA* PowerPack* enables you to discover GCP resources at either the Organization or Project hierarchy level.

What are GCP Regions and Zones?

GCP resources are hosted in data centers around the globe. A GCP region is an individual data center located in a specific geographic locale. Regional resources are accessible by any resource within the same region. Examples of regional resources include IP addresses, subnets, regionally managed instance groups, and regional operations.

Each region consists of one or more zones. Zone-specific resources are unique to that zone and accessible only by other resources in the same zone. Examples of zone resources include instances, disks, machine types, zone-managed instance groups, and per-zone operations.

Zone names include the region name combined with a letter identifier. For example, "zone a" in the East Asia region is named "asia-east1-a".

The Dynamic Applications in the *Google Cloud Platform *BETA** PowerPack create a "region" component device for each discovered data center region and a "zone" component device for each zone discovered under those regions.

The PowerPack supports the following GCP regions and zones:

- Generic Region/Zone
- Asia East 1 (Taiwan)
- Asia East 2 (Hong Kong)
- Asia Northeast 1 (Tokyo)
- Asia Northeast 2 (Osaka)
- Asia South (Mumbai)
- Asia Southeast 1 (Singapore)
- Australia Southeast 1 (Sydney)
- Europe North 1 (Finland)
- Europe West 1 (Belgium)
- Europe West 2 (London)
- Europe West 3 (Frankfurt)
- Europe West 4 (Eemshaven)
- Europe West 6 (Zurich)
- North America Northeast 1 (Montréal)
- South America East 1 (São Paulo)
- US Central 1 (Iowa)
- US East 1 (South Carolina)
- US East 4 (Northern Virginia)
- US West 1 (Oregon)
- US West 2 (Los Angeles)

- Multi-Region Asia
- Multi-Region Europe North 1 / Europe West 4
- Multi-Region European Union
- Multi-Region United States
- Multi-Region US Central 1 / US East 1


Installing the Google Cloud Platform *BETA* PowerPack

Before completing the steps in this manual, you must import and install the latest version of the *Google Cloud Platform *BETA* PowerPack*.

NOTE: ScienceLogic does not recommend upgrading to version 101 of the *Google Cloud Platform *BETA* PowerPack* from version 100. Instead, if you are currently using version 100, ScienceLogic recommends that you delete your previously discovered Google Cloud devices as well as version 100 of the *Google Cloud Platform *BETA* PowerPack*, and then install version 101. However, be advised that doing so will cause you to lose historical device data. It will also require you to create a new Google Cloud credential and re-discover your Google Cloud account.

To download and install a PowerPack:

1. Download the PowerPack from the [ScienceLogic Customer Portal](#).
2. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
3. In the **PowerPack Manager** page, click the **[Actions]** button, then select *Import PowerPack*.
4. The **Import PowerPack** dialog box appears:



5. Click the **[Browse]** button and navigate to the PowerPack file.
6. When the **PowerPack Installer** modal page appears, click the **[Install]** button to install the PowerPack.

NOTE: If you exit the **PowerPack Installer** modal page without installing the imported PowerPack, the imported PowerPack will not appear in the **PowerPack Manager** page. However, the imported PowerPack will appear in the **Imported PowerPacks** modal page. This page appears when you click the **[Actions]** menu and select *Install PowerPack*.

Configuring Google Cloud Platform for Monitoring

Overview

The following sections describe how to configure Google Cloud Platform resources for monitoring by SL1 using the *Google Cloud Platform *BETA* PowerPack*:

<i>Creating a Google Cloud Platform Service Account</i>	7
<i>Enabling Google Cloud Platform APIs</i>	10
<i>Creating a SOAP/XML Credential for Google Cloud Platform</i>	11
<i>Creating a Google Cloud Platform Virtual Device</i>	13
<i>Aligning the Google Cloud Platform Dynamic Applications</i>	13
<i>Discovering Google Cloud Platform Component Devices</i>	14
<i>Viewing Google Cloud Platform Component Devices</i>	16
<i>Relationships Between Component Devices</i>	18

Creating a Google Cloud Platform Service Account

To monitor Google Cloud Platform (GCP) resources with SL1, you must first create a GCP **service account** for SL1 in the GCP Console. This service account belongs to SL1 instead of an individual end user, and enables SL1 to communicate with Google APIs when monitoring your GCP resources.

