

Q&As

Professional Cloud Architect on Google Cloud Platform

Pass Google PROFESSIONAL-CLOUD-ARCHITECT Exam with 100% Guarantee

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

<https://www.pass2lead.com/professional-cloud-architect.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

TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour. How should you design the data ingestion?

- A. Vehicles write data directly to GCS
- B. Vehicles write data directly to Google Cloud Pub/Sub
- C. Vehicles stream data directly to Google BigQuery
- D. Vehicles continue to write data using the existing system (FTP)

Correct Answer: B

Streamed data is available for real-time analysis within a few seconds of the first streaming insertion into a table.

Instead of using a job to load data into BigQuery, you can choose to stream your data into BigQuery one record at a time by using the `tabledata().insertAll()` method. This approach enables querying data without the delay of running a load job.

References: <https://cloud.google.com/bigquery/streaming-data-into-bigquery>

QUESTION 2

Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others. Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier.

How should you configure the network?

- A. Add each tier to a different subnetwork
- B. Set up software based firewalls on individual VMs
- C. Add tags to each tier and set up routes to allow the desired traffic flow
- D. Add tags to each tier and set up firewall rules to allow the desired traffic flow

Correct Answer: D

Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.

References:

<https://cloud.google.com/docs/compare/openstack/>

<https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

QUESTION 3

You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

- A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
- B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
- C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
- D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Correct Answer: B

QUESTION 4

At Dress4Win, an operations engineer wants to create a low-cost solution to remotely archive copies of database backup files.

The database files are compressed tar files stored in their current data center.

How should he proceed?

- A. Create a cron script using gsutil to copy the files to a Coldline Storage bucket.
- B. Create a cron script using gsutil to copy the files to a Regional Storage bucket.
- C. Create a Cloud Storage Transfer Service Job to copy the files to a Coldline Storage bucket.
- D. Create a Cloud Storage Transfer Service job to copy the files to a Regional Storage bucket.

Correct Answer: C

As per the latest case study on google cloud website , they have DB storage of 1 PB out of which 600 TB is used. So you get the size of the data. These are the thumb rules as per GCP documentation Transferring from another cloud storage provider Use Storage Transfer Service

Transferring less than 1 TB from on-premises Use gsutil Transferring more than 1 TB from on-premises Use Transfer service for on-premises data <https://cloud.google.com/storage-transfer/docs/overview>

QUESTION 5

Your company has an application running on Compute Engine that allows users to play their favorite music. There are a fixed number of instances. Files are stored in Cloud Storage, and data is streamed directly to users. Users are reporting that they sometimes need to attempt to play popular songs multiple times before they are successful. You need to improve the performance of the application. What should you do?

- A. 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances.
- 2. Serve music files directly from the backend Compute Engine instance.

B. 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances.

2.

Download popular songs in Cloud Filestore.

3.

Serve music files directly from the backend Compute Engine instance.

C. 1. Copy popular songs into CloudSQL as a blob.

2. Update application code to retrieve data from CloudSQL when Cloud Storage is overloaded.

D. 1. Create a managed instance group with Compute Engine instances.

2.

Create a global load balancer and configure it with two backends:

Managed instance group

Cloud Storage bucket

3.

Enable Cloud CDN on the bucket backend.

Correct Answer: D

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)