



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^{Q&As}

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 company runs applications in Google Kubernetes Engine (GKE). Several applications rely on ephemeral volumes. You noticed some applications were unstable due to the DiskPressure node condition on the worker nodes. You need to identify which Pods are causing the issue, but you do not have execute access to workloads and nodes. What should you do?

- A. Check the node/ephemeral_storage/used_bytes metric by using Metrics Explorer.
- B. Check the container/ephemeral_storage/used_bytes metric by using Metrics Explorer.
- C. Locate all the Pods with emptyDir volumes. Use the `df -h` command to measure volume disk usage.
- D. Locate all the Pods with emptyDir volumes. Use the `df -sh *` command to measure volume disk usage.

Correct Answer: A

node/ephemeral_storage/used_bytes GA Ephemeral storage usage GAUGE, INT64, By k8s_node Local ephemeral storage bytes used by the node. Sampled every 60 seconds.
https://cloud.google.com/monitoring/api/metrics_kubernetes

QUESTION 2

You support an application that stores product information in cached memory. For every cache miss, an entry is logged in Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?

- A. Link Stackdriver Logging as a source in Google Data Studio. Filter the logs on the cache misses.
- B. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
- C. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
- D. Configure BigQuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table.

Correct Answer: C

<https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric>

QUESTION 3

You are managing the production deployment to a set of Google Kubernetes Engine (GKE) clusters. You want to make sure only images which are successfully built by your trusted CI/CD pipeline are deployed to production. What should you do?

- A. Enable Cloud Security Scanner on the clusters.
- B. Enable Vulnerability Analysis on the Container Registry.
- C. Set up the Kubernetes Engine clusters as private clusters.



D. Set up the Kubernetes Engine clusters with Binary Authorization.

Correct Answer: D

<https://cloud.google.com/binary-authorization/docs/overview>

QUESTION 4

Your company runs an ecommerce website built with JVM-based applications and microservice architecture in Google Kubernetes Engine (GKE). The application load increases during the day and decreases during the night. Your operations team has configured the application to run enough Pods to handle the evening peak load. You want to automate scaling by only running enough Pods and nodes for the load. What should you do?

- A. Configure the Vertical Pod Autoscaler, but keep the node pool size static.
- B. Configure the Vertical Pod Autoscaler, and enable the cluster autoscaler.
- C. Configure the Horizontal Pod Autoscaler, but keep the node pool size static.
- D. Configure the Horizontal Pod Autoscaler, and enable the cluster autoscaler.

Correct Answer: D

QUESTION 5

You support a service with a well-defined Service Level Objective (SLO). Over the previous 6 months, your service has consistently met its SLO and customer satisfaction has been consistently high. Most of your service's operations tasks are automated and few repetitive tasks occur frequently. You want to optimize the balance between reliability and deployment velocity while following site reliability engineering best practices. What should you do? (Choose two.)

- A. Make the service's SLO more strict.
- B. Increase the service's deployment velocity and/or risk.
- C. Shift engineering time to other services that need more reliability.
- D. Get the product team to prioritize reliability work over new features.
- E. Change the implementation of your Service Level Indicators (SLIs) to increase coverage.

Correct Answer: BC

<https://sre.google/workbook/implementing-slos/#slo-decision-matrix>

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Practice Test](#)

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

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