

Google Cloud Professional Cloud Developer 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	17

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. When migrating from an on-premises Hadoop environment, which approach minimizes changes to existing jobs and architecture?**
 - A. Migrate your data stored in Hadoop to BigQuery.**
 - B. Create Compute Engine instances with HDD instead of SSD.**
 - C. Create a Cloud Dataproc cluster on Google Cloud Platform.**
 - D. Create a Cloud Dataproc cluster and move HDFS data to larger disks.**
- 2. When testing Cloud Functions, what is a crucial action to perform with your edited code?**
 - A. Deploy it directly to production**
 - B. Run it on a separate server for load testing**
 - C. Send mock requests to evaluate functionality**
 - D. Only review it in code comments**
- 3. What is the best approach to enhance the resilience of a MySQL deployment?**
 - A. Use single instance MySQL on Compute Engine with read-only servers.**
 - B. Replicate data from single instance MySQL to Cloud SQL.**
 - C. Replace with Cloud SQL and configure for high availability.**
 - D. Switch to Cloud SQL for redundancy without extra configuration.**
- 4. Does Migrate for Anthos simply put a VM into a container?**
 - A. True**
 - B. False**
- 5. What is the appropriate action to secure a multi-tenant platform on Google Kubernetes Engine?**
 - A. Enable Application-layer Secrets on the GKE cluster**
 - B. Deploy a namespace per tenant and use Network Policies in each blog deployment**
 - C. Use GKE Audit Logging to identify malicious containers**
 - D. Build a custom image of the blogging software and use Binary Authorization**

- 6. How should you secure a Cloud Function that accesses other Google Cloud resources with the least privilege?**
- A. Create a service account with Editor authority for resource access.**
 - B. Create a new service account with a custom IAM role for resource access.**
 - C. Create a service account with Editor authority and allow the deployer to act on its behalf.**
 - D. Create a custom IAM service account and allow the deployer to act on its behalf.**
- 7. Which API is designed to transcribe audio into text?**
- A. Speech-to-Text API**
 - B. