



Monitoring Google Cloud Platform

Google Cloud Platform *BETA* PowerPack version 100

Table of Contents

Introduction	3
What is Google Cloud Platform?	3
What Does the Google Cloud Platform *BETA* PowerPack Monitor?	4
What are GCP Regions and Zones?	4
Installing the Google Cloud Platform *BETA* PowerPack	5
Configuring Google Cloud Platform for Monitoring	7
Creating a Google Cloud Platform Service Account	7
Enabling Google Cloud Platform APIs	9
Creating an SSH/Key Credential for Google Cloud Platform	10
Creating a Google Cloud Platform Virtual Device	12
Aligning the Google Cloud Platform Dynamic Applications	13
Discovering Google Cloud Platform Component Devices	13
Viewing Google Cloud Platform Component Devices	15
Relationships Between Component Devices	17

Chapter 1

Introduction

Overview

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

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

- [What is Google Cloud Platform? 3](#)
- [What Does the Google Cloud Platform *BETA* PowerPack Monitor? 4](#)
 - [What are GCP Regions and Zones? 4](#)
- [Installing the Google Cloud Platform *BETA* PowerPack 5](#)

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 the ScienceLogic platform, 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:

- An example Credential that you can use as a template to create SSH/Key credentials to monitor GCP resources
- Dynamic Applications to discover, model, and monitor the following GCP resources:
 - Compute instances
 - Storage persistent disks
 - VPC networks
- Device Classes for each GCP region and zone, plus all of the GCP resources the ScienceLogic platform monitors
- Event Policies and corresponding alerts that are triggered when GCP resources meet certain status criteria

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 Northeast 1 (Tokyo)
- Asia South (Mumbai)

- Asia Southeast 1 (Singapore)
- Australia Southeast 1 (Sydney)
- Europe West 1 (Belgium)
- Europe West 2 (London)
- Europe West 3 (Frankfurt)
- Europe West 4 (Eemshaven)
- 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)

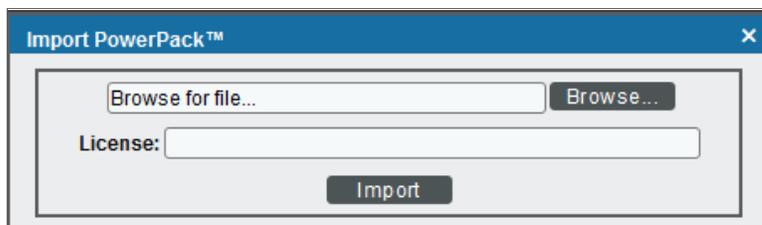
Installing the Google Cloud Platform *BETA* PowerPack

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

TIP: By default, installing a new version of a PowerPack overwrites all content in that PowerPack that has already been installed on the target system. You can use the **Enable Selective PowerPack Field Protection** setting in the **Behavior Settings** page (System > Settings > Behavior) to prevent new PowerPacks from overwriting local changes for some commonly customized fields. (For more information, see the **System Administration** manual.)

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 the ScienceLogic platform using the *Google Cloud Platform *BETA** PowerPack:

<i>Creating a Google Cloud Platform Service Account</i>	7
<i>Enabling Google Cloud Platform APIs</i>	9
<i>Creating an SSH/Key Credential for Google Cloud Platform</i>	10
<i>Creating a Google Cloud Platform Virtual Device</i>	12
<i>Aligning the Google Cloud Platform Dynamic Applications</i>	13
<i>Discovering Google Cloud Platform Component Devices</i>	13
<i>Viewing Google Cloud Platform Component Devices</i>	15
<i>Relationships Between Component Devices</i>	17

