

GCP Integrations

Introduction

This document shows information related to GCP Integration.

The Google Cloud integration collects and parses Google Cloud Audit Logs, VPC Flow Logs, Firewall Rules Logs and Cloud DNS Logs that have been exported from Cloud Logging to a Google Pub/Sub topic sink and collects Google Cloud metrics and metadata from Google Cloud Monitoring.

Requirements

To use this Google Cloud Platform (GCP) integration, the following needs to be set up:

- Service Account with a Role. (Section 3.1.1)
- Service Account Key to access data on your GCP project. (Section 3.1.2)

Service Accounts

First, please create a Service Account. A Service Account (SA) is a particular type of Google account intended to represent a non-human user who needs to access the GCP resources.

The Log Collector uses the SA (Service Account) to access data on Google Cloud Platform using the Google APIs.

Here is a reference from Google related to Service Accounts:

<https://cloud.google.com/iam/docs/best-practices-for-securing-service-accounts>

Service Account with a Role

You need to grant your Service Account (SA) access to Google Cloud Platform resources by assigning a role to the account. In order to assign minimal privileges, create a custom role that has only the privileges required by the Log Collector. Those privileges are below:

- compute.instances.list (required for GCP Compute instance metadata collection) (**2)
- monitoring.metricDescriptors.list
- monitoring.timeSeries.list
- pubsub.subscriptions.consume
- pubsub.subscriptions.create (*1)
- pubsub.subscriptions.get
- pubsub.topics.attachSubscription (*1)

*1 Only required if Agent is expected to create a new subscription. If you create the subscriptions yourself, you may omit these privileges.

**2 Only required if corresponding collection will be enabled.

After you have created the custom role, assign the role to your service account.

Service Account Key

Next, with the Service Account (SA) with access to Google Cloud Platform (GCP) resources setup in Section 3.1.1, you need some credentials to associate with it: a Service Account Key.

From the list of SA (Service Accounts):

1. Click the one you just created to open the detailed view.
2. From the Keys section, click "Add key" > "Create new key" and select JSON as the type.
3. Download and store the generated private key securely (remember, the private key can't be recovered from GCP if lost).

GCP Integrations Procedures

1. GCP Audit Logs

The audit dataset collects audit logs of administrative activities and accesses within your Google Cloud resources.

Procedures

The "Project Id" and either the "Credentials File" or "Credentials JSON" will need to be provided in the integration UI when adding the Google Cloud Platform integration.

Please provide the following information to CyTech:

- **Project ID**

The Project ID is the Google Cloud project ID where your resources exist.

- **Credentials File vs JSON**

Based on your preference, specify the information in either the Credentials File OR the Credentials JSON field.

OPTION 1: CREDENTIALS FILE

Save the JSON file with the private key in a secure location of the file system, and make sure that the Elastic Agent has at least read-only privileges to this file.

Specify the file path in the Elastic Agent integration UI in the "Credentials File" field. For example: /home/ubuntu/credentials.json.

OPTION 2: CREDENTIALS JSON

Specify the content of the JSON file you downloaded from Google Cloud Platform directly in the Credentials JSON field in the Elastic Agent integration.

2. GCP DNS Logs

The dns dataset collects queries that name servers resolve for your Virtual Private Cloud (VPC) networks, as well as queries from an external entity directly to a public zone.

Procedures

The "Project Id" and either the "Credentials File" or "Credentials JSON" will need to be provided in the integration UI when adding the Google Cloud Platform integration.

Please provide the following information to CyTech:

- **Project ID**

The Project ID is the Google Cloud project ID where your resources exist.

- **Credentials File vs JSON**

Based on your preference, specify the information in either the Credentials File OR the Credentials JSON field.

OPTION 1: CREDENTIALS FILE

Save the JSON file with the private key in a secure location of the file system, and make sure that the Elastic Agent has at least read-only privileges to this file.

Specify the file path in the Elastic Agent integration UI in the "Credentials File" field. For example: /home/ubuntu/credentials.json.

OPTION 2: CREDENTIALS JSON

Specify the content of the JSON file you downloaded from Google Cloud Platform directly in the Credentials JSON field in the Elastic Agent integration.

3. GCP Firewall Logs

The firewall dataset collects logs from Firewall Rules in your Virtual Private Cloud (VPC) networks.

Procedures

The "Project Id" and either the "Credentials File" or "Credentials JSON" will need to be provided in the integration UI when adding the Google Cloud Platform integration.

Please provide the following information to CyTech:

- **Project ID**

The Project ID is the Google Cloud project ID where your resources exist.

- **Credentials File vs JSON**

Based on your preference, specify the information in either the Credentials File OR the Credentials JSON field.

OPTION 1: CREDENTIALS FILE

Save the JSON file with the private key in a secure location of the file system, and make sure that the Elastic Agent has at least read-only privileges to this file.

Specify the file path in the Elastic Agent integration UI in the "Credentials File" field. For example: /home/ubuntu/credentials.json.

OPTION 2: CREDENTIALS JSON

Specify the content of the JSON file you downloaded from Google Cloud Platform directly in the Credentials JSON field in the Elastic Agent integration.

4. GCP VPC Flow Logs

The vpcflow dataset collects logs sent from and received by VM instances, including instances used as GKE nodes.

Procedures

The "Project Id" and either the "Credentials File" or "Credentials JSON" will need to be provided in the integration UI when adding the Google Cloud Platform integration.

Please provide the following information to CyTech:

- **Project ID**

The Project ID is the Google Cloud project ID where your resources exist.

- **Credentials File vs JSON**

Based on your preference, specify the information in either the Credentials File OR the Credentials JSON field.

OPTION 1: CREDENTIALS FILE

Save the JSON file with the private key in a secure location of the file system, and make sure that the Elastic Agent has at least read-only privileges to this file.

Specify the file path in the Elastic Agent integration UI in the "Credentials File" field. For example: `/home/ubuntu/credentials.json`.

OPTION 2: CREDENTIALS JSON

Specify the content of the JSON file you downloaded from Google Cloud Platform directly in the Credentials JSON field in the Elastic Agent integration.

Revision #3

Created 23 April 2024 12:59:48

Updated 19 June 2024 06:54:01