

# PROFESSIONAL-CLOUD-DEVOPS- ENGINEER<sup>Q&As</sup>

Professional Cloud DevOps Engineer

**Pass Google PROFESSIONAL-CLOUD-DEVOPS-  
ENGINEER Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/professional-cloud-devops-engineer.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Google  
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



**QUESTION 1**

Your application images are built and pushed to Google Container Registry (GCR). You want to build an automated pipeline that deploys the application when the image is updated while minimizing the development effort. What should you do?

- A. Use Cloud Build to trigger a Spinnaker pipeline.
- B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.
- C. Use a custom builder in Cloud Build to trigger Jenkins pipeline.
- D. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).

Correct Answer: B

[https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud#triggering\\_a\\_spinnaker\\_pipeline\\_when\\_a\\_docker\\_image\\_is\\_pushed\\_to\\_container\\_registry](https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud#triggering_a_spinnaker_pipeline_when_a_docker_image_is_pushed_to_container_registry)

---

**QUESTION 2**

You are running an experiment to see whether your users like a new feature of a web application. Shortly after deploying the feature as a canary release, you receive a spike in the number of 500 errors sent to users, and your monitoring reports show increased latency. You want to quickly minimize the negative impact on users. What should you do first?

- A. Roll back the experimental canary release.
- B. Start monitoring latency, traffic, errors, and saturation.
- C. Record data for the postmortem document of the incident.
- D. Trace the origin of 500 errors and the root cause of increased latency.

Correct Answer: A

<https://www.examttopics.com/discussions/google/1/>

---

**QUESTION 3**

You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need measure application reliability from a user perspective without making any engineering changes to it. What should you do? (Choose two.)

- A. Review current application metrics and add new ones as needed.
- B. Modify the code to capture additional information for user interaction.
- C. Analyze the web proxy logs only and capture response time of each request.
- D. Create new synthetic clients to simulate a user journey using the application.

E. Use current and historic Request Logs to trace customer interaction with the application.

Correct Answer: DE

[https://cloud.google.com/architecture/adopting-slos#choosing\\_a\\_measurement\\_method](https://cloud.google.com/architecture/adopting-slos#choosing_a_measurement_method)

---

## QUESTION 4

You currently store the virtual machine (VM) utilization logs in Stackdriver. You need to provide an easy-to-share interactive VM utilization dashboard that is updated in real time and contains information aggregated on a quarterly basis. You want to use Google Cloud Platform solutions. What should you do?

A. 1. Export VM utilization logs from Stackdriver to BigQuery.

2.

Create a dashboard in Data Studio.

3.

Share the dashboard with your stakeholders.

B. 1. Export VM utilization logs from Stackdriver to Cloud Pub/Sub.

2.

From Cloud Pub/Sub, send the logs to a Security Information and Event Management (SIEM) system.

3.

Build the dashboards in the SIEM system and share with your stakeholders.

C. 1. Export VM utilization logs from Stackdriver to BigQuery.

2.

From BigQuery, export the logs to a CSV file.

3.

Import the CSV file into Google Sheets.

4.

Build a dashboard in Google Sheets and share it with your stakeholders.

D. 1. Export VM utilization logs from Stackdriver to a Cloud Storage bucket.

2.

Enable the Cloud Storage API to pull the logs programmatically.

3.

Build a custom data visualization application.

4.

Display the pulled logs in a custom dashboard.

Correct Answer: A

---

## QUESTION 5

You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
- C. Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- D. Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

Correct Answer: A

"The filter\_record\_transformer filter plugin mutates/transforms incoming event streams in a versatile manner. If there is a need to add/delete/modify events, this plugin is the first filter to try. It is included in the Fluentd's core."

[https://docs.fluentd.org/filter/record\\_transformer](https://docs.fluentd.org/filter/record_transformer)

[Latest PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Dumps](#)

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Practice Test](#)

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Study Guide](#)