



**Vendor:** Google

**Exam Code:** Professional-Cloud-Database-Engineer

**Exam Name:** Google Professional Cloud Database  
Engineer

**Version:** DEMO

### QUESTION 1

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**

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/replication/configure-replication-from-external>

### QUESTION 2

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**

### QUESTION 3

Your application follows a microservices architecture and uses a single large Cloud SQL instance, which is starting to have performance issues as your application grows. In the Cloud Monitoring dashboard, the CPU utilization looks normal. You want to follow Google-recommended practices to resolve and prevent these performance issues while avoiding any major refactoring. What should you do?

- A. Use Cloud Spanner instead of Cloud SQL.
- B. Increase the number of CPUs for your instance.
- C. Increase the storage size for the instance.
- D. Use many smaller Cloud SQL instances.

**Answer: D**

### QUESTION 4

You need to perform a one-time migration of data from a running Cloud SQL for MySQL instance in the us-central1 region to a new Cloud SQL for MySQL instance in the us-east1 region. You want to follow Google-recommended practices to minimize performance impact on the currently running instance. What should you do?

- A. Create and run a Dataflow job that uses JdbcIO to copy data from one Cloud SQL instance to another.
- B. Create two Datastream connection profiles, and use them to create a stream from one Cloud SQL instance to another.
- C. Create a SQL dump file in Cloud Storage using a temporary instance, and then use that file to import into a new instance.
- D. Create a CSV file by running the SQL statement SELECT...INTO OUTFILE, copy the file to a Cloud Storage bucket, and import it into a new instance.

**Answer: C**

#### QUESTION 5

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?

- 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: A**

#### QUESTION 6

Your organization deployed a new version of a critical application that uses Cloud SQL for MySQL with high availability (HA) and binary logging enabled to store transactional information. The latest release of the application had an error that caused massive data corruption in your Cloud SQL for MySQL database. You need to minimize data loss. What should you do?

- A. Open the Google Cloud Console, navigate to SQL > Backups, and select the last version of the automated backup before the corruption.
- B. Reload the Cloud SQL for MySQL database using the LOAD DATA command to load data from CSV files that were used to initialize the instance.
- C. Perform a point-in-time recovery of your Cloud SQL for MySQL database, selecting a date and time before the data was corrupted.
- D. Fail over to the Cloud SQL for MySQL HA instance. Use that instance to recover the transactions that occurred before the corruption.

**Answer: C**

#### QUESTION 7

You plan to use Database Migration Service to migrate data from a PostgreSQL on-premises instance to Cloud SQL. You need to identify the prerequisites for creating and automating the task. What should you do? (Choose two.)

- A. Drop or disable all users except database administration users.
- B. Disable all foreign key constraints on the source PostgreSQL database.
- C. Ensure that all PostgreSQL tables have a primary key.
- D. Shut down the database before the Data Migration Service task is started.
- E. Ensure that pglogical is installed on the source PostgreSQL database.

**Answer:** CE

**Explanation:**

<https://cloud.google.com/database-migration/docs/postgres/faq>

#### QUESTION 8

You are using Compute Engine on Google Cloud and your data center to manage a set of MySQL databases in a hybrid configuration. You need to create replicas to scale reads and to offload part of the management operation. What should you do?

- A. Use external server replication.
- B. Use Data Migration Service.
- C. Use Cloud SQL for MySQL external replica.
- D. Use the mysqldump utility and binary logs.

**Answer:** C

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/replication/external-server>

#### QUESTION 9

Your company is shutting down their data center and migrating several MySQL and PostgreSQL databases to Google Cloud. Your database operations team is severely constrained by ongoing production releases and the lack of capacity for additional on-premises backups. You want to ensure that the scheduled migrations happen with minimal downtime and that the Google Cloud databases stay in sync with the on-premises data changes until the applications can cut over. What should you do? (Choose two.)

- A. Use Database Migration Service to migrate the databases to Cloud SQL.
- B. Use a cross-region read replica to migrate the databases to Cloud SQL.
- C. Use replication from an external server to migrate the databases to Cloud SQL.
- D. Use an external read replica to migrate the databases to Cloud SQL.
- E. Use a read replica to migrate the databases to Cloud SQL.

**Answer:** AC

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#basic-steps-to-change-parallel-replication-flags>

#### QUESTION 10

Your company is migrating the existing infrastructure for a highly transactional application to Google Cloud. You have several databases in a MySQL database instance and need to decide how to transfer the data to Cloud SQL. You need to minimize the downtime for the migration of your 500 GB instance. What should you do?

