

Google Cloud DevOps Certification Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	9
Explanations	11
Next Steps	19

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 - 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

- 1. For setting up CI/CD pipeline load testing before promoting to production GKE, how should you utilize Binary Authorization?**
 - A. Create an attestation for builds that pass the load test using a service account JSON key.**
 - B. Create an attestation for builds that pass the load test authenticated through Workload Identity.**
 - C. Create an attestation for builds that can be manually approved by an operator.**
 - D. Create an attestation for builds that pass tests stored in a public repository.**
- 2. What steps should you take to customize error reporting in a Python application running on the App Engine flexible environment?**
 - A. Install the Cloud Error Reporting library for Python and execute your code on a Compute Engine VM.**
 - B. Install the Cloud Error Reporting library for Python and run your code on Google Kubernetes Engine.**
 - C. Install the Cloud Error Reporting library for Python and run your code on the App Engine flexible environment.**
 - D. Utilize the Cloud Error Reporting API to direct errors from your application to ReportedErrorEvent.**
- 3. Which tool should you use for validating and enforcing security policies on container images?**
 - A. Cloud Build service account permissions.**
 - B. Kritis for image scanning.**
 - C. Cloud Security Command Center.**
 - D. Binary Authorization in GKE clusters.**
- 4. If you are not seeing expected network traffic information in VPC Flow Logs, what should you check first?**
 - A. Verify if you are filtering by the TCP protocol.**
 - B. Check if the traffic you are filtering for is of very low volume.**
 - C. Consider lowering the sample rate in the VPC Flow Log settings.**
 - D. Investigate if the traffic is originating from a VM outside your VPC.**

5. What should you confirm if your application's logs are not appearing on Cloud's operations suite dashboard?

- A. Confirm that the Cloud agent has been installed in the hosting virtual machine.**
- B. Confirm that your account has the proper permissions to use the Cloud dashboard.**
- C. Confirm that port 25 has been opened in the firewall for messages to Cloud's operations suite.**
- D. Confirm that the application is using the required client library and the service account key has permissions.**

6. What is the best way to define a Service Level Indicator (SLI) for scaling an NGINX-based application deployed in GKE?

- A. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.**
- B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.**
- C. Install the Cloud custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.**
- D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.**

7. What is the most effective way to prevent personally identifiable information (PII) from being written into log entries on Google Cloud?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.**
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.**
- C. Wait for the application developers to patch the application before monitoring the log entries.**
- D. Stage log entries to Cloud Storage and trigger a function to clean the entries.**

8. What is the best approach to apply policy parameters to GKE clusters automatically from a GitHub repository?

- A. Set up a GitHub action to trigger Cloud Build for changes.**
- B. Use a webhook to send updates to Anthos Service Mesh.**
- C. Configure Anthos Config Management with the GitHub repository.**
- D. Configure Config Connector with the GitHub repository.**

9. What deployment and testing strategy should you choose for a CI/CD pipeline that aims to mitigate deployment complexity and rollback duration?

- A. Recreate deployment and canary testing.**
- B. Blue/green deployment and canary testing.**
- C. Rolling update deployment and A/B testing.**
- D. Rolling update deployment and shadow testing.**

10. What is a primary benefit of using Cluster Autoscaler in GKE?

- A. It automatically resizes Pods based on resource requests**
- B. It adjusts node pool size in response to workload demands**
- C. It monitors application performance metrics continuously**
- D. It simplifies service discovery in Kubernetes**

SAMPLE

Answers

SAMPLE

1. B
2. C
3. D
4. B
5. A
6. C
7. A
8. C
9. B
10. B

SAMPLE

Explanations

SAMPLE

1. For setting up CI/CD pipeline load testing before promoting to production GKE, how should you utilize Binary Authorization?

