

Professional-Cloud-Database-Engineer Dumps

Google Cloud Certified - Professional Cloud Database Engineer

<https://www.certleader.com/Professional-Cloud-Database-Engineer-dumps.html>



NEW QUESTION 1

Your organization has an existing app that just went viral. The app uses a Cloud SQL for MySQL backend database that is experiencing slow disk performance while using hard disk drives (HDDs). You need to improve performance and reduce disk I/O wait times. What should you do?

- A. Export the data from the existing instance, and import the data into a new instance with solid-state drives (SSDs).
- B. Edit the instance to change the storage type from HDD to SSD.
- C. Create a high availability (HA) failover instance with SSDs, and perform a failover to the new instance.
- D. Create a read replica of the instance with SSDs, and perform a failover to the new instance

Answer: A

Explanation:

<https://stackoverflow.com/questions/72034607/can-i-change-storage-type-from-hdd-to-ssd-on-cloud-sql-after-creating-an-instanc>

NEW QUESTION 2

Your organization works with sensitive data that requires you to manage your own encryption keys. You are working on a project that stores that data in a Cloud SQL database. You need to ensure that stored data is encrypted with your keys. What should you do?

- A. Export data periodically to a Cloud Storage bucket protected by Customer-Supplied Encryption Keys.
- B. Use Cloud SQL Auth proxy.
- C. Connect to Cloud SQL using a connection that has SSL encryption.
- D. Use customer-managed encryption keys with Cloud SQL.

Answer: D

NEW QUESTION 3

Your company is migrating all legacy applications to Google Cloud. All on-premises applications are using legacy Oracle 12c databases with Oracle Real Application Cluster (RAC) for high availability (HA) and Oracle Data Guard for disaster recovery. You need a solution that requires minimal code changes, provides the same high availability you have today on-premises, and supports a low latency network for migrated legacy applications. What should you do?

- A. Migrate the databases to Cloud Spanner.
- B. Migrate the databases to Cloud SQL, and enable a standby database.
- C. Migrate the databases to Compute Engine using regional persistent disks.
- D. Migrate the databases to Bare Metal Solution for Oracle.

Answer: D

Explanation:

BMS is the only Google database service which supports Oracle aside from GCVE. It allows you to use all native Oracle features including RAC. Since GCVE isn't mentioned, it has to be D - Bare Metal Solution.

NEW QUESTION 4

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: AE

Explanation:

https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online_versus_offline_transfer

NEW QUESTION 5

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?

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

Answer: A

Explanation:

Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL dialects: Google Standard SQL (ANSI 2011 with extensions) and PostgreSQL.

NEW QUESTION 6

You need to migrate a 1 TB PostgreSQL database from a Compute Engine VM to Cloud SQL for PostgreSQL. You want to ensure that there is minimal downtime during the migration. What should you do?

- A. Export the data from the existing database, and load the data into a new Cloud SQL database.
- B. Use Migrate for Compute Engine to complete the migration.
- C. Use Datastream to complete the migration.
- D. Use Database Migration Service to complete the migration.

Answer: D

Explanation:

<https://www.cloudskillsboost.google/focuses/22792?parent=catalog>

NEW QUESTION 7

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: CE

Explanation:

<https://cloud.google.com/database-migration/docs/postgres/faq>

NEW QUESTION 8

You are working on a new centralized inventory management system to track items available in 200 stores, which each have 500 GB of data. You are planning a gradual rollout of the system to a few stores each week. You need to design an SQL database architecture that minimizes costs and user disruption during each regional rollout and can scale up or down on nights and holidays. What should you do?

- A. Use Oracle Real Application Cluster (RAC) databases on Bare Metal Solution for Oracle.
- B. Use sharded Cloud SQL instances with one or more stores per database instance.
- C. Use a Bigtable cluster with autoscaling.
- D. Use Cloud Spanner with a custom autoscaling solution.