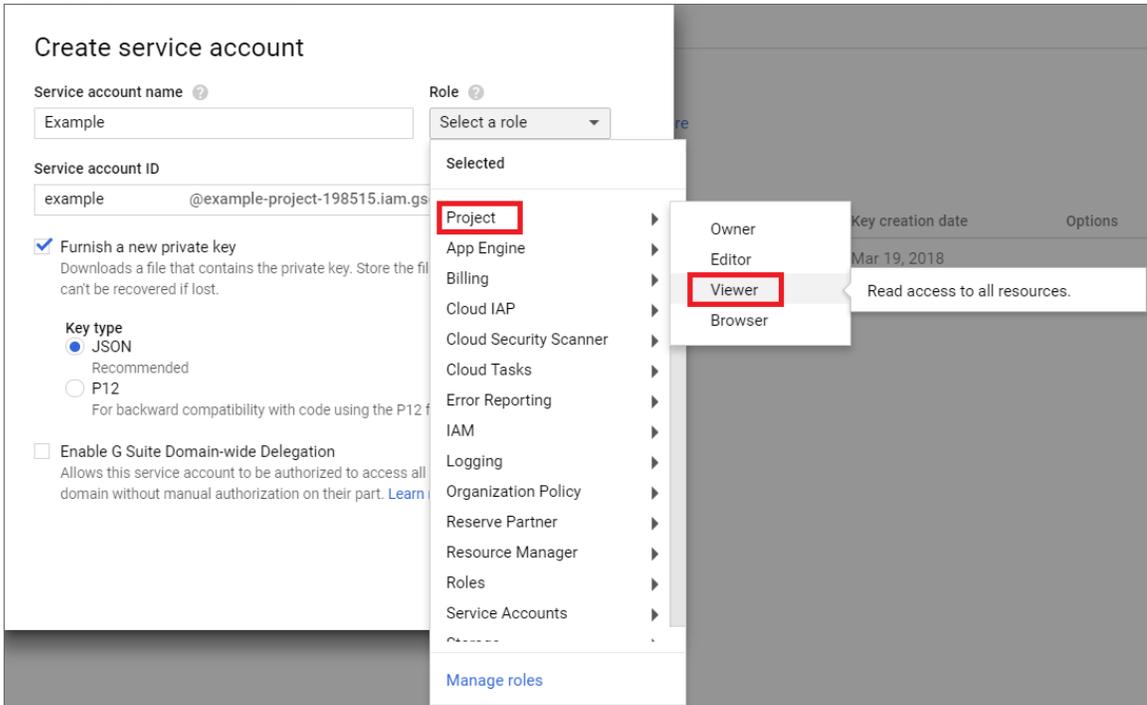
Creating a Google Cloud Platform Service Account

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

This service account's credentials will include a unique email address and a public/private key pair. You will include this email address and key information when you create the SSH/Key credential that enables the platform to monitor your GCP resources.

To create a GCP service account:

1. Log in to the Google API 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:



- **Service account name.** Type a name for the service account.
- **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.

- **Service account ID.** This field auto-populates with a service account ID that is based on your **Service Account Name**.
 - **Furnish a new private key.** Select this checkbox.
 - **Key type.** Select *JSON*.
4. Click **[Create]**. Your service account is created and the private key JSON file is saved to your computer.
 5. Click **[Close]**.

6. 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. When copied, the key also includes end line characters ("\n") that **must be removed** before you paste the value into the SSH/Key credential. ScienceLogic recommends that you first paste the private key into a text editor and search for "\n" (without the quotation marks) to remove all of these characters.

