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

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

1 / 8

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

2 / 8

QUESTION 1

You are managing two different applications: Order Management and Sales Reporting. Both applications interact with the same Cloud SQL for MySQL database. The Order Management application reads and writes to the database 24/7, but the Sales Reporting application is read-only. Both applications need the latest data. You need to ensure that the Performance of the Order Management application is not affected by the Sales Reporting application. What should you do?

A. Create a read replica for the Sales Reporting application.

B. Create two separate databases in the instance, and perform dual writes from the Order Management application.

C. Use a Cloud SQL federated query for the Sales Reporting application.

D. Queue up all the requested reports in PubSub, and execute the reports at night.

Correct Answer: C

QUESTION 2

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

QUESTION 3

Your application uses Cloud SQL for MySQL. Your users run reports on data that relies on near-real time; however, the additional analytics caused excessive load on the primary database. You created a read replica for the analytics workloads, but now your users are complaining about the lag in data changes and that their reports are still slow. You need to improve the report performance and shorten the lag in data replication without making changes to the current reports. Which two approaches should you implement? (Choose two.)

A. Create secondary indexes on the replica.

B. Create additional read replicas, and partition your analytics users to use different read replicas.

C. Disable replication on the read replica, and set the flag for parallel replication on the read replica. Re-enable replication and optimize performance by setting flags on the primary instance.

D. Disable replication on the primary instance, and set the flag for parallel replication on the primary instance. Re-enable replication and optimize performance by setting flags on the read replica.

E. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

Correct Answer: BE

QUESTION 4

You are responsible for designing a new database for an airline ticketing application in Google Cloud. This application must be able to:

Work with transactions and offer strong consistency.

Work with structured and semi-structured (JSON) data.

Scale transparently to multiple regions globally as the operation grows.

You need a Google Cloud database that meets all the requirements of the application.

What should you do?

A. Use Cloud SQL for PostgreSQL with both cross-region read replicas.

B. Use Cloud Spanner in a multi-region configuration.

C. Use Firestore in Datastore mode.

D. Use a Bigtable instance with clusters in multiple regions.

Correct Answer: A

QUESTION 5

Your team is running a Cloud SQL for MySQL instance with a 5 TB database that must be available 24/7. You need to save database backups on object storage with minimal operational overhead or risk to your production workloads. What should you do?

A. Use Cloud SQL serverless exports.

B. Create a read replica, and then use the mysqldump utility to export each table.

C. Clone the Cloud SQL instance, and then use the mysqldump utlity to export the data.

D. Use the mysqldump utility on the primary database instance to export the backup.

Correct Answer: C

QUESTION 6

You are the DBA of an online tutoring application that runs on a Cloud SQL for PostgreSQL database. You are testing the implementation of the cross-regional failover configuration. The database in region R1 fails over successfully to region R2, and the database becomes available for the application to process data. During testing, certain scenarios of the application work as expected in region R2, but a few scenarios fail with database errors. The application-related database queries, when executed in isolation from Cloud SQL for PostgreSQL in region R2, work as expected. The application performs completely as expected when the database fails back to region R1. You need to identify the cause of the database errors in region R2. What should you do?

A. Determine whether the versions of Cloud SQL for PostgreSQL in regions R1 and R2 are different.

B. Determine whether the database patches of Cloud SQl for PostgreSQL in regions R1 and R2 are different.

C. Determine whether the failover of Cloud SQL for PostgreSQL from region R1 to region R2 is in progress or has completed successfully.

D. Determine whether Cloud SQL for PostgreSQL in region R2 is a near-real-time copy of region R1 but not an exact copy.

Correct Answer: B

QUESTION 7

You are migrating your data center to Google Cloud. You plan to migrate your applications to Compute Engine and your Oracle databases to Bare Metal Solution for Oracle. You must ensure that the applications in different projects can communicate securely and efficiently with the Oracle databases. What should you do?

A. Set up a Shared VPC, configure multiple service projects, and create firewall rules.

B. Set up Serverless VPC Access.

C. Set up Private Service Connect.

D. Set up Traffic Director.

Correct Answer: A

QUESTION 8

Your organization has a production Cloud SQL for MySQL instance. Your instance is configured with 16 vCPUs and 104 GB of RAM that is running between 90% and 100% CPU utilization for most of the day. You need to scale up the database and add vCPUs with minimal interruption and effort. What should you do?

A. Issue a gcloud sql instances patch command to increase the number of vCPUs.

B. Update a MySQL database flag to increase the number of vCPUs.

C. Issue a gcloud compute instances update command to increase the number of vCPUs.

D. Back up the database, create an instance with additional vCPUs, and restore the database.

Correct Answer: A

QUESTION 9

You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SLA. You want to follow Google- recommended practices. You need the database to be performant and highly available while you keep costs low. What should you do?

A. Manually scale down the number of nodes after the peak period has passed.

B. Use interleaving to co-locate parent and child rows.

C. Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.

D. Use granular instance sizing in Cloud Spanner and Autoscaler.

Correct Answer: C

QUESTION 10

Your company uses Cloud Spanner for a mission-critical inventory management system that is globally available. You recently loaded stock keeping unit (SKU) and product catalog data from a company acquisition and observed hot-spots in the Cloud Spanner database. You want to follow Google-recommended schema design practices to avoid performance degradation. What should you do? (Choose two.)

A. Use an auto-incrementing value as the primary key.

B. Normalize the data model.

C. Promote low-cardinality attributes in multi-attribute primary keys.

D. Promote high-cardinality attributes in multi-attribute primary keys.

E. Use bit-reverse sequential value as the primary key.

Correct Answer: AD

QUESTION 11

You are choosing a new database backend for an existing application. The current database is running PostgreSQL on an on-premises VM and is managed by a database administrator and operations team. The application data is relational and has light traffic. You want to minimize costs and the migration effort for this application. What should you do?

A. Migrate the existing database to Firestore.

B. Migrate the existing database to Cloud SQL for PostgreSQL.

C. Migrate the existing database to Cloud Spanner.

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

6 / 8

D. Migrate the existing database to PostgreSQL running on Compute Engine.

Correct Answer: B

QUESTION 12

You manage a production MySQL database running on Cloud SQL at a retail company. You perform routine maintenance on Sunday at midnight when traffic is slow, but you want to skip routine maintenance during the year-end holiday shopping season. You need to ensure that your production system is available 24/7 during the holidays. What should you do?

A. Define a maintenance window on Sundays between 12 AM and 1 AM, and deny maintenance periods between November 1 and January 15.

B. Define a maintenance window on Sundays between 12 AM and 5 AM, and deny maintenance periods between November 1 and February 15.

C. Build a Cloud Composer job to start a maintenance window on Sundays between 12 AM and 1AM, and deny maintenance periods between November 1 and January 15.

D. Create a Cloud Scheduler job to start maintenance at 12 AM on Sundays. Pause the Cloud Scheduler job between November 1 and January 15.

Correct Answer: B

QUESTION 13

Your organization is running a critical production database on a virtual machine (VM) on Compute Engine. The VM has an ext4-formatted persistent disk for data files. The database will soon run out of storage space. You need to implement a solution that avoids downtime. What should you do?

A. In the Google Cloud Console, increase the size of the persistent disk, and use the resize2fs command to extend the disk.

B. In the Google Cloud Console, increase the size of the persistent disk, and use the fdisk command to verify that the new space is ready to use

C. In the Google Cloud Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

D. In the Google Cloud Console, create a new persistent disk attached to the VM, and configure the database service to move the files to the new disk.

Correct Answer: D

QUESTION 14

You are migrating a telehealth care company\\'s on-premises data center to Google Cloud. The migration plan specifies:

PostgreSQL databases must be migrated to a multi-region backup configuration with cross- region replicas to allow

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test |
PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

7 / 8

restore and failover in multiple scenarios. MySQL databases handle personally identifiable information (PII) and require data residency compliance at the regional level.

You want to set up the environment with minimal administrative effort. What should you do?

A. Set up Cloud Logging and Cloud Monitoring with Cloud Functions to send an alert every time a new database instance is created, and manually validate the region.

B. Set up different organizations for each database type, and apply policy constraints at the organization level.

C. Set up Pub/Sub to ingest data from Cloud Logging, send an alert every time a new database instance is created, and manually validate the region.

D. Set up different projects for PostgreSQL and MySQL databases, and apply organizational policy constraints at a project level.

Correct Answer: B

QUESTION 15

Your organization has a ticketing system that needs an online marketing analytics and reporting application. You need to select a relational database that can manage hundreds of terabytes of data to support this new application. Which database should you use?

A. Cloud SQL

B. BigQuery

C. Cloud Spanner

D. Bigtable

Correct Answer: B

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps

PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test

PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

Latest PROFESSIONAL-CLOUD-DATABASE-ENGINEER Dumps | PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test | PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps

8 / 8