Answer: D

Explanation:

<https://cloud.google.com/spanner/docs/autoscaling-overview>

* 1. CloudSQL max out at 64TB, so unable to hold 100TB of data. https://cloud.google.com/sql/docs/quotas#metrics_collection_limit 2. Scale is done manually on SQL Cloud

NEW QUESTION 9

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 1 AM, and deny maintenance periods between November 1 and January 15.
- D. Create a Cloud Scheduler job to start maintenance at 12 AM on Sunday
- E. Pause the Cloud Scheduler job between November 1 and January 15.

Answer: A

Explanation:

"Deny maintenance period. A block of days in which Cloud SQL does not schedule maintenance. Deny maintenance periods can be up to 90 days long."
<https://cloud.google.com/sql/docs/mysql/maintenance>

NEW QUESTION 10

Your online delivery business that primarily serves retail customers uses Cloud SQL for MySQL for its inventory and scheduling application. The required recovery time objective (RTO) and recovery point objective (RPO) must be in minutes rather than hours as a part of your high availability and disaster recovery design. You need a high availability configuration that can recover without data loss during a zonal or a regional failure. What should you do?

- A. Set up all read replicas in a different region using asynchronous replication.
- B. Set up all read replicas in the same region as the primary instance with synchronous replication.
- C. Set up read replicas in different zones of the same region as the primary instance with synchronous replication, and set up read replicas in different regions with asynchronous replication.
- D. Set up read replicas in different zones of the same region as the primary instance with asynchronous replication, and set up read replicas in different regions with synchronous replication.

Answer: C

Explanation:

This answer meets the RTO and RPO requirements by using synchronous replication within the same region, which ensures that all writes made to the primary instance are replicated to disks in both zones before a transaction is reported as committed. This minimizes data loss and downtime in case of a zonal or an instance failure, and allows for a quick failover to the standby instance.

This answer also meets the high availability and disaster recovery requirements by using asynchronous replication across different regions, which ensures that the data changes made to the primary instance are replicated to the read replicas in other regions with minimal delay. This provides additional redundancy and backup in case of a regional failure, and allows for a manual failover to the read replica in another region.

NEW QUESTION 10

You manage a meeting booking application that uses Cloud SQL. During an important launch, the Cloud SQL instance went through a maintenance event that resulted in a downtime of more than 5 minutes and adversely affected your production application. You need to immediately address the maintenance issue to prevent any unplanned events in the future. What should you do?

- A. Set your production instance's maintenance window to non-business hours.
- B. Migrate the Cloud SQL instance to Cloud Spanner to avoid any future disruptions due to maintenance.
- C. Contact Support to understand why your Cloud SQL instance had a downtime of more than 5 minutes.
- D. Use Cloud Scheduler to schedule a maintenance window of no longer than 5 minutes.

Answer: A

NEW QUESTION 13

You are running an instance of Cloud Spanner as the backend of your ecommerce website. You learn that the quality assurance (QA) team has doubled the number of their test cases. You need to create a copy of your Cloud Spanner database in a new test environment to accommodate the additional test cases. You want to follow Google-recommended practices. What should you do?

- A. Use Cloud Functions to run the export in Avro format.
- B. Use Cloud Functions to run the export in text format.
- C. Use Dataflow to run the export in Avro format.
- D. Use Dataflow to run the export in text format.

Answer: C

Explanation:

<https://cloud.google.com/spanner/docs/import-export-overview#file-format>

NEW QUESTION 14

Your DevOps team is using Terraform to deploy applications and Cloud SQL databases. After every new application change is rolled out, the environment is torn down and recreated, and the persistent database layer is lost. You need to prevent the database from being dropped. What should you do?

- A. Set Terraform `deletion_protection` to true.
- B. Rerun terraform apply.
- C. Create a read replica.
- D. Use point-in-time-recovery (PITR) to recover the database.

