

Regular Free Updates Professional-Cloud-Database-Engineer Dumps Real Exam Questions Test Engine Mar 13, 2023 [Q19-Q43]



Regular Free Updates Professional-Cloud-Database-Engineer Dumps Real Exam Questions Test Engine Mar 13, 2023 Practice Test Questions Verified Answers As Experienced in the Actual Test!

Google Professional-Cloud-Database-Engineer Exam Syllabus Topics:

TopicDetailsTopic 1- Apply concepts to implement highly scalable and available databases in Google Cloud- Given a scenario, define maintenance windows and notifications based on application availability requirementsTopic 2- Manage database users, including authentication and access- Continuously assess and optimize the cost of running a database solutionTopic 3- Differentiate between managed and unmanaged database services- Analyze the cost of running database solutions in Google CloudTopic 4- Provision high availability database solutions in Google Cloud- Design scalable and highly available cloud database solutionsTopic 5- Distinguish between SQL and NoSQL business requirements- Evaluate tradeoffs between multi-region, region, and zonal database deployment strategiesTopic 6- Determine the correct database migration tools for a given scenario- Size database compute and storage based on performance requirementsTopic 7- Deploy scalable and highly available databases in Google Cloud- Determine database connectivity and access management considerationsTopic 8- Plan and perform database migration, including fallback plans and schema conversion- Test high availability and disaster recovery strategies periodicallyTopic 9- Design for recovery time objective (RTO) and recovery point objective (RPO)- Assess slow running queries and database locking and identify missing indexesTopic 10- Justify the use of session pooler services- Given a scenario, perform solution sizing based on current environment workload metrics and future requirements

Topic 11- Automate database instance provisioning- Determine how applications will connect to the database

Q19. Your organization is running a critical production database on a virtual machine (VM) on Compute Engine. The VM has an ext4-formatted persistent disk for data files. The database will soon run out of storage space. You need to implement a solution that avoids downtime. What should you do?

- * In the Google Cloud Console, increase the size of the persistent disk, and use the `resize2fs` command to extend the disk.
- * In the Google Cloud Console, increase the size of the persistent disk, and use the `fdisk` command to verify that the new space is ready to use
- * In the Google Cloud Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.
- * In the Google Cloud Console, create a new persistent disk attached to the VM, and configure the database service to move the files to the new disk.

Q20. 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.

Q21. Your team is running a Cloud SQL for MySQL instance with a 5 TB database that must be available 24/7. You need to save database backups on object storage with minimal operational overhead or risk to your production workloads. What should you do?

- * Use Cloud SQL serverless exports.
- * Create a read replica, and then use the `mysqldump` utility to export each table.
- * Clone the Cloud SQL instance, and then use the `mysqldump` utility to export the data.
- * Use the `mysqldump` utility on the primary database instance to export the backup.

Q22. 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.)

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

Q23. You are choosing a database backend for a new application. The application will ingest data points from IoT sensors. You need to ensure that the application can scale up to millions of requests per second with sub-10ms latency and store up to 100 TB of history. What should you do?

- * Use Cloud SQL with read replicas for throughput.
- * Use Firestore, and rely on automatic serverless scaling.
- * Use Memorystore for Memcached, and add nodes as necessary to achieve the required throughput.
- * Use Bigtable, and add nodes as necessary to achieve the required throughput.

Q24. You are managing multiple applications connecting to a database on Cloud SQL for PostgreSQL. You need to be able to

monitor database performance to easily identify applications with long-running and resource-intensive queries. What should you do?

- * Use log messages produced by Cloud SQL.
- * Use Query Insights for Cloud SQL.
- * Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- * Use Cloud SQL instance monitoring in the Google Cloud Console.

Q25. 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.

Q26. Your retail organization is preparing for the holiday season. Use of catalog services is increasing, and your DevOps team is supporting the Cloud SQL databases that power a microservices-based application. The DevOps team has added instrumentation through Sqlcommenter. You need to identify the root cause of why certain microservice calls are failing. What should you do?

- * Watch Query Insights for long running queries.
- * Watch the Cloud SQL instance monitor for CPU utilization metrics.
- * Watch the Cloud SQL recommenders for overprovisioned instances.
- * Watch Cloud Trace for application requests that are failing.

Q27. Your organization is migrating 50 TB Oracle databases to Bare Metal Solution for Oracle. Database backups must be available for quick restore. You also need to have backups available for 5 years. You need to design a cost-effective architecture that meets a recovery time objective (RTO) of 2 hours and recovery point objective (RPO) of 15 minutes. What should you do?

- * Create the database on a Bare Metal Solution server with the database running on flash storage.

Keep a local backup copy on all flash storage.

Keep backups older than one day stored in Actifio OnVault storage.

- * Create the database on a Bare Metal Solution server with the database running on flash storage.

Keep a local backup copy on standard storage.

