



VCE & PDF

PassApply.com

<https://www.passapply.com/professional-cloud-architect.html>

2024 Latest passapply PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE dumps Download

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.passapply.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



VCE & PDF

PassApply.com

<https://www.passapply.com/professional-cloud-architect.html>

2024 Latest passapply PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE dumps Download

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





QUESTION 1

You need to deploy an application to Google Cloud. The application receives traffic via TCP and reads and writes data to the filesystem. The application does not support horizontal scaling. The application process requires full control over the data on the file system because concurrent access causes corruption. The business is willing to accept a downtime when an incident occurs, but the application must be available 24/7 to support their business operations. You need to design the architecture of this application on Google Cloud.

What should you do?

- A. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use an HTTP load balancer in front of the instances.
- B. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use a network load balancer in front of the instances.
- C. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use an HTTP load balancer in front of the instances.
- D. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

Correct Answer: D

Reference: <https://cloud.google.com/compute/docs/instance-groups>

QUESTION 2

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

- A. Write the log concurrently in the cloud and on premises.
- B. Use a SQL database and limit who can modify the log table.
- C. Digitally sign each timestamp and log entry and store the signature.
- D. Create a JSON dump of each log entry and store it in Google Cloud Storage.

Correct Answer: C

<https://cloud.google.com/storage/docs/access-logs> <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

QUESTION 3

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform. Which two steps should be part of their migration plan? (Choose two.)

- A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.



- B. Write a schema migration plan to denormalize data for better performance in BigQuery.
- C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
- D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.
- E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Correct Answer: AB

QUESTION 4

Your company just finished a rapid lift and shift to Google Compute Engine for your compute needs. You have another 9 months to design and deploy a more cloud-native solution. Specifically, you want a system that is no-ops and auto-scaling. Which two compute products should you choose? Choose 2 answers

- A. Compute Engine with containers
- B. Google Kubernetes Engine with containers
- C. Google App Engine Standard Environment
- D. Compute Engine with custom instance types
- E. Compute Engine with managed instance groups

Correct Answer: BC

B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.

Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.

C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps. App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.

Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.

References: <https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%99s-NoOps-claim>

QUESTION 5

You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

- A. Install the latest BigQuery API client library for Python
- B. Run your script on a new virtual machine with the BigQuery access scope enabled



C. Create a new service account with BigQuery access and execute your script with that user

D. Install the bq component for gcloud with the command gcloud components install bq.

Correct Answer: B

The error is most likely caused by the access scope issue. When creating a new instance, you have the default Compute engine default service account but most services access including BigQuery is not enabled. Create an instance. Most access is not enabled by default. You have a default service account but don't have the permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you run your script on a new virtual machine with the BigQuery access scope enabled, it also works.

<https://cloud.google.com/compute/docs/access/service-accounts>

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)