

# 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.leads4pass.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 have an application that sends banking events to Bigtable cluster-a in us-east. You decide to add cluster-b in us-central1. Cluster-a replicates data to cluster-b. You need to ensure that Bigtable continues to accept read and write requests if one of the clusters becomes unavailable and that requests are routed automatically to the other cluster. What deployment strategy should you use?

- A. Use the default app profile with single-cluster routing.
- B. Use the default app profile with multi-cluster routing.
- C. Create a custom app profile with multi-cluster routing.
- D. Create a custom app profile with single-cluster routing.

Correct Answer: C

<https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile> The question states that a single cluster existed first, then a second cluster was added. Google's documentation states, "if you created the instance with one cluster, the default app profile uses single-cluster routing. This ensures that adding additional clusters later does not change the behavior of your existing applications". Simply adding a second cluster does not change the default profile from single-cluster routing to multi-cluster routing. Since you need multi-cluster routing, you're going to need a custom app profile. So C is correct. <https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile>

---

## QUESTION 2

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: A

## QUESTION 3

You are managing a Cloud SQL for PostgreSQL instance in Google Cloud. You have a primary instance in region 1 and a read replica in region 2. After a failure of region 1, you need to make the Cloud SQL instance available again. You want to minimize data loss and follow Google-recommended practices. What should you do?

- A. Restore the Cloud SQL instance from the automatic backups in region 3.
- B. Restore the Cloud SQL instance from the automatic backups in another zone in region 1.

C. Check "Lag Bytes" in the monitoring dashboard for the primary instance in the read replica instance. Check the replication status using `pg_catalog.pg_last_wal_receive_lsn()`. Then, fail over to region 2 by promoting the read replica instance.

D. Check your instance operational log for the automatic failover status. Look for time, type, and status of the operations. If the failover operation is successful, no action is necessary. Otherwise, manually perform `gcloud sql instances failover`.

Correct Answer: C

[https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#disaster\\_recovery](https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#disaster_recovery)

---

## QUESTION 4

You work for a large retail and ecommerce company that is starting to extend their business globally. Your company plans to migrate to Google Cloud. You want to use platforms that will scale easily, handle transactions with the least amount of latency, and provide a reliable customer experience. You need a storage layer for sales transactions and current inventory levels. You want to retain the same relational schema that your existing platform uses. What should you do?

A. Store your data in Firestore in a multi-region location, and place your compute resources in one of the constituent regions.

B. Deploy Cloud Spanner using a multi-region instance, and place your compute resources close to the default leader region.

C. Build an in-memory cache in Memorystore, and deploy to the specific geographic regions where your application resides.

D. Deploy a Bigtable instance with a cluster in one region and a replica cluster in another geographic region.

Correct Answer: B

## QUESTION 5

You are running an instance of Cloud Spanner as the backend of your ecommerce website. You learn that the quality assurance (QA) team has doubled the number of their test cases. You need to create a copy of your Cloud Spanner database in a new test environment to accommodate the additional test cases. You want to follow Google-recommended practices. What should you do?

A. Use Cloud Functions to run the export in Avro format.

B. Use Cloud Functions to run the export in text format.

C. Use Dataflow to run the export in Avro format.

D. Use Dataflow to run the export in text format.

Correct Answer: C

<https://cloud.google.com/spanner/docs/import-export-overview#file-format>

---

**QUESTION 6**

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: D

Granular instance is available in Public Preview. With this feature, you can run workloads on Spanner at as low as 1/10th the cost of regular instances, <https://cloud.google.com/blog/products/databases/get-more-out-of-spanner-with-granularinstance-sizing>

---

**QUESTION 7**

Your DevOps team is using Terraform to deploy applications and Cloud SQL databases. After every new application change is rolled out, the environment is torn down and recreated, and the persistent database layer is lost. You need to prevent the database from being dropped. What should you do?

- A. Set Terraform `deletion_protection` to true.
- B. Rerun terraform apply.
- C. Create a read replica.
- D. Use point-in-time-recovery (PITR) to recover the database.

Correct Answer: A

From Google's documentation, "For stateful resources, such as databases, ensure that deletion protection is enabled. The syntax is: `lifecycle { prevent_destroy = true }`" <https://cloud.google.com/docs/terraform/best-practices-forterraform#stateful-resources>

---

**QUESTION 8**

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 9

You need to migrate a 1 TB PostgreSQL database from a Compute Engine VM to Cloud SQL for PostgreSQL. You want to ensure that there is minimal downtime during the migration. What should you do?

