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

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam Questions |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

1 / 4

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam Questions |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

2 / 4

QUESTION 1

You are designing a database architecture for a global application that stores information about public parks worldwide. The application uses the database for read-only purposes, and a centralized batch job updates the database nightly. You want to select an open source, SQL-compliant database. What should you do?

A. Use Bigtable with multi-region clusters.

B. Use Memorystore for Redis with multi-zones within a region.

C. Use Cloud SQL for PostgreSQL with cross-region replicas.

D. Use Cloud Spanner with multi-region configuration.

Correct Answer: C

QUESTION 2

You are configuring a brand new Cloud SQL for PostgreSQL database instance in Google Cloud. Your application team wants you to deploy one primary instance, one standby instance, and one read replica instance. You need to ensure that you are following Google- recommended practices for high availability. What should you do?

A. Configure the primary instance in zone A, the standby instance in zone C, and the read replica in zone B, all in the same region.

B. Configure the primary and standby instances in zone A and the read replica in zone B, all in the same region.

C. Configure the primary instance in one region, the standby instance in a second region, and the read replica in a third region.

D. Configure the primary, standby, and read replica instances in zone A, all in the same region.

Correct Answer: B

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: 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.

Correct Answer: C

QUESTION 5

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?

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

Correct Answer: B

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps | PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam Questions | PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

4 / 4