This service account's credentials will include a unique email address and a secret JSON key. You will include this email address and key information when you create the SOAP/XML credential that enables SL1 to monitor your GCP resources.


To create a GCP service account:

1. Log in to the GCP Console and go to the **Service accounts** page. If prompted, select a project.
2. Click the **[CREATE SERVICE ACCOUNT]** button.
3. Complete the following fields on the **Create service account** page:

The screenshot shows the Google Cloud Platform console interface for creating a service account. The top navigation bar includes the Google Cloud Platform logo, the current project name 'Example Project', and a search icon. The left sidebar is titled 'IAM & admin' and lists various IAM-related options, with 'Service accounts' highlighted. The main content area is titled 'Create service account' and features a breadcrumb trail: '1 Service account details' — '2 Grant this service account access to project (optional)' — '3 Grant users access to this service account (optional)'. The 'Service account details' section contains three input fields: 'Service account name', 'Service account ID' (pre-filled with '@example-project-198515.iam.gserviceaccount.com'), and 'Service account description'. Below the fields are 'CREATE' and 'CANCEL' buttons.

- **Service account name.** Type a name for the service account.
 - **Service account ID.** This field auto-populates with a service account ID that is based on your **Service account name**.
 - **Service account description.** Type a description for the service account.
4. Click **[Create]**. Your service account is created, and the **Service account permissions** page displays.

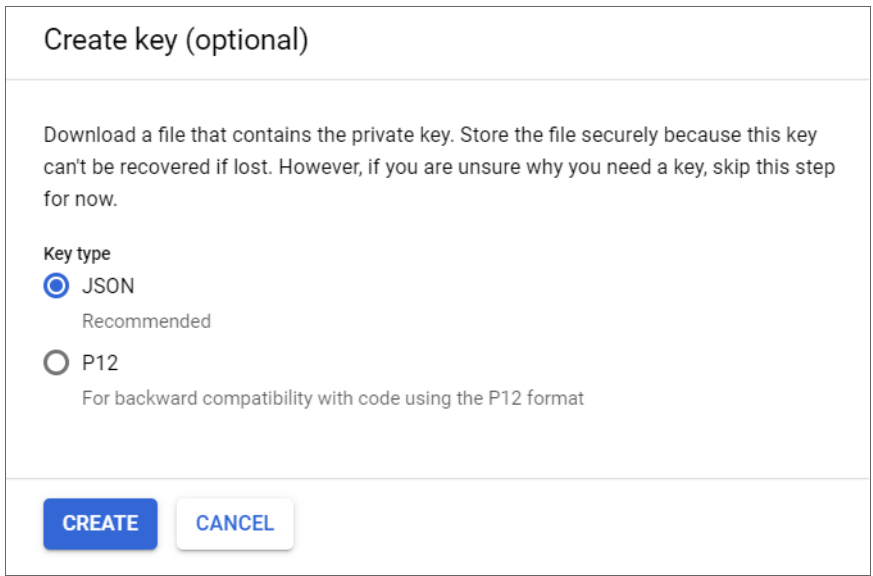
5. Complete the following fields on the **Service account permissions** page:



- **Role.** Select *Project > Viewer*.

NOTE: At a minimum, the service account must have a role of "Project" with "Viewer" permissions for the GCP service that you want to monitor.

6. Click **[Continue]**. The **Grant users access to this service account** page displays.
7. Click **[Create Key]**. The **Create key** pane appears.
8. On the **Create key** pane, select the JSON radio button and then click **[Create]**. The private JSON key is saved to your computer.



9. Click **[Close]**, and then click **[Done]**.
10. Open the JSON file that was downloaded to your computer and copy the following information:
 - client_email
 - private_key

TIP: When you copy the private key from the JSON file, it must include the "BEGIN PRIVATE KEY" and "END PRIVATE KEY" lines, including all leading and ending dashes.

If you are discovering GCP resources at the Project level, then you can skip the following steps and continue on to the [Enabling Google Cloud APIs](#) section.

However, *if you are discovering GCP resources at the Organization level*, then you must also do the following:

11. In the GCP Console, go to the **IAM** page and select your organization.
12. Click **[Add]**.
13. Add your service account as a member of the organization, and then add the following mandatory roles:
 - Role > Project > Viewer
 - Role > Resource Manager > Folder Viewer
 - Role > Resource Manager > Organization Viewer
14. When you are finished, click **[Save]**.

Enabling Google Cloud Platform APIs


Before SL1 can monitor GCP, you must also enable two APIs in the GCP portal:

- Cloud Resource Manager API
- Compute Engine API


To enable these GCP APIs:

1. Log in to the GCP Console for your project and go to the **API & Services Dashboard** page.

2. Click **[ENABLE APIS AND SERVICES]**. The **API Library** page appears.



3. In the search bar, type "Cloud Resource Manager API". The page will filter search results while you type.
4. Click the **Cloud Resource Manager API** box.



5. On the **Cloud Resource Manager API** page, click the **[Enable]** button.
6. Click **[Dashboard]** on the **API & Services** left menu and then repeat steps 2-5 to enable the **Compute Engine API**.

Creating a SOAP/XML Credential for Google Cloud Platform

To configure SL1 to monitor GCP, you must create a SOAP/XML credential that allows the Dynamic Applications in the *Google Cloud Platform *BETA** PowerPack to connect with your GCP service. An example SOAP/XML credential that you can edit for your own use is included in the *Google Cloud Platform *BETA** PowerPack.

To create a SOAP/XML credential to access GCP:

1. Go to the **Credential Management** page (System > Manage > Credentials).

2. Locate the **GCP SOAP Credential** and then click its wrench icon (🔧). The **Edit SOAP/XML Credential** modal page appears.
3. Complete the following fields:

Basic Settings

- **Profile Name.** Type a new name for the credential.

SOAP Options

- **Embedded Password [%P].** Paste the "private_key" value from the private key JSON file.
- **Embed Value [%1].** Type the "client_email" value from the private key JSON file. For example: myprojectid@myaccount.iam.gserviceaccount.com.

TIP: When you copy the "private_key" from the JSON file, it must include the "BEGIN PRIVATE KEY" and "END PRIVATE KEY" lines, including all leading and ending dashes.

4. For all remaining fields, use the default values.
5. Click the **[Save As]** button, and then click **[OK]**.

Creating a Google Cloud Platform Virtual Device

Because the GCP service does not have a static IP address, you cannot discover GCP devices using a regular discovery session. Instead, you must create a **virtual device** that represents the GCP service. A virtual device is a user-defined container that represents a device or service that cannot be discovered by SL1. You can use the virtual device to store information gathered by policies or Dynamic Applications.

To create a virtual device that represents your GCP service:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Click the **[Actions]** button and select *Create Virtual Device* from the menu. The **Virtual Device** modal page appears.
3. Complete the following fields:

The screenshot shows a modal window titled "Virtual Device" with a close button (X) in the top right. Inside the modal, there is a sub-header "Create Virtual Device" and a "Reset" button. Below this, there are four input fields: "Device Name" (a text box), "Organization" (a dropdown menu with "GCP_ViewerProj1" selected), "Device Class" (a dropdown menu with "GCP | Service" selected), and "Collector" (a dropdown menu with "CUG2" selected). At the bottom of the modal is an "Add" button.

