



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 are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

A. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.

2.

Stop the on-premises application.

3.

Create a mysqldump of the on-premises MySQL server.

4.

Upload the dump to a Cloud Storage bucket.

5.

Import the dump into Cloud SQL.

6.

Modify the source code of the application to write queries to both databases and read from its local database.

7.

Start the Compute Engine application.

8.

Stop the on-premises application.

B. 1. Set up Cloud SQL proxy and MySQL proxy.

2.

Create a mysqldump of the on-premises MySQL server.

3.

Upload the dump to a Cloud Storage bucket.

4.

Import the dump into Cloud SQL.

5.

Stop the on-premises application.



6.

Start the Compute Engine application.

C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.

2.

Stop the on-premises application.

3.

Start the Compute Engine application, configured to read and write to the on-premises MySQL server.

4.

Create the replication configuration in Cloud SQL.

5.

Configure the source database server to accept connections from the Cloud SQL replica.

6.

Finalize the Cloud SQL replica configuration.

7.

When replication has been completed, stop the Compute Engine application.

8.

Promote the Cloud SQL replica to a standalone instance.

9.

Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.

D. 1. Stop the on-premises application.

2.

Create a mysqldump of the on-premises MySQL server.

3.

Upload the dump to a Cloud Storage bucket.

4.

Import the dump into Cloud SQL.

5.

Start the application on Compute Engine.



Correct Answer: C

External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.

<https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept>

The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration. - https://cloud.google.com/architecture/migrating-mysql-to-cloudsqlconcept#external_replica_promotion_migration

QUESTION 2

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

- A. Deploy each microservice as a Deployment. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
- B. Deploy each microservice as a Deployment. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
- C. Deploy each microservice as a Pod. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- D. Deploy each microservice as a Pod. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Correct Answer: A

<https://kubernetes.io/docs/concepts/services-networking/ingress/>

QUESTION 3

Your organization has stored sensitive data in a Cloud Storage bucket. For regulatory reasons, your company must be able to rotate the encryption key used to encrypt the data in the bucket. The data will be processed in Dataproc. You want to follow Google-recommended practices for security. What should you do?

- A. Create a key with Cloud Key Management Service (KMS). Encrypt the data using the encrypt method of Cloud KMS.
- B. Create a key with Cloud Key Management Service (KMS). Set the encryption key on the bucket to the Cloud KMS key.
- C. Generate a GPG key pair. Encrypt the data using the GPG key. Upload the encrypted data to the bucket.



D. Generate an AES-256 encryption key. Encrypt the data in the bucket using the customer-supplied encryption keys feature.

Correct Answer: B

<https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key>

QUESTION 4

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

- A. Use a persistent disk for each instance.
- B. Use a regional persistent disk for each instance.
- C. Create a Cloud Filestore instance and mount it in each instance.
- D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Correct Answer: C

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

<https://cloud.google.com/storage/docs/gcs-fuse#notes> Reference: <https://cloud.google.com/storage/docs/gcs-fuse>

QUESTION 5

Mountkirk Games\' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

- A. Verify that the database is online
- B. Verify that the project quota hasn\'t been exceeded
- C. Verify that the new feature code did not introduce any performance bugs
- D. Verify that the load-testing team is not running their tool against production

Correct Answer: B

503 is service unavailable error. If the database was online everyone would get the 503 error.