

QUESTIONS & ANSWERS

Kill your exam at first Attempt



Google

Google-PCDE

Google Cloud Certified - Professional Cloud Database Engineer

<https://killexams.com/pass4sure/exam-detail/Google-PCDE>



Question: 30

Your company has PostgreSQL databases on-premises and on Amazon Web Services (AWS). You are planning multiple database migrations to Cloud SQL in an effort to reduce costs and downtime. You want to follow Google-recommended practices and use Google native data migration tools. You also want to closely monitor the migrations as part of the cutover strategy .

What should you do?

- A. Use Database Migration Service to migrate all databases to Cloud SQL.
- B. Use Database Migration Service for one-time migrations, and use third-party or partner tools for change data capture (CDC) style migrations.
- C. Use data replication tools and CDC tools to enable migration.
- D. Use a combination of Database Migration Service and partner tools to support the data migration strategy.

Answer: A

Question: 31

Your organization operates in a highly regulated industry. Separation of concerns (SoC) and security principle of least privilege (PoLP) are critical. The operations team consists of: Person A is a database administrator.

Person B is an analyst who generates metric reports.

Application C is responsible for automatic backups.

You need to assign roles to team members for Cloud Spanner .

Which roles should you assign?

- A. roles/spanner.databaseAdmin for Person A
roles/spanner.databaseReader for Person B
roles/spanner.backupWriter for Application C
- B. roles/spanner.databaseAdmin for Person A
roles/spanner.databaseReader for Person B
roles/spanner.backupAdmin for Application C
- C. roles/spanner.databaseAdmin for Person A
roles/spanner.databaseUser for Person B
roles/spanner.databaseReader for Application C
- D. roles/spanner.databaseAdmin for Person A
roles/spanner.databaseUser for Person B
roles/spanner.backupWriter for Application C

Answer: B

Question: 32

You are setting up a Bare Metal Solution environment. You need to update the operating system to the latest version. You need to connect the Bare Metal Solution environment to the internet so you can receive software updates .

What should you do?

- A. Setup a static external IP address in your VPC network.
- B. Set up bring your own IP (BYOIP) in your VPC.
- C. Set up a Cloud NAT gateway on the Compute Engine VM.
- D. Set up Cloud NAT service.

Answer: C

Question: 33

Your organization has a critical business app that is running with a Cloud SQL for MySQL backend database. Your company wants to build the most fault-tolerant and highly available solution possible. You need to ensure that the application database can survive a zonal and regional failure with a primary region of us-central1 and the backup region of us-east1 .

What should you do?

- A. Provision a Cloud SQL for MySQL instance in us-central1-a.
Create a multiple-zone instance in us-west1-b.
Create a read replica in us-east1-c.
- B. Provision a Cloud SQL for MySQL instance in us-central1-a.
Create a multiple-zone instance in us-central1-b.
Create a read replica in us-east1-b.
- C. Provision a Cloud SQL for MySQL instance in us-central1-a.
Create a multiple-zone instance in us-east-b.
Create a read replica in us-east1-c.
- D. Provision a Cloud SQL for MySQL instance in us-central1-a.
Create a multiple-zone instance in us-east1-b.
Create a read replica in us-central1-b.

Answer: B

Question: 34

Your customer is running a MySQL database on-premises with read replicas. The nightly incremental backups are expensive and add maintenance overhead. You want to follow Google-recommended practices to migrate the database to Google Cloud, and you need to ensure minimal downtime .

What should you do?

- A. Create a Google Kubernetes Engine (GKE) cluster, install MySQL on the cluster, and then import the dump file.
- B. Use the mysqldump utility to take a backup of the existing on-premises database, and then import it into Cloud SQL.
- C. Create a Compute Engine VM, install MySQL on the VM, and then import the dump file.
- D. Create an external replica, and use Cloud SQL to synchronize the data to the replica.

Answer: D

Question: 35

You are managing multiple applications connecting to a database on Cloud SQL for PostgreSQL. You need to be able to monitor database performance to easily identify applications with long-running and resource-intensive queries .

What should you do?

- A. Use log messages produced by Cloud SQL.
- B. Use Query Insights for Cloud SQL.
- C. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- D. Use Cloud SQL instance monitoring in the Google Cloud Console.

Answer: B

Question: 36

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.

Answer: C

Question: 37

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.

Answer: B

Question: 38

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.

Answer: B, C

Question: 39

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 SL

- A. 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.

Answer: C

Question: 40

Your team uses thousands of connected IoT devices to collect device maintenance data for your oil and gas customers in real time. You want to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by these devices. You need a managed solution that is highly scalable, supports a multi-cloud strategy, and offers low latency for these IoT devices .

What should you do?

- A. Use Firestore with Looker.
- B. Use Cloud Spanner with Data Studio.
- C. Use MongoDB Atlas with Charts.
- D. Use Bigtable with Looker.

Answer: C

For More exams visit <https://killexams.com/vendors-exam-list>