- A. Create an attestation for builds that pass the load test using a service account JSON key.**
- B. Create an attestation for builds that pass the load test authenticated through Workload Identity.**
- C. Create an attestation for builds that can be manually approved by an operator.**
- D. Create an attestation for builds that pass tests stored in a public repository.**

Using Binary Authorization in the context of setting up a CI/CD pipeline for load testing before promoting to production in Google Kubernetes Engine (GKE) involves ensuring that only builds that meet specific criteria can be deployed. The correct choice involves creating an attestation for builds that pass the load test authenticated through Workload Identity. Workload Identity allows GKE applications to authenticate to Google Cloud services without needing to manage service account keys. By leveraging Workload Identity, you can securely associate Kubernetes service accounts with Google Cloud service accounts, which enhances security by eliminating the need for service account JSON keys that can be exploited if leaked. This method ensures that the process of authentication is more secure and aligns with modern best practices. In this scenario, establishing an attestation for builds authenticated through Workload Identity means using a structured, secure approach to verify that the builds have met predetermined quality and performance standards before promoting them to production. This method not only reinforces security but also simplifies management by utilizing IAM features inherent to Google Cloud. The other options, while they involve attestation processes, do not incorporate the same level of security and integration with GKE's capabilities as Workload Identity does. For instance, using a service account JSON key introduces additional management overhead and potential security risks. Manual approvals

SAMPLE

2. What steps should you take to customize error reporting in a Python application running on the App Engine flexible environment?

- A. Install the Cloud Error Reporting library for Python and execute your code on a Compute Engine VM.**
- B. Install the Cloud Error Reporting library for Python and run your code on Google Kubernetes Engine.**
- C. Install the Cloud Error Reporting library for Python and run your code on the App Engine flexible environment.**
- D. Utilize the Cloud Error Reporting API to direct errors from your application to ReportedErrorEvent.**

To customize error reporting in a Python application running on the App Engine flexible environment, installing the Cloud Error Reporting library for Python and running your code on App Engine is essential. This approach allows your application to record and report exceptions and errors directly, leveraging the integrated capabilities of the App Engine platform for managing logs and error reports seamlessly. The App Engine flexible environment is specifically designed to handle applications and services built using popular programming languages and frameworks. This means that by utilizing the Cloud Error Reporting library within this environment, you can effectively implement error tracking and gain insights into performance issues or exceptions that occur in real time. Using the Cloud Error Reporting library enables you to automatically log errors with rich context, which will be sent to Google Cloud's error reporting services. This is highly beneficial for monitoring application health, debugging issues, and managing errors without having to set up additional infrastructure components, as would be the case in alternative environments like Compute Engine or Google Kubernetes Engine. In summary, the correct answer reflects the intended use of Google Cloud services tailored to run in the App Engine flexible environment, specifically leveraging the capabilities offered by Cloud Error Reporting for effective application monitoring and error handling.

SAMPLE

3. Which tool should you use for validating and enforcing security policies on container images?

- A. Cloud Build service account permissions.
- B. Kritis for image scanning.
- C. Cloud Security Command Center.
- D. Binary Authorization in GKE clusters.**

Binary Authorization in GKE clusters is a powerful tool specifically designed to validate and enforce security policies on container images. It acts as a deployment safety mechanism that ensures only trusted container images are deployed to your Google Kubernetes Engine (GKE) clusters. By defining policies that require certain criteria to be met—such as having undergone security scans, being signed by trusted authorities, or meeting compliance standards—Binary Authorization helps prevent vulnerabilities introduced by malicious or unverified images. The use of Binary Authorization offers a flexible yet robust approach to maintaining a secure deployment process within Kubernetes, as it integrates seamlessly with the CI/CD pipelines. This tool effectively enhances the security posture of your applications running in GKE by ensuring that only those images that have passed predefined security checks are allowed to be deployed. Other choices, while important in their own capacities, do not specifically enforce security policies during the deployment phase of container images. For instance, the Cloud Build service account permissions focus more on access control rather than on validation of images themselves. Kritis, which provides image scanning and compliance checks, is relevant to image security but does not enforce deployment policies directly. Meanwhile, Cloud Security Command Center is a broader security management tool designed for monitoring and managing security across Google Cloud resources rather than specifically validating container images.

4. If you are not seeing expected network traffic information in VPC Flow Logs, what should you check first?

- A. Verify if you are filtering by the TCP protocol.
- B. Check if the traffic you are filtering for is of very low volume.**
- C. Consider lowering the sample rate in the VPC Flow Log settings.
- D. Investigate if the traffic is originating from a VM outside your VPC.

