

DUMPS ARENA

Google Cloud Certified - Professional Cloud Database Engineer

Google Professional-Cloud-Database-Engineer

Version Demo

Total Demo Questions: 10

Total Premium Questions: 120

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

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.

ANSWER: C**QUESTION NO: 2**

Your digital-native business runs its database workloads on Cloud SQL. Your website must be globally accessible 24/7. You need to prepare your Cloud SQL instance for high availability (HA). You want to follow Google-recommended practices. What should you do? (Choose two.)

- A. Set up manual backups.
- B. Create a PostgreSQL database on-premises as the HA option.
- C. Configure single zone availability for automated backups.
- D. Enable point-in-time recovery.
- E. Schedule automated backups.

ANSWER: B D**QUESTION NO: 3**

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

E. Use a read replica to migrate the databases to Cloud SQL.

ANSWER: C E

QUESTION NO: 4

You use Python scripts to generate weekly SQL reports to assess the state of your databases and determine whether you need to reorganize tables or run statistics. You want to automate this report but need to minimize operational costs and overhead. What should you do?

- A. Create a VM in Compute Engine, and run a cron job.
- B. Create a Cloud Composer instance, and create a directed acyclic graph (DAG).
- C. Create a Cloud Function, and call the Cloud Function using Cloud Scheduler.
- D. Create a Cloud Function, and call the Cloud Function from a Cloud Tasks queue.

ANSWER: C

QUESTION NO: 5

Your organization has a ticketing system that needs an online marketing analytics and reporting application. You need to select a relational database that can manage hundreds of terabytes of data to support this new application. Which database should you use?

- A. Cloud SQL
- B. BigQuery
- C. Cloud Spanner
- D. Bigtable

ANSWER: B

QUESTION NO: 6

Your application uses Cloud SQL for MySQL. Your users run reports on data that relies on near-real time; however, the additional analytics caused excessive load on the primary database. You created a read replica for the analytics workloads, but now your users are complaining about the lag in data changes and that their reports are still slow. You need to improve the report performance and shorten the lag in data replication without making changes to the current reports. Which two approaches should you implement? (Choose two.)

- A. Create secondary indexes on the replica.
- B. Create additional read replicas, and partition your analytics users to use different read replicas.

- C. Disable replication on the read replica, and set the flag for parallel replication on the read replica. Re-enable replication and optimize performance by setting flags on the primary instance.
- D. Disable replication on the primary instance, and set the flag for parallel replication on the primary instance. Re-enable replication and optimize performance by setting flags on the read replica.
- E. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

ANSWER: B E

QUESTION NO: 7

You released a popular mobile game and are using a 50 TB Cloud Spanner instance to store game data in a PITR-enabled production environment. When you analyzed the game statistics, you realized that some players are exploiting a loophole to gather more points to get on the leaderboard. Another DBA accidentally ran an emergency bugfix script that corrupted some of the data in the production environment. You need to determine the extent of the data corruption and restore the production environment. What should you do? (Choose two.)

- A. If the corruption is significant, use backup and restore, and specify a recovery timestamp.
- B. If the corruption is significant, perform a stale read and specify a recovery timestamp. Write the results back.
- C. If the corruption is significant, use import and export.
- D. If the corruption is insignificant, use backup and restore, and specify a recovery timestamp.
- E. If the corruption is insignificant, perform a stale read and specify a recovery timestamp. Write the results back.

ANSWER: B E

QUESTION NO: 8

You plan to use Database Migration Service to migrate data from a PostgreSQL on-premises instance to Cloud SQL. You need to identify the prerequisites for creating and automating the task. What should you do? (Choose two.)

- A. Drop or disable all users except database administration users.
- B. Disable all foreign key constraints on the source PostgreSQL database.
- C. Ensure that all PostgreSQL tables have a primary key.
- D. Shut down the database before the Data Migration Service task is started.
- E. Ensure that pglogical is installed on the source PostgreSQL database.

ANSWER: B E

QUESTION NO: 9

You are configuring a new application that has access to an existing Cloud Spanner database. The new application reads from this database to gather statistics for a dashboard. You want to follow Google-recommended practices when granting Identity and Access Management (IAM) permissions. What should you do?

- A. Reuse the existing service account that populates this database.
- B. Create a new service account, and grant it the Cloud Spanner Database Admin role.
- C. Create a new service account, and grant it the Cloud Spanner Database Reader role.
- D. Create a new service account, and grant it the spanner.databases.select permission.

ANSWER: B

QUESTION NO: 10

Your organization has an existing app that just went viral. The app uses a Cloud SQL for MySQL backend database that is experiencing slow disk performance while using hard disk drives (HDDs). You need to improve performance and reduce disk I/O wait times. What should you do?

- A. Export the data from the existing instance, and import the data into a new instance with solid-state drives (SSDs).
- B. Edit the instance to change the storage type from HDD to SSD.
- C. Create a high availability (HA) failover instance with SSDs, and perform a failover to the new instance.
- D. Create a read replica of the instance with SSDs, and perform a failover to the new instance

ANSWER: C