

PROFESSIONAL-CLOUD- DEVELOPER^{Q&As}

Professional Cloud Developer

**Pass Google PROFESSIONAL-CLOUD-DEVELOPER
Exam with 100% Guarantee**

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

<https://www.leads4pass.com/professional-cloud-developer.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 recently joined a new team that has a Cloud Spanner database instance running in production. Your manager has asked you to optimize the Spanner instance to reduce cost while maintaining high reliability and availability of the database. What should you do?

- A. Use Cloud Logging to check for error logs, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- B. Use Cloud Trace to monitor the requests per sec of incoming requests to Spanner, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- C. Use Cloud Monitoring to monitor the CPU utilization, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- D. Use Snapshot Debugger to check for application errors, and reduce Spanner processing units by small increments until you find the minimum capacity required.

Correct Answer: C

https://cloud.google.com/spanner/docs/computecapacity#increasing_and_decreasing_compute_capacity

QUESTION 2

Your company's product team has a new requirement based on customer demand to autoscale your stateless and distributed service running in a Google Kubernetes Engine (GKE) cluster. You want to find a solution that minimizes changes because this feature will go live in two weeks. What should you do?

- A. Deploy a Vertical Pod Autoscaler, and scale based on the CPU load.
- B. Deploy a Vertical Pod Autoscaler, and scale based on a custom metric.
- C. Deploy a Horizontal Pod Autoscaler, and scale based on the CPU load.
- D. Deploy a Horizontal Pod Autoscaler, and scale based on a custom metric.

Correct Answer: C

<https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler>

The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption, or in response to custom metrics reported from within Kubernetes or external metrics from sources outside of your cluster.

QUESTION 3

You need to redesign the ingestion of audit events from your authentication service to allow it to handle a large increase in traffic. Currently, the audit service and the authentication system run in the same Compute Engine virtual machine.

You plan to use the following Google Cloud tools in the new architecture:

Multiple Compute Engine machines, each running an instance of the authentication service

Multiple Compute Engine machines, each running an instance of the audit service

Pub/Sub to send the events from the authentication services.

How should you set up the topics and subscriptions to ensure that the system can handle a large volume of messages and can scale efficiently?

- A. Create one Pub/Sub topic. Create one pull subscription to allow the audit services to share the messages.
- B. Create one Pub/Sub topic. Create one pull subscription per audit service instance to allow the services to share the messages.
- C. Create one Pub/Sub topic. Create one push subscription with the endpoint pointing to a load balancer in front of the audit services.
- D. Create one Pub/Sub topic per authentication service. Create one pull subscription per topic to be used by one audit service.
- E. Create one Pub/Sub topic per authentication service. Create one push subscription per topic, with the endpoint pointing to one audit service.

Correct Answer: A

Explanation: <https://cloud.google.com/pubsub/docs/subscriber> "Multiple subscribers can make pull calls to the same "shared" subscription. Each subscriber will receive a subset of the messages."

QUESTION 4

You are a developer at a large organization. You are deploying a web application to Google Kubernetes Engine (GKE). The DevOps team has built a CI/CD pipeline that uses Cloud Deploy to deploy the application to Dev Test, and Prod clusters in GKE. After Cloud Deploy successfully deploys the application to the Dev cluster you want to automatically promote it to the Test Cluster. How should you configure this process following Google-recommended best practices?

- A. 1 Create a Cloud Build trigger that listens for SUCCEEDED Pub/Sub messages from the cloudeploy-operations topic. 2 Configure Cloud Build to include a step that promotes the application to the Test cluster
- B. 1 Create a Cloud Function that calls the Google Cloud Deploy API to promote the application to the Test cluster 2 Configure this function to be triggered by SUCCEEDED Pub/Sub messages from the cloud-builds topic
- C. 1 Create a Cloud Function that calls the Google Cloud Deploy API to promote the application to the Test cluster 2 Configure this function to be triggered by SUCCEEDED Pub/Sub messages from the cloudeploy operations topic
- D. 1 Create a Cloud Build pipeline that uses the gke-deploy builder 2 Create a Cloud Build trigger that listens to SUCCEEDED Pub/Sub messages from the cloud-builds topic 3 Configure this pipeline to run a deployment step to the Test cluster

Correct Answer: B

QUESTION 5

HipLocal is configuring their access controls.

Which firewall configuration should they implement?

- A. Block all traffic on port 443.
- B. Allow all traffic into the network.
- C. Allow traffic on port 443 for a specific tag.
- D. Allow all traffic on port 443 into the network.

Correct Answer: D

[PROFESSIONAL-CLOUD-DEVELOPER Practice Test](#)

[PROFESSIONAL-CLOUD-DEVELOPER Exam Questions](#)

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)