Answer: A

Explanation:

From Google's documentation, "For stateful resources, such as databases, ensure that deletion protection is enabled. The syntax is: `lifecycle { prevent_destroy = true }`" <https://cloud.google.com/docs/terraform/best-practices-for-terraform#stateful-resources>

NEW QUESTION 15

You are managing a set of Cloud SQL databases in Google Cloud. Regulations require that database backups reside in the region where the database is created. You want to minimize operational costs and administrative effort. What should you do?

- A. Configure the automated backups to use a regional Cloud Storage bucket as a custom location.
- B. Use the default configuration for the automated backups location.
- C. Disable automated backups, and create an on-demand backup routine to a regional Cloud Storage bucket.
- D. Disable automated backups, and configure serverless exports to a regional Cloud Storage bucket.

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/mysql/backup-recovery/backing-up#locationbackups> You can use a custom location for on-demand and automatic backups. For a complete list of valid location values, see the Instance locations.

NEW QUESTION 19

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?

- A. Reuse the existing service account that populates this database.
- B. Create a new service account, and grant it the Cloud Spanner Database Admin role.
- C. Create a new service account, and grant it the Cloud Spanner Database Reader role.
- D. Create a new service account, and grant it the `spanner.databases.select` permission.

Answer: C

Explanation:

<https://cloud.google.com/iam/docs/overview>

NEW QUESTION 21

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.)

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

Answer: DE

Explanation:

D. Enable point-in-time recovery - This feature allows you to restore your database to a specific point in time. It helps protect against data loss and can be used in the event of data corruption or accidental data deletion. E. Schedule automated backups - Automated backups allow you to take regular backups of your database without manual intervention. You can use these backups to restore your database in the event of data loss or corruption.

NEW QUESTION 23

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.
- D. Restore the copy from two hours ago to a new database instance.
- E. Use the Cloud SQL database import feature.
- F. Import last week's dump file from Cloud Storage.

Answer: C

Explanation:

Using import/export from last week is slow for large scale databases and will restore database from last week.

NEW QUESTION 25

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?

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

Answer: A

Explanation:

<https://cloud.google.com/blog/topics/developers-practitioners/scheduling-cloud-sql-exports-using-cloud-functions-and-cloud-scheduler>

NEW QUESTION 30

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?

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

Answer: A

Explanation:

Cloud Trace doesn't support Cloud SQL. Eliminate D. Cloud SQL recommenders for overprovisioned instances would tell you about Cloud SQL instances which are too large for their workload. Eliminate C. Monitoring CPU utilization wouldn't tell you why microservice calls are failing. Eliminate B. SQLcommenter integrates with Query Insights. So A is the best answer. <https://cloud.google.com/blog/topics/developers-practitioners/introducing-sqlcommenter-open-source-orm-auto-instrumentation-library>

NEW QUESTION 31

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?

- A. Check the `cloudsql.googleapis.com/postgres.log` instance log.
- B. Perform the `gcloud sql operations list` command.
- C. Use Cloud Audit Logs to verify the status.
- D. Use the Google Cloud Console.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres/backup-recovery/backups#troubleshooting-backups> Under Troubleshooting: Issue: "You can't see the current operation's status." The Google Cloud console reports only success or failure when the operation is done. It isn't designed to show warnings or other updates. Run the `gcloud sql operations list` command to list all operations for the given Cloud SQL instance.

NEW QUESTION 32

Your company is migrating their MySQL database to Cloud SQL and cannot afford any planned downtime during the month of December. The company is also concerned with cost, so you need the most cost-effective solution. What should you do?

