



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.passapply.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

PassApply.com

<https://www.passapply.com/professional-cloud-database-engineer.html>
2024 Latest passapply 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 are troubleshooting a connection issue with a newly deployed Cloud SQL instance on Google Cloud. While investigating the Cloud SQL Proxy logs, you see the message Error 403: Access Not Configured. What should you do?

- A. Check the app.yaml value cloud_sql_instances for a misspelled or incorrect instance connection name.
- B. Check whether your service account has cloudsql.instances.connect permission.
- C. Enable the Cloud SQL Admin API.
- D. Ensure that you are using an external (public) IP address interface.

Correct Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-auth-proxy#troubleshooting> C because in docs it says "Make sure to enable the Cloud SQL Admin API. If it is not, you see output like Error 403: Access Not Configured in your Cloud SQL

QUESTION 2

You are the primary DBA of a Cloud SQL for PostgreSQL database that supports 6 enterprise applications in production. You used Cloud SQL Insights to identify inefficient queries and now need to identify the application that is originating the inefficient queries. You want to follow Google-recommended practices. What should you do?

- A. Shut down and restart each application.
- B. Write a utility to scan database query logs.
- C. Write a utility to scan application logs.
- D. Use query tags to add application-centric database monitoring.

Correct Answer: D

https://cloud.google.com/sql/docs/postgres/using-query-insights#filter_by_query_tags

QUESTION 3

You are designing for a write-heavy application. During testing, you discover that the write workloads are performant in a regional Cloud Spanner instance but slow down by an order of magnitude in a multi-regional instance. You want to make the write workloads faster in a multi-regional instance. What should you do?

- A. Place the bulk of the read and write workloads closer to the default leader region.
- B. Use staleness of at least 15 seconds.
- C. Add more read-write replicas.
- D. Keep the total CPU utilization under 45% in each region.

Correct Answer: A



<https://cloud.google.com/spanner/docs/instance-configurations#multi-region-best-practices> Best practices For optimal performance, follow these best practices: Design a schema that prevents hotspots and other performance issues. For optimal write latency, place compute resources for write-heavy workloads within or close to the default leader region. For optimal read performance outside of the default leader region, use staleness of at least 15 seconds. To avoid single-region dependency for your workloads, place critical compute resources in at least two regions. A good option is to place them next to the two different read-write regions so that any single region outage will not impact all of your application. Provision enough compute capacity to keep high priority total CPU utilization under 45% in each region.

QUESTION 4

Your company is developing a new global transactional application that must be ACID-compliant and have 99.999% availability. You are responsible for selecting the appropriate Google Cloud database to serve as a datastore for this new application. What should you do?

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

Correct Answer: B

QUESTION 5

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>

[PROFESSIONAL-CLOUD-](#) [PROFESSIONAL-CLOUD-](#) [PROFESSIONAL-CLOUD-](#)



VCE & PDF

PassApply.com

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

[DATABASE-ENGINEER
PDF Dumps](#)

[DATABASE-ENGINEER
Practice Test](#)

[DATABASE-ENGINEER
Study Guide](#)