

# PROFESSIONAL-CLOUD- DEVELOPER<sup>Q&As</sup>

Professional Cloud Developer

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

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

<https://www.pass2lead.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

Your company wants to expand their users outside the United States for their popular application. The company wants to ensure 99.999% availability of the database for their application and also wants to minimize the read latency for their users across the globe.

Which two actions should they take? (Choose two.)

- A. Create a multi-regional Cloud Spanner instance with "nam-asia-eur1" configuration.
- B. Create a multi-regional Cloud Spanner instance with "nam3" configuration.
- C. Create a cluster with at least 3 Spanner nodes.
- D. Create a cluster with at least 1 Spanner node.
- E. Create a minimum of two Cloud Spanner instances in separate regions with at least one node.
- F. Create a Cloud Dataflow pipeline to replicate data across different databases.

Correct Answer: BF

---

### QUESTION 2

Your development team has been tasked with maintaining a .NET legacy application. The application incurs occasional changes and was recently updated. Your goal is to ensure that the application provides consistent results while moving through the CI/CD pipeline from environment to environment. You want to minimize the cost of deployment while making sure that external factors and dependencies between hosting environments are not problematic. Containers are not yet approved in your organization. What should you do?

- A. Rewrite the application using .NET Core, and deploy to Cloud Run. Use revisions to separate the environments.
- B. Use Cloud Build to deploy the application as a new Compute Engine image for each build. Use this image in each environment.
- C. Deploy the application using MS Web Deploy, and make sure to always use the latest, patched MS Windows Server base image in Compute Engine.
- D. Use Cloud Build to package the application, and deploy to a Google Kubernetes Engine cluster. Use namespaces to separate the environments.

Correct Answer: B

[https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase\\_1\\_rehost\\_in\\_the\\_cloud](https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase_1_rehost_in_the_cloud) <https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud>

---

### QUESTION 3

Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

- A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.
- C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Correct Answer: C

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

---

#### QUESTION 4

Your application is logging to Stackdriver. You want to get the count of all requests on all /api/alpha/\* endpoints.

What should you do?

- A. Add a Stackdriver counter metric for path:/api/alpha/.
- B. Add a Stackdriver counter metric for endpoint:/api/alpha/\*.
- C. Export the logs to Cloud Storage and count lines matching /api/alpha.
- D. Export the logs to Cloud Pub/Sub and count lines matching /api/alpha.

Correct Answer: C

---

#### QUESTION 5

In order for HipLocal to store application state and meet their stated business requirements, which database service should they migrate to?

- A. Cloud Spanner
- B. Cloud Datastore
- C. Cloud Memorystore as a cache
- D. Separate Cloud SQL clusters for each region

Correct Answer: D

---

[Dumps](#)