



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^{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.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

Your team develops services that run on Google Cloud. You need to build a data processing service and will use Cloud Functions. The data to be processed by the function is sensitive. You need to ensure that invocations can only happen from authorized services and follow Google-recommended best practices for securing functions. What should you do?

- A. Enable Identity-Aware Proxy in your project. Secure function access using its permissions.
- B. Create a service account with the Cloud Functions Viewer role. Use that service account to invoke the function.
- C. Create a service account with the Cloud Functions Invoker role. Use that service account to invoke the function.
- D. Create an OAuth 2.0 client ID for your calling service in the same project as the function you want to secure. Use those credentials to invoke the function.

Correct Answer: C

Reference: <https://medium.com/google-cloud/how-to-securely-invoke-a-cloud-function- from-google-kubernetes-engine-running-on-another-gcp-79797ec2b2c6>

QUESTION 2

You support an application that uses the Cloud Storage API. You review the logs and discover multiple HTTP 503 Service Unavailable error responses from the API. Your application logs the error and does not take any further action. You want to implement Google-recommended retry logic to improve success rates. Which approach should you take?

- A. Retry the failures in batch after a set number of failures is logged.
- B. Retry each failure at a set time interval up to a maximum number of times.
- C. Retry each failure at increasing time intervals up to a maximum number of tries.
- D. Retry each failure at decreasing time intervals up to a maximum number of tries.

Correct Answer: C

Explanation: <https://cloud.google.com/storage/docs/retry-strategy>

QUESTION 3

You want to view the memory usage of your application deployed on Compute Engine. What should you do?

- A. Install the Stackdriver Client Library.
- B. Install the Stackdriver Monitoring Agent.
- C. Use the Stackdriver Metrics Explorer.
- D. Use the Google Cloud Platform Console.

Correct Answer: C



Reference: <https://stackoverflow.com/questions/43991246/google-cloud-platform-how-to-monitor-memory-usage-of-vm-instances>

QUESTION 4

Your teammate has asked you to review the code below. Its purpose is to efficiently add a large number of small rows to a BigQuery table.

```
BigQuery service = BigQueryOptions.newBuilder().build().getService();

public void writeToBigQuery(Collection<Map<String, String>> rows){
    for(Map<String, String> row : rows) {
        InsertAllRequest insertRequest = InsertAllRequest.newBuilder(
            "datasetId", "tableId",
            InsertAllRequest.RowToInsert.of(row)).build();
        service.insertAll(insertRequest);
    }
}
```

Which improvement should you suggest your teammate make?

- A. Include multiple rows with each request.
- B. Perform the inserts in parallel by creating multiple threads.
- C. Write each row to a Cloud Storage object, then load into BigQuery.
- D. Write each row to a Cloud Storage object in parallel, then load into BigQuery.

Correct Answer: B

QUESTION 5

Your team has created an application that is hosted on a Google Kubernetes Engine (GKE) cluster. You need to connect the application to a legacy REST service that is deployed in two GKE clusters in two different regions. You want to connect your application to the legacy service in a way that is resilient and requires the fewest number of steps. You also want to be able to run probe-based health checks on the legacy service on a separate port. How should you set up the connection?

- A. Use Traffic Director with a sidecar proxy to connect the application to the service.
- B. Use a proxyless Traffic Director configuration to connect the application to the service.
- C. Configure the legacy service's firewall to allow health checks originating from the proxy.
- D. Configure the legacy service's firewall to allow health checks originating from the application.
- E. Configure the legacy service's firewall to allow health checks originating from the Traffic Director control plane.

Correct Answer: AC



Explanation: <https://cloud.google.com/traffic-director/docs/advanced-setup#routing-rule-maps>

<https://cloud.google.com/traffic-director/docs/advanced-setup>

A. Using Traffic Director with a sidecar proxy can provide resilience for your application by allowing for failover to the secondary region in the event of an outage. The sidecar proxy can route traffic to the legacy service in either of the two GKE clusters, ensuring high availability. C. Configuring the legacy service's firewall to allow health checks originating from the proxy allows the proxy to periodically check the health of the legacy service and ensure that it is functioning properly. This helps to ensure that traffic is only routed to healthy instances of the legacy service, further improving the resilience of the setup.

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

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

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