- **Device Name.** Type a name for the device.
 - **Organization.** Select the organization for this device. The organization you associate with the device limits the users that will be able to view and edit the device. Typically, only members of the organization will be able to view and edit the device.
 - **Device Class.** Select *GCP | Service*.
 - **Collector.** Select the collector group that will monitor the device.
4. Click **[Add]** to create the virtual device.

Aligning the Google Cloud Platform Dynamic Applications

The Dynamic Applications in the *Google Cloud Platform *BETA** PowerPack are divided into the following types:

- **Discovery.** These Dynamic Applications poll GCP for new instances of services or changes to existing instances of services.
- **Configuration.** These Dynamic Applications retrieve configuration information about each service instance and retrieve any changes to that configuration information.
- **Performance.** These Dynamic Applications poll GCP for performance metrics.


When configuring SL1 to monitor GCP services, you must manually align Dynamic Applications to discover GCP component devices.

Discovering Google Cloud Platform Component Devices


To discover all the components of your GCP service, you must manually align two Dynamic Applications with the GCP virtual device. The specific Dynamic Applications that you must align to the virtual device vary based on whether you are discovering GCP resources from the Organization level or the Project level.

- If you are discovering an Organization, you must align the following Dynamic Applications:
 - GCP: Token
 - GCP: Organization Discovery
- If you are discovering GCP resources from the Project level, you must align the following Dynamic Applications:
 - GCP: Token
 - GCP: Project Discovery

To manually align these Dynamic Applications:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Click the wrench icon () for your GCP virtual device.
3. In the **Device Administration** panel, click the **[Collections]** tab. The **Dynamic Application Collections** page appears.
4. Click the **[Actions]** button and select *Add Dynamic Application* from the menu.

5. In the **Dynamic Application Alignment** modal:



- In the **Dynamic Applications** field, select *GCP Token*.
 - In the **Credentials** field, select the credential you created for your GCP service.
6. Click [**Save**] to align the Dynamic Application with the GCP virtual device.
 7. Repeat steps 2-6 to align the "GCP Project Discovery" or "GCP Project Discovery" Dynamic Application, depending on whether you are discovering an Organization or a Project.

NOTE: You must align the "GCP: Token" Dynamic Application **before** you align the "GCP: Organization Discovery" or "GCP: Project Discovery" Dynamic Application.

When you align the Dynamic Applications with the virtual device representing your GCP service, SL1 creates a component device representing your GCP Organization or Project.

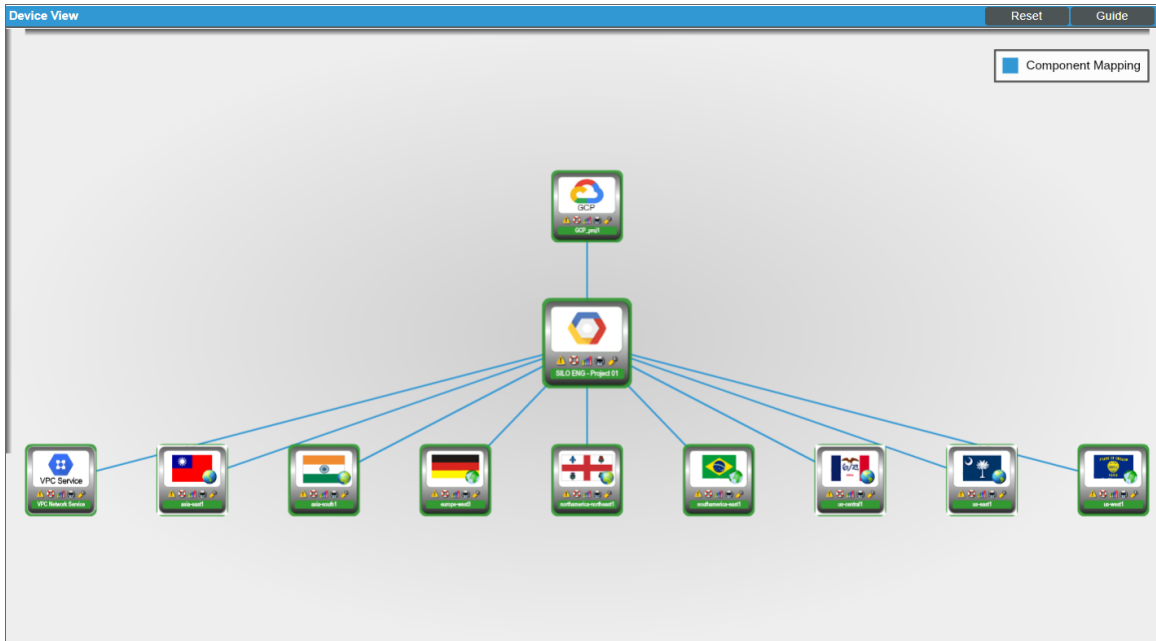
SL1 then automatically aligns several other Dynamic Applications to that component device. These Dynamic Applications discover and create additional component devices representing your GCP resources.