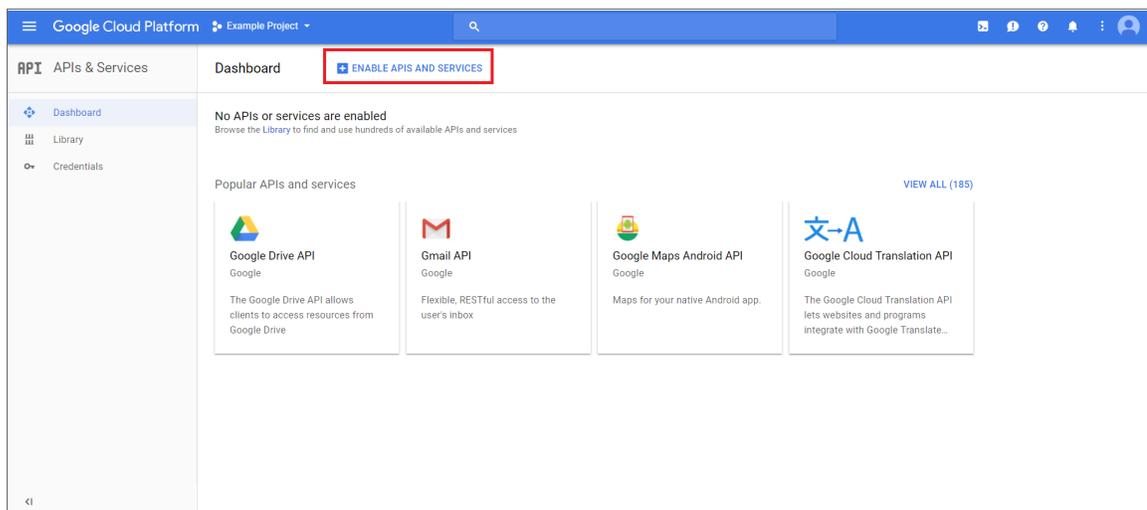
Enabling Google Cloud Platform APIs

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

- Cloud Resource Manager API
- Stackdriver Monitoring API

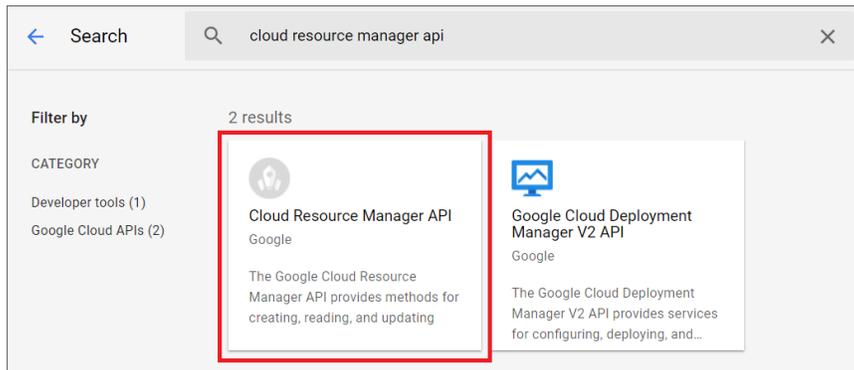
To enable these GCP APIs:

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 **Stackdriver Monitoring API**.

Creating an SSH/Key Credential for Google Cloud Platform

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

To create an SSH/Key credential to access GCP:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Locate the **GCP Sample Credential** and then click its wrench icon (🔧). The **Edit SSH/Key Credential** modal page appears.

3. Complete the following fields:

The screenshot shows a window titled "Credential Editor [79]" with a sub-header "Edit SSH/Key Credential #79". The window contains a "Basic Settings" section with the following fields:

- Credential Name:** A text box containing "GCP Sample Credential".
- Hostname/IP:** A text box containing "<CLIENT_EMAIL>".
- Port:** A text box containing "22".
- Timeout(ms):** A text box containing "0".
- Username:** An empty text box.
- Password:** An empty text box.
- Private Key (PEM Format):** A large empty text area.

At the top right of the window are "New" and "Reset" buttons. At the bottom are "Save" and "Save As" buttons.

- **Credential Name.** Type a new name for the credential.
- **Hostname/IP.** Type the "client_email" value from the private key JSON file. For example, "myprojectid@myaccount.iam.gserviceaccount.com".

NOTE: The account that you enter in the **Hostname/IP** field must at a minimum have a role of "Project" with "Viewer" permissions for the GCP service that you want to monitor.

- **Port.** Type "22".
- **Timeout(ms).** Type the time, in milliseconds, after which the ScienceLogic platform will stop trying to communicate with GCP.
- **Username.** Leave this field blank.
- **Password.** Leave this field blank.

- **Private Key (PEM Format)**. Paste the "private_key" value from the private key JSON file.

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. When copied, the key also includes end line characters ("\n") that **must be removed** for the SSH/Key credential to work properly. ScienceLogic recommends that you first paste the private key into a text editor and search for "\n" (without the quotation marks) to remove all of these characters, and then copy and paste the remaining value from the text editor into the **Private Key (PEM Format)** field.

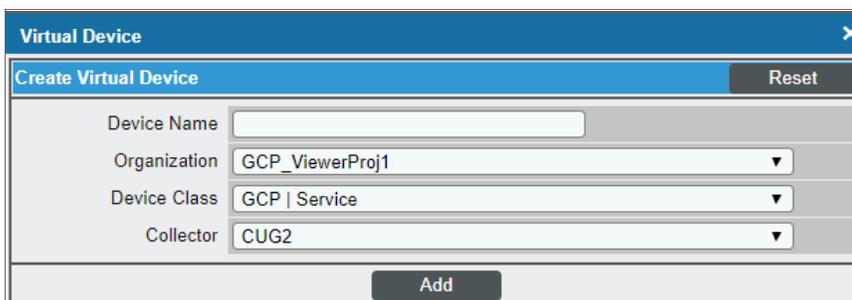
4. Click **[Save As]**.
5. When the confirmation message appears, click **[OK]**.

