

KillTest

質量更高 服務更好



練習題

<https://www.killtest.net>

一年免費更新服務

**Exam : Professional Cloud
Database Engineer**

**Title : Google Cloud Certified -
Professional Cloud
Database Engineer**

Version : DEMO

1.You are developing a new application on a VM that is on your corporate network. The application will use Java Database Connectivity (JDBC) to connect to Cloud SQL for PostgreSQL. Your Cloud SQL instance is configured with IP address 192.168.3.48, and SSL is disabled. You want to ensure that your application can access your database instance without requiring configuration changes to your database. What should you do?

- A. Define a connection string using your Google username and password to point to the external (public) IP address of your Cloud SQL instance.
- B. Define a connection string using a database username and password to point to the internal (private) IP address of your Cloud SQL instance.
- C. Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the internal (private) IP address of your Cloud SQL instance.
- D. Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the external (public) IP address of your Cloud SQL instance.

Answer: C

Explanation:

The Cloud SQL connectors are libraries that provide encryption and IAM-based authorization when connecting to a Cloud SQL instance. They can't provide a network path to a Cloud SQL instance if one is not already present. Other ways to connect to a Cloud SQL instance include using a database client or the Cloud SQL Auth proxy. <https://cloud.google.com/sql/docs/postgres/connect-connectors>
<https://github.com/GoogleCloudPlatform/cloud-sql-jdbc-socket-factory/blob/main/docs/jdbc-postgres.md>

2.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.

Answer: DE

Explanation:

D. Enable point-in-time recovery - This feature allows you to restore your database to a specific point in time. It helps protect against data loss and can be used in the event of data corruption or accidental data deletion.

E. Schedule automated backups - Automated backups allow you to take regular backups of your database without manual intervention. You can use these backups to restore your database in the event of data loss or corruption.

3.Your company wants to move to Google Cloud. Your current data center is closing in six months. You are running a large, highly transactional Oracle application footprint on VMWare. You need to design a solution with minimal disruption to the current architecture and provide ease of migration to Google Cloud.

What should you do?

- A. Migrate applications and Oracle databases to Google Cloud VMware Engine (VMware Engine).
- B. Migrate applications and Oracle databases to Compute Engine.
- C. Migrate applications to Cloud SQL.
- D. Migrate applications and Oracle databases to Google Kubernetes Engine (GKE).

Answer: A

Explanation:

<https://cloud.google.com/blog/products/databases/migrate-databases-to-google-cloud-vmware-engine-gcve>

4. Your customer has a global chat application that uses a multi-regional Cloud Spanner instance. The application has recently experienced degraded performance after a new version of the application was launched. Your customer asked you for assistance. During initial troubleshooting, you observed high read latency.

What should you do?

- A. Use query parameters to speed up frequently executed queries.
- B. Change the Cloud Spanner configuration from multi-region to single region.
- C. Use SQL statements to analyze SPANNER_SYS.READ_STATS* tables.
- D. Use SQL statements to analyze SPANNER_SYS.QUERY_STATS* tables.

Answer: C

Explanation:

To troubleshoot high read latency, you can use SQL statements to analyze the SPANNER_SYS.READ_STATS* tables. These tables contain statistics about read operations in Cloud Spanner, including the number of reads, read latency, and the number of read errors. By analyzing these tables, you can identify the cause of the high read latency and take appropriate action to resolve the issue. Other options, such as using query parameters to speed up frequently executed queries or changing the Cloud Spanner configuration from multi-region to single region, may not be directly related to the issue of high read latency. Similarly, analyzing the SPANNER_SYS.QUERY_STATS* tables, which contain statistics about query operations, may not be relevant to the issue of high read latency.

5. Your company has PostgreSQL databases on-premises and on Amazon Web Services (AWS). You are planning multiple database migrations to Cloud SQL in an effort to reduce costs and downtime. You want to follow Google-recommended practices and use Google native data migration tools. You also want to closely monitor the migrations as part of the cutover strategy.

What should you do?

- A. Use Database Migration Service to migrate all databases to Cloud SQL.
- B. Use Database Migration Service for one-time migrations, and use third-party or partner tools for change data capture (CDC) style migrations.
- C. Use data replication tools and CDC tools to enable migration.
- D. Use a combination of Database Migration Service and partner tools to support the data migration strategy.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/databases/tips-for-migrating-across-compatible-database-engines>