- A. Open a support ticket in Google Cloud to prevent any maintenance in that MySQL instance during the month of December.
- B. Use Cloud SQL maintenance settings to prevent any maintenance during the month of December.
- C. Create MySQL read replicas in different zones so that, if any downtime occurs, the read replicas will act as the primary instance during the month of December.
- D. Create a MySQL regional instance so that, if any downtime occurs, the standby instance will act as the primary instance during the month of December.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/maintenance?hl=fr>

NEW QUESTION 35

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
- C. 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

Answer: D

Explanation:

<https://cloud.google.com/bigtable/docs/replication-settings#regional-failover>

NEW QUESTION 38

You are using Compute Engine on Google Cloud and your data center to manage a set of MySQL databases in a hybrid configuration. You need to create replicas to scale reads and to offload part of the management operation. What should you do?

- A. Use external server replication.
- B. Use Data Migration Service.
- C. Use Cloud SQL for MySQL external replica.
- D. Use the `mysqldump` utility and binary logs.

Answer: C

Explanation:

An external replica is a method that allows you to create a read-only copy of your Cloud SQL instance on an external server, such as a Compute Engine instance or an on-premises database server¹. An external replica can help you scale reads and offload management operations from your data center to Google Cloud. You can also use an external replica for disaster recovery, migration, or reporting purposes¹.

To create an external replica, you need to configure a Cloud SQL instance that replicates to one or more replicas external to Cloud SQL, and a source representation instance that

represents the source database server in Cloud SQL¹. You also need to enable access on the Cloud SQL instance for the IP address of the external replica, create a replication user, and export and import the data from the source database server to the external replica¹.

NEW QUESTION 40

You are designing a payments processing application on Google Cloud. The application must continue to serve requests and avoid any user disruption if a regional failure occurs. You need to use AES-256 to encrypt data in the database, and you want to control where you store the encryption key. What should you do?

- A. Use Cloud Spanner with a customer-managed encryption key (CMEK).
- B. Use Cloud Spanner with default encryption.
- C. Use Cloud SQL with a customer-managed encryption key (CMEK).
- D. Use Bigtable with default encryption.

Answer: A

Explanation:

Yes default encryption comes with AES-256 but the question states that you need to control where you store the encryption keys. that can be achieved by CMEK.

NEW QUESTION 44

You are starting a large CSV import into a Cloud SQL for MySQL instance that has many open connections. You checked memory and CPU usage, and sufficient resources are available. You want to follow Google-recommended practices to ensure that the import will not time out. What should you do?

- A. Close idle connections or restart the instance before beginning the import operation.
- B. Increase the amount of memory allocated to your instance.
- C. Ensure that the service account has the Storage Admin role.
- D. Increase the number of CPUs for the instance to ensure that it can handle the additional import operation.

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/mysql/import-export#troubleshooting>

NEW QUESTION 48

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: B

Explanation:

Financial data, such as transaction histories, stock prices, and currency exchange rates.

<https://cloud.google.com/bigtable/docs/overview#what-its-good-for>

With SSD:

Reads - up to 10,000 rows per second Writes - up to 10,000 rows per second Scans - up to 220 MB/s

<https://cloud.google.com/bigtable/docs/performance#typical-workloads>

NEW QUESTION 50

You need to issue a new server certificate because your old one is expiring. You need to avoid a restart of your Cloud SQL for MySQL instance. What should you do in your Cloud SQL instance?

- A. Issue a rollback, and download your server certificate.
- B. Create a new client certificate, and download it.
- C. Create a new server certificate, and download it.
- D. Reset your SSL configuration, and download your server certificate.

Answer: C

Explanation:

<https://cloud.google.com/sql/docs/sqlserver/configure-ssl-instance#server-certs>

NEW QUESTION 53

Your team recently released a new version of a highly consumed application to accommodate additional user traffic. Shortly after the release, you received an alert from your production monitoring team that there is consistently high replication lag between your primary instance and the read replicas of your Cloud SQL for MySQL instances. You need to resolve the replication lag. What should you do?

- A. Identify and optimize slow running queries, or set parallel replication flags.
- B. Stop all running queries, and re-create the replicas.
- C. Edit the primary instance to upgrade to a larger disk, and increase vCPU count.
- D. Edit the primary instance to add additional memory.