When troubleshooting issues with VPC Flow Logs and unexpected network traffic information, the first thing to check should indeed be the volume of traffic you are filtering for. If the traffic of interest is very low, it may not appear in the logs due to the sampling method employed by VPC Flow Logs. VPC Flow Logs can be set up to sample data, which means that they may only capture a portion of the network traffic, especially when dealing with low-volume situations. Therefore, if there is very little traffic passing through a particular interface or to specific destinations, it is possible that this traffic might fall beneath the threshold that is required for it to be logged, leading to an absence of expected entries in your flow logs. Recognizing the volume of traffic is crucial, as low-volume data may not be reliably logged, especially compared to higher-volume traffic that is more likely to be captured. Once you confirm that the traffic is indeed of low volume, you can adjust your monitoring and logging strategies appropriately, perhaps by changing the sample rate or looking for different traffic patterns. The importance of this initial check lies in efficiently narrowing down the cause of the issue before diving into more complex diagnostics. This helps in maintaining effective resource use and time management when working with VPC network

5. What should you confirm if your application's logs are not appearing on Cloud's operations suite dashboard?

- A. Confirm that the Cloud agent has been installed in the hosting virtual machine.**
- B. Confirm that your account has the proper permissions to use the Cloud dashboard.**
- C. Confirm that port 25 has been opened in the firewall for messages to Cloud's operations suite.**
- D. Confirm that the application is using the required client library and the service account key has permissions.**

To ensure that application logs are appearing on Cloud's operations suite dashboard, confirming that the Cloud agent has been installed in the hosting virtual machine is crucial because the Cloud agent is responsible for collecting and sending logs to the operations suite. Without this agent, there would be no mechanism to gather and transmit the application's logs, resulting in their absence from the dashboard. The Cloud agent serves as a connection between your application and Google Cloud's logging services. It is essential for monitoring, logging, and resource management, thereby making it a critical component for ensuring that logs are visible and accessible. If the agent is missing or not functioning properly, the application logs cannot be collected, meaning they will not appear in the dashboard as expected. While permissions, firewall settings, and client libraries are important aspects of configuring a cloud application, they do not specifically address the capability to collect and send logs to the operations suite. Therefore, the installation of the Cloud agent is the first line of action to resolve issues related to log visibility on the dashboard.

SAMPLE

6. What is the best way to define a Service Level Indicator (SLI) for scaling an NGINX-based application deployed in GKE?

- A. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.**
- B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.**
- C. Install the Cloud custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.**
- D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.**

The best way to define a Service Level Indicator (SLI) for scaling an NGINX-based application deployed in Google Kubernetes Engine (GKE) is to install the Cloud custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the Google Cloud Load Balancer (GCLB). This approach is effective because it allows dynamic scaling of application pods based on actual user demand. By leveraging the number of incoming requests being handled by the load balancer, the autoscaler can make more informed decisions about scaling up or down the number of pods. This ensures that the application remains responsive and available under varying load conditions, directly aligning with performance metrics that are critical for maintaining high availability and service quality. Additionally, the Cloud custom metrics adapter enables the horizontal pod autoscaler to utilize custom metrics, which are often better representatives of system load, as opposed to default metrics like CPU or memory usage, especially for web applications that might not have a clear correlation between resource usage and performance under load. In contrast, utilizing the average response time from liveness and readiness probes, as suggested in one of the options, may not reflect the actual load on the application or provide actionable data for scaling. This is particularly important as liveness

7. What is the most effective way to prevent personally identifiable information (PII) from being written into log entries on Google Cloud?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.**
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.**
- C. Wait for the application developers to patch the application before monitoring the log entries.**
- D. Stage log entries to Cloud Storage and trigger a function to clean the entries.**

Using the filter-record-transformer plugin in Fluentd is the most effective method to prevent personally identifiable information (PII) from being captured in log entries as they are generated. This plugin enables real-time modification of log messages, allowing for the removal of specific fields that contain sensitive information before they reach their destination, such as Cloud Logging. By applying transformations directly in-flight, you can ensure that any PII is stripped from the logs at the source, significantly reducing the risk of exposure and maintaining compliance with privacy regulations. This proactive approach is essential in a DevOps culture that emphasizes continuous integration and delivery since it integrates seamlessly into the logging pipeline and prevents sensitive information from being logged in the first place, rather than having to deal with it after the fact. Options that suggest waiting for developers to patch applications or staging entries for later cleanup introduce unnecessary delays and risks. Optioning for output plugins that modify entries instead of filters might not address PII at the most meaningful point in the logging process. Therefore, using the filter-record-transformer plugin provides not only immediate remediation but also streamlined logging practices that bolster security and compliance.

