



VCE & PDF

PassApply.com

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

---

# PROFESSIONAL-CLOUD-DEVOPS- ENGINEER<sup>Q&As</sup>

Professional Cloud DevOps Engineer

**Pass Google PROFESSIONAL-CLOUD-DEVOPS-  
ENGINEER Exam with 100% Guarantee**

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

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

---

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





### QUESTION 1

Your team is designing a new application for deployment into Google Kubernetes Engine (GKE). You need to set up monitoring to collect and aggregate various application-level metrics in a centralized location. You want to use Google Cloud Platform services while minimizing the amount of work required to set up monitoring. What should you do?

- A. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.
- B. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.
- C. Install the OpenTelemetry client libraries in the application, configure Stackdriver as the export destination for the metrics, and then observe the application's metrics in Stackdriver.
- D. Emit all metrics in the form of application-specific log messages, pass these messages from the containers to the Stackdriver logging collector, and then observe metrics in Stackdriver.

Correct Answer: A

<https://cloud.google.com/trace/docs/setup>

---

### QUESTION 2

You need to reduce the cost of virtual machines (VM) for your organization. After reviewing different options, you decide to leverage preemptible VM instances. Which application is suitable for preemptible VMs?

- A. A scalable in-memory caching system.
- B. The organization's public-facing website.
- C. A distributed, eventually consistent NoSQL database cluster with sufficient quorum.
- D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket.

Correct Answer: D

Reference: <https://cloud.google.com/preemptible-vm>

---

### QUESTION 3

You need to deploy a new service to production. The service needs to automatically scale using a Managed Instance Group (MIG) and should be deployed over multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity. What should you do?

- A. Use the n1-highcpu-96 machine type in the configuration of the MIG.
- B. Monitor results of Stackdriver Trace to determine the required amount of resources.
- C. Validate that the resource requirements are within the available quota limits of each region.



D. Deploy the service in one region and use a global load balancer to route traffic to this region.

Correct Answer: C

Validate that the resource requirements are within the available quota limits of each region. It is important to ensure that the resource requirements are within the available quota limits in each region before deploying the service, to avoid exceeding the limits and causing problems. This is essential to ensure that the service is deployed correctly and has the necessary capacity to handle the load.

---

#### QUESTION 4

You are responsible for creating and modifying the Terraform templates that define your Infrastructure. Because two new engineers will also be working on the same code, you need to define a process and adopt a tool that will prevent you from overwriting each other's code. You also want to ensure that you capture all updates in the latest version. What should you do?

- A. Store your code in a Git-based version control system. Establish a process that allows developers to merge their own changes at the end of each day. Package and upload code to a versioned Cloud Storage basket as the latest master version.
- B. Store your code in a Git-based version control system. Establish a process that includes code reviews by peers and unit testing to ensure integrity and functionality before integration of code. Establish a process where the fully integrated code in the repository becomes the latest master version.
- C. Store your code as text files in Google Drive in a defined folder structure that organizes the files. At the end of each day, confirm that all changes have been captured in the files within the folder structure. Rename the folder structure with a predefined naming convention that increments the version.
- D. Store your code as text files in Google Drive in a defined folder structure that organizes the files. At the end of each day, confirm that all changes have been captured in the files within the folder structure and create a new .zip archive with a predefined naming convention. Upload the .zip archive to a versioned Cloud Storage bucket and accept it as the latest version.

Correct Answer: B

Using a Git-based version control system such as GitHub or GitLab is a best practice for managing code in a collaborative environment. It provides a central repository where all changes are tracked and versioned, and it also allows for concurrent development by multiple team members.

Establishing a process that includes code reviews by peers and unit testing before merging changes ensures the integrity and functionality of the code, and it also helps to prevent conflicts and errors.

Once the changes are fully integrated and tested, the latest version of the code in the repository should be considered the master version, and this should be the version that is used for deployment.

---

#### QUESTION 5

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to `/var/log/app_messages.log`, and you want to send these log entries to Stackdriver Logging. What should you do?



- A. Use the default Stackdriver Kubernetes Engine Monitoring agent configuration.
- B. Deploy a Fluentd daemonset to GKE. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- C. Install Kubernetes on Google Compute Engine (GCE) and redeploy your applications. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Write a script to tail the log file within the pod and write entries to standard output. Run the script as a sidecar container with the application's pod. Configure a shared volume between the containers to allow the script to have read access to /var/log in the application container.

Correct Answer: B

Reference: <https://cloud.google.com/solutions/customizing-stackdriver-logs-fluentd>

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

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Study Guide](#)

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