Free Questions for Professional-Cloud-Database-Engineer

Shared by Rush on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are designing for a write-heavy application. During testing, you discover that the write workloads are performant in a regional Cloud Spanner instance but slow down by an order of magnitude in a multi-regional instance. You want to make the write workloads faster in a multi-regional instance. What should you do?

Options:

A- Place the bulk of the read and write workloads closer to the default leader region.
B- Use staleness of at least 15 seconds.
C- Add more read-write replicas.
D- Keep the total CPU utilization under 45% in each region.

Answer:

A

Explanation:

https://cloud.google.com/spanner/docs/instance-configurations#multi-region-best-practices Best practices For optimal performance, follow these best practices: Design a schema that prevents hotspots and other performance issues. For optimal write latency, place compute resources for write-heavy workloads within or close to the default leader region. For optimal read performance outside of the default leader region, use staleness of at least 15 seconds. To avoid single-region dependency for your workloads, place critical compute resources in at least two regions. A good option is to place them next to the two different read-write regions so that any single region outage will not impact all of your application. Provision enough compute capacity to keep high priority total CPU utilization under 45% in each region.

Question 2

Question Type: MultipleChoice

Your company is using Cloud SQL for MySQL with an internal (private) IP address and wants to replicate some tables into BigQuery in near-real time for analytics and machine learning. You need to ensure that replication is fast and reliable and uses Google-managed services. What should you do?

Options:

A- Develop a custom data replication service to send data into BigQuery.
B- Use Cloud SQL federated queries.
C- Use Database Migration Service to replicate tables into BigQuery.
D- Use Datastream to capture changes, and use Dataflow to write those changes to BigQuery.

Answer:

D

Explanation:

''Datastream is a serverless and easy-to-use Change Data Capture (CDC) and replication service that allows you to synchronize data across heterogeneous databases, storage systems, and applications reliably and with minimal latency. Datastream supports change data streaming from Oracle and MySQL databases to Google Cloud Storage (GCS). The service offers streamlined integration with Dataflow templates to power up to date materialized views in BigQuery for analytics, replicate their databases into Cloud SQL or Cloud Spanner for database synchronization, or leverage the event stream directly from GCS to realize event-driven architectures.''

Question 3

Question Type: MultipleChoice

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?

Options:

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

Question 4

Question Type: MultipleChoice

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?

Options:

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

Question 5

Question Type: MultipleChoice

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?

Options:

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

Answer:

B

Explanation:

This answer is correct because it meets the requirements of having an optimal and highly available environment with automatic failover.According to the Google Cloud documentation1, a regional Cloud SQL instance is an instance that has a primary server in one zone and a standby server in another zone within the same region. The primary and standby servers are kept in sync using synchronous replication, which ensures zero data loss and minimal downtime in case of a zonal outage or an instance failure.If the primary server becomes unavailable, Cloud SQL automatically fails over to the standby server, which becomes the new primary server1.

Question 6

Question Type: MultipleChoice

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?

Options:

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

Answer:

C

Explanation:

Recovery Process: Once Zone-B becomes available again, Cloud SQL will initiate the recovery process for the impacted read replica. The recovery process involves the following steps: 1. Synchronization: Cloud SQL will compare the data in the recovered read replica with the primary instance in Zone-A. If there is any data divergence due to the unavailability period, Cloud SQL will synchronize the read replica with the primary instance to ensure data consistency. 2. Catch-up Replication: The recovered read replica will start catching up on the changes that occurred on the primary instance during its unavailability. It will apply the necessary updates from the primary instance's binary logs (binlogs) to bring the replica up to date. 3. Resuming Read Traffic: Once the synchronization and catch-up replication processes are complete, the read replica in Zone-B will resume its normal operation. It will be able to serve read traffic and stay updated with subsequent changes from the primary instance.

Question 7

Question Type: MultipleChoice

You are designing a database architecture for a global application that stores information about public parks worldwide. The application uses the database for read-only purposes, and a centralized batch job updates the database nightly. You want to select an open source, SQL-compliant database. What should you do?

Options:

A- Use Bigtable with multi-region clusters.
B- Use Memorystore for Redis with multi-zones within a region.

C- Use Cloud SQL for PostgreSQL with cross-region replicas.
D- Use Cloud Spanner with multi-region configuration.

Answer:

C

Question 8

Question Type: MultipleChoice

You want to migrate your PostgreSQL database from another cloud provider to Cloud SQL. You plan on using Database Migration Service and need to assess the impact of any known limitations. What should you do? (Choose two.)

Options:

A- Identify whether the database has over 512 tables.
B- Identify all tables that do not have a primary key.
C- Identity all tables that do not have at least one foreign key.
D- Identify whether the source database is encrypted using pgcrypto extension.
E- Identify whether the source database uses customer-managed encryption keys (CMEK).

Answer:

C, E

Question 9

Question Type: MultipleChoice

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 dat

a. What should you do?

Options:

A- Use Identity and Access Management (IAM) for Bigtable access control.
B- Use VPC Service Controls to create a trusted network for the Bigtable service.

C- Use customer-managed encryption keys (CMEK).
D- Use Google Cloud Armor to add IP addresses to an allowlist.

Answer:

B

Explanation:

''Users can define a security perimeter around Google Cloud resources such as Cloud Storage buckets, Bigtable instances, and BigQuery datasets to constrain data within a VPC and control the flow of data.'' https://cloud.google.com/vpc-service-controls

Question 10

Question Type: MultipleChoice

You are managing a Cloud SQL for PostgreSQL instance in Google Cloud. You have a primary instance in region 1 and a read replica in region 2. After a failure of region 1, you need to make the Cloud SQL instance available again. You want to minimize data loss and follow Google-recommended practices. What should you do?

Options:

A- Restore the Cloud SQL instance from the automatic backups in region 3.
B- Restore the Cloud SQL instance from the automatic backups in another zone in region 1.
C- Check 'Lag Bytes' in the monitoring dashboard for the primary instance in the read replica instance. Check the replication status using pg_catalog.pg_last_wal_receive_lsn(). Then, fail over to region 2 by promoting the read replica instance.
D- Check your instance operational log for the automatic failover status. Look for time, type, and status of the operations. If the failover operation is successful, no action is necessary. Otherwise, manually perform gcloud sql instances failover .

Answer:

C

Explanation:

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