

PROFESSIONAL-CLOUD-DEVOPS- ENGINEER^{Q&As}

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

As a Site Reliability Engineer, you support an application written in Go that runs on Google Kubernetes Engine (GKE) in production. After releasing a new version of the application, you notice the application runs for about 15 minutes and then restarts. You decide to add Cloud Profiler to your application and now notice that the heap usage grows constantly until the application restarts. What should you do?

- A. Increase the CPU limit in the application deployment.
- B. Add high memory compute nodes to the cluster.
- C. Increase the memory limit in the application deployment.
- D. Add Cloud Trace to the application, and redeploy.

Correct Answer: C

QUESTION 2

You are investigating issues in your production application that runs on Google Kubernetes Engine (GKE). You determined that the source of the issue is a recently updated container image, although the exact change in code was not identified. The deployment is currently pointing to the latest tag. You need to update your cluster to run a version of the container that functions as intended. What should you do?

- A. Create a new tag called stable that points to the previously working container, and change the deployment to point to the new tag.
- B. Alter the deployment to point to the sha256 digest of the previously working container.
- C. Build a new container from a previous Git tag, and do a rolling update on the deployment to the new container.
- D. Apply the latest tag to the previous container image, and do a rolling update on the deployment.

Correct Answer: B

<https://cloud.google.com/kubernetes-engine/docs/concepts/about-container-images>

QUESTION 3

You support an application deployed on Compute Engine. The application connects to a Cloud SQL instance to store and retrieve data. After an update to the application, users report errors showing database timeout messages. The number of concurrent active users remained stable. You need to find the most probable cause of the database timeout. What should you do?

- A. Check the serial port logs of the Compute Engine instance.
- B. Use Stackdriver Profiler to visualize the resources utilization throughout the application.
- C. Determine whether there is an increased number of connections to the Cloud SQL instance.
- D. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Correct Answer: B

QUESTION 4

You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how much it costs to run each of the systems. What should you do?

- A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.
- B. Assign all instances a label specific to the system they run. Configure BigQuery billing export and query costs per label.
- C. Enrich all instances with metadata specific to the system they run. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.
- D. Name each virtual machine (VM) after the system it runs. Set up a usage report export to a Cloud Storage bucket. Configure the bucket as a source in BigQuery to query costs based on VM name.

Correct Answer: B

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

QUESTION 5

You are creating a CI/CD pipeline in Cloud Build to build an application container image. The application code is stored in GitHub. Your company requires that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch. You want the image build to be as automated as possible. What should you do? (Choose two.)

- A. Create a trigger on the Cloud Build job. Set the repository event setting to 'Pull request'.
- B. Add the OWNERS file to the Included files filter on the trigger.
- C. Create a trigger on the Cloud Build job. Set the repository event setting to 'Push to a branch'.
- D. Configure a branch protection rule for the main branch on the repository.
- E. Enable the Approval option on the trigger.

Correct Answer: CD

Option C: Setting the repository event setting to 'Push to a branch' will trigger the Cloud Build job whenever a push is made to any branch in the repository. This is necessary because you want the image build to be triggered when a push is

made to the main branch.

Option D: Configuring a branch protection rule for the main branch on the repository will require that all pushes to the main branch be approved by the change control team. This is necessary to ensure that only approved changes are made to

the main branch, which will then trigger the image build.

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

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

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