



Free Questions for **Professional-Cloud-Database-Engineer** by **certscare**

Shared by **Cortez** on **15-04-2024**

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You work for a financial services company that wants to use fully managed database services. Traffic volume for your consumer services products has increased annually at a constant rate with occasional spikes around holidays. You frequently need to upgrade the capacity of your database. You want to use Cloud Spanner and include an automated method to increase your hardware capacity to support a higher level of concurrency. What should you do?

Options:

- A- Use linear scaling to implement the Autoscaler-based architecture
- B- Use direct scaling to implement the Autoscaler-based architecture.
- C- Upgrade the Cloud Spanner instance on a periodic basis during the scheduled maintenance window.
- D- Set up alerts that are triggered when Cloud Spanner utilization metrics breach the threshold, and then schedule an upgrade during the scheduled maintenance window.

Answer:

A

Explanation:

Linear scaling is best used with load patterns that change more gradually or have a few large peaks. The method calculates the minimum number of nodes or processing units required to keep utilization below the scaling threshold. The number of nodes or processing units added or removed in each scaling event is not limited to a fixed step amount.

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

Question 2

Question Type: MultipleChoice

Your hotel booking company is expanding into Country A, where personally identifiable information (PII) must comply with regional data residency requirements and audits. You need to isolate customer data in Country A from the rest of the customer dat

a. You want to design a multi-tenancy strategy to efficiently manage costs and operations. What should you do?

Options:

- A- Apply a schema data management pattern.
- B- Apply an instance data management pattern.
- C- Apply a table data management pattern.

D- Apply a database data management pattern.

Answer:

B

Explanation:

<https://cloud.google.com/solutions/implementing-multi-tenancy-cloud-spanner#multi-tenancy-data-management-patterns>

<https://cloud.google.com/solutions/implementing-multi-tenancy-cloud-spanner>

Question 3

Question Type: MultipleChoice

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?

Options:

- 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 committed¹. This minimizes data loss and downtime in case of a zonal or an instance failure, and allows for a quick failover to the standby instance¹.

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 delay². This provides additional redundancy and backup in case of a regional failure, and allows for a manual failover to the read replica in another region².

Question 4

Question Type: MultipleChoice

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?

Options:

- 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>

Question 5

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 6

Question Type: MultipleChoice

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?

Options:

- 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>

Question 7

Question Type: MultipleChoice

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?

Options:

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>

**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>

