

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 are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster.

Your application can scale horizontally, and each instance of your application needs to have a stable network identity and its own persistent disk.

Which GKE object should you use?

- A. Deployment
- B. StatefulSet
- C. ReplicaSet
- D. ReplicaController

Correct Answer: B

Reference: <https://livebook.manning.com/book/kubernetes-in-action/chapter-10/46>

QUESTION 2

Your website is deployed on Compute Engine. Your marketing team wants to test conversion rates between 3 different website designs. Which approach should you use?

- A. Deploy the website on App Engine and use traffic splitting.
- B. Deploy the website on App Engine as three separate services.
- C. Deploy the website on Cloud Functions and use traffic splitting.
- D. Deploy the website on Cloud Functions as three separate functions.

Correct Answer: A

Reference: <https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

QUESTION 3

Your team develops services that run on Google Cloud. You want to process messages sent to a Pub/Sub topic, and then store them. Each message must be processed exactly once to avoid duplication of data and any data conflicts. You need to use the cheapest and most simple solution. What should you do?

- A. Process the messages with a Dataproc job, and write the output to storage.
- B. Process the messages with a Dataflow streaming pipeline using Apache Beam's PubSubIO package, and write the output to storage.
- C. Process the messages with a Cloud Function, and write the results to a BigQuery location where you can run a job to deduplicate the data.

D. Retrieve the messages with a Dataflow streaming pipeline, store them in Cloud Bigtable, and use another Dataflow streaming pipeline to deduplicate messages.

Correct Answer: B

<https://cloud.google.com/dataflow/docs/concepts/streaming-with-cloud-pubsub>

QUESTION 4

You manage your company's ecommerce platform's payment system, which runs on Google Cloud. Your company must retain user logs for 1 year for internal auditing purposes and for 3 years to meet compliance requirements. You need to

store new user logs on Google Cloud to minimize on-premises storage usage and ensure that they are easily searchable. You want to minimize effort while ensuring that the logs are stored correctly.

What should you do?

- A. Store the logs in a Cloud Storage bucket with bucket lock turned on.
- B. Store the logs in a Cloud Storage bucket with a 3-year retention period.
- C. Store the logs in Cloud Logging as custom logs with a custom retention period.
- D. Store the logs in a Cloud Storage bucket with a 1-year retention period. After 1 year, move the logs to another bucket with a 2-year retention period.

Correct Answer: C

<https://cloud.google.com/logging/docs/buckets#custom-retention>

QUESTION 5

Which database should HipLocal use for storing user activity?

- A. BigQuery
- B. Cloud SQL
- C. Cloud Spanner
- D. Cloud Datastore

Correct Answer: A

[PROFESSIONAL-CLOUD-DEVELOPER VCE Dumps](#)

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

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