


[Jan-2023 Free Professional-Cloud-Database-Engineer Exam Questions Professional-Cloud-Database-Engineer Actual Free Exam Questions [Q20-Q43]



[Jan-2023] Free Professional-Cloud-Database-Engineer Exam Questions Professional-Cloud-Database-Engineer Actual Free Exam Questions [Q20-Q43]

[Jan-2023] Free Professional-Cloud-Database-Engineer Exam Questions Professional-Cloud-Database-Engineer Actual Free Exam Questions

Verified Professional-Cloud-Database-Engineer dumps and 120 unique questions

NO.20 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.21 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.)

- * Use a change data capture (CDC) migration strategy.
- * Move the physical database servers from on-premises to Google Cloud.

- * Keep the network bandwidth at 1 Gbps, and then perform an offline data migration.
- * Increase the network bandwidth to 2 Gbps, and then perform an offline data migration.
- * Increase the network bandwidth to 10 Gbps, and then perform an offline data migration.

NO.22 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.

NO.23 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.24 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?

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

NO.25 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?

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

NO.26 You are managing a Cloud SQL for MySQL environment in Google Cloud. You have deployed a primary instance in Zone A and a read replica instance in Zone B, both in the same region. You are notified that the replica instance in Zone B was unavailable for 10 minutes. You need to ensure that the read replica instance is still working. What should you do?

- * Use the Google Cloud Console or gcloud CLI to manually create a new clone database.
- * Use the Google Cloud Console or gcloud CLI to manually create a new failover replica from backup.
- * Verify that the new replica is created automatically.
- * Start the original primary instance and resume replication.

NO.27 During an internal audit, you realized that one of your Cloud SQL for MySQL instances does not have high availability (HA) enabled. You want to follow Google-recommended practices to enable HA on your existing instance. What should you do?

- * Create a new Cloud SQL for MySQL instance, enable HA, and use the export and import option to migrate your data.

- * Create a new Cloud SQL for MySQL instance, enable HA, and use Cloud Data Fusion to migrate your data.
- * Use the `gcloud instances patch` command to update your existing Cloud SQL for MySQL instance.
- * Shut down your existing Cloud SQL for MySQL instance, and enable HA.

NO.28 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?

- * Set the automated backups to occur every other day to lower the frequency of backups.
- * Change the storage tier of the automated backups from solid-state drive (SSD) to hard disk drive (HDD).
- * Select a different region to store your backups.
- * Reduce the number of automated backups that are retained to two (2).

NO.29 Your company is developing a new global transactional application that must be ACID-compliant and have 99.999% availability. You are responsible for selecting the appropriate Google Cloud database to serve as a datastore for this new application. What should you do?

- * Use Firestore.
- * Use Cloud Spanner.
- * Use Cloud SQL.
- * Use Bigtable.

NO.30 You are the primary DBA of a Cloud SQL for PostgreSQL database that supports 6 enterprise applications in production. You used Cloud SQL Insights to identify inefficient queries and now need to identify the application that is originating the inefficient queries. You want to follow Google-recommended practices. What should you do?

- * Shut down and restart each application.
- * Write a utility to scan database query logs.
- * Write a utility to scan application logs.
- * Use query tags to add application-centric database monitoring.

NO.31 You are managing a small Cloud SQL instance for developers to do testing. The instance is not critical and has a recovery point objective (RPO) of several days. You want to minimize ongoing costs for this instance. What should you do?

- * Take no backups, and turn off transaction log retention.
- * Take one manual backup per day, and turn off transaction log retention.
- * Turn on automated backup, and turn off transaction log retention.
- * Turn on automated backup, and turn on transaction log retention.

NO.32 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.)

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

NO.33 You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SLA.

You want to follow Google-recommended practices. You need the database to be performant and highly available while you keep

costs low. What should you do?

- * Manually scale down the number of nodes after the peak period has passed.
- * Use interleaving to co-locate parent and child rows.
- * Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.
- * Use granular instance sizing in Cloud Spanner and Autoscaler.

NO.34 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?

- * Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis.
- * 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.
- * Use Memorystore to handle your low-latency requirements and for real-time analytics.
- * Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

NO.35 Your organization has a busy transactional Cloud SQL for MySQL instance. Your analytics team needs access to the data so they can build monthly sales reports. You need to provide data access to the analytics team without adversely affecting performance. What should you do?

- * Create a read replica of the database, provide the database IP address, username, and password to the analytics team, and grant read access to required tables to the team.
- * Create a read replica of the database, enable the `cloudsql.iam_authentication` flag on the replica, and grant read access to required tables to the analytics team.
- * Enable the `cloudsql.iam_authentication` flag on the primary database instance, and grant read access to required tables to the analytics team.
- * Provide the database IP address, username, and password of the primary database instance to the analytics, team, and grant read access to required tables to the team.

NO.36 Your organization stores marketing data such as customer preferences and purchase history on Bigtable. The consumers of this database are predominantly data analysts and operations users. You receive a service ticket from the database operations department citing poor database performance between 9 AM-10 AM every day. The application team has confirmed no latency from their logs. A new cohort of pilot users that is testing a dataset loaded from a third-party data provider is experiencing poor database performance. Other users are not affected. You need to troubleshoot the issue. What should you do?

- * Isolate the data analysts and operations user groups to use different Bigtable instances.
- * Check the Cloud Monitoring table/bytes_used metric from Bigtable.
- * Use Key Visualizer for Bigtable.
- * Add more nodes to the Bigtable cluster.

NO.37 You are troubleshooting a connection issue with a newly deployed Cloud SQL instance on Google Cloud. While investigating the Cloud SQL Proxy logs, you see the message Error 403: Access Not Configured. What should you do?

- * Check the app.yaml value `cloud_sql_instances` for a misspelled or incorrect instance connection name.
- * Check whether your service account has `cloudsql.instances.connect` permission.
- * Enable the Cloud SQL Admin API.
- * Ensure that you are using an external (public) IP address interface.

NO.38 Your organization has a production Cloud SQL for MySQL instance. Your instance is configured with 16 vCPUs and 104 GB of RAM that is running between 90% and 100% CPU utilization for most of the day. You need to scale up the database and add vCPUs with minimal interruption and effort. What should you do?

- * Issue a `gcloud sql instances patch` command to increase the number of vCPUs.
- * Update a MySQL database flag to increase the number of vCPUs.
- * Issue a `gcloud compute instances update` command to increase the number of vCPUs.

- * Back up the database, create an instance with additional vCPUs, and restore the database.

NO.39 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.)

- * Use an auto-incrementing value as the primary key.
- * Normalize the data model.
- * Promote low-cardinality attributes in multi-attribute primary keys.
- * Promote high-cardinality attributes in multi-attribute primary keys.
- * Use bit-reverse sequential value as the primary key.

NO.40 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?

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

NO.41 You are building an application that allows users to customize their website and mobile experiences. The application will capture user information and preferences. User profiles have a dynamic schema, and users can add or delete information from their profile. You need to ensure that user changes automatically trigger updates to your downstream BigQuery data warehouse. What should you do?

- * Store your data in Bigtable, and use the user identifier as the key. Use one column family to store user profile data, and use another column family to store user preferences.
- * Use Cloud SQL, and create different tables for user profile data and user preferences from your recommendations model. Use SQL to join the user profile data and preferences
- * Use Firestore in Native mode, and store user profile data as a document. Update the user profile with preferences specific to that user and use the user identifier to query.
- * Use Firestore in Datastore mode, and store user profile data as a document. Update the user profile with preferences specific to that user and use the user identifier to query.

NO.42 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.43 Your team is building a new inventory management application that will require read and write database instances in multiple Google Cloud regions around the globe. Your database solution requires 99.99% availability and global transactional consistency. You need a fully managed backend relational database to store inventory changes. What should you do?

- * Use Bigtable.
- * Use Firestore.
- * Use Cloud SQL for MySQL
- * Use Cloud Spanner.

Latest 100% Passing Guarantee - Brilliant Professional-Cloud-Database-Engineer Exam Questions PDF:
<https://www.actualtests4sure.com/Professional-Cloud-Database-Engineer-test-questions.html>