

# PROFESSIONAL-CLOUD-DATABASE- ENGINEER<sup>Q&As</sup>

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.certbus.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

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



### QUESTION 1

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` 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.

Correct Answer: B

<https://cloud.google.com/sql/docs/postgres/backup-recovery/backups#troubleshooting-backups> Under Troubleshooting: Issue: "You can't see the current operation's status." The Google Cloud console reports only success or failure when the operation is done. It isn't designed to show warnings or other updates. Run the `gcloud sql operations list` command to list all operations for the given Cloud SQL instance.

---

### QUESTION 2

You are migrating an on-premises application to Compute Engine and Cloud SQL. The application VMs will live in their own project, separate from the Cloud SQL instances which have their own project. What should you do to configure the networks?

- A. Create a new VPC network in each project, and use VPC Network Peering to connect the two together.
- B. Create a Shared VPC that both the application VMs and Cloud SQL instances will use.
- C. Use the default networks, and leverage Cloud VPN to connect the two together.
- D. Place both the application VMs and the Cloud SQL instances in the default network of each project.

Correct Answer: B

[https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5\\_HPXytY?pli=1](https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5_HPXytY?pli=1)

---

### QUESTION 3

You want to migrate an on-premises 100 TB Microsoft SQL Server database to Google Cloud over a 1 Gbps network link. You have 48 hours allowed downtime to migrate this database. What should you do? (Choose two.)

- A. Use a change data capture (CDC) migration strategy.
- B. Move the physical database servers from on-premises to Google Cloud.
- C. Keep the network bandwidth at 1 Gbps, and then perform an offline data migration.
- D. Increase the network bandwidth to 2 Gbps, and then perform an offline data migration.

E. Increase the network bandwidth to 10 Gbps, and then perform an offline data migration.

Correct Answer: AE

[https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online\\_versus\\_offline\\_transfer](https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online_versus_offline_transfer)

#### QUESTION 4

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.

Correct Answer: AE

<https://cloud.google.com/spanner/docs/pitr#ways-to-recover> To recover the entire database, backup or export the database specifying a timestamp in the past and then restore or import it to a new database. This is typically used to recover from data corruption issues when you have to revert the entire database to a point-in-time before the corruption occurred. This part describes significant corruption -A

To recover a portion of the database, perform a stale read specifying a query-condition and timestamp in the past, and then write the results back into the live database. This is typically used for surgical operations on a live database. For example, if you accidentally delete a particular row or incorrectly update a subset of data, you can recover it with this method. This describes insignificant corruption case ?E <https://cloud.google.com/spanner/docs/pitr>  
<https://cloud.google.com/spanner/docs/backup/restore-backup>

#### QUESTION 5

You have a large Cloud SQL for PostgreSQL instance. The database instance is not mission-critical, and you want to minimize operational costs. What should you do to lower the cost of backups in this environment?

- A. Set the automated backups to occur every other day to lower the frequency of backups.
- B. Change the storage tier of the automated backups from solid-state drive (SSD) to hard disk drive (HDD).
- C. Select a different region to store your backups.
- D. Reduce the number of automated backups that are retained to two (2).

Correct Answer: D

By default, for each instance, Cloud SQL retains seven automated backups, in addition to on-demand backups. You can configure how many automated backups to retain (from 1 to 365). We charge a lower rate for backup storage than for other types of instances. <https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>

---

#### QUESTION 6

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 7

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: A

"Deny maintenance period. A block of days in which Cloud SQL does not schedule maintenance. Deny maintenance periods can be up to 90 days long. " <https://cloud.google.com/sql/docs/mysql/maintenance>

---

#### QUESTION 8

You are designing a database strategy for a new web application in one region. You need to minimize write latency.

What should you do?

- A. Use Cloud SQL with cross-region replicas.
- B. Use high availability (HA) Cloud SQL with multiple zones.
- C. Use zonal Cloud SQL without high availability (HA).
- D. Use Cloud Spanner in a regional configuration.

Correct Answer: D

<https://docs.google.com/forms/d/e/1FAIpQLSfZ77ZnuUL0NpU-bOtO5QUkC0cnRCe5YKMiubLXwfv3abBqkg/viewform>

---

### QUESTION 9

You need to migrate existing databases from Microsoft SQL Server 2016 Standard Edition on a single Windows Server 2019 Datacenter Edition to a single Cloud SQL for SQL Server instance. During the discovery phase of your project, you notice that your on-premises server peaks at around 25,000 read IOPS. You need to ensure that your Cloud SQL instance is sized appropriately to maximize read performance. What should you do?

