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.passapply.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 are building a data warehouse on BigQuery. Sources of data include several MySQL databases located on-premises.

You need to transfer data from these databases into BigQuery for analytics. You want to use a managed solution that has low latency and is easy to set up. What should you do?

A. Create extracts from your on-premises databases periodically, and push these extracts to Cloud Storage. Upload the changes into BigQuery, and merge them with existing tables.

B. Use Cloud Data Fusion and scheduled workflows to extract data from MySQL. Transform this data into the appropriate schema, and load this data into your BigQuery database.

C. Use Datastream to connect to your on-premises database and create a stream. Have Datastream write to Cloud Storage. Then use Dataflow to process the data into BigQuery.

D. Use Database Migration Service to replicate data to a Cloud SQL for MySQL instance. Create federated tables in BigQuery on top of the replicated instances to transform and load the data into your BigQuery database.

Correct Answer: C

QUESTION 2

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?

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.

Correct Answer: A

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 3

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?

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.

Correct Answer: C

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

QUESTION 4

You have a Cloud SQL instance (DB-1) with two cross-region read replicas (DB-2 and DB-3). During a business continuity test, the primary instance (DB-1) was taken offline and a replica (DB-2) was promoted. The test has concluded and you want to return to the pre-test configuration. What should you do?

A. Bring DB-1 back online.

B. Delete DB-1, and re-create DB-1 as a read replica in the same region as DB-1.

C. Delete DB-2 so that DB-1 automatically reverts to the primary instance.

D. Create DB-4 as a read replica in the same region as DB-1, and promote DB-4 to primary.

Correct Answer: D

If you need to have the primary instance in the zone that had the outage, you can do a failback. A failback performs the same steps as the failover, only in the opposite direction, to reroute traffic back to the original instance. To perform a failback, use the procedure in Initiating failover. https://cloud.google.com/sql/docs/mysql/high-availability#failback

QUESTION 5

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?

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.

Correct Answer: A

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%20th e%20CPU%20utilization%20of%20the%20instance%20is%20above%20the%20recommen

ded%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#reviewschema