Google Professional-Cloud-Database-Engineer Practice Questions

Shared by Mays on 24-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are designing a highly available (HA) Cloud SQL for PostgreSQL instance that will be used by 100 databases. Each database contains 80 tables that were migrated from your on-premises environment to Google Cloud. The applications that use these databases are located in multiple regions in the US, and you need to ensure that read and write operations have low latency. What should you do?

Options:

A- Deploy 2 Cloud SQL instances in the us-central1 region with HA enabled, and create read replicas in us-east1 and us-west1.
B- Deploy 2 Cloud SQL instances in the us-central1 region, and create read replicas in us-east1 and us-west1.
C- Deploy 4 Cloud SQL instances in the us-central1 region with HA enabled, and create read replicas in us-central1, us-east1, and us-west1.
D- Deploy 4 Cloud SQL instances in the us-central1 region, and create read replicas in us-central1, us-east1 and us-west1.

Answer:

A

Explanation:

https://cloud.google.com/sql/docs/mysql/quotas#table_limit

Question 2

Question Type: MultipleChoice

Your company wants to migrate an Oracle-based application to Google Cloud. The application team currently uses Oracle Recovery Manager (RMAN) to back up the database to tape for long-term retention (LTR). You need a cost-effective backup and restore solution that meets a 2-hour recovery time objective (RTO) and a 15-minute recovery point objective (RPO). What should you do?

Options:

A- Migrate the Oracle databases to Bare Metal Solution for Oracle, and store backups on tapes on-premises.
B- Migrate the Oracle databases to Bare Metal Solution for Oracle, and use Actifio to store backup files on Cloud Storage using the Nearline Storage class.
C- Migrate the Oracle databases to Bare Metal Solution for Oracle, and back up the Oracle databases to Cloud Storage using the Standard Storage class.
D- Migrate the Oracle databases to Compute Engine, and store backups on tapes on-premises.

Answer:

B

Explanation:

https://www.actifio.com/solutions/cloud/google/

Question 3

Question Type: MultipleChoice

Your organization has a security policy to ensure that all Cloud SQL for PostgreSQL databases are secure. You want to protect sensitive data by using a key that meets specific locality or residency requirements. Your organization needs to control the key's lifecycle activities. You need to ensure that data is encrypted at rest and in transit. What should you do?

Options:

A- Create the database with Google-managed encryption keys.
B- Create the database with customer-managed encryption keys.
C- Create the database persistent disk with Google-managed encryption keys.
D- Create the database persistent disk with customer-managed encryption keys.

Answer:

B

Explanation:

https://cloud.google.com/sql/docs/postgres/configure-cmek#createcmekinstance

Question 4

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 5

Question Type: MultipleChoice

Your organization stores marketing data such as customer preferences and purchase history on Bigtable. The consumers of this database are predominantly data analysts and operations users. You receive a service ticket from the database operations department citing poor database performance between 9 AM-10 AM every day. The application team has confirmed no latency from their logs. A new cohort of pilot users that is testing a dataset loaded from a third-party data provider is experiencing poor database performance. Other users are not affected. You need to troubleshoot the issue. What should you do?

Options:

A- Isolate the data analysts and operations user groups to use different Bigtable instances.
B- Check the Cloud Monitoring table/bytes_used metric from Bigtable.
C- Use Key Visualizer for Bigtable.
D- Add more nodes to the Bigtable cluster.

Answer:

C

Explanation:

https://cloud.google.com/bigtable/docs/performance#troubleshooting

Question 6

Question Type: MultipleChoice

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?

Options:

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

Question 7

Question Type: MultipleChoice

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?

Options:

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.

Question 8

Question Type: MultipleChoice

You need to redesign the architecture of an application that currently uses Cloud SQL for PostgreSQL. The users of the application complain about slow query response times. You want to enhance your application architecture to offer sub-millisecond query latency. What should you do?

Options:

A- Configure Firestore, and modify your application to offload queries.
B- Configure Bigtable, and modify your application to offload queries.
C- Configure Cloud SQL for PostgreSQL read replicas to offload queries.
D- Configure Memorystore, and modify your application to offload queries.

Answer:

D

Explanation:

'sub-millisecond latency' always involves Memorystore. Furthermore, as we are talking about a relational DB (Cloud SQL), BigTable is not a solution to be considered.

Question 9

Question Type: MultipleChoice

You are migrating an on-premises application to Compute Engine and Cloud SQL. The application VMs will live in their own project, separate from the Cloud SQL instances which have their own project. What should you do to configure the networks?

Options:

A- Create a new VPC network in each project, and use VPC Network Peering to connect the two together.
B- Create a Shared VPC that both the application VMs and Cloud SQL instances will use.
C- Use the default networks, and leverage Cloud VPN to connect the two together.
D- Place both the application VMs and the Cloud SQL instances in the default network of each

project.

Answer:

B

Explanation:

https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5_HPXytY?pli=1

P2P exams