

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.pass2lead.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

You are deploying a Cloud Build job that deploys Terraform code when a Git branch is updated. While testing, you noticed that the job fails. You see the following error in the build logs:

Initializing the backend...

Error: Failed to get existing workspaces: querying Cloud Storage failed: googleapi: Error 403

You need to resolve the issue by following Google-recommended practices. What should you do?

- A. Change the Terraform code to use local state.
- B. Create a storage bucket with the name specified in the Terraform configuration.
- C. Grant the roles/owner Identity and Access Management (IAM) role to the Cloud Build service account on the project.
- D. Grant the roles/storage.objectAdmin Identity and Access Management (IAM) role to the Cloud Build service account on the state file bucket.

Correct Answer: D

QUESTION 2

Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk: Mean Time to Detect (MTTD) in minutes Mean Time to Repair (MTTR) in minutes Mean Time Between Failure (MTBF) in days User Impact Percentage

The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

- A. MTTD: 5 MTTR: 10 MTBF: 90 Impact: 33%
- B. MTTD: 5 MTTR: 20 MTBF: 90 Impact: 33%
- C. MTTD: 5 MTTR: 10 MTBF: 90 Impact: 50%
- D. MTTD: 5 MTTR: 20 MTBF: 90 Impact: 50%

Correct Answer: B

QUESTION 3

You are the Operations Lead for an ongoing incident with one of your services. The service usually runs at around 70% capacity. You notice that one node is returning 5xx errors for all requests. There has also been a noticeable increase in

support cases from customers. You need to remove the offending node from the load balancer pool so that you can isolate and investigate the node. You want to follow Google-recommended practices to manage the incident and reduce the impact on users. What should you do?

A. 1. Communicate your intent to the incident team.

2.

Perform a load analysis to determine if the remaining nodes can handle the increase in traffic offloaded from the removed node, and scale appropriately.

3.

When any new nodes report healthy, drain traffic from the unhealthy node, and remove the unhealthy node from service.

B. 1. Communicate your intent to the incident team.

2.

Add a new node to the pool, and wait for the new node to report as healthy.

3.

When traffic is being served on the new node, drain traffic from the unhealthy node, and remove the old node from service.

C. 1. Drain traffic from the unhealthy node and remove the node from service.

2.

Monitor traffic to ensure that the error is resolved and that the other nodes in the pool are handling the traffic appropriately.

3.

Scale the pool as necessary to handle the new load.

4.

Communicate your actions to the incident team.

D. 1. Drain traffic from the unhealthy node and remove the old node from service.

2.

Add a new node to the pool, wait for the new node to report as healthy, and then serve traffic to the new node.

3.

Monitor traffic to ensure that the pool is healthy and is handling traffic appropriately.

4.

Communicate your actions to the incident team.

Correct Answer: A

QUESTION 4

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 5

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do?

- A. Use Stackdriver Kubernetes Engine Monitoring.
- B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
- C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.
- D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Correct Answer: A

*

<https://cloud.google.com/anthos/clusters/docs/on-prem> GKE on-prem is also called Anthos clusters on VMware

*

<https://cloud.google.com/anthos/clusters/docs/on-prem/concepts/logging-and-monitoring> You have several logging and monitoring options for your Anthos clusters on VMware:

+

Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with Anthos clusters on VMware.

+

Prometheus and Grafana, disabled by default.

+

Validated configurations with third-party solutions.

=> it means, if not a special situation, the correct should be using the first option: Logging and Monitoring. In this case, we want metrics, so Monitoring (aka. Cloud Monitoring, Stackdriver Monitoring) should be used. We are talking about GKE, so we will use Kubernetes Engine Monitoring (<https://cloud.google.com/kubernetes-engine-monitoring>).

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

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

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Braindumps](#)