



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-developer.html>
2024 Latest geekcert PROFESSIONAL-CLOUD-DEVELOPER PDF and VCE
dumps Download

PROFESSIONAL-CLOUD- DEVELOPER^{Q&As}

Professional Cloud Developer

**Pass Google PROFESSIONAL-CLOUD-DEVELOPER
Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/professional-cloud-developer.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-developer.html>

2024 Latest geekcert PROFESSIONAL-CLOUD-DEVELOPER PDF and VCE dumps Download

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





QUESTION 1

You have two tables in an ANSI-SQL compliant database with identical columns that you need to quickly combine into a single table, removing duplicate rows from the result set.

What should you do?

- A. Use the JOIN operator in SQL to combine the tables.
- B. Use nested WITH statements to combine the tables.
- C. Use the UNION operator in SQL to combine the tables.
- D. Use the UNION ALL operator in SQL to combine the tables.

Correct Answer: C

Reference: https://www.techonthenet.com/sql/union_all.php

QUESTION 2

Your team develops services that run on Google Kubernetes Engine. You need to standardize their log data using Google-recommended practices and make the data more useful in the fewest number of steps. What should you do? (Choose two.)

- A. Create aggregated exports on application logs to BigQuery to facilitate log analytics.
- B. Create aggregated exports on application logs to Cloud Storage to facilitate log analytics.
- C. Write log output to standard output (stdout) as single-line JSON to be ingested into Cloud Logging as structured logs.
- D. Mandate the use of the Logging API in the application code to write structured logs to Cloud Logging.
- E. Mandate the use of the Pub/Sub API to write structured data to Pub/Sub and create a Dataflow streaming pipeline to normalize logs and write them to BigQuery for analytics.

Correct Answer: AC

Explanation: https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best_practices

QUESTION 3

You are working on a new application that is deployed on Cloud Run and uses Cloud Functions. Each time new features are added, new Cloud Functions and Cloud Run services are deployed. You use ENV variables to keep track of the services and enable interservice communication but the maintenance of the ENV variables has become difficult. You want to implement dynamic discovery in a scalable way. What should you do?

- A. Create a Service Directory Namespace. Use API calls to register the services during deployment, and query during runtime.



- B. Configure your microservices to use the Cloud Run Admin and Cloud Functions APIs to query for deployed Cloud Run services and Cloud Functions in the Google Cloud project.
- C. Deploy Hashicorp Consul on a single Compute Engine Instance Register the services with Consul during deployment and query during runtime
- D. Rename the Cloud Functions and Cloud Run services endpoints using a well- documented naming convention

Correct Answer: B

QUESTION 4

You are developing an ecommerce application that stores customer, order, and inventory data as relational tables inside Cloud Spanner. During a recent load test, you discover that Spanner performance is not scaling linearly as expected. Which of the following is the cause?

- A. The use of 64-bit numeric types for 32-bit numbers.
- B. The use of the STRING data type for arbitrary-precision values.
- C. The use of Version 1 UUIDs as primary keys that increase monotonically.
- D. The use of LIKE instead of STARTS_WITH keyword for parameterized SQL queries.

Correct Answer: C

QUESTION 5

In order to meet their business requirements, how should HipLocal store their application state?

- A. Use local SSDs to store state.
- B. Put a memcache layer in front of MySQL.
- C. Move the state storage to Cloud Spanner.
- D. Replace the MySQL instance with Cloud SQL.

Correct Answer: B

[PROFESSIONAL-CLOUD-DEVELOPER VCE Dumps](#)

[PROFESSIONAL-CLOUD-DEVELOPER Exam Questions](#)

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)