



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.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 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.

Correct Answer: A

<https://firebase.google.com/docs/firestore>

QUESTION 2

You are managing a Cloud SQL for MySQL environment in Google Cloud. You have deployed a primary instance in Zone A and a read replica instance in Zone B, both in the same region. You are notified that the replica instance in Zone B was unavailable for 10 minutes. You need to ensure that the read replica instance is still working. What should you do?

- A. Use the Google Cloud Console or gcloud CLI to manually create a new clone database.
- B. Use the Google Cloud Console or gcloud CLI to manually create a new failover replica from backup.
- C. Verify that the new replica is created automatically.
- D. Start the original primary instance and resume replication.

Correct Answer: C

Recovery Process: Once Zone-B becomes available again, Cloud SQL will initiate the recovery process for the impacted read replica. The recovery process involves the following steps: 1. Synchronization: Cloud SQL will compare the data in the recovered read replica with the primary instance in Zone-A. If there is any data divergence due to the unavailability period, Cloud SQL will synchronize the read replica with the primary instance to ensure data consistency. 2. Catch-up Replication: The recovered read replica will start catching up on the changes that occurred on the primary instance during its unavailability. It will apply the necessary updates from the primary instance's binary logs (binlogs) to bring the replica up to date. 3. Resuming Read Traffic: Once the synchronization and catch-up replication processes are complete, the read replica in Zone-B will resume its normal operation. It will be able to serve read traffic and stay updated with subsequent changes from the primary instance.

QUESTION 3

Your team uses thousands of connected IoT devices to collect device maintenance data for your oil and gas customers in real time. You want to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by these devices. You need a managed solution that is highly scalable, supports a



multi-cloud strategy, and offers low latency for these IoT devices. What should you do?

- A. Use Firestore with Looker.
- B. Use Cloud Spanner with Data Studio.
- C. Use MongoDB Atlas with Charts.
- D. Use Bigtable with Looker.

Correct Answer: C

This scenario has BigTable written all over it -large amounts of data from many devices to be analysed in realtime. I would even argue it could qualify as a multicloud solution, given the links to HBASE. BUT it does not support SQL queries and is not therefore compatible (on its own) with Looker. Firestore + Looker has the same problem. Spanner + Data Studio is at least a compatible pairing, but I agree with others that it doesn't fit this use-case -not least because it's Google-native. By contrast, MongoDB Atlas is a managed solution (just not by Google) which is compatible with the proposed reporting tool (Mongo's own Charts), it's specifically designed for this type of solution and of course it can run on any cloud.

QUESTION 4

Your customer is running a MySQL database on-premises with read replicas. The nightly incremental backups are expensive and add maintenance overhead. You want to follow Google-recommended practices to migrate the database to Google Cloud, and you need to ensure minimal downtime. What should you do?

- A. Create a Google Kubernetes Engine (GKE) cluster, install MySQL on the cluster, and then import the dump file.
- B. Use the mysqldump utility to take a backup of the existing on-premises database, and then import it into Cloud SQL.
- C. Create a Compute Engine VM, install MySQL on the VM, and then import the dump file.
- D. Create an external replica, and use Cloud SQL to synchronize the data to the replica.

Correct Answer: D

<https://cloud.google.com/sql/docs/mysql/replication/configure-replication-from-external>

QUESTION 5

Your team is building an application that stores and analyzes streaming time series financial data. You need a database solution that can perform time series-based scans with sub-second latency. The solution must scale into the hundreds of terabytes and be able to write up to 10k records per second and read up to 200 MB per second. What should you do?

- A. Use Firestore.
- B. Use Bigtable
- C. Use BigQuery.
- D. Use Cloud Spanner.



Correct Answer: B

Financial data, such as transaction histories, stock prices, and currency exchange rates.

<https://cloud.google.com/bigtable/docs/overview#what-its-good-for>

With SSD: Reads -up to 10,000 rows per second Writes -up to 10,000 rows per second Scans -up to 220 MB/s

<https://cloud.google.com/bigtable/docs/performance#typical-workloads>

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

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

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