NOTE: SL1 might take several minutes to align these Dynamic Applications and create the component devices in your GCP service.

Viewing Google Cloud Platform Component Devices

In addition to the **Device Manager** page (Registry > Devices > Device Manager), you can view the GCP service and all associated component devices in the following places in the user interface:


- The **Device View** modal page (click the bar-graph icon [img alt="bar-graph icon"] for a device, then click the **Topology** tab) displays a map of a particular device and all of the devices with which it has parent-child relationships. Double-clicking any of the devices listed reloads the page to make the selected device the primary device:



- The **Device Components** page (Registry > Devices > Device Components) displays a list of all root devices and component devices discovered by SL1 in an indented view, so you can easily view the hierarchy and relationships between child devices, parent devices, and root devices. To view the component devices associated with the GCP service, find the GCP virtual device and click its plus icon (+):

Device Name	IP Address	Device Category	Device Class Sub-class	DID	Organization	Current State	Collection Group	Collection State
GCP_proj1	--	Service	GCP Service	1	GCP_ViewerProj1	Healthy	CUG_Automation	Active
SILO ENG - Project 01	--	Service	GCP Project	2	GCP_ViewerProj1	Notice	CUG_Automation	Active
asia-east1	--	Region	GCP Region Asia East 1 (Taiwan)	8	GCP_ViewerProj1	Healthy	CUG_Automation	Active
asia-east1-a	--	AvailabilityZone	GCP Zone Asia East 1 (Taiwan)	30	GCP_ViewerProj1	Healthy	CUG_Automation	Active
asia-east1-a Compute Service	--	Service	GCP Compute Service	55	GCP_ViewerProj1	Notice	CUG_Automation	Active
asia-east1-a trmp-delme-jpen	--	Compute	GCP Compute Instance	141	GCP_ViewerProj1	Major	CUG_Automation	Unavailable
asia-east1-a vmmwithmoveddisks	--	Compute	GCP Compute Instance	91	GCP_ViewerProj1	Healthy	CUG_Automation	Active
asia-east1-a Persistent Disk Service	--	Service	GCP Persistent Disk Service	56	GCP_ViewerProj1	Notice	CUG_Automation	Active
asia-east1-a qppersistentdiskstandard	--	Storage	GCP Storage Persistent Disk	96	GCP_ViewerProj1	Healthy	CUG_Automation	Active
asia-east1-a template-disk-ssd-moved	--	Storage	GCP Storage Persistent Disk	100	GCP_ViewerProj1	Healthy	CUG_Automation	Active

- The **Component Map** page (Views > Device Maps > Components) allows you to view devices by root node and view the relationships between root nodes, parent components, and child components in a map. This makes it easy to visualize and manage root nodes and their components. SL1 automatically updates the **Component Map** as new component devices are discovered. The platform also updates each map with the latest status and event information. To view the map for the GCP service, go to the **Component Map** page and select the map from the list in the left NavBar. To learn more about the **Component Map** page, see the **Views** manual.



