



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-developer.html>

2024 Latest geekcert PROFESSIONAL-CLOUD-DEVELOPER PDF and VCE dumps Download

---

# 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.geekcert.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



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-developer.html>

2024 Latest geekcert PROFESSIONAL-CLOUD-DEVELOPER PDF and VCE dumps Download

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





### QUESTION 1

You are a developer at a large organization. You have an application written in Go running in a production Google Kubernetes Engine (GKE) cluster. You need to add a new feature that requires access to BigQuery. You want to grant BigQuery access to your GKE cluster following Google-recommended best practices. What should you do?

- A. Create a Google service account with BigQuery access. Add the JSON key to Secret Manager, and use the Go client library to access the JSON key.
- B. Create a Google service account with BigQuery access. Add the Google service account JSON key as a Kubernetes secret, and configure the application to use this secret.
- C. Create a Google service account with BigQuery access. Add the Google service account JSON key to Secret Manager, and use an init container to access the secret for the application to use.
- D. Create a Google service account and a Kubernetes service account. Configure Workload Identity on the GKE cluster, and reference the Kubernetes service account on the application Deployment.

Correct Answer: D

[https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity#what\\_is](https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity#what_is) Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs. Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

---

### QUESTION 2

HipLocal\\s.net-based auth service fails under intermittent load. What should they do?

- A. Use App Engine for autoscaling.
- B. Use Cloud Functions for autoscaling.
- C. Use a Compute Engine cluster for the service.
- D. Use a dedicated Compute Engine virtual machine instance for the service.

Correct Answer: D

Reference: <https://www.qwiklabs.com/focuses/611?parent=catalog>

---

### QUESTION 3

Your team is writing a backend application to implement the business logic for an interactive voice response (IVR) system that will support a payroll application. The IVR system has the following technical characteristics:

?Each customer phone call is associated with a unique IVR session.

?The IVR system creates a separate persistent gRPC connection to the backend for each session.



?If the connection is interrupted, the IVR system establishes a new connection, causing a slight latency for that call.

You need to determine which compute environment should be used to deploy the backend application. Using current call data, you determine that:

?Call duration ranges from 1 to 30 minutes.

?Calls are typically made during business hours.

?There are significant spikes of calls around certain known dates (e.g., pay days), or when large payroll changes occur.

You want to minimize cost, effort, and operational overhead. Where should you deploy the backend application?

- A. Compute Engine
- B. Google Kubernetes Engine cluster in Standard mode
- C. Cloud Functions
- D. Cloud Run

Correct Answer: D

This page shows Cloud Run-specific details for developers who want to use gRPC to connect a Cloud Run service with other services, for example, to provide simple, high performance communication between internal microservices. You can

use all gRPC types, streaming or unary, with Cloud Run.

Possible use cases include:

Communication between internal microservices.

High loads of data (gRPC uses protocol buffers, which are up to seven times faster than REST calls).

Only a simple service definition is needed, you don't want to write a full client library. Use streaming gRPCs in your gRPC server to build more responsive applications and APIs.

<https://cloud.google.com/run/docs/tutorials/secure-services#:~:text=The%20backend%20service%20is%20private,Google%20Cloud%20except%20where%20necessary.>

#### QUESTION 4

You recently developed an application that monitors a large number of stock prices. You need to configure Pub/Sub to receive a high volume messages and update the current stock price in a single large in-memory database. The downstream service needs only the most up-to-date prices in the in-memory database to perform stock trading transactions. Each message contains three pieces of information: ?Stock symbol ?Stock price ?Timestamp for the update. How should you set up your Pub/Sub subscription?

- A. Create a pull subscription with both ordering and exactly-once delivery turned off
- B. Create a pull subscription with exactly-once delivery enabled
- C. Create a push subscription with exactly-once delivery enabled



---

D. Create a push subscription with both ordering and exactly-once delivery turned off

Correct Answer: C

---

### QUESTION 5

You are developing a flower ordering application. Currently you have three microservices.

?Order Service (receives the orders).

?Order Fulfillment Service (processes the orders).

?Notification Service (notifies the customer when the order is filled).

You need to determine how the services will communicate with each other. You want incoming orders to be processed quickly and you need to collect order information for fulfillment. You also want to make sure orders are not lost between

your services and are able to communicate asynchronously. How should the requests be processed?

A.

B.

C.

D.

Correct Answer: A

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

[PROFESSIONAL-CLOUD-DEVELOPER Study Guide](#)

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)