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 the ScienceLogic platform. 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:



- **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 the ScienceLogic platform to monitor GCP services, you must manually align Dynamic Applications to discover GCP component devices.

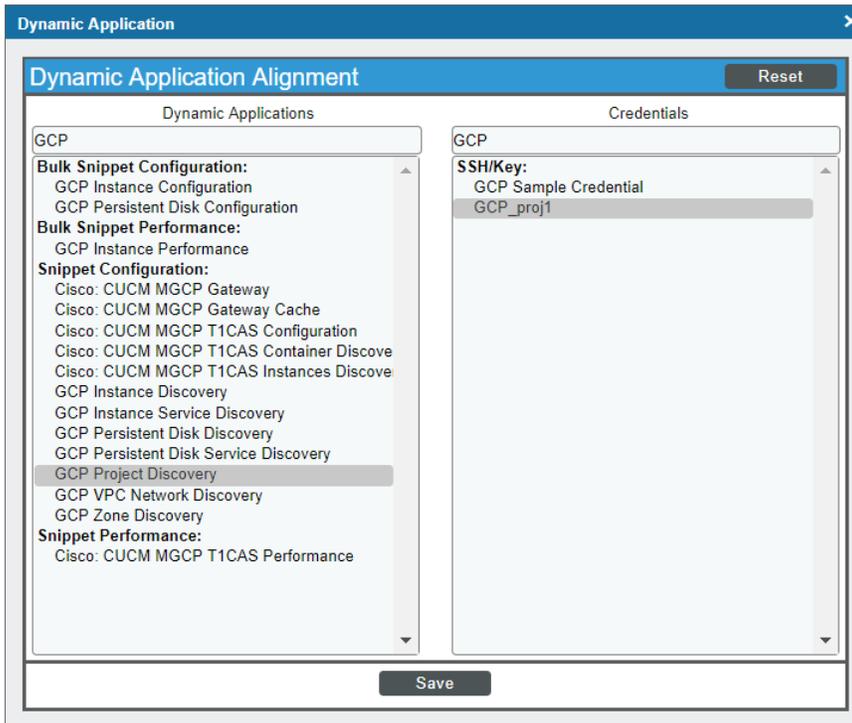
Discovering Google Cloud Platform Component Devices

To discover all the components of your GCP service, you must manually align the "GCP Project Discovery" Dynamic Application with the GCP virtual device.

To manually align the "GCP Project Discovery" Dynamic Application:

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 Project Discovery*.
- 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.

When you align the "GCP Project Discovery" Dynamic Application with the virtual device representing your GCP service, the ScienceLogic platform creates a component device representing your GCP project.

The platform then automatically aligns several other Dynamic Applications to that component device. These Dynamic Applications discover and create additional component devices for each region and zone used by the GCP project, plus the following components:

