Community Tutorials

COMMUNITY HOME (/COMMUNITY) SEARCH TUTORIALS (/DOCS/TUTORIALS)

EDIT ON G
(HTTPS://GITHUB.COM/GOOGLECLOUDPLATFORM/COMMUNITY/EDIT/MASTER/TUTORIALS/EXPO
STACKDRIVER-ELASTICCLOUD/INDI

REPORT ISSUE
(HTTPS://GITHUB.COM/GOOGLECLOUDPLATFORM/COMMUNITY/ISSUES/NEW?
TITLE=ISSUE%20WITH%20TUTORIALS/EXPORTING-STACKDRIVER-
ELASTICCLOUD/INDEX.MD&BODY=ISSUE%20DESCRIPTION)

PAG
(HTTPS://GITHUB.COM/GOOGLECLOUDPLATFORM/COMMUNITY/COMMITS/MASTER/TUTORIALS/
STACKDRIVER-ELASTICCLOUD

Exporting Stackdriver logs to Elastic Cloud

Author(s): @twenny (https://github.com/twenny), Published: 2019-03-20

★ Google Cloud Community tutorials submitted from the community do not represent official Google Cloud product documentation.

Overview

This tutorial explains how to export Stackdriver logs to the Elastic Cloud Elasticsearch SaaS platform to perform log analytics. Elastic Cloud is a SaaS offering, which saves time by not needing to build and manage the Elasticsearch infrastructure.

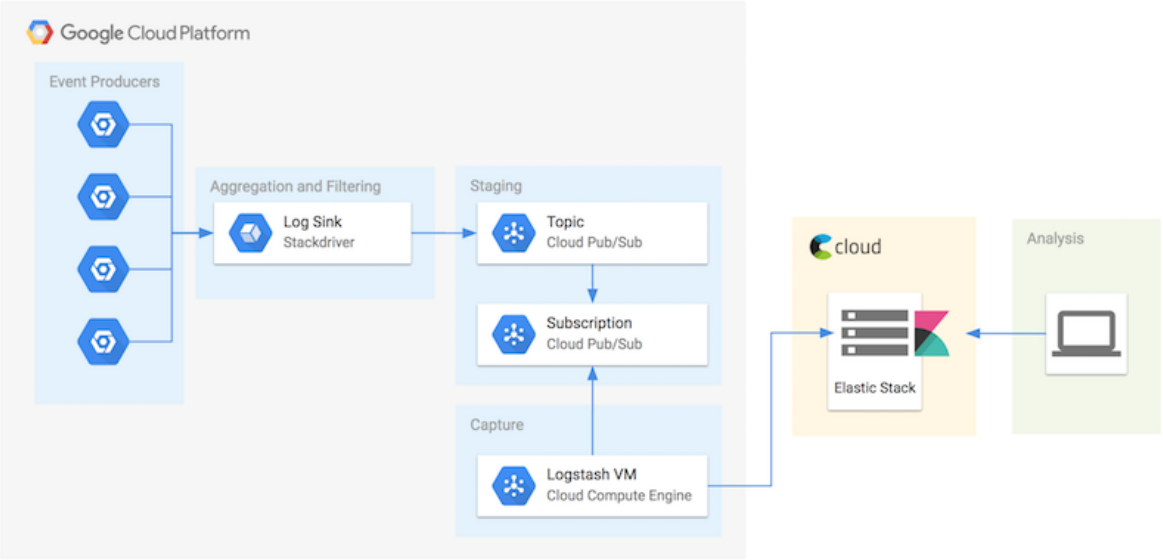

Costs

This tutorial uses billable components of Google Cloud Platform (GCP), including Compute Engine.

New GCP users might be eligible for a free trial (/free-trial).

Configure GCP resources

The high-level steps in this section:

1. Create a user-managed service account
2. Create a VM for Logstash
3. Create a Cloud Pub/Sub topic
4. Create a Stackdriver log sink and subscribe it to the Cloud Pub/Sub topic

Enable APIs

Log in or sign up for Google Cloud Platform (https://cloud.google.com), then open the Cloud Console (https://console.cloud.google.com).

The examples in this document use the `gcloud` command-line inteface. GCP APIs must be enabled via the Services and APIs page (https://console.cloud.google.com/apis/dashboard) in the console before they can be used with `gcloud`. To perform the steps in this tutorial, enable the following APIs:

- Compute Engine
- Cloud Pub/Sub
- Identity and Access Management (IAM)
- Stackdriver

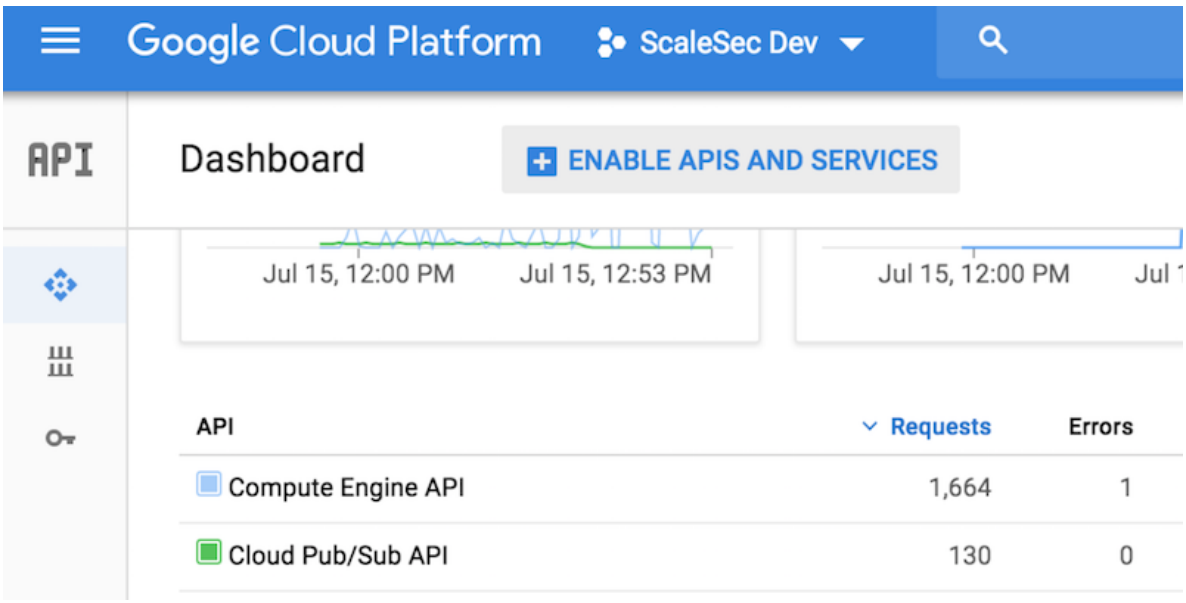

Activate Google Cloud Shell

The GCP Console provides an interactive shell that includes the `gcloud` command-line interface. At the top right corner of the page, click the **Activate Google Cloud Shell** button.

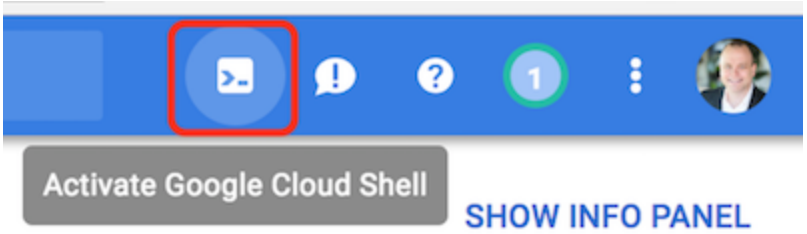

Create a service account

GCP best practices (/vpc/docs/firewalls#service-accounts-vs-tags) suggest using a service account to configure security controls to a VM. A service account is useful for a VM to determine which other GCP resources can be accessed by the VM and its applications, and which firewall rules should be applied to the VM.

While credentials can be created to be used by a service account, this step is not necessary when the service account is attached to a VM running on Google Compute Engine. Google manages the keys, and applications can retrieve the credentials securely (/compute/docs/access/create-enable-service-accounts-for-instances#authenticating_applications_using_service_account_credentials) with the metadata service.

1. Create a service account to attach to the VM:

```
gcloud iam service-accounts create logstash \
    --display-name="Logstash to Stackdriver"
```

Expected response:

```
Created service account [logstash].
```

2. Provide IAM permissions allowing the new service account to access Cloud Pub/Sub using the `pubsub.subscriber` role.

```
gcloud projects add-iam-policy-binding scalesec-dev \
--member serviceAccount:logstash@scalesec-dev.iam.gserviceaccount.com
--role roles/pubsub.subscriber
```

Excerpt of expected response:

```
Updated IAM policy for project [scalesec-dev].
[...]
```

```
- members:
  - serviceAccount:logstash@scalesec-dev.iam.gserviceaccount.com
  role: roles/pubsub.subscriber
[...]
etag: BwWEjM0909E=
version: 1
```

Create a Cloud Pub/Sub topic and subscription

1. Create a Cloud Pub/Sub topic where Stackdriver will send events to be picked up by Logstash:

   ```
   gcloud pubsub topics create stackdriver-topic
   ```

   Expected response:

   ```
   Created topic [projects/scalesec-dev/topics/stackdriver-topic].
   ```

   Next, create a subscription:

   ```
   gcloud pubsub subscriptions create logstash-sub --topic=stackdriver-t
   ```

   Expected response:

   ```
   Created subscription [projects/scalesec-dev/subscriptions/logstash-su
   ```

Create a Stackdriver log sink

1. Create a log sink to be used to export Stackdriver logs to the new Cloud Pub/Sub topic.

```
gcloud logging sinks create logstash-sink pubsub.googleapis.com/proje
--log-filter='resource.type="project"'
```

Expected response:

```
Created [https://logging.googleapis.com/v2/projects/scalesec-dev/sink
Please remember to grant `serviceAccount:p352005273005-058743@gcp-sa-
Publisher role to the topic.
More information about sinks can be found at /logging/docs/export/
```

The filter specified above will produce events associated with changes to IAM, which is a typical area to be monitored closely. Stackdriver supports monitoring activities for vpn_gateway and other resource types. See the documentation (/logging/docs/view/overview) for more filter ideas.

The second part of the output is a reminder to verify that the service account used by Stackdriver has permissions to publish events to the Cloud Pub/Sub topic. The beta version of `gcloud` CLI supports permissions management for Cloud Pub/Sub.

```
gcloud beta pubsub topics add-iam-policy-binding stackdriver-topic \
--member serviceAccount:p352005273005-776084@gcp-sa-logging.iam.gserv
--role roles/pubsub.publisher
```

Expected response:

```
Updated IAM policy for topic [stackdriver-topic].
bindings:
- members:
  - serviceAccount:p352005273005-776084@gcp-sa-logging.iam.gserviceac
  role: roles/pubsub.publisher
etag: BwWEi9uEM1A=
```

Create the Logstash VM

Note: Some system responses are omitted in this section for brevity.

1. Create a VM to run `logstash` to pull logs from the Cloud Pub/Sub logging sink and send them to ElasticSearch:

```
gcloud compute --project=scalesec-dev instances create logstash \
--zone=us-west1-a \
--machine-type=n1-standard-1 \
--subnet=default \
--service-account=logstash@scalesec-dev.iam.gserviceaccount.com \
--scopes="https://www.googleapis.com/auth/cloud-platform" \
--image-family=ubuntu-1804-lts \
--image-project=ubuntu-os-cloud \
--boot-disk-size=10GB \
--boot-disk-type=pd-ssd \
--boot-disk-device-name=logstash
```

Expected response:

```
Created [https://www.googleapis.com/compute/beta/projects/scalesec-de
NAME      ZONE        MACHINE_TYPE   PREEMPTIBLE  INTERNAL_IP  EXTERN
logstash  us-west1-a  n1-standard-1               10.138.0.3   35.233
```

Create Elastic Cloud deployment

1. Go to https://cloud.elastic.co/login. A trial account provides suitable service to complete this tutorial.

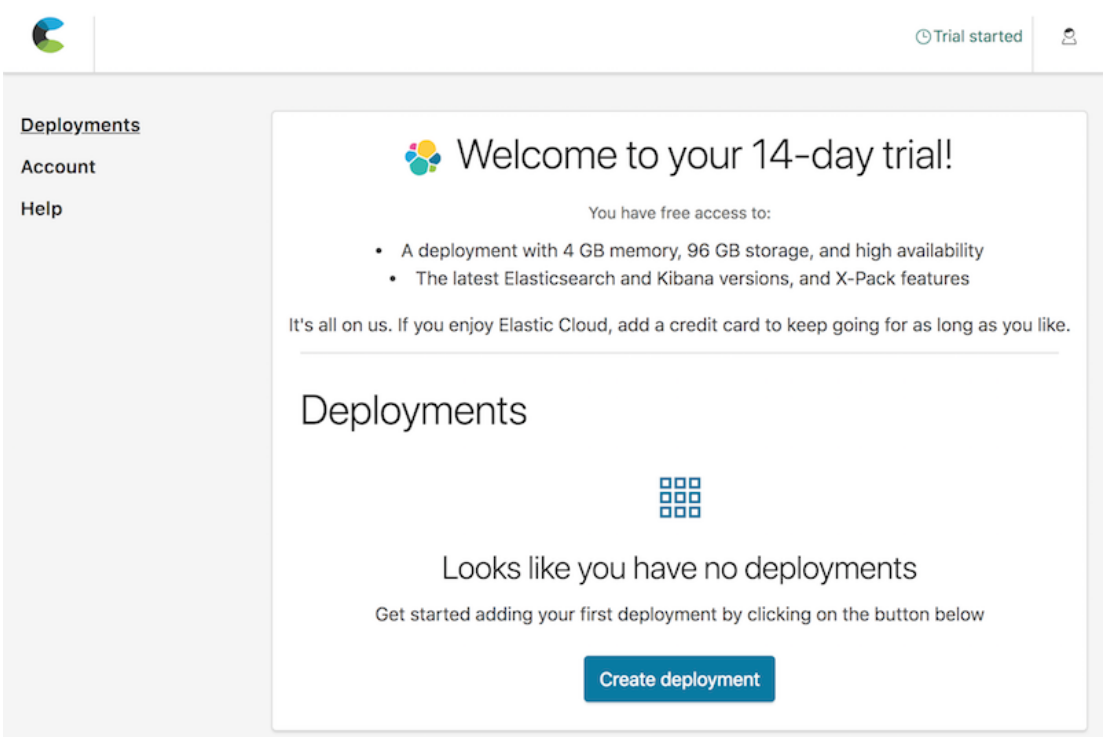

2. Create an Elasticsearch deployment. This example is deployed on GCP in us-west1.

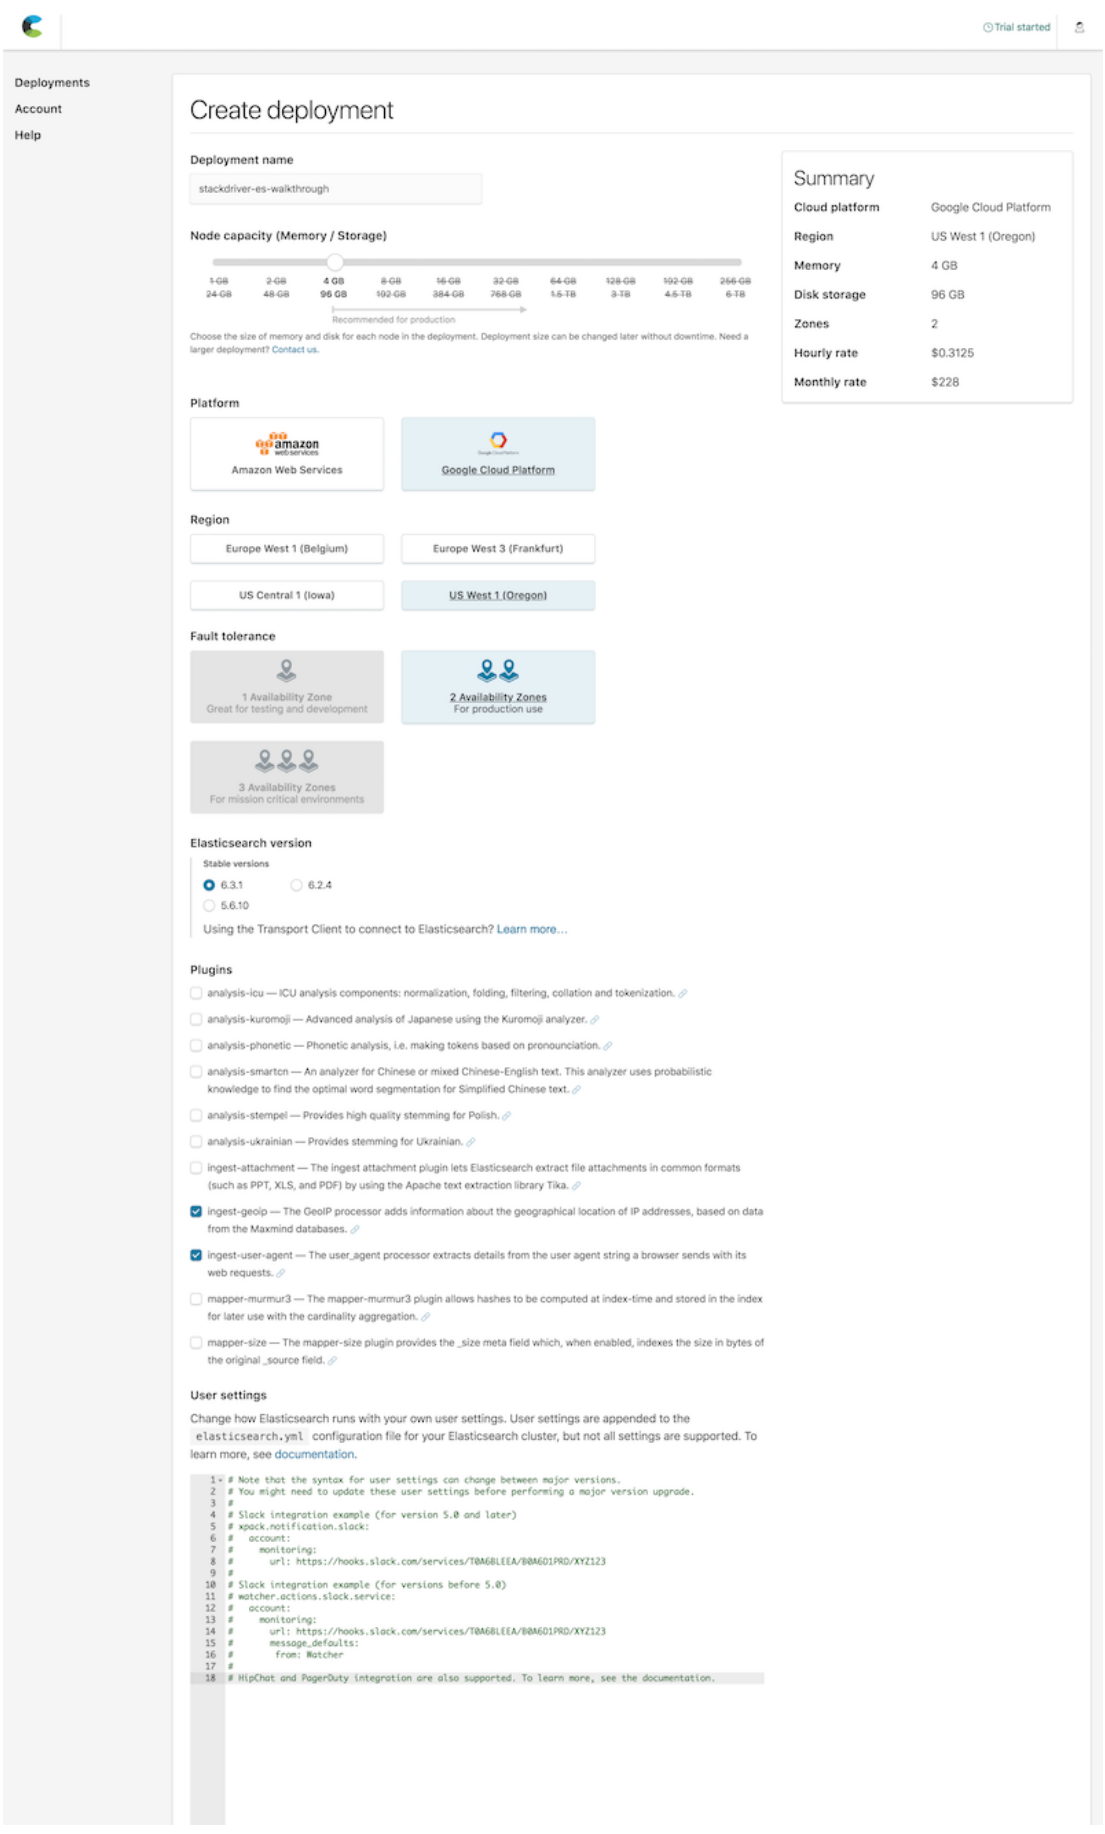
Trial started
Deployments
Account
Help
Create deployment
Deployment name
stackdriver-es-walkthrough
Node capacity (Memory / Storage)
1 GB 24 GB
2 GB 48 GB
4 GB 96 GB
8 GB 192 GB
16 GB 384 GB
32 GB 768 GB
64 GB 1.5 TB
128 GB 3 TB
192 GB 4.5 TB
256 GB 6 TB
Recommended for production
Choose the size of memory and disk for each node in the deployment. Deployment size can be changed later without downtime. Need a larger deployment? Contact us.
Summary
Cloud platform Google Cloud Platform
Region US West 1 (Oregon)
Memory 4 GB
Disk storage 96 GB
Zones 2
Hourly rate $0.3125
Monthly rate $228
Platform
Amazon Web Services
Google Cloud Platform
Region
Europe West 1 (Belgium)
Europe West 3 (Frankfurt)
US Central 1 (Iowa)
US West 1 (Oregon)
Fault tolerance
1 Availability Zone
Great for testing and development
2 Availability Zones
For production use
3 Availability Zones
For mission critical environments
Elasticsearch version
Stable versions
6.3.1
6.2.4
5.6.10
Using the Transport Client to connect to Elasticsearch? Learn more…
Plugins
analysis-icu — ICU analysis components: normalization, folding, filtering, collation and tokenization.
analysis-kuromoji — Advanced analysis of Japanese using the Kuromoji analyzer.
analysis-phonetic — Phonetic analysis, i.e. making tokens based on pronounciation.
analysis-smartcn — An analyzer for Chinese or mixed Chinese-English text. This analyzer uses probabilistic knowledge to find the optimal word segmentation for Simplified Chinese text.
analysis-stempel — Provides high quality stemming for Polish.
analysis-ukrainian — Provides stemming for Ukrainian.
ingest-attachment — The ingest attachment plugin lets Elasticsearch extract file attachments in common formats (such as PPT, XLS, and PDF) by using the Apache text extraction library Tika.
ingest-geoip — The GeoIP processor adds information about the geographical location of IP addresses, based on data from the Maxmind databases.
ingest-user-agent — The user_agent processor extracts details from the user agent string a browser sends with its web requests.
mapper-murmur3 — The mapper-murmur3 plugin allows hashes to be computed at index-time and stored in the index for later use with the cardinality aggregation.
mapper-size — The mapper-size plugin provides the _size meta field which, when enabled, indexes the size in bytes of the original _source field.
User settings
Change how Elasticsearch runs with your own user settings. User settings are appended to the elasticsearch.yml configuration file for your Elasticsearch cluster, but not all settings are supported. To learn more, see documentation.
Note that the syntax for user settings can change between major versions.
You might need to update these user settings before performing a major version upgrade.
#
Slack integration example (for version 5.0 and later)
xpack.notification.slack:
account:
monitoring:
url: https://hooks.slack.com/services/T0A6BLEEA/B0A6D1PRD/XYZ123
#
Slack integration example (for versions before 5.0)
watcher.actions.slack.service:
account:
monitoring:
url: https://hooks.slack.com/services/T0A6BLEEA/B0A6D1PRD/XYZ123
message_defaults:
from: Watcher
#
HipChat and PagerDuty integration are also supported. To learn more, see the documentation.

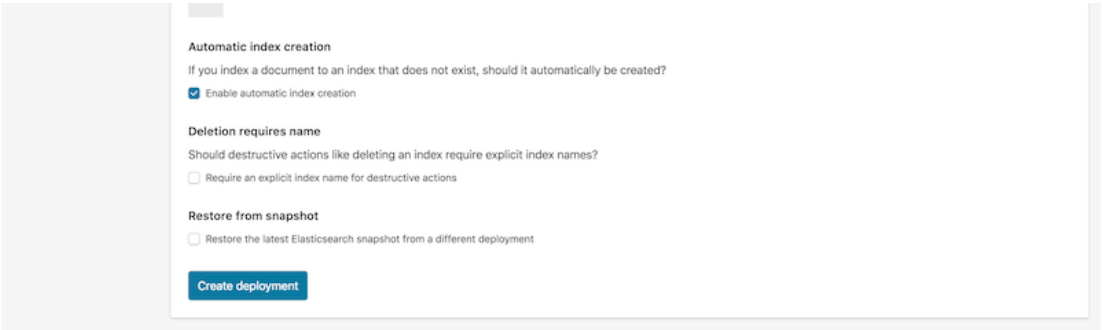

3. While the deployment is finishing up, make sure to capture the credentials and store them in a safe place. While the Cloud ID can be viewed from the deployment page, this is the only time the password for the elastic user is available. Visit the **Security** page to reset the password if needed. When considering production environments, create new Elasticsearch credentials with tighter permissions and avoid using the `elastic` user. As documented (https://www.elastic.co/guide/en/cloud/master/ec-cloud-id.html): "On a production system, you should adapt these examples by creating a user that can write to and access only the minimally required indices."

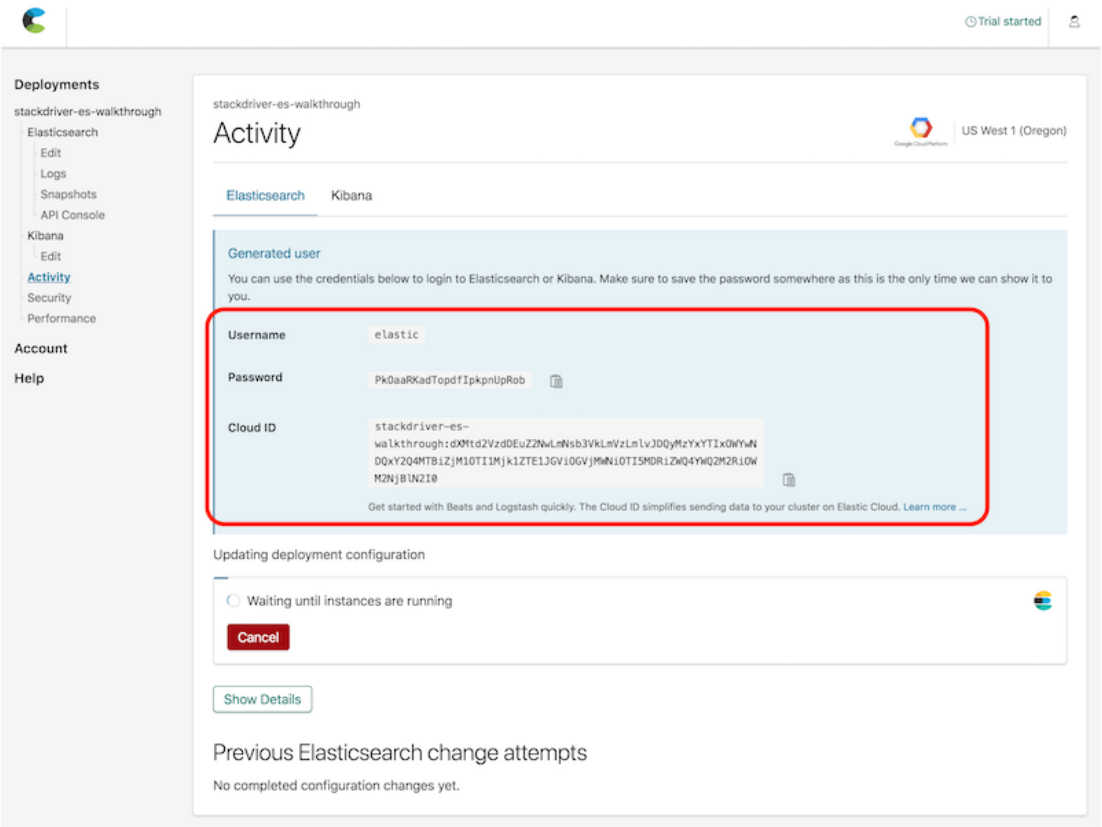

4. Obtain the URI of the Elasticsearch endpoint that has been provisioned. A link to this endpoint can be copied from the **Deployments** page. This value will be needed to configure Logstash output plugin configuration.

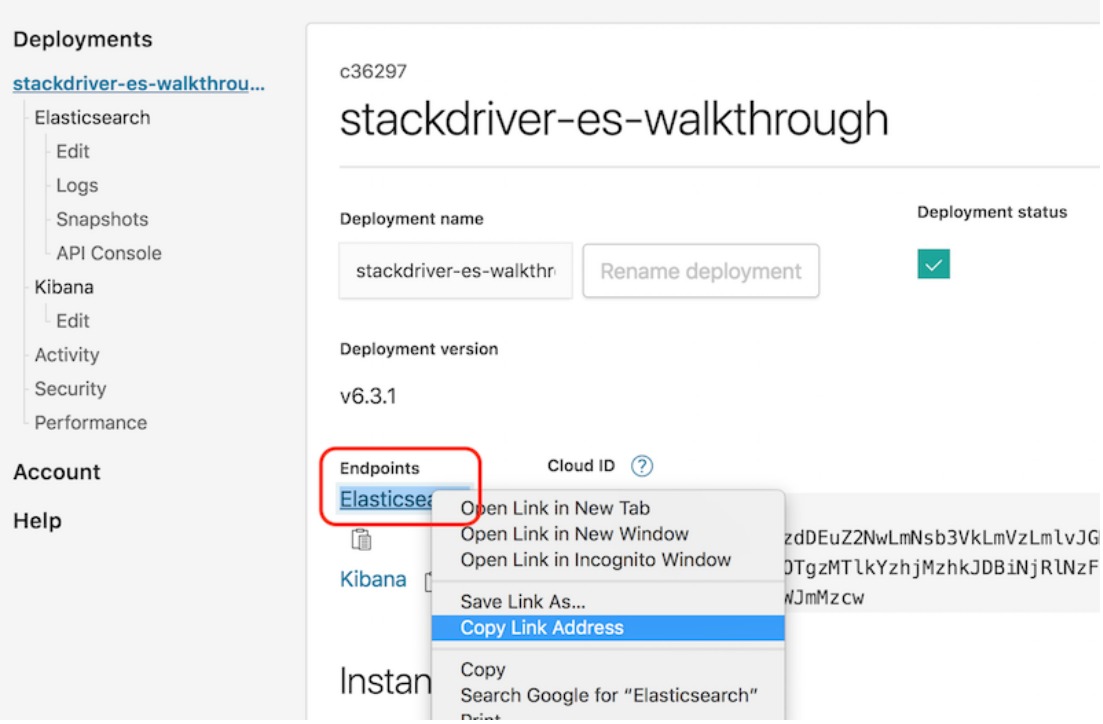

The next section provides steps to complete the setup to send events to the new Elasticsearch deployment.

Configure the Logstash VM

1. Compute Engine supports several ways (/compute/docs/instances/connecting-to-instance) to access your VM. You can use the `gcloud` command in Cloud Shell to leverage `oslogin` to connect to the `logstash` VM via SSH, noting the zone from the VM creation step above.

```
gcloud compute ssh logstash --zone us-west1-a
```

2. Perform typical system updates and install OpenJDK:

```
sudo apt-get update
sudo apt-get -y upgrade
sudo apt -y install openjdk-8-jre-headless
echo "export JAVA_HOME=\"/usr/lib/jvm/java-8-openjdk-amd64\"" >> ~/.p
sudo reboot
```

After a few moments, the VM will complete its reboot and can be accessed again via `gcloud`.

```
gcloud compute ssh logstash --zone us-west1-a
```

Install Logstash

1. Install logstash from Elastic.

```
wget -qO - https://artifacts.elastic.co/GPG-KEY-elasticsearch | sudo
echo "deb https://artifacts.elastic.co/packages/6.x/apt stable main"
sudo apt-get update
sudo apt-get install logstash
```

2. Install the Logstash Plugin for Cloud Pub/Sub.

```
cd /usr/share/logstash
sudo -u root sudo -u logstash bin/logstash-plugin install logstash-ir
```

Expected response:

```
Validating logstash-input-google_pubsub
Installing logstash-input-google_pubsub
Installation successful
```

Configure Logstash

Logstash comes with no default configuration.

1. Create a new file `/etc/logstash/conf.d/logstash.conf` with these contents, modifying values as needed:

```
input
{
    google_pubsub {
        project_id => "scalesec-dev"
        topic => "stackdriver-topic"
        subscription => "logstash-sub"
        include_metadata => true
        codec => "json"
    }
    # optional, but helpful to generate the ES index and test the plu
    heartbeat {
        interval => 10
        type => "heartbeat"
    }
}
filter {
    # don't modify logstash heartbeat events
    if [type] != "heartbeat" {
        mutate {
            add_field => { "messageId" => "%{[@metadata][pubsub_messa
        }
    }
}
output
{
    stdout { codec => rubydebug }
    elasticsearch
    {
        hosts => ["https://c36297ebbc024cd4b29c98319dc8c38d.us-west1.
        user => "elastic"
        password => "NTmWdNJXkzMWL4kkIcIzY8O6"
        index => "logstash-%{+YYYY.MM.dd}"
    }
}
```

Start Logstash

1. Start Logstash:

```
sudo service logstash start
```

2. Monitor the startup logs closely for issues:

```
sudo tail -f /var/log/syslog
```

3. Review log messages. It may take a few moments for events to begin flowing.

   Log messages like these indicate that Logstash is working internally:

```
Jul 15 20:43:09 logstash logstash[2537]: {
Jul 15 20:43:09 logstash logstash[2537]:          "type" => "heartbe
Jul 15 20:43:09 logstash logstash[2537]:     "messageId" => "%{[@met
Jul 15 20:43:09 logstash logstash[2537]:       "message" => "ok",
Jul 15 20:43:09 logstash logstash[2537]:    "@timestamp" => 2018-07-
Jul 15 20:43:09 logstash logstash[2537]:      "@version" => "1",
Jul 15 20:43:09 logstash logstash[2537]:          "host" => "logstas
Jul 15 20:43:09 logstash logstash[2537]: }
```

   Log messages like these indicate that Logstash is pulling events from Cloud Pub/Sub. Actual message content will differ.

```
Jul 17 20:58:13 logstash logstash[15198]:             "logName" => "
Jul 17 20:58:13 logstash logstash[15198]:          "resource" => {
Jul 17 20:58:13 logstash logstash[15198]:        "labels" => {
Jul 17 20:58:13 logstash logstash[15198]:            "project_id" =>
Jul 17 20:58:13 logstash logstash[15198]:               "region" =>
Jul 17 20:58:13 logstash logstash[15198]:            "gateway_id" =>
Jul 17 20:58:13 logstash logstash[15198]:        },
Jul 17 20:58:13 logstash logstash[15198]:          "type" => "vpn_ga
Jul 17 20:58:13 logstash logstash[15198]:    },
Jul 17 20:58:13 logstash logstash[15198]:            "severity" => "
Jul 17 20:58:13 logstash logstash[15198]:          "@timestamp" => 2
Jul 17 20:58:13 logstash logstash[15198]:         "textPayload" => "
Jul 17 20:58:13 logstash logstash[15198]:            "insertId" => "
Jul 17 20:58:13 logstash logstash[15198]:          "timestamp" => "
Jul 17 20:58:13 logstash logstash[15198]:            "@version" => "
```

```
Jul 17 20:58:13 logstash logstash[15198]:              "labels" => {
Jul 17 20:58:13 logstash logstash[15198]:          "tunnel_id" => "109
Jul 17 20:58:13 logstash logstash[15198]:      },
Jul 17 20:58:13 logstash logstash[15198]:           "messageId" => "
Jul 17 20:58:13 logstash logstash[15198]:      "receiveTimestamp" => "
```

Configure Kibana

Kibana is a powerful graphical user interface that uses the underlying Elasticsearch data. This is the main console to monitor and triage security events and perform searches and investigations.

1. Return to the Elasticsearch deployment page and click the link to Kibana.

2. Log in as the **`elastic`** user.

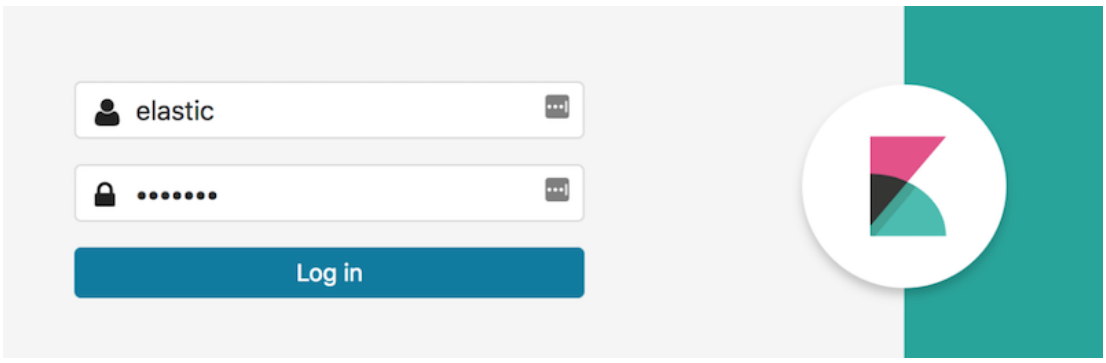

3. Navigate to the **Management** page to set up index patterns for Kibana.

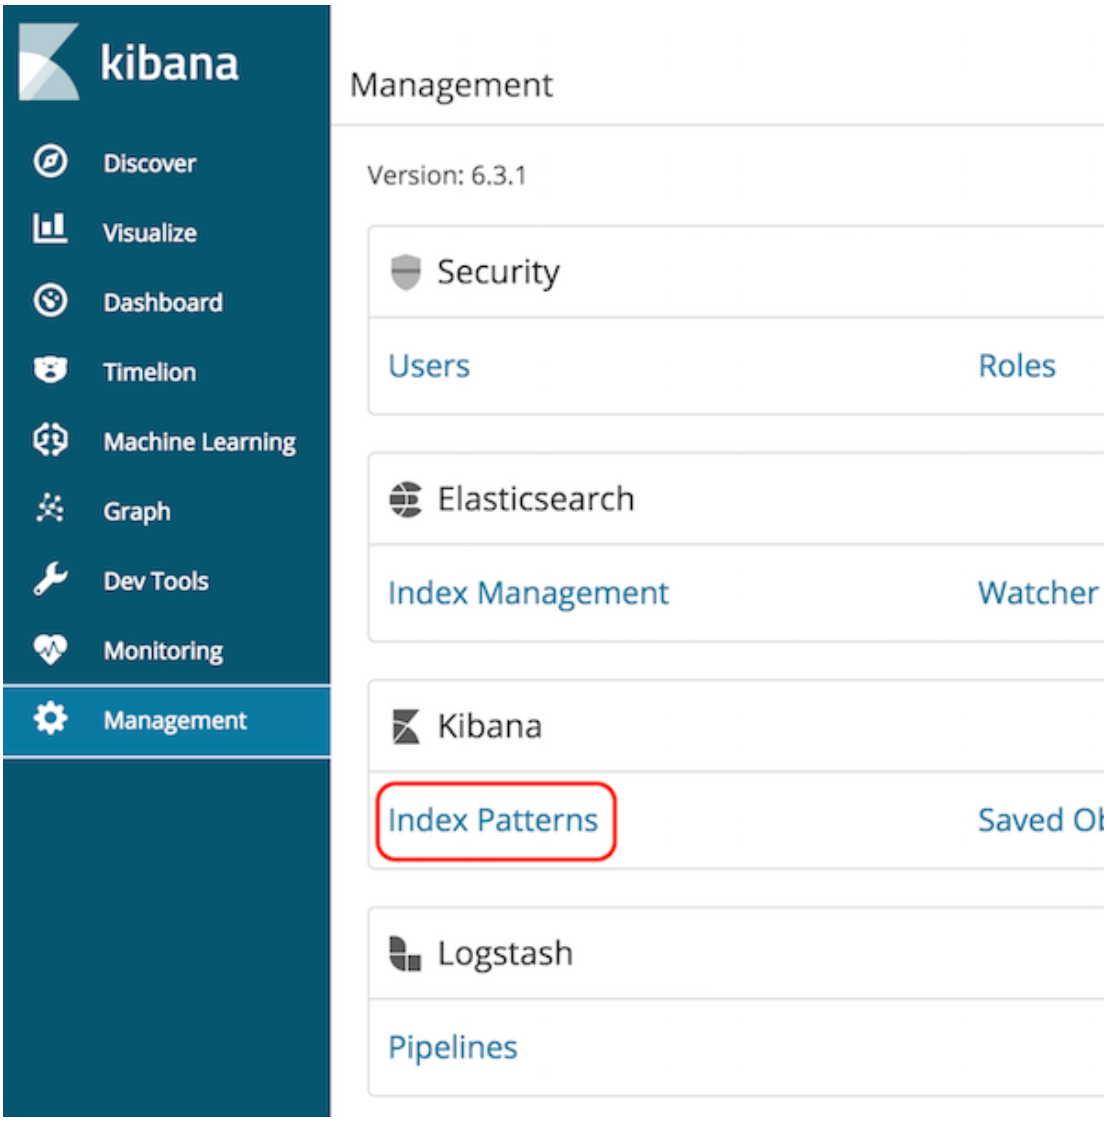

4. Enter `logstash-*` for the index pattern.

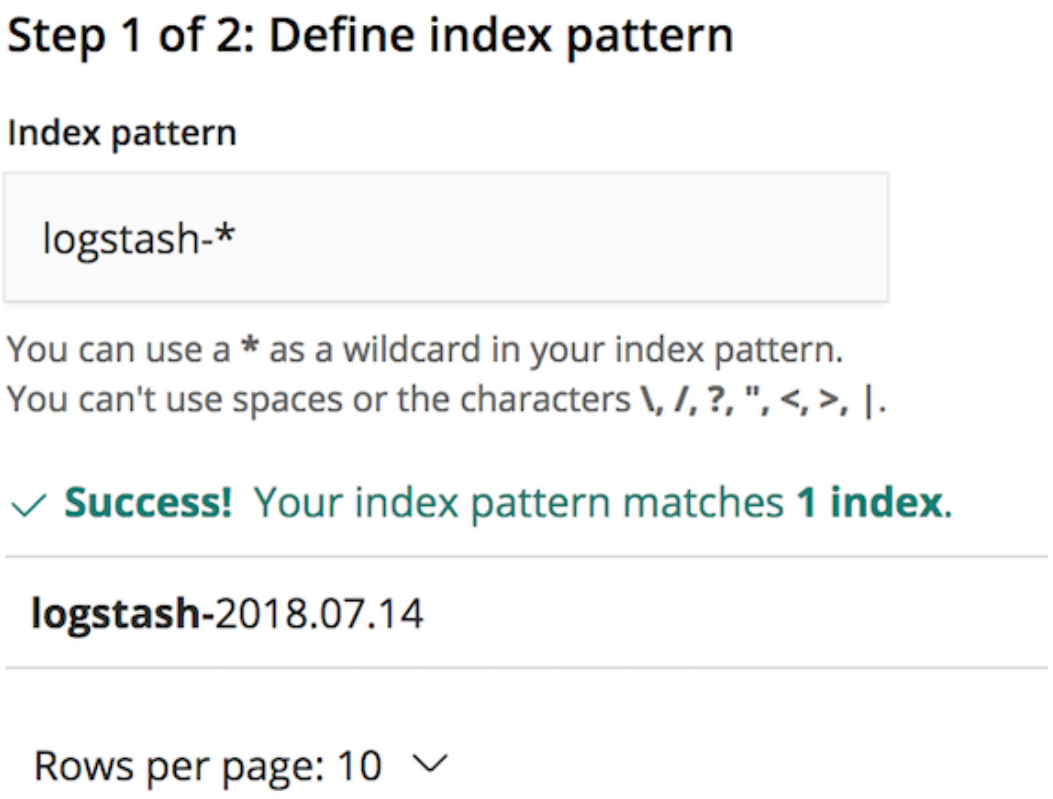

5. Use `@timestamp` for the time field.

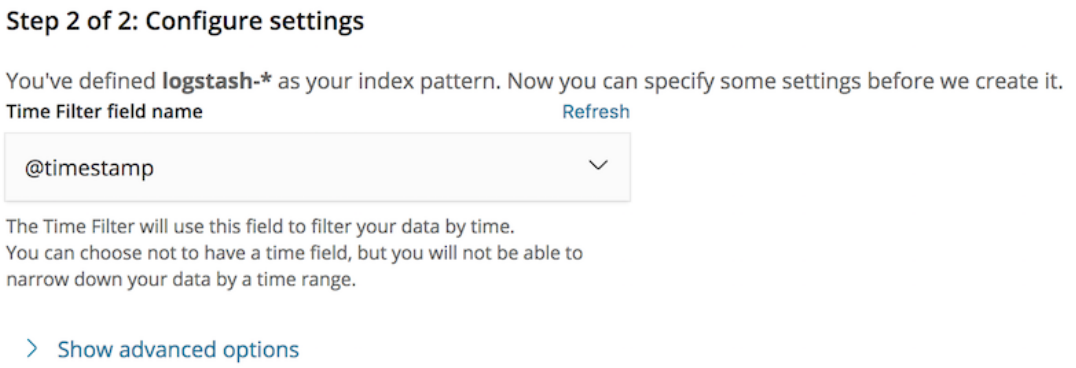

Verify log flow

Return to the main Kibana dashboard (shown as **Discover** in the navigation menu). The Kibana dashboard should display Stackdriver events similar to those shown below:

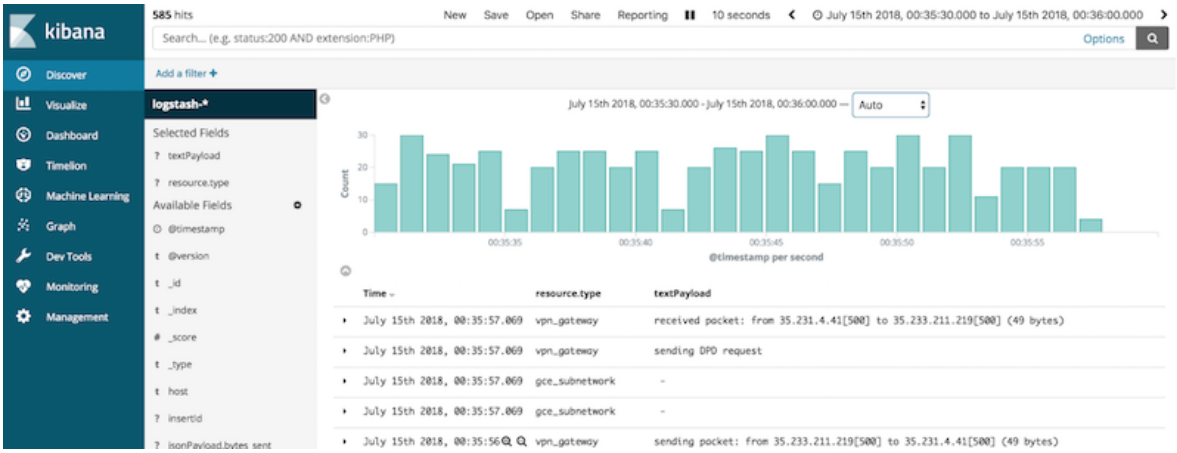

Submit a Tutorial

Share step-by-step guides

SUBMIT A TUTORIAL (/COMMUNITY/TUTORIALS/WRITE)

Request a Tutorial

Ask for community help

SUBMIT A REQUEST (HTTPS://GITHUB.COM/GOOGLECLOUDPLATFORM/COMMUNITY/ISSUES?Q=IS%3AOPEN+IS%3AISSUE+LABEL%3A%22TUTORIAL+REQUEST%22)

GCP Tutorials

Tutorials published by GCP

VIEW TUTORIALS (/DOCS/TUTORIALS)

(/community) (/docs/tutorials)