8. What is the best approach to apply policy parameters to GKE clusters automatically from a GitHub repository?

- A. Set up a GitHub action to trigger Cloud Build for changes.**
- B. Use a webhook to send updates to Anthos Service Mesh.**
- C. Configure Anthos Config Management with the GitHub repository.**
- D. Configure Config Connector with the GitHub repository.**

Using Anthos Config Management with a GitHub repository is the best approach to automatically apply policy parameters to GKE clusters. Anthos Config Management provides a way to manage Kubernetes applications and configurations using a GitOps methodology. By integrating with a GitHub repository, you can store your configurations as code, allowing for version control, automated deployment, and policy enforcement across multiple GKE clusters. When you commit changes to the GitHub repository, Anthos Config Management monitors the repository for updates and automatically applies those changes to the GKE clusters. This ensures that the configuration in your clusters is always in sync with what is specified in your Git repository. This capability allows organizations to manage infrastructure as code, providing better visibility, reproducibility, and compliance with organizational policies. Other options may involve various integration mechanisms, but they do not provide the same level of streamlined and consistent policy application mechanism that Anthos Config Management offers when paired with GitHub.

9. What deployment and testing strategy should you choose for a CI/CD pipeline that aims to mitigate deployment complexity and rollback duration?

- A. Recreate deployment and canary testing.
- B. Blue/green deployment and canary testing.**
- C. Rolling update deployment and A/B testing.
- D. Rolling update deployment and shadow testing.

Choosing a blue/green deployment strategy combined with canary testing is effective for mitigating deployment complexity and minimizing rollback duration. In a blue/green deployment, two identical environments are maintained: one (the blue environment) serves the live production traffic, while the other (the green environment) is used for staging the new version of the application. Once the new version is fully tested in the green environment, the traffic can be switched to it with minimal downtime and risk. If any issues arise after the switch, rolling back to the previous version is straightforward and quick, as the blue environment is still intact and can be reactivated. Canary testing further enhances this approach by allowing you to release the new version to a small subset of users before fully transitioning to it. This provides an opportunity to monitor the new version's performance and gather feedback without impacting the entire user base. If the canary release encounters issues, it can be rolled back without affecting all users, thus reducing the complexity and risk associated with deployments. This combination ensures that deployments are safer and more manageable while allowing for prompt responses to any issues that may arise, ultimately leading to lower rollback durations and less disruption.

10. What is a primary benefit of using Cluster Autoscaler in GKE?

- A. It automatically resizes Pods based on resource requests
- B. It adjusts node pool size in response to workload demands**
- C. It monitors application performance metrics continuously
- D. It simplifies service discovery in Kubernetes

The primary benefit of using Cluster Autoscaler in Google Kubernetes Engine (GKE) is its ability to dynamically adjust the size of the node pool based on the demands of workloads. This means that when there are more workloads requiring resources—for example, if additional Pods are scheduled or existing Pods need more resources—the Cluster Autoscaler can increase the number of nodes in the cluster to accommodate these needs. Conversely, when the workloads decrease, it can also reduce the number of nodes, helping to optimize resource utilization and cost. This ability to respond to workload demands ensures that applications have the necessary resources available to run efficiently without manual intervention, which is particularly valuable in environments with fluctuating workloads. By automating the scaling process, organizations can improve operational efficiency and better manage costs associated with cloud resources. In contrast, resizing Pods based on resource requests does not fall under the capabilities of Cluster Autoscaler, as that responsibility lies with Kubernetes's built-in features. Continuous monitoring of application performance metrics is related to other tools and practices, such as Prometheus or Google's Operations Suite. Lastly, simplifying service discovery pertains to how services within a Kubernetes cluster communicate with each other, which is handled by the Kubernetes networking model rather than the function of the Cluster Autoscaler.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

<https://googleclouddevops.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE