



# PROFESSIONAL-CLOUD-DATABASE-ENGINEER<sup>Q&As</sup>

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.passapply.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

PassApply.com

<https://www.passapply.com/professional-cloud-database-engineer.html>  
2024 Latest passapply 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

You need to migrate a 1 TB PostgreSQL database from a Compute Engine VM to Cloud SQL for PostgreSQL. You want to ensure that there is minimal downtime during the migration. What should you do?

- A. Export the data from the existing database, and load the data into a new Cloud SQL database.
- B. Use Migrate for Compute Engine to complete the migration.
- C. Use Datastream to complete the migration.
- D. Use Database Migration Service to complete the migration.

Correct Answer: A

---

### QUESTION 2

You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SLA. You want to follow Google- recommended practices. You need the database to be performant and highly available while you keep costs low. What should you do?

- A. Manually scale down the number of nodes after the peak period has passed.
- B. Use interleaving to co-locate parent and child rows.
- C. Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.
- D. Use granular instance sizing in Cloud Spanner and Autoscaler.

Correct Answer: C

---

### QUESTION 3

You are the DBA of an online tutoring application that runs on a Cloud SQL for PostgreSQL database. You are testing the implementation of the cross-regional failover configuration. The database in region R1 fails over successfully to region R2, and the database becomes available for the application to process data. During testing, certain scenarios of the application work as expected in region R2, but a few scenarios fail with database errors. The application-related database queries, when executed in isolation from Cloud SQL for PostgreSQL in region R2, work as expected. The application performs completely as expected when the database fails back to region R1. You need to identify the cause of the database errors in region R2. What should you do?

- A. Determine whether the versions of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- B. Determine whether the database patches of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- C. Determine whether the failover of Cloud SQL for PostgreSQL from region R1 to region R2 is in progress or has completed successfully.
- D. Determine whether Cloud SQL for PostgreSQL in region R2 is a near-real-time copy of region R1 but not an exact



copy.

Correct Answer: B

---

#### QUESTION 4

Your application follows a microservices architecture and uses a single large Cloud SQL instance, which is starting to have performance issues as your application grows. In the Cloud Monitoring dashboard, the CPU utilization looks normal. You want to follow Google-recommended practices to resolve and prevent these performance issues while avoiding any major refactoring. What should you do?

- A. Use Cloud Spanner instead of Cloud SQL.
- B. Increase the number of CPUs for your instance.
- C. Increase the storage size for the instance.
- D. Use many smaller Cloud SQL instances.

Correct Answer: A

---

#### QUESTION 5

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

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

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER Study Guide](#)

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