Configuring a GCS Endpoint

Google Cloud Storage (GCS) provides a bit more flexibility when it comes to configuring your bucket in preparation for log streaming from StackPath. A few examples of this include having the ability to choose between using one of Google Cloud's built-in Roles or defining your own, as well as the ability to grant StackPath access to either the entire bucket or a subdirectory within a bucket.

You can use this document to learn how to configure your GCS account to enable StackPath to stream logs to your GCS bucket.

At a high level, you will:

- Create and configure a GCS bucket, enabling it to receive logs from StackPath.
    - Create a GCS IAM Role
    - Grant permissions to StackPath's Log Streaming Service Account
    - Restrict access to a subdirectory (optional)

Support

Your Account or Sales representative can assist you in enabling any of the fields listed in the **Available Fields** section here, or you may reach out to Support at hi@stackpath.com for assistance with this as well.

Support requires that you provide the following information when requesting assistance with enabling logs or additional fields:

- Site name
- Desired fields
- Bucket name (e.g *my-stackpath-logs-bucket*)
- Subdirectory to use (e.g if you want to use my-stackpath-logs-bucket/cdn/accesslogs, the subdirectory is cdn/accesslogs)

For more information, see the official documentation for GCS Identity and Access Management.

GCS Configuration

This section will explain the steps required before StackPath can start sending logs to your GCS bucket. Follow the steps below to configure the sink you would like your logs sent to.

Step 1: Select a Role

You can either use Google Cloud's pre-made roles, or you can define your own.

Google Cloud Roles

If using Google Cloud's roles, use Storage Object Creator and Storage Object Viewer.

- Storage Object Creator is for writing logs.
- Storage Object Viewer is for GCS's endpoint healthcheck.

The permissions these roles include can be reviewed in the official docs.

Custom Role

If you prefer to grant narrower permissions, a custom role can be used instead. This must include the following permissions: `storage.objects.create` and `storage.objects.list`

- `storage.objects.create` is the exact permission for writing logs
- `storage.objects.list` is the exact permission for GCS's endpoint healthcheck

Note that `storage.objects.list` is not strictly necessary for streaming logs, however, excluding this permission will make all healthchecks fail, which makes validating configuration and access more difficult.

To create a custom role, please see the official docs on Creating Custom Roles.

Example using `gsutil`:

bash

```
gcloud iam roles create role-id --project=project-id \
--title=role-title --description=role-description \
 --permissions="storage.objects.create,storage.objects.list"
```

Replace `project-id` with the ID of the project your bucket is in, and give the role a descriptive `role-title`, `role-description`, and `role-id`. Record the `role-id` and save it for the next step.

Step 2: Grant Permissions to StackPath's Log Streaming Service Account

Using either the two built-in roles (Storage Object Creator and Storage Object Viewer), or the custom role you created, grant permission to the StackPath Log Streaming Service Account for the role(s) on the bucket you want logs written to.

The serviceAccount to grant permission to is `stackpath-log-streaming-gcs@sp-log-streaming.iam.gserviceaccount.com`.

To do this, please see the official docs on Adding a principal to a bucket-level policy.

Example using gsutil and the built-in roles:

```
bash
gsutil iam ch \
serviceAccount:stackpath-log-streaming-gcs@sp-log-streaming.iam.gserviceaccount.com:objectViewer \
gs://BUCKET_NAME

gsutil iam ch \
serviceAccount:stackpath-log-streaming-gcs@sp-log-streaming.iam.gserviceaccount.com:objectCreator \
gs://BUCKET_NAME
```

Be sure to replace **BUCKET_NAME** with the name of your bucket.

Example using gsutil and a custom role:

```
bash
gsutil iam ch \
serviceAccount:stackpath-log-streaming-gcs@sp-log-streaming.iam.gserviceaccount.com:projects/PROJECT_ID/roles/ROLE_ID \
gs://BUCKET_NAME
```

Be sure to replace **PROJECT_ID** with your project's ID, **ROLE_ID** with the ID of the role created in Step 1, and **BUCKET_NAME** with the name of your bucket.

Special Case: Subdirectories

If you need to restrict access to only a specific subdirectory, you can, however, the behavior is different.

Specifically, the storage.objects.list permission only works at the bucket-level, so granting list at subdirectory levels effectively does nothing. Writing logs to the bucket subdirectory will still succeed, however, the healthcheck endpoint will always fail, which makes validating configuration and access more difficult.

If restriction to a subdirectory is needed, you can grant this by adding a Condition to the policy.

Please see the docs on how to Use IAM Conditions on buckets.

First, determine the full resource path you want to restrict access to. The format will appear as follows:

```
projects/_/buckets/BUCKET_NAME/objects/PATH_TO_SUBDIRECTORY/
```

The Condition will be:

"Name Starts with projects/_/buckets/BUCKET_NAME/objects/PATH_TO_SUBDIRECTORY/"

For example, given a bucket named *my-logs-bucket*, with a subdirectory path of logs/cdn/stackpath, the Condition will be:

"Name Starts with projects/_/buckets/my-logs-bucket/objects/logs/cdn/stackpath/"

Adding this in the Google Cloud Console is easiest, however if following the docs' gsutil example, the condition for this example would be:

```
json
"condition": {
"title": "some-descriptive-title",
"description": "some-descriptive-description",
"expression": "resource.name.startsWith(\"projects/_/buckets/my-logs-bucket/objects/logs/cdn/stackpath/\")"
}
```

For more information on log streaming, please see Enabling Log Streaming.