Answer: A

Explanation:

https://cloud.google.com/sql/docs/mysql/replication/replication-lag#optimize_queries_and_schema

NEW QUESTION 57

Your team uses thousands of connected IoT devices to collect device maintenance data for your oil and gas customers in real time. You want to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by these devices. You need a managed solution that is highly scalable, supports a multi-cloud strategy, and offers low latency for these IoT devices. What should you do?

- A. Use Firestore with Looker.
- B. Use Cloud Spanner with Data Studio.
- C. Use MongoDB Atlas with Charts.
- D. Use Bigtable with Looker.

Answer: C

Explanation:

This scenario has BigTable written all over it - large amounts of data from many devices to be analysed in realtime. I would even argue it could qualify as a multicloud solution, given the links to HBASE. BUT it does not support SQL queries and is not therefore compatible (on its own) with Looker. Firestore + Looker has the same problem. Spanner + Data Studio is at least a compatible pairing, but I agree with others that it doesn't fit this use-case - not least because it's Google-native. By contrast, MongoDB Atlas is a managed solution (just not by Google) which is compatible with the proposed reporting tool (Mongo's own Charts), it's specifically designed for this type of solution and of course it can run on any cloud.

NEW QUESTION 61

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?

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

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/using-query-insights#introduction>

NEW QUESTION 65

You work for a large retail and ecommerce company that is starting to extend their business globally. Your company plans to migrate to Google Cloud. You want to use platforms that will scale easily, handle transactions with the least amount of latency, and provide a reliable customer experience. You need a storage layer for sales transactions and current inventory levels. You want to retain the same relational schema that your existing platform uses. What should you do?

- A. Store your data in Firestore in a multi-region location, and place your compute resources in one of the constituent regions.
- B. Deploy Cloud Spanner using a multi-region instance, and place your compute resources close to the default leader region.
- C. Build an in-memory cache in Memorystore, and deploy to the specific geographic regions where your application resides.
- D. Deploy a Bigtable instance with a cluster in one region and a replica cluster in another geographic region.

Answer: B

NEW QUESTION 70

You are running a mission-critical application on a Cloud SQL for PostgreSQL database with a multi-zonal setup. The primary and read replica instances are in the same region but in different zones. You need to ensure that you split the application load between both instances. What should you do?

- A. Use Cloud Load Balancing for load balancing between the Cloud SQL primary and read replica instances.
- B. Use PgBouncer to set up database connection pooling between the Cloud SQL primary and read replica instances.
- C. Use HTTP(S) Load Balancing for database connection pooling between the Cloud SQL primary and read replica instances.
- D. Use the Cloud SQL Auth proxy for database connection pooling between the Cloud SQL primary and read replica instances.

Answer: B

Explanation:

<https://severalnines.com/blog/how-achieve-postgresql-high-availability-pgbouncer/>

<https://cloud.google.com/blog/products/databases/using-haproxy-to-scale-read-only-workloads-on-cloud-sql-for-postgresql>

This answer is correct because PgBouncer is a lightweight connection pooler for PostgreSQL that can help you distribute read requests between the Cloud SQL primary and read replica instances¹. PgBouncer can also improve performance and scalability by reducing the overhead of creating new connections and reusing existing ones¹. You can install PgBouncer on a Compute Engine instance and configure it to connect to the Cloud SQL instances using private IP addresses or the Cloud SQL Auth proxy².

NEW QUESTION 75

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?

- A. roles/spanner.databaseAdmin for Person A roles/spanner.databaseReader for Person B roles/spanner.backupWriter for Application C
- B. roles/spanner.databaseAdmin for Person A roles/spanner.databaseReader for Person B roles/spanner.backupAdmin for Application C
- C. roles/spanner.databaseAdmin for Person A roles/spanner.databaseUser for Person B roles/spanner.databaseReader for Application C
- D. roles/spanner.databaseAdmin for Person A roles/spanner.databaseUser for Person B roles/spanner.backupWriter for Application C

Answer: A

Explanation:

<https://cloud.google.com/spanner/docs/iam#spanner.backupWriter>

NEW QUESTION 78

You support a consumer inventory application that runs on a multi-region instance of Cloud Spanner. A customer opened a support ticket to complain about slow response times. You notice a Cloud Monitoring alert about high CPU utilization. You want to follow Google- recommended practices to address the CPU performance issue. What should you do first?