Relationships Between Component Devices

In addition to parent/child relationships between component devices, SL1 also creates relationships between the following component devices:

- Compute Instances and Storage Persistent Disks
- Compute Instances and Subnets
- Compute Instances and VPC Networks
- Load Balancing Global HTTPS and Backend Buckets
- Load Balancing Global HTTPS and Backend Services
- Load Balancing Global HTTPS and Default Backend Services
- Load Balancing Global SSL Proxy and Backend Services
- Load Balancing Global TCP Proxy and Backend Services
- Load Balancing Regional Network TCP/UDP and Compute Instances
- VPC Subnets and VPC Networks

NOTE: If an instance is configured in GCP to automatically delete any associated read-write persistent disks when the instance is deleted, then that behavior will also occur in SL1: If the instance is deleted, its related persistent disks will also be deleted. This behavior is controlled in GCP on the **VM Instances** page by the *Delete boot disk when instance is deleted* checkbox for boot disks and the **When deleting instance** field for additional disks.

Additionally, SL1 can also build relationships between GCP VM Instances and Kubernetes Nodes, for users who also have the *Kubernetes PowerPack* installed.

Google Cloud Platform Dashboards

Overview


The following sections describe the device dashboards that are included in the *Google Cloud Platform *BETA* PowerPack*:

Device Dashboards	19
Google Backend Service	20
Google Compute Instance	21
Google Compute Instance Service	22
Google Global HTTP(S) Load Balancer	23
Google Persistent Disk	24
Google Persistent Disk Service	25
Google Project	26
Google Regional TCP/UDP Load Balancer	27
Google Storage Bucket	28
Google Storage Bucket Service	29

Device Dashboards

The *Google Cloud Platform *BETA* PowerPack* includes device dashboards that provide summary information for GCP component devices. Each of the device dashboards in the *Google Cloud Platform *BETA* PowerPack* is set as the default device dashboard for the equivalent device class.

Google Backend Service



The **Google Backend Service** dashboard displays the following information:

- The basic information about the device
- Six instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Inbound traffic
 - Outbound traffic
 - New connections per second
 - Closed connections per second
 - Frontend RTT
 - Open connections

Google Compute Instance




3

The **Google Compute Instance** dashboard displays the following information:

- The basic information about the device
- A list of events and tickets for the device
- Nine instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - CPU utilization
 - Availability
 - Reserved cores
 - Disk IOPS
 - Dropped packets
 - Packets sent and received
 - Disk read/write bytes count
 - Dropped bytes
 - Bytes sent/received

Google Compute Instance Service



The **Google Compute Instance Service** dashboard displays the following information:

- The basic information about the device
- Three instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Top 10 instances by CPU utilization
 - Top 10 instances by disk write IOPS
 - Top 10 instances by packets received

