



Q&As

Professional Cloud Architect on Google Cloud Platform

Pass Google PROFESSIONAL-CLOUD-ARCHITECT Exam with 100% Guarantee

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

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

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

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





QUESTION 1

Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time. You want to use Google-recommended practices to detect anomalies in your company data. What should you do?

- A. Upload your files into Cloud Storage. Use Cloud Datalab to explore and clean your data.
- B. Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.
- C. Connect Cloud Datalab to your on-premises systems. Use Cloud Datalab to explore and clean your data.
- D. Connect Cloud Dataprep to your on-premises systems. Use Cloud Dataprep to explore and clean your data.

Correct Answer: B

<https://cloud.google.com/dataprep/>

QUESTION 2

Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table. Any logs older than 45 days should be removed. You want to optimize storage and follow Google-recommended practices. What should you do?

- A. Configure the expiration time for your tables at 45 days
- B. Make the tables time-partitioned, and configure the partition expiration at 45 days
- C. Rely on BigQuery's default behavior to prune application logs older than 45 days
- D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Correct Answer: B

Reference: <https://cloud.google.com/bigquery/docs/managing-tables>

QUESTION 3

Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter. You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change.

Which product should you use?

- A. Google Cloud Dataflow
- B. Google Cloud Dataproc
- C. Google Compute Engine



D. Google Kubernetes Engine

Correct Answer: B

Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.

References: <https://cloud.google.com/dataproc/docs/resources/faq>

QUESTION 4

An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs.

What should you do?

- A. Direct them to download and install the Google StackDriver logging agent
- B. Send them a list of online resources about logging best practices
- C. Help them define their requirements and assess viable logging tools
- D. Help them upgrade their current tool to take advantage of any new features

Correct Answer: A

The Stackdriver Logging agent streams logs from your VM instances and from selected third party software packages to Stackdriver Logging. Using the agent is optional but we recommend it. The agent runs under both Linux and Microsoft Windows.

Note: Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source. Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time.

References: <https://cloud.google.com/logging/docs/agent/installation>

QUESTION 5

Your operations team has asked you to help diagnose a performance issue in a production application that runs on Compute Engine. The application is dropping requests that reach it when under heavy load. The process list for affected instances shows a single application process that is consuming all available CPU, and autoscaling has reached the upper limit of instances. There is no abnormal load on any other related systems, including the database. You want to allow production traffic to be served again as quickly as possible. Which action should you recommend?

- A. Change the autoscaling metric to `agent.googleapis.com/memory/percent_used`.
- B. Restart the affected instances on a staggered schedule.
- C. SSH to each instance and restart the application process.



D. Increase the maximum number of instances in the autoscaling group.

Correct Answer: D

Reference: <https://cloud.google.com/blog/products/sap-google-cloud/best-practices-for-sap-app-server-autoscaling-on-google-cloud>

[Latest PROFESSIONAL-CLOUD-ARCHITECT Dumps](#)

[PROFESSIONAL-CLOUD-ARCHITECT PDF Dumps](#)

[PROFESSIONAL-CLOUD-ARCHITECT Practice Test](#)