



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.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 organization has strict policies on tracking rollouts to production and periodically shares this information with external auditors to meet compliance requirements. You need to enable auditing on several Cloud Spanner databases. What should you do?

- A. Use replication to roll out changes to higher environments.
- B. Use backup and restore to roll out changes to higher environments.
- C. Use Liquibase to roll out changes to higher environments.
- D. Manually capture detailed DBA audit logs when changes are rolled out to higher environments.

Correct Answer: C

To satisfy audit reporting you would need a way to record what was changed and when. The best answer is one which uses some kind of source code control system (SCCS). That rules out A and B. Any mention of anything manual in a cloud environment should look suspicious, which leave option C. As it happens, Liquibase is an SCCS and can be integrated with Spanner. <https://cloud.google.com/spanner/docs/use-liquibase>

QUESTION 2

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 3

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 4

Your organization has hundreds of Cloud SQL for MySQL instances. You want to follow Google-recommended practices to optimize platform costs. What should you do?

- A. Use Query Insights to identify idle instances.
- B. Remove inactive user accounts.
- C. Run the Recommender API to identify overprovisioned instances.
- D. Build indexes on heavily accessed tables.

Correct Answer: C

The Cloud SQL overprovisioned instance recommender helps you detect instances that are unnecessarily large for a given workload. It then provides recommendations on how to resize such instances and reduce cost. This page describes how this recommender works and how to use it.

QUESTION 5

You are using Compute Engine on Google Cloud and your data center to manage a set of MySQL databases in a hybrid configuration. You need to create replicas to scale reads and to offload part of the management operation. What should you do?

- A. Use external server replication.
- B. Use Data Migration Service.
- C. Use Cloud SQL for MySQL external replica.
- D. Use the mysqldump utility and binary logs.

Correct Answer: C

An external replica is a method that allows you to create a read-only copy of your Cloud SQL instance on an external server, such as a Compute Engine instance or an on-premises database server¹. An external replica can help you scale reads and offload management operations from your data center to Google Cloud. You can also use an external replica for disaster recovery, migration, or reporting purposes¹. To create an external replica, you need to configure a Cloud



SQL instance that replicates to one or more replicas external to Cloud SQL, and a source representation instance that represents the source database server in Cloud SQL1. You also need to enable access on the Cloud SQL instance for the IP address of the external replica, create a replication user, and export and import the data from the source database server to the external replica1.

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

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

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