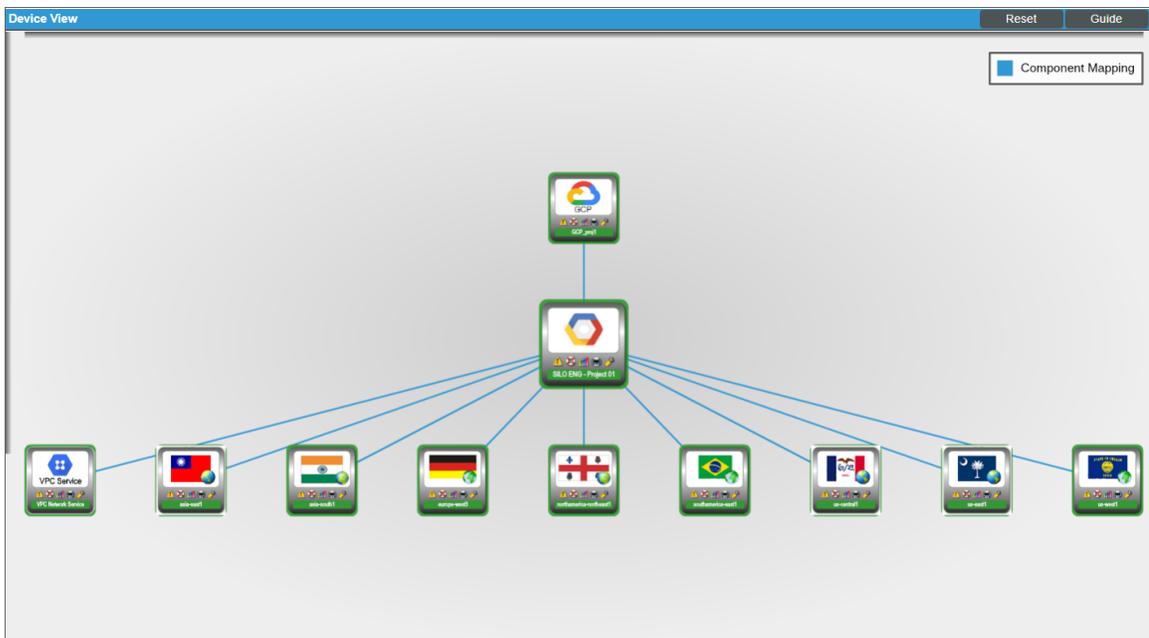
- Compute Services
 - Compute Instances
- Persistent Disk Services
 - Storage Persistent Disks
- VPC Network Services
 - VPC Networks

NOTE: The ScienceLogic platform 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:

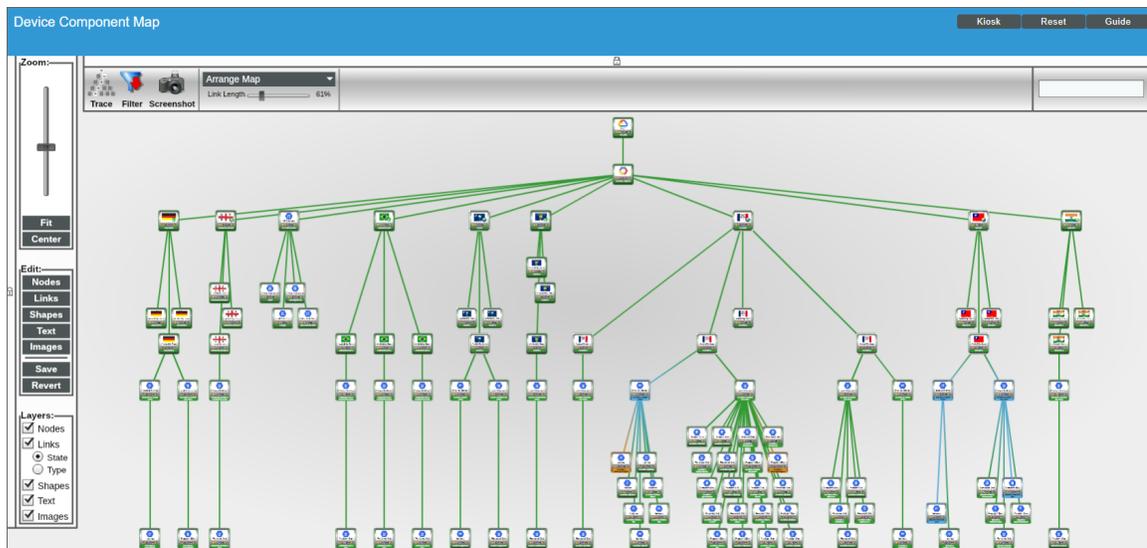
- The **Device View** modal page (click the bar-graph icon  for a device, then click the **Topology** tab) displays a map of a particular device and all of the devices with which it has parent-child relationships. Double-clicking any of the devices listed reloads the page to make the selected device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



- The **Device Components** page (Registry > Devices > Device Components) displays a list of all root devices and component devices discovered by the ScienceLogic platform 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-dsime-jpen	--	Compute	GCP Compute Instance	141	GCP_ViewerProj1	Major	CUG_Automation	Unavailable
asia-east1-a vrmwithmoweddskis	--	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 qpersistentdiskstandard	--	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. The ScienceLogic platform 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, the ScienceLogic platform also creates relationships between the following component devices:

- Compute instances and storage persistent disks
- Compute instances 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 the ScienceLogic platform: 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.

© 2003 - 2018, 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