



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-devops-engineer.html>  
2024 Latest pass4itsure 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.pass4itsure.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

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-devops-engineer.html>

2024 Latest pass4itsure 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

You are building and running client applications in Cloud Run and Cloud Functions. Your client requires that all logs must be available for one year so that the client can import the logs into their logging service. You must minimize required code changes. What should you do?

- A. Deploy Falco or Twistlock on GKE to monitor for vulnerabilities on your running Pods.
- B. Configure Identity and Access Management (IAM) policies to create a least privilege model on your GKE clusters.
- C. Use Binary Authorization to attest images during your CI/CD pipeline.
- D. Enable Container Analysis in Artifact Registry, and check for common vulnerabilities and exposures (CVEs) in your container images.

Correct Answer: A

---

### QUESTION 2

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 3

Your company operates in a highly regulated domain. Your security team requires that only trusted container images can be deployed to Google Kubernetes Engine (GKE). You need to implement a solution that meets the requirements of the security team while minimizing management overhead. What should you do?

- A. Configure Binary Authorization in your GKE clusters to enforce deploy-time security policies.
- B. Grant the roles/artifactregistry.writer role to the Cloud Build service account. Confirm that no employee has Artifact Registry write permission.
- C. Use Cloud Run to write and deploy a custom validator. Enable an Eventarc trigger to perform validations when new images are uploaded.
- D. Configure Kritis to run in your GKE clusters to enforce deploy-time security policies.

Correct Answer: A

---



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

#### QUESTION 4

You use Terraform to manage an application deployed to a Google Cloud environment. The application runs on instances deployed by a managed instance group. The Terraform code is deployed by using a CI/CD pipeline. When you change the machine type on the instance template used by the managed instance group, the pipeline fails at the terraform apply stage with the following error message:

```
Error waiting for Deleting Instance Template: The instance_template resource 'project/my-project/global/instanceTemplates/my-it-2022010101010000000001' is already being used by 'projects/my-project/regions/us-central1/instanceGroupManagers/my-mig'
```

You need to update the instance template and minimize disruption to the application and the number of pipeline runs. What should you do?

- A. Delete the managed instance group, and recreate it after updating the instance template.
- B. Add a new instance template, update the managed instance group to use the new instance template, and delete the old instance template.
- C. Remove the managed instance group from the Terraform state file, update the instance template, and reimport the managed instance group.
- D. Set the create\_before\_destroy meta-argument to true in the lifecycle block on the instance template.

Correct Answer: D

#### QUESTION 5

As a Site Reliability Engineer, you support an application written in Go that runs on Google Kubernetes Engine (GKE) in production. After releasing a new version of the application, you notice the application runs for about 15 minutes and then restarts. You decide to add Cloud Profiler to your application and now notice that the heap usage grows constantly until the application restarts. What should you do?

- A. Increase the CPU limit in the application deployment.
- B. Add high memory compute nodes to the cluster.
- C. Increase the memory limit in the application deployment.
- D. Add Cloud Trace to the application, and redeploy.

Correct Answer: C

#### QUESTION 6



You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality – a data intensive reporting feature – is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

- A. As the I/O wait times aggregated across all report generation backends
- B. As the proportion of report generation requests that result in a successful response
- C. As the application's report generation queue size compared to a known-good threshold
- D. As the reporting backend PD throughout capacity compared to a known-good threshold

Correct Answer: B

As the proportion of report generation requests that result in a successful response is a valid availability Service Level Indicator (SLI) for the report generation feature. This indicator measures the percentage of requests for report generation that are successfully completed, and it is an indicator of the service availability. This SLI provides a clear and measurable way to track the availability of the report generation feature. It is a simple and easy to understand metric, that it can be easily monitored and reported. It can be used to track the performance of the report generation feature over time and detect any potential issues that may cause it to become unavailable. It also allows you to detect and diagnose issues in the system quickly and take appropriate action to mitigate them. Additionally, it aligns well with the customer's expectation of the report generation feature as they want to see a high percentage of successful report generation requests, which indicates that the feature is working correctly.

## QUESTION 7

You are deploying an application that needs to access sensitive information. You need to ensure that this information is encrypted and the risk of exposure is minimal if a breach occurs. What should you do?

- A. Store the encryption keys in Cloud Key Management Service (KMS) and rotate the keys frequently
- B. Inject the secret at the time of instance creation via an encrypted configuration management system.
- C. Integrate the application with a Single sign-on (SSO) system and do not expose secrets to the application.
- D. Leverage a continuous build pipeline that produces multiple versions of the secret for each instance of the application.

Correct Answer: A

Store the encryption keys in Cloud Key Management Service (KMS) and rotate the keys frequently. This ensures that the sensitive information is encrypted at rest and in transit, and that the encryption keys are regularly rotated to minimize the risk of exposure in the event of a breach.

## QUESTION 8

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What



should you do?

- A. Disable the CI pipeline and revert to manually building and pushing the artifacts.
- B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
- D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Correct Answer: D

Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

This option allows you to compare the previous and current Cloud Build Configuration files and find what has been changed and understand what could be causing the issue. After identifying the problem, you can fix it and ensure that the

pipeline is working correctly again. This approach is based on the SRE practices of identifying and resolving issues quickly and effectively.

---

#### QUESTION 9

You need to create a Cloud Monitoring SLO for a service that will be published soon. You want to verify that requests to the service will be addressed in fewer than 300 ms at least 90% of the time per calendar month. You need to identify the metric and evaluation method to use. What should you do?

- A. Select a latency metric for a request-based method of evaluation.
- B. Select a latency metric for a window-based method of evaluation.
- C. Select an availability metric for a request-based method of evaluation.
- D. Select an availability metric for a window-based method of evaluation.

Correct Answer: D

---

#### QUESTION 10

Your Cloud Run application writes unstructured logs as text strings to Cloud Logging. You want to convert the unstructured logs to JSON-based structured logs. What should you do?

- A. Modify the application to use Cloud Logging software development kit (SDK), and send log entries with a jsonPayload field.
- B. Install a Fluent Bit sidecar container, and use a JSON parser.
- C. Install the log agent in the Cloud Run container image, and use the log agent to forward logs to Cloud Logging.
- D. Configure the log agent to convert log text payload to JSON payload.

Correct Answer: D

---



<https://cloud.google.com/logging/docs/agent/logging/configuration#process-payload>

### QUESTION 11

You support a popular mobile game application deployed on Google Kubernetes Engine (GKE) across several Google Cloud regions. Each region has multiple Kubernetes clusters. You receive a report that none of the users in a specific region can connect to the application. You want to resolve the incident while following Site Reliability Engineering practices. What should you do first?

- A. Reroute the user traffic from the affected region to other regions that don't report issues.
- B. Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
- C. Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
- D. Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.

Correct Answer: A

Reference: <https://cloud.google.com/error-reporting/docs/viewing-errors>

### QUESTION 12

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do?

- A. Use Stackdriver Kubernetes Engine Monitoring.
- B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
- C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.
- D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Correct Answer: A

\*

<https://cloud.google.com/anthos/clusters/docs/on-prem> GKE on-prem is also called Anthos clusters on VMware

\*

<https://cloud.google.com/anthos/clusters/docs/on-prem/concepts/logging-and-monitoring> You have several logging and monitoring options for your Anthos clusters on VMware:

+

Cloud Logging and Cloud Monitoring, enabled by in-cluster agents deployed with Anthos clusters on VMware.

+





Prometheus and Grafana, disabled by default.

+

Validated configurations with third-party solutions.

=> it means, if not a special situation, the correct should be using the first option: Logging and Monitoring. In this case, we want metrics, so Monitoring (aka. Cloud Monitoring, Stackdriver Monitoring) should be used. We are talking about GKE, so we will use Kubernetes Engine Monitoring (<https://cloud.google.com/kubernetes-engine-monitoring>).

---

### QUESTION 13

You are running an experiment to see whether your users like a new feature of a web application. Shortly after deploying the feature as a canary release, you receive a spike in the number of 500 errors sent to users, and your monitoring reports show increased latency. You want to quickly minimize the negative impact on users. What should you do first?

- A. Roll back the experimental canary release.
- B. Start monitoring latency, traffic, errors, and saturation.
- C. Record data for the postmortem document of the incident.
- D. Trace the origin of 500 errors and the root cause of increased latency.

Correct Answer: A

<https://www.examtopycs.com/discussions/google/1/>

---

### QUESTION 14

You encountered a major service outage that affected all users of the service for multiple hours. After several hours of incident management, the service returned to normal, and user access was restored. You need to provide an incident summary to relevant stakeholders following the Site Reliability Engineering recommended practices. What should you do first?

- A. Call individual stakeholders to explain what happened.
- B. Develop a post-mortem to be distributed to stakeholders.
- C. Send the Incident State Document to all the stakeholders.
- D. Require the engineer responsible to write an apology email to all stakeholders.

Correct Answer: B

<https://sre.google/sre-book/postmortem-culture/>

---

### QUESTION 15

You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need measure application





reliability from a user perspective without making any engineering changes to it. What should you do? (Choose two.)

- A. Review current application metrics and add new ones as needed.
- B. Modify the code to capture additional information for user interaction.
- C. Analyze the web proxy logs only and capture response time of each request.
- D. Create new synthetic clients to simulate a user journey using the application.
- E. Use current and historic Request Logs to trace customer interaction with the application.

Correct Answer: DE

[https://cloud.google.com/architecture/adopting-slos#choosing\\_a\\_measurement\\_method](https://cloud.google.com/architecture/adopting-slos#choosing_a_measurement_method)

[Latest PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Dumps](#)   [PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Study Guide](#)   [PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam Questions](#)