Keep backups older than one day stored in Actifio OnVault storage.

- * Create the database on a Bare Metal Solution server with the database running on flash storage.

Keep a local backup copy on standard storage.

Use the Oracle Recovery Manager (RMAN) backup utility to move backups older than one day to a Coldline Storage bucket.

- * Create the database on a Bare Metal Solution server with the database running on flash storage.

Keep a local backup copy on all flash storage.

Use the Oracle Recovery Manager (RMAN) backup utility to move backups older than one day to an Archive Storage bucket.

Q28. Your organization operates in a highly regulated industry. Separation of concerns (SoC) and security principle of least privilege (PoLP) are critical. The operations team consists of:

Person A is a database administrator.

Person B is an analyst who generates metric reports.

Application C is responsible for automatic backups.

You need to assign roles to team members for Cloud Spanner. Which roles should you assign?

* roles/spanner.databaseAdmin for Person A

roles/spanner.databaseReader for Person B

roles/spanner.backupWriter for Application C

* roles/spanner.databaseAdmin for Person A

roles/spanner.databaseReader for Person B

roles/spanner.backupAdmin for Application C

* roles/spanner.databaseAdmin for Person A

roles/spanner.databaseUser for Person B

roles/spanner.databaseReader for Application C

* roles/spanner.databaseAdmin for Person A

roles/spanner.databaseUser for Person B

roles/spanner.backupWriter for Application C

Q29. You are migrating a telehealth care company's on-premises data center to Google Cloud. The migration plan specifies:

PostgreSQL databases must be migrated to a multi-region backup configuration with cross-region replicas to allow restore and failover in multiple scenarios.

MySQL databases handle personally identifiable information (PII) and require data residency compliance at the regional level.

You want to set up the environment with minimal administrative effort. What should you do?

* Set up Cloud Logging and Cloud Monitoring with Cloud Functions to send an alert every time a new database instance is created, and manually validate the region.

* Set up different organizations for each database type, and apply policy constraints at the organization level.

* Set up Pub/Sub to ingest data from Cloud Logging, send an alert every time a new database instance is created, and manually validate the region.

* Set up different projects for PostgreSQL and MySQL databases, and apply organizational policy constraints at a project level.

Q30. You are designing a database strategy for a new web application. You plan to start with a small pilot in one country and eventually expand to millions of users in a global audience. You need to ensure that the application can run 24/7 with minimal downtime for maintenance. What should you do?

* Use Cloud Spanner in a regional configuration.

* Use Cloud Spanner in a multi-region configuration.

* Use Cloud SQL with cross-region replicas.

* Use highly available Cloud SQL with multiple zones.

Q31. 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.

Q32. You finished migrating an on-premises MySQL database to Cloud SQL. You want to ensure that the daily export of a table, which was previously a cron job running on the database server, continues. You want the solution to minimize cost and operations overhead. What should you do?

- * Use Cloud Scheduler and Cloud Functions to run the daily export.
- * Create a streaming Datatlow job to export the table.
- * Set up Cloud Composer, and create a task to export the table daily.
- * Run the cron job on a Compute Engine instance to continue the export.

Q33. 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.

Q34. You are deploying a new Cloud SQL instance on Google Cloud using the Cloud SQL Auth proxy. You have identified snippets of application code that need to access the new Cloud SQL instance. The snippets reside and execute on an application server running on a Compute Engine machine. You want to follow Google-recommended practices to set up Identity and Access Management (IAM) as quickly and securely as possible. What should you do?

- * For each application code, set up a common shared user account.
- * For each application code, set up a dedicated user account.
- * For the application server, set up a service account.
- * For the application server, set up a common shared user account.

Q35. You are writing an application that will run on Cloud Run and require a database running in the Cloud SQL managed service. You want to secure this instance so that it only receives connections from applications running in your VPC environment in Google Cloud. What should you do?

- * Create your instance with a specified external (public) IP address.

Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.

Use Cloud SQL Auth proxy to connect to the instance.

- * Create your instance with a specified external (public) IP address.

Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.

Connect to the instance using a connection pool to best manage connections to the instance.

- * Create your instance with a specified internal (private) IP address.

Choose the VPC with private service connection configured.

Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.

Use Cloud SQL Auth proxy to connect to the instance.

* Create your instance with a specified internal (private) IP address.

Choose the VPC with private service connection configured.

Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.

Connect to the instance using a connection pool to best manage connections to the instance.

Q36. You are migrating an on-premises application to Google Cloud. The application requires a high availability (HA) PostgreSQL database to support business-critical functions. Your company's disaster recovery strategy requires a recovery time objective (RTO) and recovery point objective (RPO) within 30 minutes of failure. You plan to use a Google Cloud managed service. What should you do to maximize uptime for your application?

* Deploy Cloud SQL for PostgreSQL in a regional configuration. Create a read replica in a different zone in the same region and a read replica in another region for disaster recovery.

* Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Take periodic backups, and use this backup to restore to a new Cloud SQL for PostgreSQL instance in another region during a disaster recovery event.

* Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Create a cross-region read replica, and promote the read replica as the primary node for disaster recovery.

* Migrate the PostgreSQL database to multi-regional Cloud Spanner so that a single region outage will not affect your application. Update the schema to support Cloud Spanner data types, and refactor the application.

Q37. You need to migrate existing databases from Microsoft SQL Server 2016 Standard Edition on a single Windows Server 2019 Datacenter Edition to a single Cloud SQL for SQL Server instance. During the discovery phase of your project, you notice that your on-premises server peaks at around 25,000 read IOPS. You need to ensure that your Cloud SQL instance is sized appropriately to maximize read performance. What should you do?

* Create a SQL Server 2019 Standard on Standard machine type with 4 vCPUs, 15 GB of RAM, and 800 GB of solid-state drive (SSD).

* Create a SQL Server 2019 Standard on High Memory machine type with at least 16 vCPUs, 104 GB of RAM, and 200 GB of SSD.

* Create a SQL Server 2019 Standard on High Memory machine type with 16 vCPUs, 104 GB of RAM, and 4 TB of SSD.

* Create a SQL Server 2019 Enterprise on High Memory machine type with 16 vCPUs, 104 GB of RAM, and 500 GB of SSD.

Q38. Your project is using Bigtable to store data that should not be accessed from the public internet under any circumstances, even if the requestor has a valid service account key. You need to secure access to this data. What should you do?

* Use Identity and Access Management (IAM) for Bigtable access control.

* Use VPC Service Controls to create a trusted network for the Bigtable service.

* Use customer-managed encryption keys (CMEK).

* Use Google Cloud Armor to add IP addresses to an allowlist.

Q39. You are managing a mission-critical Cloud SQL for PostgreSQL instance. Your application team is running important transactions on the database when another DBA starts an on-demand backup. You want to verify the status of the backup. What should you do?

* Check the `cloudsql.googleapis.com/postgres.log` instance log.

* Perform the `gcloud sql operations list` command.

* Use Cloud Audit Logs to verify the status.

* Use the Google Cloud Console.

Q40. 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.

Q41. Your organization needs to migrate a critical, on-premises MySQL database to Cloud SQL for MySQL. The on-premises database is on a version of MySQL that is supported by Cloud SQL and uses the InnoDB storage engine. You need to migrate the database while preserving transactions and minimizing downtime. What should you do?

- * Use Database Migration Service to connect to your on-premises database, and choose continuous replication.

After the on-premises database is migrated, promote the Cloud SQL for MySQL instance, and connect applications to your Cloud SQL instance.

- * Build a Cloud Data Fusion pipeline for each table to migrate data from the on-premises MySQL database to Cloud SQL for MySQL.

Schedule downtime to run each Cloud Data Fusion pipeline.

Verify that the migration was successful.

Re-point the applications to the Cloud SQL for MySQL instance.

- * Pause the on-premises applications.

Use the mysqldump utility to dump the database content in compressed format.

Run gsutil -m to move the dump file to Cloud Storage.

Use the Cloud SQL for MySQL import option.

After the import operation is complete, re-point the applications to the Cloud SQL for MySQL instance.

- * Pause the on-premises applications.

Use the mysqldump utility to dump the database content in CSV format.

Run gsutil -m to move the dump file to Cloud Storage.

Use the Cloud SQL for MySQL import option.

After the import operation is complete, re-point the applications to the Cloud SQL for MySQL instance.

Q42. You are designing a new gaming application that uses a highly transactional relational database to store player authentication and inventory data in Google Cloud. You want to launch the game in multiple regions. What should you do?

- * Use Cloud Spanner to deploy the database.
- * Use Bigtable with clusters in multiple regions to deploy the database
- * Use BigQuery to deploy the database
- * Use Cloud SQL with a regional read replica to deploy the database.

Q43. Your organization is running a low-latency reporting application on Microsoft SQL Server. In addition to the database engine,

you are using SQL Server Analysis Services (SSAS), SQL Server Reporting Services (SSRS), and SQL Server Integration Services (SSIS) in your on-premises environment. You want to migrate your Microsoft SQL Server database instances to Google Cloud. You need to ensure minimal disruption to the existing architecture during migration. What should you do?

- * Migrate to Cloud SQL for SQL Server.
- * Migrate to Cloud SQL for PostgreSQL.
- * Migrate to Compute Engine.
- * Migrate to Google Kubernetes Engine (GKE).

Pass Google Professional-Cloud-Database-Engineer Exam in First Attempt Easily:

<https://www.examlabs.com/Google/Google-Cloud-Certified/best-Professional-Cloud-Database-Engineer-exam-dumps.html>