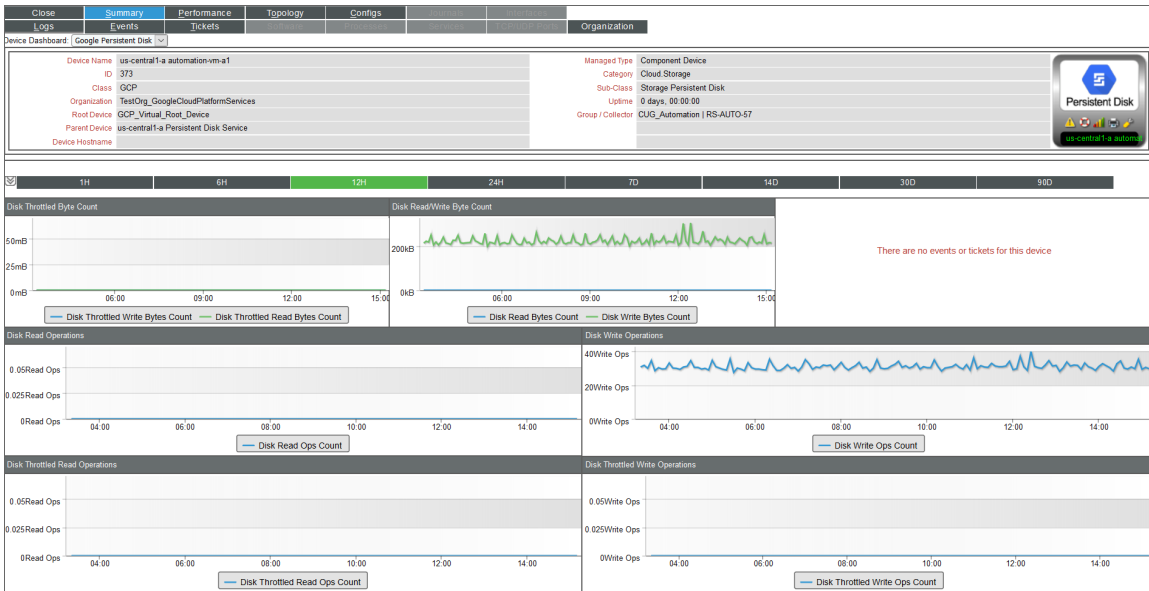
Google Global HTTP(S) Load Balancer



The **Google Global HTTP(S) Load Balancer** dashboard displays the following information:

- The basic information about the device
- A list of events and tickets for the device
- Five instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Latency
 - Request counts
 - Frontend RTT
 - Request bytes
 - Response bytes

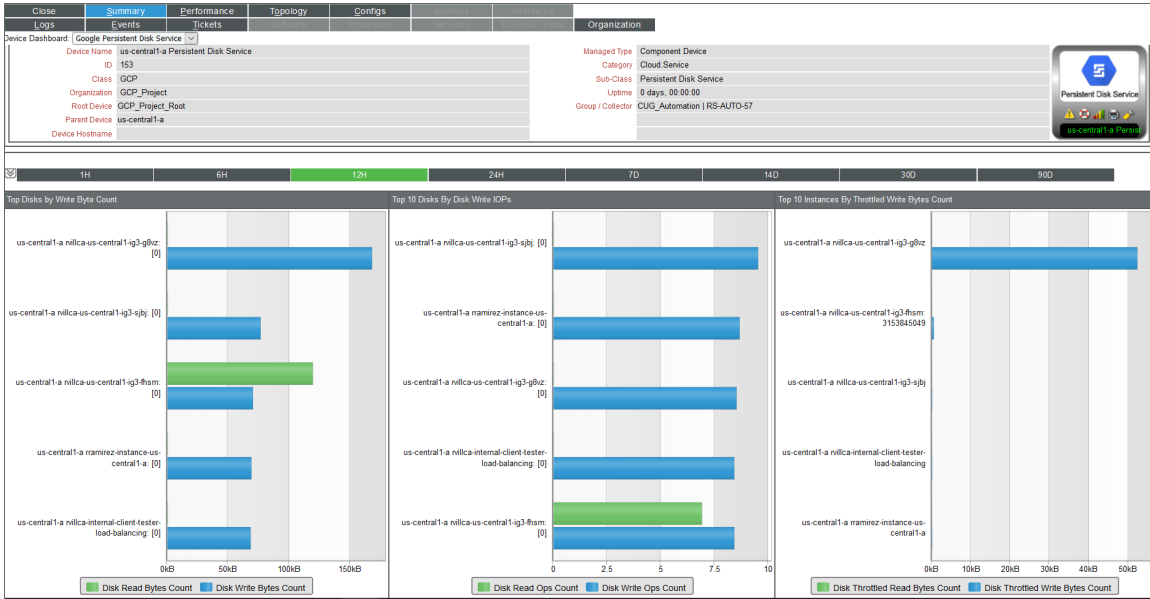
Google Persistent Disk



The **Google Persistent Disk** dashboard displays the following information:

