

This tutorial describes how to customize [Fluentd](https://www.fluentd.org/) logging for a [Google Kubernetes Engine](/kubernetes-engine/) cluster. You'll learn how to host your own configurable Fluentd daemonset to send logs to Stackdriver, instead of selecting the cloud logging option when creating the Google Kubernetes Engine (GKE) cluster, which does not allow configuration of the Fluentd daemon.

- Deploy your own Fluentd daemonset on a Google Kubernetes Engine cluster, configured to log data to [Stackdriver](/stackdriver/). We assume that you are already familiar with [Kubernetes](https://kubernetes.io/docs/home/).
- Customize GKE logging to remove sensitive data from the Stackdriver logs.
- Customize GKE logging to add node-level events to to the Stackdriver logs

This tutorial uses billable components of Cloud Platform, including:

- A three-node [Google Kubernetes Engine](/kubernetes-engine/) cluster.

The [Pricing Calculator](/products/calculator/#id=38ec76f1-971f-41b5-8aec-a04e732129cc) estimates the cost of this environment at around \$1.14 for 8 hours.

1. [Sign in](https://accounts.google.com/Login) to your Google Account.

If you don't already have one, [sign up for a new account](https://accounts.google.com/SignUp).

2. In the Cloud Console, on the project selector page, select or create a Cloud project.

★ **Note:** If you don't plan to keep the resources that you create in this procedure, create a project instead of selecting an existing project. After you finish these steps, you can delete the project, removing all resources associated with the project.

[Go to the project selector page](https://console.cloud.google.com/projectselector2/home/dashboard) (https://console.cloud.google.com/projectselector2/home/dashboard)

3. Make sure that billing is enabled for your Google Cloud project. [Learn how to confirm billing is enabled for your project](#) (/billing/docs/how-to/modify-project).
4. Enable the Google Kubernetes Engine, Compute Engine APIs.

[Enable the APIs](https://console.cloud.google.com/flows/enableapi?apiid=container,compute.googleapis.com) (https://console.cloud.google.com/flows/enableapi?apiid=container,compute.googleapis.com)

You must define several variables that control where elements of the infrastructure are deployed.

1. Using a text editor, edit the following script, substituting your project ID for [YOUR\_PROJECT\_ID]. The script sets the region to us-east-1. If you make any changes to the script, make sure that the zone values reference the region you specify.

2. Go to Cloud Shell.

[Open Cloud Shell](https://console.cloud.google.com/?cloudshell=true) (https://console.cloud.google.com/?cloudshell=true)

3. Copy the script into your Cloud Shell window and run it.
4. Run the following commands to set the default zone and project ID so you don't have to specify these values in every subsequent command:

Unless otherwise noted, you enter all the commands for this tutorial at the command line of your computer or in Cloud Shell.

1. Clone the sample repository. The sample repository includes the Kubernetes manifests for the Fluentd daemonset and a test logging program that you will deploy:
2. Change your working directory to the cloned repository:
3. Create the GKE cluster without cloud logging turned on:

By default, the sample application that you deploy continuously emits random logging statements. The Docker container it uses is available at `gcr.io/cloud-solutions-images/test-logger`, and its source code is included in the `test-logger` subdirectory.

1. Deploy the `test-logger` application to the GKE cluster:
2. View the status of the `test-logger` pods:
3. Repeat this command until the output looks like the following, with all three `test-logger` pods running:

NAME	READY	STATUS	RESTARTS	AGE
test-logger-1704239063-cc01k	1/1	Running	0	55s
test-logger-1704239063-j2tjd	1/1	Running	0	55s
test-logger-1704239063-svtl3	1/1	Running	0	55s

Next you will configure and deploy the Fluentd daemonset.

**Note:** The Kubernetes manifests for Fluentd that you deploy in this procedure are modified versions of the ones available from the Kubernetes site for [logging using Stackdriver](https://kubernetes.io/docs/tasks/debug-application-cluster/logging-stackdriver/) (<https://kubernetes.io/docs/tasks/debug-application-cluster/logging-stackdriver/>) and for [watching changes to Docker log files](https://kubernetes.io/docs/concepts/cluster-administration/logging/) (<https://kubernetes.io/docs/concepts/cluster-administration/logging/>).

