



# 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.pass4itsure.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

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-database-engineer.html>

2024 Latest pass4itsure 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 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 an external read replica to migrate the databases to Cloud SQL.
- B. Use a read replica to migrate the databases to Cloud SQL.
- C. Use Database Migration Service to migrate the databases to Cloud SQL.
- D. Use a cross-region read replica to migrate the databases to Cloud SQL.
- E. Use replication from an external server to migrate the databases to Cloud SQL.

Correct Answer: CE

---

**QUESTION 2**

Your company wants to migrate its MySQL, PostgreSQL, and Microsoft SQL Server on-premises databases to Google Cloud. You need a solution that provides near-zero downtime, requires no application changes, and supports change data

capture (CDC).

What should you do?

- A. Use the native export and import functionality of the source database.
- B. Create a database on Google Cloud, and use database links to perform the migration.
- C. Create a database on Google Cloud, and use Dataflow for database migration.
- D. Use Database Migration Service.

Correct Answer: D

Simplify migrations to the cloud. Available now for MySQL and PostgreSQL, with SQL Server and Oracle migrations in preview.

1.

Migrate to Cloud SQL and AlloyDB for PostgreSQL from on-premises, Google Cloud, or other clouds

2.



Replicate data continuously for minimal downtime migrations

3.

Serverless and easy to set up

---

### QUESTION 3

You host an application in Google Cloud. The application is located in a single region and uses Cloud SQL for transactional data. Most of your users are located in the same time zone and expect the application to be available 7 days a week, from 6 AM to 10 PM. You want to ensure regular maintenance updates to your Cloud SQL instance without creating downtime for your users. What should you do?

- A. Configure a maintenance window during a period when no users will be on the system. Control the order of update by setting non-production instances to earlier and production instances to later.
- B. Create your database with one primary node and one read replica in the region.
- C. Enable maintenance notifications for users, and reschedule maintenance activities to a specific time after notifications have been sent.
- D. Configure your Cloud SQL instance with high availability enabled.

Correct Answer: A

Configure a maintenance window during a period when no users will be on the system. Control the order of update by setting non-production instances to earlier and production instances to later.

---

### QUESTION 4

Your customer has a global chat application that uses a multi-regional Cloud Spanner instance. The application has recently experienced degraded performance after a new version of the application was launched. Your customer asked you for assistance. During initial troubleshooting, you observed high read latency. What should you do?

- A. Use query parameters to speed up frequently executed queries.
- B. Change the Cloud Spanner configuration from multi-region to single region.
- C. Use SQL statements to analyze SPANNER\_SYS.READ\_STATS\* tables.
- D. Use SQL statements to analyze SPANNER\_SYS.QUERY\_STATS\* tables.

Correct Answer: C

To troubleshoot high read latency, you can use SQL statements to analyze the SPANNER\_SYS.READ\_STATS\* tables. These tables contain statistics about read operations in Cloud Spanner, including the number of reads, read latency, and the number of read errors. By analyzing these tables, you can identify the cause of the high read latency and take appropriate action to resolve the issue. Other options, such as using query parameters to speed up frequently executed queries or changing the Cloud Spanner configuration from multi-region to single region, may not be directly related to the issue of high read latency. Similarly, analyzing the SPANNER\_SYS.QUERY\_STATS\* tables, which contain statistics about query operations, may not be relevant to the issue of high read latency.

---

**QUESTION 5**

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?

- 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.

Correct Answer: C

<https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile> The question states that a single cluster existed first, then a second cluster was added. Google's documentation states, "if you created the instance with one cluster, the default app profile uses single-cluster routing. This ensures that adding additional clusters later does not change the behavior of your existing applications". Simply adding a second cluster does not change the default profile from single-cluster routing to multi-cluster routing. Since you need multi-cluster routing, you're going to need a custom app profile. So C is correct. <https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile>

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

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

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps](#)