



PROFESSIONAL-CLOUD-DATABASE-ENGINEER^{Q&As}

Google Cloud Certified - Professional Cloud Database Engineer

Pass Google PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam with 100% Guarantee

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

<https://www.geekcert.com/professional-cloud-database-engineer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-database-engineer.html>
2024 Latest geekcert PROFESSIONAL-CLOUD-DATABASE-ENGINEER PDF
and VCE dumps Download

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





QUESTION 1

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. Create a Cloud SQL for MySQL instance for your databases, and configure Datastream to stream your database changes to Cloud SQL. 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. Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.
- B. Create migration job using Database Migration Service. 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. Wait for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete. Update your application connections to the new instance.
- C. Create migration job using Database Migration Service. Set the migration job type to One-time, and perform this migration during a maintenance window. 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. Update your application connections to the new instance.
- D. Use the mysqldump utility to manually initiate a backup of MySQL during the application maintenance window. Move the files to Cloud Storage, and import each database into your Cloud SQL instance. Continue to dump each database until all the databases are migrated. Update your application connections to the new instance.

Correct Answer: B

<https://cloud.google.com/datastream/docs/overview>.

QUESTION 2

You are managing a Cloud SQL for PostgreSQL instance in Google Cloud. You have a primary instance in region 1 and a read replica in region 2. After a failure of region 1, you need to make the Cloud SQL instance available again. You want to minimize data loss and follow Google-recommended practices. What should you do?

- A. Restore the Cloud SQL instance from the automatic backups in region 3.
- B. Restore the Cloud SQL instance from the automatic backups in another zone in region 1.
- C. Check "Lag Bytes" in the monitoring dashboard for the primary instance in the read replica instance. Check the replication status using `pg_catalog.pg_last_wal_receive_lsn()`. Then, fail over to region 2 by promoting the read replica instance.
- D. Check your instance operational log for the automatic failover status. Look for time, type, and status of the operations. If the failover operation is successful, no action is necessary. Otherwise, manually perform `gcloud sql instances failover`.

Correct Answer: C

https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#disaster_recovery



QUESTION 3

Your organization deployed a new version of a critical application that uses Cloud SQL for MySQL with high availability (HA) and binary logging enabled to store transactional information. The latest release of the application had an error that caused massive data corruption in your Cloud SQL for MySQL database. You need to minimize data loss. What should you do?

- A. Open the Google Cloud Console, navigate to SQL > Backups, and select the last version of the automated backup before the corruption.
- B. Reload the Cloud SQL for MySQL database using the LOAD DATA command to load data from CSV files that were used to initialize the instance.
- C. Perform a point-in-time recovery of your Cloud SQL for MySQL database, selecting a date and time before the data was corrupted.
- D. Fail over to the Cloud SQL for MySQL HA instance. Use that instance to recover the transactions that occurred before the corruption.

Correct Answer: C

Binary Logging enabled, with that you can identify the point of time the data was good and recover from that point time.
https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#perform_the_point-in-time_recovery_using_binary_log_positions

QUESTION 4

You are the primary DBA of a Cloud SQL for PostgreSQL database that supports 6 enterprise applications in production. You used Cloud SQL Insights to identify inefficient queries and now need to identify the application that is originating the inefficient queries. You want to follow Google-recommended practices. What should you do?

- A. Shut down and restart each application.
- B. Write a utility to scan database query logs.
- C. Write a utility to scan application logs.
- D. Use query tags to add application-centric database monitoring.

Correct Answer: D

https://cloud.google.com/sql/docs/postgres/using-query-insights#filter_by_query_tags

QUESTION 5

You are designing a physician portal app in Node.js. This application will be used in hospitals and clinics that might have intermittent internet connectivity. If a connectivity failure occurs, the app should be able to query the cached data. You need to ensure that the application has scalability, strong consistency, and multi-region replication. What should you do?

- A. Use Firestore and ensure that the PersistenceEnabled option is set to true.
- B. Use Memorystore for Memcached.



C. Use Pub/Sub to synchronize the changes from the application to Cloud Spanner.

D. Use Table.read with the exactStaleness option to perform a read of rows in Cloud Spanner.

Correct Answer: A

<https://firebase.google.com/docs/firestore/manage-data/enable-offline>

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER PDF Dumps](#)

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER VCE Dumps](#)

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam Questions](#)