- A. Create a SQL Server 2019 Standard on Standard machine type with 4 vCPUs, 15 GB of RAM, and 800 GB of solid-state drive (SSD).
- B. Create a SQL Server 2019 Standard on High Memory machine type with at least 16 vCPUs, 104 GB of RAM, and 200 GB of SSD.
- C. Create a SQL Server 2019 Standard on High Memory machine type with 16 vCPUs, 104 GB of RAM, and 4 TB of SSD.
- D. Create a SQL Server 2019 Enterprise on High Memory machine type with 16 vCPUs, 104 GB of RAM, and 500 GB of SSD.

Correct Answer: C

Given that Google SSD performance is related to the size of the disk in an order of 30 IOPS for each GB, it would require at least 833 GB to handle 25000 IOPS, the only answer that exceeds this value is C.

<https://cloud.google.com/compute/docs/disks/performance>

---

### QUESTION 10

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.

Correct Answer: C

<https://cloud.google.com/sql/docs/mysql/import-export#serverless>

---

### QUESTION 11

An analytics team needs to read data out of Cloud SQL for SQL Server and update a table in Cloud Spanner. You need to create a service account and grant least privilege access using predefined roles. What roles should you assign to the service account?

- A. roles/cloudsql.viewer and roles/spanner.databaseUser
- B. roles/cloudsql.editor and roles/spanner.admin
- C. roles/cloudsql.client and roles/spanner.databaseReader
- D. roles/cloudsql.instanceUser and roles/spanner.databaseUser

Correct Answer: A

To read data out of Cloud SQL for SQL Server, you need to use a service account with the roles/cloudsql.viewer role on the Cloud SQL instance. This role grants the service account permission to read data from the instance. Whereas roles/cloudsql.instanceUser will only allow to login to cloud SQL instance. No resource will be allowed to view.

---

### QUESTION 12

You manage a meeting booking application that uses Cloud SQL. During an important launch, the Cloud SQL instance went through a maintenance event that resulted in a downtime of more than 5 minutes and adversely affected your production application. You need to immediately address the maintenance issue to prevent any unplanned events in the future. What should you do?

- A. Set your production instance's maintenance window to non-business hours.
- B. Migrate the Cloud SQL instance to Cloud Spanner to avoid any future disruptions due to maintenance.
- C. Contact Support to understand why your Cloud SQL instance had a downtime of more than 5 minutes.
- D. Use Cloud Scheduler to schedule a maintenance window of no longer than 5 minutes.

Correct Answer: A

---

### QUESTION 13

Your organization needs to migrate a critical, on-premises MySQL database to Cloud SQL for MySQL. The on-premises database is on a version of MySQL that is supported by Cloud SQL and uses the InnoDB storage engine. You need to migrate the database while preserving transactions and minimizing downtime. What should you do?

- A. Use Database Migration Service to connect to your on-premises database, and choose continuous replication. After

the on-premises database is migrated, promote the Cloud SQL for MySQL instance, and connect applications to your Cloud SQL instance.

B. Build a Cloud Data Fusion pipeline for each table to migrate data from the on-premises MySQL database to Cloud SQL for MySQL. Schedule downtime to run each Cloud Data Fusion pipeline. Verify that the migration was successful. Re-point the applications to the Cloud SQL for MySQL instance.

C. Pause the on-premises applications. Use the mysqldump utility to dump the database content in compressed format. Run gsutil

Correct Answer: A

<https://cloud.google.com/database-migration/docs/mysql/configure-source-database>

To migrate the database while preserving transactions and minimizing downtime, you should use Database Migration Service. This service will allow you to migrate the database in a way that is transparent to your users and applications. It will also allow you to test the migration before you make it live, so that you can be sure that everything will work as expected.

#### QUESTION 14

Your company is evaluating Google Cloud database options for a mission-critical global payments gateway application. The application must be available 24/7 to users worldwide, horizontally scalable, and support open source databases. You need to select an automatically shardable, fully managed database with 99.999% availability and strong transactional consistency. What should you do?

- A. Select Bare Metal Solution for Oracle.
- B. Select Cloud SQL.
- C. Select Bigtable.
- D. Select Cloud Spanner.

Correct Answer: D

The application must be available 24/7 to users worldwide, horizontally scalable, and support open source databases.

#### QUESTION 15

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.

[Latest PROFESSIONAL-CL  
OUD-DATABASE-  
ENGINEER Dumps](#)

[PROFESSIONAL-CLOUD-  
DATABASE-ENGINEER  
PDF Dumps](#)

[PROFESSIONAL-CLOUD-  
DATABASE-ENGINEER  
Study Guide](#)