1. Deploy the Fluentd configuration:

2. Deploy the Fluentd daemonset:

3. Check that the Fluentd pods have started:

If they're running, you see output like the following:

NAME	READY	STATUS	RESTARTS	AGE
fluentd-gcp-v3.2.0-1r69v	2/2	Running	0	12s
fluentd-gcp-v3.2.0-9b6x7	2/2	Running	0	12s
fluentd-gcp-v3.2.0-smgtn	2/2	Running	0	12s

4. Verify that you're seeing logs in Stackdriver. In the console, in the left-hand menu click **Stackdriver** > **Logging** > **Logs** and select **Kubernetes Container** in the list.

The screenshot shows the Stackdriver Logging console. The left sidebar has a menu with 'Logs' selected. The main area displays a list of logs for a 'Kubernetes Container'. The logs are filtered to show the last hour ending at 9:27 AM (CEST). The log entries are as follows:

Timestamp	Log Level	Message
2019-04-30 09:27:15.951 CEST	Info	Something happened..with social 222334444
2019-04-30 09:27:16.406 CEST	Info	Users email is joker@batman.com
2019-04-30 09:27:16.632 CEST	Info	Users email is joker@batman.com
2019-04-30 09:27:17.951 CEST	Info	error happened with social security number 111-22-3333
2019-04-30 09:27:18.407 CEST	Info	Something happened..with social 222334444
2019-04-30 09:27:18.633 CEST	Info	Users email is joker@batman.com
2019-04-30 09:27:19.951 CEST	Info	Processing credit card 1234 5678 9012 3456
2019-04-30 09:27:20.407 CEST	Info	Processing credit card 1234 5678 9012 3456
2019-04-30 09:27:20.633 CEST	Info	Processing credit card 1234 5678 9012 3456
2019-04-30 09:27:21.951 CEST	Info	Users email is joker@batman.com
2019-04-30 09:27:22.407 CEST	Info	error happened with social security number 111-22-3333
2019-04-30 09:27:22.633 CEST	Info	Something happened..with social 222334444
2019-04-30 09:27:23.952 CEST	Info	Something happened..with social 222334444
2019-04-30 09:27:24.407 CEST	Info	Processing credit card 1234 5678 9012 3456
2019-04-30 09:27:24.634 CEST	Info	Something happened..with social 222334444
2019-04-30 09:27:25.952 CEST	Info	Processing credit card 1234 5678 9012 3456
2019-04-30 09:27:26.408 CEST	Info	Processing credit card 1234 5678 9012 3456
2019-04-30 09:27:26.634 CEST	Info	error happened with social security number 111-22-3333
2019-04-30 09:27:27.952 CEST	Info	Users email is joker@batman.com

The next step is to specify that Fluentd should filter certain data so that it is not logged. For this tutorial, you filter out the Social Security numbers, credit card numbers, and email addresses. To make this update, you change the daemonset to use a different ConfigMap that contains these filters. You use Kubernetes rolling updates feature and preserve the old version of the ConfigMap.

1. Open the `kubernetes/fluentd-configmap.yaml` file in an editor.
2. Uncomment the lines between and not including the lines `### sample log scrubbing filters` and `### end sample log scrubbing filters`:

```
kubernetes/fluentd-configmap.yaml
(https://github.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-configmap.yaml)
```

`.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-configmap.yaml)`

3. Change the name of the ConfigMap from `fluentd-gcp-config` to `fluentd-gcp-config-filtered` by editing the `metadata.name` field:

[kubernetes/fluentd-configmap.yaml](https://github.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-configmap.yaml)

(<https://github.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-configmap.yaml>)

.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-configmap.yaml)

4. Save and close the file.

Now you change `kubernetes/fluentd-daemonset.yaml` to mount the ConfigMap `fluentd-gcp-config-filtered` instead of `fluentd-gcp-config`.

