DUMPSQUEEN

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

dumpsqueen.com

QUESTION NO: 1

Your company uses Cloud Spanner for a mission-critical inventory management system that is globally available. You recently loaded stock keeping unit (SKU) and product catalog data from a company acquisition and observed hot-spots in the Cloud Spanner database. You want to follow Google-recommended schema design practices to avoid performance degradation. What should you do? (Choose two.)

A. Use an auto-incrementing value as the primary key.

B. Normalize the data model.

C. Promote low-cardinality attributes in multi-attribute primary keys.

D. Promote high-cardinality attributes in multi-attribute primary keys.

E. Use bit-reverse sequential value as the primary key.

ANSWER: A D

QUESTION NO: 2

You want to migrate your PostgreSQL database from another cloud provider to Cloud SQL. You plan on using Database Migration Service and need to assess the impact of any known limitations. What should you do? (Choose two.)

A. Identify whether the database has over 512 tables.

B. Identify all tables that do not have a primary key.

C. Identity all tables that do not have at least one foreign key.

D. Identify whether the source database is encrypted using pgcrypto extension.

E. Identify whether the source database uses customer-managed encryption keys (CMEK).

ANSWER: C E

QUESTION NO: 3

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.

ANSWER: B

QUESTION NO: 4

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: 5

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.

ANSWER: A

QUESTION NO: 6

You want to migrate an on-premises 100 TB Microsoft SQL Server database to Google Cloud over a 1 Gbps network link. You have 48 hours allowed downtime to migrate this database. What should you do? (Choose two.)

A. Use a change data capture (CDC) migration strategy.

B. Move the physical database servers from on-premises to Google Cloud.

C. Keep the network bandwidth at 1 Gbps, and then perform an offline data migration.

D. Increase the network bandwidth to 2 Gbps, and then perform an offline data migration.

E. Increase the network bandwidth to 10 Gbps, and then perform an offline data migration.

ANSWER: A D

QUESTION NO: 7

You are evaluating Cloud SQL for PostgreSQL as a possible destination for your on-premises PostgreSQL instances. Geography is becoming increasingly relevant to customer privacy worldwide. Your solution must support data residency requirements and include a strategy to:

configure where data is stored

control where the encryption keys are stored

govern the access to data

What should you do?

A. Replicate Cloud SQL databases across different zones.

B. Create a Cloud SQL for PostgreSQL instance on Google Cloud for the data that does not need to adhere to data residency requirements. Keep the data that must adhere to data residency requirements on-premises. Make application changes to support both databases.

C. Allow application access to data only if the users are in the same region as the Google Cloud region for the Cloud SQL for PostgreSQL database.

D. Use features like customer-managed encryption keys (CMEK), VPC Service Controls, and Identity and Access Management (IAM) policies.

ANSWER: C

QUESTION NO: 8

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.

ANSWER: C

QUESTION NO: 9

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: 10

You have a large Cloud SQL for PostgreSQL instance. The database instance is not mission-critical, and you want to minimize operational costs. What should you do to lower the cost of backups in this environment?

A. Set the automated backups to occur every other day to lower the frequency of backups.

B. Change the storage tier of the automated backups from solid-state drive (SSD) to hard disk drive (HDD).

C. Select a different region to store your backups.

D. Reduce the number of automated backups that are retained to two (2).

ANSWER: A