

## May-2023 Realistic Professional-Cloud-Database-Engineer Exam Dumps with Accurate & Updated Questions [Q29-Q43]



# May-2023 Realistic Professional-Cloud-Database-Engineer Exam Dumps with Accurate & Updated Questions [Q29-Q43]

May-2023 Realistic Professional-Cloud-Database-Engineer Exam Dumps with Accurate & Updated Questions  
Professional-Cloud-Database-Engineer Exam Dumps - PDF Questions and Testing Engine

**NO.29** 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?

- \* roles/cloudsql.viewer and roles/spanner.databaseUser
- \* roles/cloudsql.editor and roles/spanner.admin
- \* roles/cloudsql.client and roles/spanner.databaseReader
- \* roles/cloudsql.instanceUser and roles/spanner.databaseUser

**NO.30** 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?

- \* Reuse the existing service account that populates this database.
- \* Create a new service account, and grant it the Cloud Spanner Database Admin role.
- \* Create a new service account, and grant it the Cloud Spanner Database Reader role.

- \* Create a new service account, and grant it the `spanner.databases.select` permission.

**NO.31** 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?

- \* Use Cloud Spanner instead of Cloud SQL.
- \* Increase the number of CPUs for your instance.
- \* Increase the storage size for the instance.
- \* Use many smaller Cloud SQL instances.

**NO.32** You are developing a new application on a VM that is on your corporate network. The application will use Java Database Connectivity (JDBC) to connect to Cloud SQL for PostgreSQL. Your Cloud SQL instance is configured with IP address 192.168.3.48, and SSL is disabled. You want to ensure that your application can access your database instance without requiring configuration changes to your database. What should you do?

- \* Define a connection string using your Google username and password to point to the external (public) IP address of your Cloud SQL instance.
- \* Define a connection string using a database username and password to point to the internal (private) IP address of your Cloud SQL instance.
- \* Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the internal (private) IP address of your Cloud SQL instance.
- \* Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the external (public) IP address of your Cloud SQL instance.

**NO.33** You are the DBA of an online tutoring application that runs on a Cloud SQL for PostgreSQL database. You are testing the implementation of the cross-regional failover configuration. The database in region R1 fails over successfully to region R2, and the database becomes available for the application to process data. During testing, certain scenarios of the application work as expected in region R2, but a few scenarios fail with database errors. The application-related database queries, when executed in isolation from Cloud SQL for PostgreSQL in region R2, work as expected. The application performs completely as expected when the database fails back to region R1. You need to identify the cause of the database errors in region R2. What should you do?

- \* Determine whether the versions of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- \* Determine whether the database patches of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- \* Determine whether the failover of Cloud SQL for PostgreSQL from region R1 to region R2 is in progress or has completed successfully.
- \* Determine whether Cloud SQL for PostgreSQL in region R2 is a near-real-time copy of region R1 but not an exact copy.

**NO.34** You need to perform a one-time migration of data from a running Cloud SQL for MySQL instance in the `us-central1` region to a new Cloud SQL for MySQL instance in the `us-east1` region. You want to follow Google-recommended practices to minimize performance impact on the currently running instance. What should you do?

- \* Create and run a Dataflow job that uses `JdbcIO` to copy data from one Cloud SQL instance to another.
- \* Create two Datastream connection profiles, and use them to create a stream from one Cloud SQL instance to another.
- \* Create a SQL dump file in Cloud Storage using a temporary instance, and then use that file to import into a new instance.
- \* Create a CSV file by running the SQL statement `SELECT * INTO OUTFILE`, copy the file to a Cloud Storage bucket, and import it into a new instance.

**NO.35** Your company wants you to migrate their Oracle, MySQL, Microsoft SQL Server, and PostgreSQL relational databases to Google Cloud. You need a fully managed, flexible database solution when possible. What should you do?

- \* Migrate all the databases to Cloud SQL.
- \* Migrate the Oracle, MySQL, and Microsoft SQL Server databases to Cloud SQL, and migrate the PostgreSQL databases to Compute Engine.

