

# Congratulation!! You have selected the best success partner to become



## Exam overview

**Exam Name:** Professional Cloud DevOps Engineer

**Length:** 2 hours

**Registration fee:** \$200 (plus tax where applicable)

**Exam format:** Multiple choice and multiple select

**Exam delivery method:**

- a. Take the [online-proctored exam](#) from a remote location
- b. Take the onsite-proctored exam at a [testing center](#)

<https://cloud.google.com/certification/cloud-devops-engineer>

---

**QUESTION NO: 1**

You support a Node.js application running on Google Kubernetes Engine (GKE) in production. The application makes several HTTP requests to dependent applications. You want to anticipate which dependent applications might cause performance issues. What should you do?

- A.**  
Instrument all applications with Cloud Profiler.
- B.**  
Instrument all applications with Cloud Trace and review inter-service HTTP requests.
- C.**  
Use Cloud Debugger to review the execution of logic within each application to instrument all applications.
- D.**  
Modify the Node.js application to log HTTP request and response times to dependent applications. Use Cloud Logging to find dependent applications that are performing poorly.

**Answer: B**

**Explanation:**

**QUESTION NO: 2**

You created a Cloud Monitoring chart for CPU utilization in a dashboard within your workspace project. You want to share the chart with your Site Reliability Engineering (SRE) team only. You want to ensure you follow the principle of least privilege. What should you do?

- A.**  
Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- B.**  
Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
- C.**  
Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- D.**  
Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the

---

Dashboard Viewer IAM role in the workspace project.

**Answer: A**

**Explanation:**

### QUESTION NO: 3

Your organization wants to implement Site Reliability Engineering (SRE) culture and principles. Recently, a service that you support had a limited outage. A manager on another team asks you to provide a formal explanation of what happened so they can action remediations. What should you do?

**A.**

Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it with the manager only.

**B.**

Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it on the engineering organization's document portal.

**C.**

Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it with the manager only.

**D.**

Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it on the engineering organization's document portal.

**Answer: B**

**Explanation:**

### QUESTION NO: 4

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Cloud Operations for GKE. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to `/var/log/app_messages.log`, and you want to send these log entries to Cloud Logging. What should you do?

---

**A.**

Use the default Cloud Operations for GKE agent configuration.

**B.**

Deploy a Fluentd daemonset to GKE. Then create a customized input and output configuration to tail the log file in the application's pods and write to Cloud Logging.

**C.**

Install Kubernetes on Google Compute Engine (GCE) and redeploy your applications. Then customize the built-in Cloud Logging configuration to tail the log file in the application's pods and write to Cloud Logging.

**D.**

Write a script to tail the log file within the pod and write entries to standard output. Run the script as a sidecar container with the application's pod. Configure a shared volume between the containers to allow the script to have read access to /var/log in the application container.

**Answer: B**

Reference: <https://cloud.google.com/solutions/customizing-stackdriver-logs-fluentd>

## QUESTION NO: 5

You are running an application in a virtual machine (VM) using a custom Debian image. The image has the Cloud Logging agent installed. The VM has the cloud-platform scope. The application is logging information via syslog. You want to use Cloud Logging in the Google Cloud Platform Console to visualize the logs. You notice that syslog is not showing up in the "All logs" dropdown list of the Logs Viewer. What is the first thing you should do?

**A.**

Look for the agent's test log entry in the Logs Viewer.

**B.**

Install the most recent version of the Cloud Logging agent.

**C.**

Verify the VM service account access scope includes the monitoring.write scope.

**D.**

SSH to the VM and execute the following commands on your VM: `ps ax | grep fluentd`.

**Answer: D**

Reference: <https://groups.google.com/g/google-stackdriver-discussion/c/FXehB9a-5Vk?pli=1>

---

**QUESTION NO: 6**

You use a multiple step Cloud Build pipeline to build and deploy your application to Google Kubernetes Engine (GKE). You want to integrate with a third-party monitoring platform by performing a HTTP POST of the build information to a webhook. You want to minimize the development effort. What should you do?

- A.**  
Add logic to each Cloud Build step to HTTP POST the build information to a webhook.
- B.**  
Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
- C.**  
Use Cloud Logging to create a logs-based metric from the Cloud Build logs. Create an Alert with a Webhook notification type.
- D.**  
Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

**Answer: D**

**Explanation:**

**QUESTION NO: 7**

You use Spinnaker to deploy your application and have created a canary deployment stage in the pipeline. Your application has an in-memory cache that loads objects at start time. You want to automate the comparison of the canary version against the production version. How should you configure the canary analysis?

- A.**  
Compare the canary with a new deployment of the current production version.
- B.**  
Compare the canary with a new deployment of the previous production version.
- C.**  
Compare the canary with the existing deployment of the current production version.

---

**D.**

Compare the canary with the average performance of a sliding window of previous production versions.

**Answer: D**

Reference: <https://cloud.google.com/solutions/automated-canary-analysis-kubernetes-engine-spinnaker>

### QUESTION NO: 8

You support a high-traffic web application and want to ensure that the home page loads in a timely manner. As a first step, you decide to implement a Service Level Indicator (SLI) to represent home page request latency with an acceptable page load time set to 100 ms. What is the Google-recommended way of calculating this SLI?

**A.**

Bucketize the request latencies into ranges, and then compute the percentile at 100 ms.

**B.**

Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.

**C.**

Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.

**D.**

Count the number of home page request that load in under 100 ms, and then divide by the total number of all web application requests.

**Answer: C**

Reference: <https://sre.google/workbook/implementing-slos/>

### QUESTION NO: 9

You deploy a new release of an internal application during a weekend maintenance window when there is minimal user traffic. After the window ends, you learn that one of the new features isn't working as expected in the production environment. After an extended outage, you roll back the new release and deploy a fix. You want to modify your release process to reduce the mean time to

---

recovery so you can avoid extended outages in the future. What should you do? (Choose two.)

**A.**

Before merging new code, require 2 different peers to review the code changes.

**B.**

Adopt the blue/green deployment strategy when releasing new code via a CD server.

**C.**

Integrate a code linting tool to validate coding standards before any code is accepted into the repository.

**D.**

Require developers to run automated integration tests on their local development environments before release.

**E.**

Configure a CI server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

**Answer: A,C**

**Explanation:**

#### **QUESTION NO: 10**

You have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on Google Cloud?

**A.**

- Deploy the Cloud Logging agent to the application servers.
- Give the developers the IAM Logs Viewer role to access Cloud Logging and view logs.

**B.**

- Deploy the Cloud Logging agent to the application servers.
- Give the developers the IAM Logs Private Logs Viewer role to access Cloud Logging and view logs.

**C.**

- Deploy the Cloud Monitoring agent to the application servers.
- Give the developers the IAM Monitoring Viewer role to access Cloud Monitoring and view metrics.