- A. Increase the number of processing units.
- B. Modify the database schema, and add additional indexes.
- C. Shard data required by the application into multiple instances.
- D. Decrease the number of processing units.

Answer: A

Explanation:

In case of high CPU utilization like, mentioned in question, refer: <https://cloud.google.com/spanner/docs/identify-latency-point#:~:text=Check%20the%20CPU%20utilization%20of%20the%20instance.%20If%20the%20CPU%20utilization%20of%20the%20instance%20is%20above%20the%20recommended%20level%2C%20you%20should%20manually%20add%20more%20nodes%2C%20or%20set%20up%20auto%20scaling.> "Check the CPU utilization of the instance. If the CPU utilization of the instance is above the recommended level, you should manually add more nodes, or set up auto scaling." Indexes and schema are reviewed post identifying query with slow performance. Refer : <https://cloud.google.com/spanner/docs/troubleshooting-performance-regressions#review-schema>

NEW QUESTION 80

Your organization is running a Firestore-backed Firebase app that serves the same top ten news stories on a daily basis to a large global audience. You want to optimize content delivery while decreasing cost and latency. What should you do?

- A. Enable serializable isolation in the Firebase app.
- B. Deploy a US multi-region Firestore location.
- C. Build a Firestore bundle, and deploy bundles to Cloud CDN.
- D. Create a Firestore index on the news story date.

Answer: C

Explanation:

A global audience strongly suggests serving content via Google's Content Delivery Network. Changing the isolation level won't decrease cost or latency

NEW QUESTION 84

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?

- A. Issue a `gcloud sql instances patch` command to increase the number of vCPUs.
- B. Update a MySQL database flag to increase the number of vCPUs.
- C. Issue a `gcloud compute instances update` command to increase the number of vCPUs.
- D. Back up the database, create an instance with additional vCPUs, and restore the database.

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/sql/instances/patch>

NEW QUESTION 85

You have deployed a Cloud SQL for SQL Server instance. In addition, you created a cross-region read replica for disaster recovery (DR) purposes. Your company requires you to maintain and monitor a recovery point objective (RPO) of less than 5 minutes. You need to verify that your cross-region read replica meets the allowed RPO. What should you do?

- A. Use Cloud SQL instance monitoring.
- B. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- C. Use Cloud SQL logs.
- D. Use the SQL Server Always On Availability Group dashboard.

Answer: D

Explanation:

Note, you cannot create a read replica in Cloud SQL for SQL Server unless you use an Enterprise Edition. Which is also a requirement for configuring SQL Server AG. That's not a coincidence. That's how Cloud SQL for SQL Server creates SQL Server read replicas. To find out about the replication, use the AG Dashboard in SSMS.

<https://cloud.google.com/sql/docs/sqlserver/replication/manage-replicas#promote-replica>

NEW QUESTION 90

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?

- A. 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.
- B. 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.
- C. Enable the `cloudsql.iam_authentication` flag on the primary database instance, and grant read access to required tables to the analytics team.
- D. 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.

Answer: B

Explanation:

"Read replicas do not have the `cloudsql.iam_authentication` flag enabled automatically when it is enabled on the primary instance."

https://cloud.google.com/sql/docs/postgres/replication/create-replica#configure_iam_replicas

NEW QUESTION 91

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-Database-Engineer Exam with Our Prep Materials Via below:

<https://www.certleader.com/Professional-Cloud-Database-Engineer-dumps.html>