



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.geekcert.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



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-database-engineer.html>
2024 Latest geekcert PROFESSIONAL-CLOUD-DATABASE-ENGINEER PDF
and VCE dumps Download

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





QUESTION 1

You want to migrate an existing on-premises application to Google Cloud. Your application supports semi-structured data ingested from 100,000 sensors, and each sensor sends 10 readings per second from manufacturing plants. You need to make this data available for real-time monitoring and analysis. What should you do?

- A. Deploy the database using Cloud SQL.
- B. Use BigQuery, and load data in batches.
- C. Deploy the database using Bigtable.
- D. Deploy the database using Cloud Spanner.

Correct Answer: C

Bigtable is a scalable, fully managed, and high-performance NoSQL database service that can handle semi-structured data and support real-time monitoring and analysis. Cloud SQL is a relational database service that does not support semi-structured data. BigQuery is a data warehouse service that is optimized for batch processing and analytics, not real-time monitoring. Cloud Spanner is a relational database service that supports semi-structured data with JSON data type, but it is more expensive and complex than Bigtable for this use case.

QUESTION 2

You are writing an application that will run on Cloud Run and require a database running in the Cloud SQL managed service. You want to secure this instance so that it only receives connections from applications running in your VPC environment in Google Cloud. What should you do?

- A. 1. Create your instance with a specified external (public) IP address.

2.

Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.

3.

Use Cloud SQL Auth proxy to connect to the instance.

- B. 1. Create your instance with a specified external (public) IP address.

2.

Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.

3.

Connect to the instance using a connection pool to best manage connections to the instance.

- C. 1. Create your instance with a specified internal (private) IP address.

2.

Choose the VPC with private service connection configured.



3.

Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.

4.

Use Cloud SQL Auth proxy to connect to the instance.

D. 1. Create your instance with a specified internal (private) IP address.

2.

Choose the VPC with private service connection configured.

3.

Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.

4.

Connect to the instance using a connection pool to best manage connections to the instance.

Correct Answer: D

<https://cloud.google.com/sql/docs/mysql/connect-run#configure> <https://cloud.google.com/sql/docs/mysql/connect-run#connection-pools>

QUESTION 3

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: A

https://cloud.google.com/compute/docs/disks/resize-persistent-disk#resize_partitions

QUESTION 4

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 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>

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

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

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER Braindumps](#)