

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.pass2lead.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

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

An analytics team needs to read data out of Cloud SQL for SQL Server and update a table in Cloud Spanner. You need to create a service account and grant least privilege access using predefined roles. What roles should you assign to the service account?

- A. roles/cloudsql.viewer and roles/spanner.databaseUser
- B. roles/cloudsql.editor and roles/spanner.admin
- C. roles/cloudsql.client and roles/spanner.databaseReader
- D. roles/cloudsql.instanceUser and roles/spanner.databaseUser

Correct Answer: C

QUESTION 2

You are the database administrator of a Cloud SQL for PostgreSQL instance that has pgaudit disabled. Users are complaining that their queries are taking longer to execute and performance has degraded over the past few months. You need to collect and analyze query performance data to help identify slow-running queries. What should you do?

- A. View Cloud SQL operations to view historical query information.
- B. Write a Logs Explorer query to identify database queries with high execution times.
- C. Review application logs to identify database calls.
- D. Use the Query Insights dashboard to identify high execution times.

Correct Answer: D

QUESTION 3

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: B

QUESTION 4

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. Identify 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).

Correct Answer: CE

QUESTION 5

You recently launched a new product to the US market. You currently have two Bigtable clusters in one US region to serve all the traffic. Your marketing team is planning an immediate expansion to APAC. You need to roll out the regional expansion while implementing high availability according to Google-recommended practices. What should you do?

- A. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone europe-west1-d cluster-c in zone asia-east1-b
- B. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central1-b cluster-c in zone us-east1-a
- C. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone australia-southeast1-a cluster-c in zone europe-west1-d cluster-d in zone asia-east1-b
- D. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central2-a cluster-c in zone asia-northeast1-b cluster-d in zone asia-east1-b

Correct Answer: D

QUESTION 6

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: B

QUESTION 7

You manage a production MySQL database running on Cloud SQL at a retail company. You perform routine maintenance on Sunday at midnight when traffic is slow, but you want to skip routine maintenance during the year-end holiday shopping season. You need to ensure that your production system is available 24/7 during the holidays. What should you do?

- A. Define a maintenance window on Sundays between 12 AM and 1 AM, and deny maintenance periods between November 1 and January 15.
- B. Define a maintenance window on Sundays between 12 AM and 5 AM, and deny maintenance periods between November 1 and February 15.
- C. Build a Cloud Composer job to start a maintenance window on Sundays between 12 AM and 1AM, and deny maintenance periods between November 1 and January 15.
- D. Create a Cloud Scheduler job to start maintenance at 12 AM on Sundays. Pause the Cloud Scheduler job between November 1 and January 15.

Correct Answer: B

QUESTION 8

You work for a financial services company that wants to use fully managed database services. Traffic volume for your consumer services products has increased annually at a constant rate with occasional spikes around holidays. You frequently need to upgrade the capacity of your database. You want to use Cloud Spanner and include an automated method to increase your hardware capacity to support a higher level of concurrency. What should you do?

- A. Use linear scaling to implement the Autoscaler-based architecture
- B. Use direct scaling to implement the Autoscaler-based architecture.
- C. Upgrade the Cloud Spanner instance on a periodic basis during the scheduled maintenance window.
- D. Set up alerts that are triggered when Cloud Spanner utilization metrics breach the threshold, and then schedule an upgrade during the scheduled maintenance window.

Correct Answer: C

QUESTION 9

You are running a transactional application on Cloud SQL for PostgreSQL in Google Cloud. The database is running in a high availability configuration within one region. You have encountered issues with data and want to restore to the last known pristine version of the database. What should you do?

- A. Create a clone database from a read replica database, and restore the clone in the same region.
- B. Create a clone database from a read replica database, and restore the clone into a different zone.
- C. Use the Cloud SQL point-in-time recovery (PITR) feature. Restore the copy from two hours ago to a new database instance.
- D. Use the Cloud SQL database import feature. Import last week's dump file from Cloud Storage.

Correct Answer: C

QUESTION 10

Your ecommerce website captures user clickstream data to analyze customer traffic patterns in real time and support personalization features on your website. You plan to analyze this data using big data tools. You need a low-latency solution that can store 8 TB of data and can scale to millions of read and write requests per second. What should you do?

- A. Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis.
- B. Deploy a Cloud SQL environment with read replicas for improved performance. Use Datastream to export data to Cloud Storage and analyze with Dataproc and the Cloud Storage connector.
- C. Use Memorystore to handle your low-latency requirements and for real-time analytics.
- D. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

Correct Answer: B

QUESTION 11

You are designing for a write-heavy application. During testing, you discover that the write workloads are performant in a regional Cloud Spanner instance but slow down by an order of magnitude in a multi-regional instance. You want to make the write workloads faster in a multi-regional instance. What should you do?

- A. Place the bulk of the read and write workloads closer to the default leader region.
- B. Use staleness of at least 15 seconds.
- C. Add more read-write replicas.
- D. Keep the total CPU utilization under 45% in each region.

Correct Answer: C

QUESTION 12

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.

Correct Answer: AD

QUESTION 13

Your company has PostgreSQL databases on-premises and on Amazon Web Services (AWS). You are planning multiple database migrations to Cloud SQL in an effort to reduce costs and downtime. You want to follow Google-recommended practices and use Google native data migration tools. You also want to closely monitor the migrations as part of the cutover strategy. What should you do?

- A. Use Database Migration Service to migrate all databases to Cloud SQL.
- B. Use Database Migration Service for one-time migrations, and use third-party or partner tools for change data capture (CDC) style migrations.
- C. Use data replication tools and CDC tools to enable migration.
- D. Use a combination of Database Migration Service and partner tools to support the data migration strategy.

Correct Answer: B

QUESTION 14

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.

Correct Answer: BE

QUESTION 15

You are designing a physician portal app in Node.js. This application will be used in hospitals and clinics that might have intermittent internet connectivity. If a connectivity failure occurs, the app should be able to query the cached data. You need to ensure that the application has scalability, strong consistency, and multi-region replication. What should you do?

- A. Use Firestore and ensure that the PersistenceEnabled option is set to true.
- B. Use Memorystore for Memcached.
- C. Use Pub/Sub to synchronize the changes from the application to Cloud Spanner.

D. Use Table.read with the exactStaleness option to perform a read of rows in Cloud Spanner.

Correct Answer: C

[Latest PROFESSIONAL-CL
OUD-DATABASE-
ENGINEER Dumps](#)

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

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