



**Free Questions for Professional-Cloud-Database-Engineer by
certsdeals**

Shared by Rodgers on 20-10-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are designing a database strategy for a new web application in one region. You need to minimize write latency. What should you do?

Options:

- A- Use Cloud SQL with cross-region replicas.
- B- Use high availability (HA) Cloud SQL with multiple zones.
- C- Use zonal Cloud SQL without high availability (HA).
- D- Use Cloud Spanner in a regional configuration.

Answer:

A

Question 2

Question Type: MultipleChoice

You are designing a new gaming application that uses a highly transactional relational database to store player authentication and inventory data in Google Cloud. You want to launch the game in multiple regions. What should you do?

Options:

- A- Use Cloud Spanner to deploy the database.
- B- Use Bigtable with clusters in multiple regions to deploy the database
- C- Use BigQuery to deploy the database
- D- Use Cloud SQL with a regional read replica to deploy the database.

Answer:

C

Question 3

Question Type: MultipleChoice

Your team is building an application that stores and analyzes streaming time series financial dat

a. You need a database solution that can perform time series-based scans with sub-second latency. The solution must scale into the hundreds of terabytes and be able to write up to 10k records per second and read up to 200 MB per second. What should you do?

Options:

A- Use Firestore.

B- Use Bigtable

C- Use BigQuery.

D- Use Cloud Spanner.

Answer:

C

Question 4

Question Type: MultipleChoice

Your organization works with sensitive data that requires you to manage your own encryption keys. You are working on a project that stores that data in a Cloud SQL database. You need to ensure that stored data is encrypted with your keys. What should you do?

Options:

- A- Export data periodically to a Cloud Storage bucket protected by Customer-Supplied Encryption Keys.
- B- Use Cloud SQL Auth proxy.
- C- Connect to Cloud SQL using a connection that has SSL encryption.
- D- Use customer-managed encryption keys with Cloud SQL.

Answer:

C

Question 5

Question Type: MultipleChoice

You have an application that sends banking events to Bigtable cluster-a in us-east. You decide to add cluster-b in us-central1. Cluster-a replicates data to cluster-b. You need to ensure that Bigtable continues to accept read and write requests if one of the clusters becomes unavailable and that requests are routed automatically to the other cluster. What deployment strategy should you use?

Options:

- A- Use the default app profile with single-cluster routing.
- B- Use the default app profile with multi-cluster routing.
- C- Create a custom app profile with multi-cluster routing.
- D- Create a custom app profile with single-cluster routing.

Answer:

A

Question 6

Question Type: MultipleChoice

You want to migrate your on-premises PostgreSQL database to Compute Engine. You need to migrate this database with the minimum downtime possible. What should you do?

Options:

- A- Perform a full backup of your on-premises PostgreSQL, and then, in the migration window, perform an incremental backup.
- B- Create a read replica on Cloud SQL, and then promote it to a read/write standalone instance.

- C-** Use Database Migration Service to migrate your database.
- D-** Create a hot standby on Compute Engine, and use PgBouncer to switch over the connections.

Answer:

B

Question 7

Question Type: MultipleChoice

Your organization is running a critical production database on a virtual machine (VM) on Compute Engine. The VM has an ext4-formatted persistent disk for data files. The database will soon run out of storage space. You need to implement a solution that avoids downtime. What should you do?

Options:

- A-** In the Google Cloud Console, increase the size of the persistent disk, and use the `resize2fs` command to extend the disk.
- B-** In the Google Cloud Console, increase the size of the persistent disk, and use the `fdisk` command to verify that the new space is ready to use
- C-** In the Google Cloud Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old

disk, mount the new disk, and restart the database service.

D- In the Google Cloud Console, create a new persistent disk attached to the VM, and configure the database service to move the files to the new disk.

Answer:

D

Question 8

Question Type: MultipleChoice

Your organization has hundreds of Cloud SQL for MySQL instances. You want to follow Google-recommended practices to optimize platform costs. What should you do?

Options:

- A-** Use Query Insights to identify idle instances.
- B-** Remove inactive user accounts.
- C-** Run the Recommender API to identify overprovisioned instances.

D- Build indexes on heavily accessed tables.

Answer:

C

Question 9

Question Type: MultipleChoice

Your company is developing a global ecommerce website on Google Cloud. Your development team is working on a shopping cart service that is durable and elastically scalable with live traffic. Business disruptions from unplanned downtime are expected to be less than 5 minutes per month. In addition, the application needs to have very low latency writes. You need a data storage solution that has high write throughput and provides 99.99% uptime. What should you do?

Options:

A- Use Cloud SQL for data storage.

B- Use Cloud Spanner for data storage.

C- Use Memorystore for data storage.

D- Use Bigtable for data storage.

Answer:

A

Question 10

Question Type: MultipleChoice

Your organization needs to migrate a critical, on-premises MySQL database to Cloud SQL for MySQL. The on-premises database is on a version of MySQL that is supported by Cloud SQL and uses the InnoDB storage engine. You need to migrate the database while preserving transactions and minimizing downtime. What should you do?

Options:

A- Use Database Migration Service to connect to your on-premises database, and choose continuous replication.

After the on-premises database is migrated, promote the Cloud SQL for MySQL instance, and connect applications to your Cloud SQL instance.

B- Build a Cloud Data Fusion pipeline for each table to migrate data from the on-premises MySQL database to Cloud SQL for MySQL. Schedule downtime to run each Cloud Data Fusion pipeline. Verify that the migration was successful.

Re-point the applications to the Cloud SQL for MySQL instance.

C- Pause the on-premises applications.

Use the mysqldump utility to dump the database content in compressed format.

Run gsutil --m to move the dump file to Cloud Storage.

Use the Cloud SQL for MySQL import option.

After the import operation is complete, re-point the applications to the Cloud SQL for MySQL instance.

D- Pause the on-premises applications.

Use the mysqldump utility to dump the database content in CSV format.

Run gsutil --m to move the dump file to Cloud Storage.

Use the Cloud SQL for MySQL import option.

After the import operation is complete, re-point the applications to the Cloud SQL for MySQL instance.

Answer:

B

Question 11

Question Type: MultipleChoice

You finished migrating an on-premises MySQL database to Cloud SQL. You want to ensure that the daily export of a table, which was previously a cron job running on the database server, continues. You want the solution to minimize cost and operations overhead. What should you do?

Options:

- A-** Use Cloud Scheduler and Cloud Functions to run the daily export.
- B-** Create a streaming Datatlow job to export the table.
- C-** Set up Cloud Composer, and create a task to export the table daily.
- D-** Run the cron job on a Compute Engine instance to continue the export.

Answer:

C

**To Get Premium Files for Professional-Cloud-Database-Engineer
Visit**

<https://www.p2pexams.com/products/professional-cloud-database-engineer>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-database-engineer>