- The basic information about the device
- A list of events and tickets for the device
- Six instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Disk throttled byte count
 - Disk read/write byte count
 - Disk read operations
 - Disk write operations
 - Disk throttled read operations
 - Disk throttled write operations

Google Persistent Disk Service

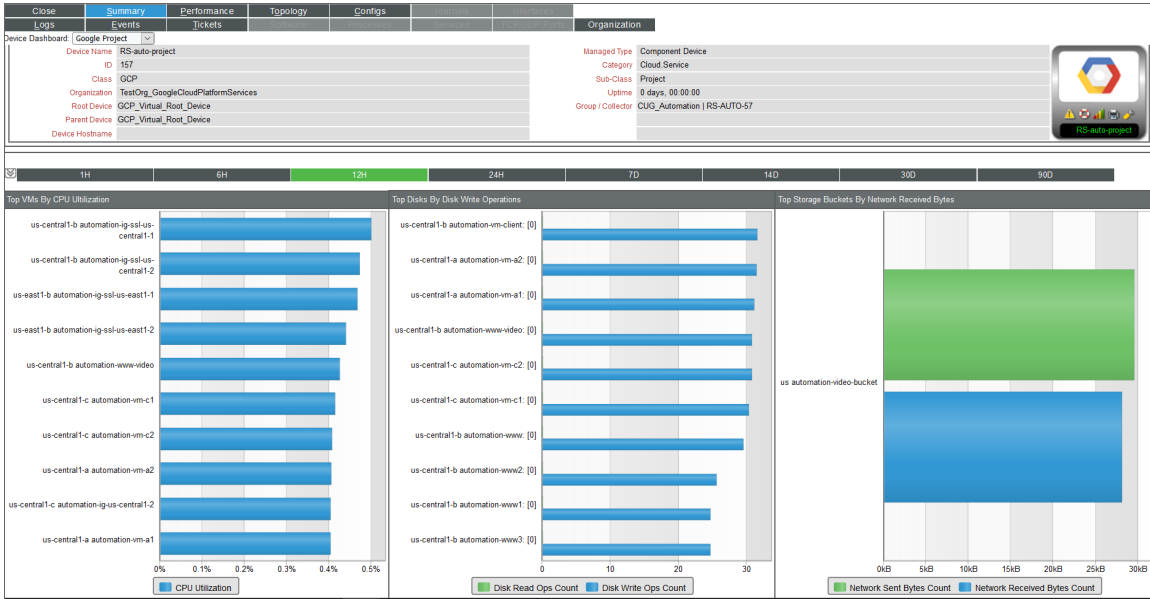


The **Google Persistent Disk Service** dashboard displays the following information:

- The basic information about the device
- Three instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Top 10 disks by write byte count
 - Top 10 disks by disk write IOPS
 - Top 10 instances by throttled write bytes count




Google Project



The **Google Project** dashboard displays the following information:

- The basic information about the device
- Three instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Top VMs by CPU utilization
 - Top disks by write operations
 - Top storage buckets by network received bytes


Google Regional TCP/UDP Load Balancer



The **Google Regional TCP/UDP Load Balancer** dashboard displays the following information:

- The basic information about the device
- A list of events and tickets for the device
- Three instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - RTT latency
 - Throughput inbound/outbound
 - Packets inbound/outbound


Google Storage Bucket



The **Google Storage Bucket** dashboard displays the following information:

- The basic information about the device
- A list of events and tickets for the device
- Five instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Network sent/received byte counts
 - Object ACL access and mutation counts
 - API request count
 - Storage object count
 - Total bytes

Google Storage Bucket Service



The **Google Storage Bucket Service** dashboard displays the following information:

- The basic information about the device
- Three instances of the Multi-series Performance Widget that display the following metrics trended over the specified period of time:
 - Top storage buckets by network sent bytes
 - Top storage buckets by size
 - Top storage buckets by ACL access count



© 2003 - 2019, 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