

Google GCP-CDE Cloud Database Engineer

- **Up to Date products, reliable and verified.**
- **Questions and Answers in PDF Format.**

Full Version Features:

- **90 Days Free Updates**
- **30 Days Money Back Guarantee**
- **Instant Download Once Purchased**
- **24 Hours Live Chat Support**

For More Information:

<https://www.testsexpert.com/>

- **Product Version**

Latest Version: 6.0

Question: 1

Your company is shutting down their data center and migrating several MySQL and PostgreSQL databases to Google Cloud. Your database operations team is severely constrained by ongoing production releases and the lack of capacity for additional on-premises backups. You want to ensure that the scheduled migrations happen with minimal downtime and that the Google Cloud databases stay in sync with the on-premises data changes until the applications can cut over. What should you do? (Choose two.)

- A. Use Database Migration Service to migrate the databases to Cloud SQL.
- B. Use a cross-region read replica to migrate the databases to Cloud SQL.
- C. Use replication from an external server to migrate the databases to Cloud SQL.
- D. Use an external read replica to migrate the databases to Cloud SQL.
- E. Use a read replica to migrate the databases to Cloud SQL.

Answer: C,E

Question: 2

Your company is migrating the existing infrastructure for a highly transactional application to Google Cloud. You have several databases in a MySQL database instance and need to decide how to transfer the data to Cloud SQL. You need to minimize the downtime for the migration of your 500 GB instance. What should you do?

- A.
 1. Create a Cloud SQL for MySQL instance for your databases, and configure Datastream to stream your database changes to Cloud SQL.
 2. Select the Backfill historical data check box on your stream configuration to initiate Datastream to backfill any data that is out of sync between the source and destination.
 3. Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.
- B.
 1. Create migration job using Database Migration Service.
 2. Set the migration job type to Continuous, and allow the databases to complete the full dump phase and start sending data in change data capture (CDC) mode.
 3. Wait for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete.
 4. Update your application connections to the new instance.
- C.
 1. Create migration job using Database Migration Service.
 2. Set the migration job type to One-time, and perform this migration during a maintenance window.
 3. Stop all write workloads to the source database and initiate the dump. Wait for the dump to be loaded into the Cloud SQL destination database and the destination database to be promoted to the primary database.

4. Update your application connections to the new instance.
- D. 1. Use the mysqldump utility to manually initiate a backup of MySQL during the application maintenance window.
2. Move the files to Cloud Storage, and import each database into your Cloud SQL instance.
3. Continue to dump each database until all the databases are migrated.
4. Update your application connections to the new instance.

Answer: C

Question: 3

Your company uses the Cloud SQL out-of-disk recommender to analyze the storage utilization trends of production databases over the last 30 days. Your database operations team uses these recommendations to proactively monitor storage utilization and implement corrective actions. You receive a recommendation that the instance is likely to run out of disk space. What should you do to address this storage alert?

- A. Normalize the database to the third normal form.
- B. Compress the data using a different compression algorithm.
- C. Manually or automatically increase the storage capacity.
- D. Create another schema to load older data.

Answer: B

Question: 4

You are managing a mission-critical Cloud SQL for PostgreSQL instance. Your application team is running important transactions on the database when another DBA starts an on-demand backup. You want to verify the status of the backup. What should you do?

- A. Check the `cloudsql.googleapis.com/postgres.log` instance log.
- B. Perform the `gcloud sql operations list` command.
- C. Use Cloud Audit Logs to verify the status.
- D. Use the Google Cloud Console.

Answer: C

Question: 5

You support a consumer inventory application that runs on a multi-region instance of Cloud Spanner. A customer opened a support ticket to complain about slow response times. You notice a Cloud

Monitoring alert about high CPU utilization. You want to follow Google-recommended practices to address the CPU performance issue. What should you do first?

- A. Increase the number of processing units.
- B. Modify the database schema, and add additional indexes.
- C. Shard data required by the application into multiple instances.
- D. Decrease the number of processing units.

Answer: A

Question: 6

Your company uses Bigtable for a user-facing application that displays a low-latency real-time dashboard. You need to recommend the optimal storage type for this read-intensive database. What should you do?

- A. Recommend solid-state drives (SSD).
- B. Recommend splitting the Bigtable instance into two instances in order to load balance the concurrent reads.
- C. Recommend hard disk drives (HDD).
- D. Recommend mixed storage types.

Answer: B

Question: 7

Your organization has a critical business app that is running with a Cloud SQL for MySQL backend database. Your company wants to build the most fault-tolerant and highly available solution possible. You need to ensure that the application database can survive a zonal and regional failure with a primary region of us-central1 and the backup region of us-east1. What should you do?

- A. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-west1-b.
3. Create a read replica in us-east1-c.
- B. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-central1-b.
3. Create a read replica in us-east1-b.
- C. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-east-b.
3. Create a read replica in us-east1-c.
- D. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-east1-b.
3. Create a read replica in us-central1-b.

Answer: B

Question: 8

You are building an Android game that needs to store data on a Google Cloud serverless database. The database will log user activity, store user preferences, and receive in-game updates. The target audience resides in developing countries that have intermittent internet connectivity. You need to ensure that the game can synchronize game data to the backend database whenever an internet network is available. What should you do?

- A. Use Firestore.
- B. Use Cloud SQL with an external (public) IP address.
- C. Use an in-app embedded database.
- D. Use Cloud Spanner.

Answer: B

Question: 9

You released a popular mobile game and are using a 50 TB Cloud Spanner instance to store game data in a PITR-enabled production environment. When you analyzed the game statistics, you realized that some players are exploiting a loophole to gather more points to get on the leaderboard. Another DBA accidentally ran an emergency bugfix script that corrupted some of the data in the production environment. You need to determine the extent of the data corruption and restore the production environment. What should you do? (Choose two.)

- A. If the corruption is significant, use backup and restore, and specify a recovery timestamp.
- B. If the corruption is significant, perform a stale read and specify a recovery timestamp. Write the results back.
- C. If the corruption is significant, use import and export.
- D. If the corruption is insignificant, use backup and restore, and specify a recovery timestamp.
- E. If the corruption is insignificant, perform a stale read and specify a recovery timestamp. Write the results back.

Answer: B,E

Question: 10

You are starting a large CSV import into a Cloud SQL for MySQL instance that has many open connections. You checked memory and CPU usage, and sufficient resources are available. You want to follow Google-recommended practices to ensure that the import will not time out. What should you do?

- A. Close idle connections or restart the instance before beginning the import operation.
- B. Increase the amount of memory allocated to your instance.
- C. Ensure that the service account has the Storage Admin role.
- D. Increase the number of CPUs for the instance to ensure that it can handle the additional import operation.

Answer: C

For More Information – Visit link below:
<https://www.testsexpert.com/>

16\$ Discount Coupon: **9M2GK4NW**

Features:

■ Money Back Guarantee.....



■ 100% Course Coverage.....



■ 90 Days Free Updates.....



■ Instant Email Delivery after Order.....

