



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



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





QUESTION 1

Your company uses Cloud Spanner for a mission-critical inventory management system that is globally available. You recently loaded stock keeping unit (SKU) and product catalog data from a company acquisition and observed hot-spots in the Cloud Spanner database. You want to follow Google-recommended schema design practices to avoid performance degradation. What should you do? (Choose two.)

- A. Use an auto-incrementing value as the primary key.
- B. Normalize the data model.
- C. Promote low-cardinality attributes in multi-attribute primary keys.
- D. Promote high-cardinality attributes in multi-attribute primary keys.
- E. Use bit-reverse sequential value as the primary key.

Correct Answer: AD

QUESTION 2

You are migrating an on-premises application to Google Cloud. The application requires a high availability (HA) PostgreSQL database to support business-critical functions. Your company's disaster recovery strategy requires a recovery time objective (RTO) and recovery point objective (RPO) within 30 minutes of failure. You plan to use a Google Cloud managed service. What should you do to maximize uptime for your application?

- A. Deploy Cloud SQL for PostgreSQL in a regional configuration. Create a read replica in a different zone in the same region and a read replica in another region for disaster recovery.
- B. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Take periodic backups, and use this backup to restore to a new Cloud SQL for PostgreSQL instance in another region during a disaster recovery event.
- C. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Create a cross-region read replica, and promote the read replica as the primary node for disaster recovery.
- D. Migrate the PostgreSQL database to multi-regional Cloud Spanner so that a single region outage will not affect your application. Update the schema to support Cloud Spanner data types, and refactor the application.

Correct Answer: C

QUESTION 3

Your ecommerce website captures user clickstream data to analyze customer traffic patterns in real time and support personalization features on your website. You plan to analyze this data using big data tools. You need a low-latency solution that can store 8 TB of data and can scale to millions of read and write requests per second. What should you do?

- A. Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis.
- B. Deploy a Cloud SQL environment with read replicas for improved performance. Use Datastream to export data to Cloud Storage and analyze with Dataproc and the Cloud Storage connector.



- C. Use Memorystore to handle your low-latency requirements and for real-time analytics.
- D. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

Correct Answer: B

QUESTION 4

Your digital-native business runs its database workloads on Cloud SQL. Your website must be globally accessible 24/7. You need to prepare your Cloud SQL instance for high availability (HA). You want to follow Google-recommended practices. What should you do? (Choose two.)

- A. Set up manual backups.
- B. Create a PostgreSQL database on-premises as the HA option.
- C. Configure single zone availability for automated backups.
- D. Enable point-in-time recovery.
- E. Schedule automated backups.

Correct Answer: BD

QUESTION 5

Your team is building a new inventory management application that will require read and write database instances in multiple Google Cloud regions around the globe. Your database solution requires 99.99% availability and global transactional consistency. You need a fully managed backend relational database to store inventory changes. What should you do?

- A. Use Bigtable.
- B. Use Firestore.
- C. Use Cloud SQL for MySQL
- D. Use Cloud Spanner.

Correct Answer: C

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

[PROFESSIONAL-CLOUD-DATABASE-ENGINEER Exam Questions](#)

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