1. Open the `kubernetes/fluentd-daemonset.yaml` file in an editor.
2. Change the name of the ConfigMap from `fluentd-gcp-config` to `fluentd-gcp-config-filtered` by editing the `configMap.name` field:

```
kubernetes/fluentd-daemonset.yaml  
(https://github.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-daemonset.yaml)
```

[com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-daemonset.yaml](https://github.com/GoogleCloudPlatform/kubernetes-engine-customize-fluentd/blob/master/kubernetes/fluentd-daemonset.yaml))

3. Deploy the new version of the ConfigMap to your cluster:

4. Roll out the new version of the daemonset:

5. Roll out the update and wait for it to complete:

```
Waiting for rollout to finish: 0 out of 3 new pods have been updated...  
Waiting for rollout to finish: 1 out of 3 new pods have been updated...  
Waiting for rollout to finish: 1 out of 3 new pods have been updated...  
Waiting for rollout to finish: 2 out of 3 new pods have been updated...  
Waiting for rollout to finish: 2 out of 3 new pods have been updated...  
Waiting for rollout to finish: 2 of 3 updated pods are available...  
daemon set "fluentd-gcp-v3.2.0" successfully rolled out
```

6. When the rollout is complete, refresh the Stackdriver logs and make sure that the Social Security number, credit card number, and email address data has been filtered out.

The screenshot shows the Stackdriver Logging interface. The left sidebar contains navigation options: Logs, Logs-based metrics, Exports, and Logs ingestion. The main panel displays a list of logs for a 'Kubernetes Container' from 'All logs' at 'Any log level' for the 'Last hour'. The logs show various events, including user logins and errors, with sensitive information like email addresses and social security numbers redacted with 'xxxx'.

Timestamp	Log Level	Message
2019-04-30 09:33:02.687 CEST	Info	Users email is user@email.tld
2019-04-30 09:33:03.985 CEST	Info	Users email is user@email.tld
2019-04-30 09:33:04.454 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:04.687 CEST	Info	Users email is user@email.tld
2019-04-30 09:33:05.985 CEST	Info	Users email is user@email.tld
2019-04-30 09:33:06.455 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:06.687 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:07.986 CEST	Info	Something happened..with social xxx-xx-xxxx
2019-04-30 09:33:08.455 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:08.688 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:09.986 CEST	Info	Something happened..with social xxx-xx-xxxx
2019-04-30 09:33:10.455 CEST	Info	Something happened..with social xxx-xx-xxxx
2019-04-30 09:33:10.688 CEST	Info	Processing credit card xxxx xxxx xxxx xxxx
2019-04-30 09:33:11.986 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:12.455 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:12.688 CEST	Info	Users email is user@email.tld
2019-04-30 09:33:13.986 CEST	Info	Users email is user@email.tld
2019-04-30 09:33:14.456 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:14.688 CEST	Info	Processing credit card xxxx xxxx xxxx xxxx
2019-04-30 09:33:15.986 CEST	Error	error happened with social security number xxx-xx-xxxx
2019-04-30 09:33:16.456 CEST	Info	Users email is user@email.tld

If you want events that happen on your GKE nodes to show up in Stackdriver as well, add the following lines to your ConfigMap and follow the instructions described in the last section:



After you've finished the tutorial, you can clean up the resources you created on GCP so you won't be billed for them in the future.

The easiest way to eliminate billing is to delete the project that you created for the tutorial.

To delete the project:


**!** **Caution:** Deleting a project has the following effects:

- **Everything in the project is deleted.** If you used an existing project for this tutorial, when you delete it, you also delete any other work you've done in the project.
- **Custom project IDs are lost.** When you created this project, you might have created a custom project ID that you want to use in the future. To preserve the URLs that use the project ID, such as an **appspot.com** URL, delete selected resources inside the project instead of deleting the whole project.

If you plan to explore multiple tutorials and quickstarts, reusing projects can help you avoid exceeding project quota limits.

1. In the Cloud Console, go to the **Manage resources** page.

[Go to the Manage resources page \(https://console.cloud.google.com/iam-admin/projects\)](https://console.cloud.google.com/iam-admin/projects)

2. In the project list, select the project you want to delete and click **Delete** .

3. In the dialog, type the project ID, and then click **Shut down** to delete the project.

If you don't want to delete the whole project, run the following command to delete the GKE cluster:

- Review [Fluentd](https://docs.fluentd.org/v1.0/articles/quickstart) (https://docs.fluentd.org/v1.0/articles/quickstart) documentation in more detail.
- Review [Google Kubernetes Engine](/kubernetes-engine/) (/kubernetes-engine/) documentation in more detail.
- Try out other Google Cloud features for yourself. Have a look at our [tutorials](/docs/tutorials) (/docs/tutorials).