- A. Export the data from the existing database, and load the data into a new Cloud SQL database.
- B. Use Migrate for Compute Engine to complete the migration.
- C. Use Datastream to complete the migration.
- D. Use Database Migration Service to complete the migration.

Correct Answer: D

<https://www.cloudskillsboost.google/focuses/22792?parent=catalog>

---

## QUESTION 10

Your organization has an existing app that just went viral. The app uses a Cloud SQL for MySQL backend database that is experiencing slow disk performance while using hard disk drives (HDDs). You need to improve performance and

reduce disk I/O wait times.

What should you do?

- A. Export the data from the existing instance, and import the data into a new instance with solid-state drives (SSDs).
- B. Edit the instance to change the storage type from HDD to SSD.
- C. Create a high availability (HA) failover instance with SSDs, and perform a failover to the new instance.
- D. Create a read replica of the instance with SSDs, and perform a failover to the new instance

Correct Answer: A

<https://stackoverflow.com/questions/72034607/can-i-change-storage-type-from-hdd-to-ssd-on-cloud-sql-after-creating-an-instanc>

---

## QUESTION 11

Your organization is running a low-latency reporting application on Microsoft SQL Server. In addition to the database engine, you are using SQL Server Analysis Services (SSAS), SQL Server Reporting Services (SSRS), and SQL Server Integration Services (SSIS) in your on-premises environment. You want to migrate your Microsoft SQL Server database instances to Google Cloud. You need to ensure minimal disruption to the existing architecture during migration. What

should you do?

- A. Migrate to Cloud SQL for SQL Server.
- B. Migrate to Cloud SQL for PostgreSQL.
- C. Migrate to Compute Engine.
- D. Migrate to Google Kubernetes Engine (GKE).

Correct Answer: C

<https://cloud.google.com/sql/docs/sqlserver/features>

---

## QUESTION 12

You are designing a database strategy for a new web application. You plan to start with a small pilot in one country and eventually expand to millions of users in a global audience. You need to ensure that the application can run 24/7 with minimal downtime for maintenance. What should you do?

- A. Use Cloud Spanner in a regional configuration.
- B. Use Cloud Spanner in a multi-region configuration.
- C. Use Cloud SQL with cross-region replicas.
- D. Use highly available Cloud SQL with multiple zones.

Correct Answer: A

<https://docs.google.com/forms/d/e/1FAIpQLSfZ77ZnuUL0NpU-bOtO5QUkC0cnRCe5YKMiubLXwfV3abBqkg/viewform>

---

## QUESTION 13

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 14

You recently launched a new product to the US market. You currently have two Bigtable clusters in one US region to serve all the traffic. Your marketing team is planning an immediate expansion to APAC. You need to roll out the regional expansion while implementing high availability according to Google-recommended practices. What should you do?

- A. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone europe-west1-d cluster-c in zone asia-east1-b
- B. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central1-b cluster-c in zone us-east1-a
- C. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone australia-southeast1-a cluster-c in zone europe-west1-d cluster-d in zone asia-east1-b
- D. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central2-a cluster-c in zone asia-northeast1-b cluster-d in zone asia-east1-b

Correct Answer: D

<https://cloud.google.com/bigtable/docs/replication-settings#regional-failover>

---

## QUESTION 15

Your company is using Cloud SQL for MySQL with an internal (private) IP address and wants to replicate some tables into BigQuery in near-real time for analytics and machine learning. You need to ensure that replication is fast and reliable and uses Google-managed services. What should you do?

- A. Develop a custom data replication service to send data into BigQuery.
- B. Use Cloud SQL federated queries.
- C. Use Database Migration Service to replicate tables into BigQuery.
- D. Use Datastream to capture changes, and use Dataflow to write those changes to BigQuery.

Correct Answer: D

"Datastream is a serverless and easy-to-use Change Data Capture (CDC) and replication service that allows you to synchronize data across heterogeneous databases, storage systems, and applications reliably and with minimal latency. Datastream supports change data streaming from Oracle and MySQL databases to Google Cloud Storage (GCS). The service offers streamlined integration with Dataflow templates to power up to date materialized views in BigQuery for analytics, replicate their databases into Cloud SQL or Cloud Spanner for database synchronization, or leverage the event stream directly from GCS to realize event-driven architectures."

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

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER Practice Test](#)

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