Free Questions for Professional-Cloud-Database-Engineer

Shared by Miles on 04-10-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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?

Options:

A- Use Cloud SQL serverless exports.

B- Create a read replica, and then use the mysqldump utility to export each table.

C- Clone the Cloud SQL instance, and then use the mysqldump utlity to export the data.

D- Use the mysqldump utility on the primary database instance to export the backup.

Answer:

A

Explanation:

https://cloud.google.com/blog/products/databases/introducing-cloud-sql-serverless-exports

Question 2

Question Type: MultipleChoice

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?

Options:

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.

Question 3

Question Type: MultipleChoice

Your company uses the Cloud SQL out-of-disk recommender to analyze the storage utilization trends of production databases over the last 30 days. Your database operations team uses these recommendations to proactively monitor storage utilization and implement corrective actions. You receive a recommendation that the instance is likely to run out of disk space. What should you do to address this storage alert?

Options:

A- Normalize the database to the third normal form.
B- Compress the data using a different compression algorithm.
C- Manually or automatically increase the storage capacity.
D- Create another schema to load older data.

Answer:

C

Explanation:

https://cloud.google.com/sql/docs/mysql/instance-settings#storage-capacity-2ndgen

Question 4

Question Type: MultipleChoice

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 dat

a. 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?

Options:
A- Determine whether the versions of Cloud SQL for PostgreSQL in regions R1 and R2 are different.

B- Determine whether the database patches of Cloud SQl for PostgreSQL in regions R1 and R2 are different.

C- Determine whether the failover of Cloud SQL for PostgreSQL from region R1 to region R2 is in progress or has completed successfully.

D- Determine whether Cloud SQL for PostgreSQL in region R2 is a near-real-time copy of region R1 but not an exact copy.

Answer:
D

Explanation:
Verify that the replica has processed all the transactions it has received from the primary. This ensures that when promoted, the replica reflects all transactions that were received before the primary became unavailable.

https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#verify_failover_criteria

Question 5

Question Type: MultipleChoice

You are managing two different applications: Order Management and Sales Reporting. Both applications interact with the same Cloud SQL for MySQL database. The Order Management application reads and writes to the database 24/7, but the Sales Reporting application is read-only. Both applications need the latest dat

a. You need to ensure that the Performance of the Order Management application is not affected by the Sales Reporting application. What should you do?

Options:

A- Create a read replica for the Sales Reporting application.

B- Create two separate databases in the instance, and perform dual writes from the Order Management application.

C- Use a Cloud SQL federated query for the Sales Reporting application.

D- Queue up all the requested reports in PubSub, and execute the reports at night.

Answer:

A

Question 6

Question Type: MultipleChoice

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?

Options:

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

Question 7

Question Type: MultipleChoice

Your company is shutting down their on-premises data center and migrating their Oracle databases using Oracle Real Application Clusters (RAC) to Google Cloud. You want minimal to no changes to the applications during the database migration. What should you do?

Options:

A- Migrate the Oracle databases to Cloud Spanner.

B- Migrate the Oracle databases to Compute Engine.

C- Migrate the Oracle databases to Cloud SQL.

D- Migrate the Oracle databases to Bare Metal Solution for Oracle.

Answer:

D

Explanation:

This answer is correct because Bare Metal Solution for Oracle is a service that provides dedicated physical servers and networking infrastructure for running Oracle databases on Google Cloud1.Bare Metal Solution for Oracle supports Oracle RAC, which is a cluster database that provides high availability, scalability, and performance for Oracle workloads2.By using Bare Metal Solution for Oracle, you can migrate your Oracle databases with minimal to no changes to the applications, and you can leverage the native Google Cloud services and interconnectivity1.

Question 8

Question Type: MultipleChoice

Your team is building an application that stores and analyzes streaming time series financial dat

a. 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?

Options:

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

Question 9

Question Type: MultipleChoice

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?

Options:

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

Question 10

Question Type: MultipleChoice

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?

Options:

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.

Answer:

C

Explanation:

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

To Get Premium Files for Professional-Cloud-Database-Engineer Visit

https://www.p2pexams.com/products/professional-cloud-database-engineer

For More Free Questions Visit

https://www.p2pexams.com/google/pdf/professional-cloud-database-engineer