PROFESSIONAL-CLOUD-DATABASE-ENGINEER Q&As

Google Cloud Certified - Professional Cloud Database Engineer

Pass Google PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/professional-cloud-database-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

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.

Correct Answer: D

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

QUESTION 2

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.

Correct Answer: B

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

QUESTION 3

Your organization has hundreds of Cloud SQL for MySQL instances. You want to follow Google-recommended practices to optimize platform costs. What should you do?

A. Use Query Insights to identify idle instances.

B. Remove inactive user accounts.

C. Run the Recommender API to identify overprovisioned instances.

D. Build indexes on heavily accessed tables.

Correct Answer: C

The Cloud SQL overprovisioned instance recommender helps you detect instances that are unnecessarily large for a given workload. It then provides recommendations on how to resize such instances and reduce cost. This page describes how this recommender works and how to use it.

QUESTION 4

You are migrating an on-premises application to Google Cloud. The application requires a high availability (HA) PostgreSQL database to support business-critical functions. Your company\\'s disaster recovery strategy requires a recovery time objective (RTO) and recovery point objective (RPO) within 30 minutes of failure. You plan to use a Google Cloud managed service. What should you do to maximize uptime for your application?

A. Deploy Cloud SQL for PostgreSQL in a regional configuration. Create a read replica in a different zone in the same region and a read replica in another region for disaster recovery.

B. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Take periodic backups, and use this backup to restore to a new Cloud SQL for PostgreSQL instance in another region during a disaster recovery event.

C. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Create a cross-region read replica, and promote the read replica as the primary node for disaster recovery.

D. Migrate the PostgreSQL database to multi-regional Cloud Spanner so that a single region outage will not affect your application. Update the schema to support Cloud Spanner data types, and refactor the application.

Correct Answer: C

The best answer is deploy an HA configuration and have a read replica you could promote to the primary in a different region

QUESTION 5

Your customer has a global chat application that uses a multi-regional Cloud Spanner instance. The application has recently experienced degraded performance after a new version of the application was launched. Your customer asked you for assistance. During initial troubleshooting, you observed high read latency. What should you do?

A. Use query parameters to speed up frequently executed queries.

B. Change the Cloud Spanner configuration from multi-region to single region.

C. Use SQL statements to analyze SPANNER_SYS.READ_STATS* tables.

D. Use SQL statements to analyze SPANNER_SYS.QUERY_STATS* tables.

Correct Answer: C

To troubleshoot high read latency, you can use SQL statements to analyze the SPANNER_SYS.READ_STATS* tables. These tables contain statistics about read operations in Cloud Spanner, including the number of reads, read latency, and the number of read errors. By analyzing these tables, you can identify the cause of the high read latency and take

appropriate action to resolve the issue. Other options, such as using query parameters to speed up frequently executed queries or changing the Cloud Spanner configuration from multi-region to single region, may not be directly related to the issue of high read latency. Similarly, analyzing the SPANNER_SYS.QUERY_STATS* tables, which contain statistics about query operations, may not be relevant to the issue of high read latency.

QUESTION 6

During an internal audit, you realized that one of your Cloud SQL for MySQL instances does not have high availability (HA) enabled. You want to follow Google-recommended practices to enable HA on your existing instance. What should you do?

A. Create a new Cloud SQL for MySQL instance, enable HA, and use the export and import option to migrate your data.

B. Create a new Cloud SQL for MySQL instance, enable HA, and use Cloud Data Fusion to migrate your data.

C. Use the gcloud instances patch command to update your existing Cloud SQL for MySQL instance.

D. Shut down your existing Cloud SQL for MySQL instance, and enable HA.

Correct Answer: C

Creating a new instance and migrating data can be time-consuming and disruptive to your application\\'s availability. Shutting down the existing instance is not a recommended approach, as it will cause downtime for your application.

The recommended approach is to use the gcloud instances patch command to enable high availability on your existing Cloud SQL for MySQL instance. This command updates the instance\\'s configuration to enable the failover replica, configure it, and enable automatic failover.

By following this approach, you can ensure minimal downtime, and your application can continue to operate during the process.

QUESTION 7

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.

Correct Answer: A

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-forterraform#stateful-resources

QUESTION 8

Your team is building a new inventory management application that will require read and write database instances in multiple Google Cloud regions around the globe. Your database solution requires 99.99% availability and global transactional consistency. You need a fully managed backend relational database to store inventory changes. What should you do?

A. Use Bigtable.

B. Use Firestore.

C. Use Cloud SQL for MySQL

D. Use Cloud Spanner.

Correct Answer: D

Spanner covers the SLA

QUESTION 9

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?

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.

Correct Answer: A

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

QUESTION 10

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.

Correct Answer: D

QUESTION 11

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.

Correct Answer: C

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 12

Your organization is running a MySQL workload in Cloud SQL. Suddenly you see a degradation in database performance. You need to identify the root cause of the performance degradation. What should you do?

A. Use Logs Explorer to analyze log data.

B. Use Cloud Monitoring to monitor CPU, memory, and storage utilization metrics.

C. Use Error Reporting to count, analyze, and aggregate the data.

D. Use Cloud Debugger to inspect the state of an application.

Correct Answer: B

QUESTION 13

You want to migrate an existing on-premises application to Google Cloud. Your application supports semi-structured data ingested from 100,000 sensors, and each sensor sends 10 readings per second from manufacturing plants. You need to make this data available for real-time monitoring and analysis. What should you do?

A. Deploy the database using Cloud SQL.

B. Use BigQuery, and load data in batches.

C. Deploy the database using Bigtable.

D. Deploy the database using Cloud Spanner.

Correct Answer: C

Bigtable is a scalable, fully managed, and high-performance NoSQL database service that can handle semi-structured data and support real-time monitoring and analysis. Cloud SQL is a relational database service that does not support semi-structured data. BigQuery is a data warehouse service that is optimized for batch processing and analytics, not real-time monitoring. Cloud Spanner is a relational database service that supports semi-structured data with JSON data type, but it is more expensive and complex than Bigtable for this use case.

QUESTION 14

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 Biglable cluster with autoscaling.

D. Use Cloud Spanner with a custom autoscaling solution.

Correct Answer: D

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

1.

CloudSQL max out at 64TB, so unable to told 100TB of data.
https://cloud.google.com/sql/docs/quotas#metrics_collection_limit

2.

Scale is done manually on SQL Cloud

QUESTION 15

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?

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.

Correct Answer: C

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