- A.
  - 1. Create a Cloud SQL for MySQL instance for your databases, and configure Datastream to stream your database changes to Cloud SQL.
  - 2. Select the Backfill historical data check box on your stream configuration to initiate Datastream to backfill any data that is out of sync between the source and destination.
  - 3. Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.
- B.
  - 1. Create migration job using Database Migration Service.
  - 2. Set the migration job type to Continuous, and allow the databases to complete the full dump phase and start sending data in change data capture (CDC) mode.
  - 3. Wait for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete.
  - 4. Update your application connections to the new instance.
- C.
  - 1. Create migration job using Database Migration Service.
  - 2. Set the migration job type to One-time, and perform this migration during a maintenance window.
  - 3. Stop all write workloads to the source database and initiate the dump. Wait for the dump to be loaded into the Cloud SQL destination database and the destination database to be promoted to the primary database.
  - 4. Update your application connections to the new instance.
- D.
  - 1. Use the mysqldump utility to manually initiate a backup of MySQL during the application maintenance window.
  - 2. Move the files to Cloud Storage, and import each database into your Cloud SQL instance.
  - 3. Continue to dump each database until all the databases are migrated.
  - 4. Update your application connections to the new instance.

**Answer: B**

#### QUESTION 11

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**

**Explanation:**

[https://cloud.google.com/sql/docs/mysql/using-ood-recommender#apply\\_recommendations](https://cloud.google.com/sql/docs/mysql/using-ood-recommender#apply_recommendations)

#### QUESTION 12

You are managing a mission-critical Cloud SQL for PostgreSQL instance. Your application team is running important transactions on the database when another DBA starts an on-demand backup. You want to verify the status of the backup. What should you do?

- A. Check the [cloudsql.googleapis.com/postgres.log](https://cloudsql.googleapis.com/postgres.log) instance log.
- B. Perform the `gcloud sql operations list` command.
- C. Use Cloud Audit Logs to verify the status.

D. Use the Google Cloud Console.

**Answer: B**

### QUESTION 13

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.

**Answer: A**

### QUESTION 14

Your company uses Bigtable for a user-facing application that displays a low-latency real-time dashboard. You need to recommend the optimal storage type for this read-intensive database. What should you do?

- A. Recommend solid-state drives (SSD).
- B. Recommend splitting the Bigtable instance into two instances in order to load balance the concurrent reads.
- C. Recommend hard disk drives (HDD).
- D. Recommend mixed storage types.

**Answer: A**

**Explanation:**

<https://cloud.google.com/bigtable/docs/choosing-ssd-hdd>

### QUESTION 15

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. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.  
2. Create a multiple-zone instance in us-west1-b.  
3. Create a read replica in us-east1-c.
- B. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.  
2. Create a multiple-zone instance in us-central1-b.  
3. Create a read replica in us-east1-b.
- C. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.  
2. Create a multiple-zone instance in us-east-b.  
3. Create a read replica in us-east1-c.
- D. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.  
2. Create a multiple-zone instance in us-east1-b.

3. Create a read replica in us-central1-b.

**Answer:** B

**QUESTION 16**

You are building an Android game that needs to store data on a Google Cloud serverless database. The database will log user activity, store user preferences, and receive in-game updates. The target audience resides in developing countries that have intermittent internet connectivity. You need to ensure that the game can synchronize game data to the backend database whenever an internet network is available. What should you do?

- A. Use Firestore.
- B. Use Cloud SQL with an external (public) IP address.
- C. Use an in-app embedded database.
- D. Use Cloud Spanner.

**Answer:** A

**Explanation:**

<https://firebase.google.com/docs/firestore>

**QUESTION 17**

You released a popular mobile game and are using a 50 TB Cloud Spanner instance to store game data in a PITR-enabled production environment. When you analyzed the game statistics, you realized that some players are exploiting a loophole to gather more points to get on the leaderboard. Another DBA accidentally ran an emergency bugfix script that corrupted some of the data in the production environment. You need to determine the extent of the data corruption and restore the production environment. What should you do? (Choose two.)

- A. If the corruption is significant, use backup and restore, and specify a recovery timestamp.
- B. If the corruption is significant, perform a stale read and specify a recovery timestamp. Write the results back.
- C. If the corruption is significant, use import and export.
- D. If the corruption is insignificant, use backup and restore, and specify a recovery timestamp.
- E. If the corruption is insignificant, perform a stale read and specify a recovery timestamp. Write the results back.

**Answer:** AE

**Explanation:**

<https://cloud.google.com/spanner/docs/pitr>

<https://cloud.google.com/spanner/docs/backup/restore-backup>

## Thank You for Trying Our Product

### Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



**10% Discount Coupon Code: ASTR14**