- \* Migrate the MySQL, Microsoft SQL Server, and PostgreSQL databases to Compute Engine, and migrate the Oracle databases to Bare Metal Solution for Oracle.
- \* Migrate the MySQL, Microsoft SQL Server, and PostgreSQL databases to Cloud SQL, and migrate the Oracle databases to Bare Metal Solution for Oracle.

**NO.36** Your organization is running a MySQL workload in Cloud SQL. Suddenly you see a degradation in database performance. You need to identify the root cause of the performance degradation. What should you do?

- \* Use Logs Explorer to analyze log data.
- \* Use Cloud Monitoring to monitor CPU, memory, and storage utilization metrics.
- \* Use Error Reporting to count, analyze, and aggregate the data.
- \* Use Cloud Debugger to inspect the state of an application.

**NO.37** 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?

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

**NO.38** Your company is migrating the existing infrastructure for a highly transactional application to Google Cloud. You have several databases in a MySQL database instance and need to decide how to transfer the data to Cloud SQL. You need to minimize the downtime for the migration of your 500 GB instance. What should you do?

- \* Create a Cloud SQL for MySQL instance for your databases, and configure Datastream to stream your database changes to Cloud SQL.

Select the Backfill historical data check box on your stream configuration to initiate Datastream to backfill any data that is out of sync between the source and destination.

Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.

- \* Create migration job using Database Migration Service.

Set the migration job type to Continuous, and allow the databases to complete the full dump phase and start sending data in change data capture (CDC) mode.

Wait for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete.

Update your application connections to the new instance.

- \* Create migration job using Database Migration Service.

Set the migration job type to One-time, and perform this migration during a maintenance window.

Stop all write workloads to the source database and initiate the dump. Wait for the dump to be loaded into the Cloud SQL destination database and the destination database to be promoted to the primary database.

Update your application connections to the new instance.

- \* Use the mysqldump utility to manually initiate a backup of MySQL during the application maintenance window.

Move the files to Cloud Storage, and import each database into your Cloud SQL instance.

Continue to dump each database until all the databases are migrated.

Update your application connections to the new instance.

**NO.39** 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.)

- \* Create secondary indexes on the replica.
- \* Create additional read replicas, and partition your analytics users to use different read replicas.
- \* 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.
- \* 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.
- \* Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

**NO.40** You are configuring a brand new PostgreSQL database instance in Cloud SQL. Your application team wants to have an optimal and highly available environment with automatic failover to avoid any unplanned outage. What should you do?

- \* Create one regional Cloud SQL instance with a read replica in another region.
- \* Create one regional Cloud SQL instance in one zone with a standby instance in another zone in the same region.
- \* Create two read-write Cloud SQL instances in two different zones with a standby instance in another region.
- \* Create two read-write Cloud SQL instances in two different regions with a standby instance in another zone.

**NO.41** 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?

- \* Use Firestore and ensure that the PersistenceEnabled option is set to true.
- \* Use Memorystore for Memcached.
- \* Use Pub/Sub to synchronize the changes from the application to Cloud Spanner.
- \* Use Table.read with the exactStaleness option to perform a read of rows in Cloud Spanner.

**NO.42** 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?

- \* Use Firestore.
- \* Use Bigtable
- \* Use BigQuery.
- \* Use Cloud Spanner.

**NO.43** You are setting up a Bare Metal Solution environment. You need to update the operating system to the latest version. You need to connect the Bare Metal Solution environment to the internet so you can receive software updates. What should you do?

- \* Setup a static external IP address in your VPC network.
- \* Set up bring your own IP (BYOIP) in your VPC.
- \* Set up a Cloud NAT gateway on the Compute Engine VM.
- \* Set up Cloud NAT service.

**Pass Google Professional-Cloud-Database-Engineer Exam Quickly With BraindumpsIT:**  
[https://www.braindumpsit.com/Professional-Cloud-Database-Engineer\\_real-exam.html](https://www.braindumpsit.com/Professional-Cloud-Database-Engineer_real-exam.html)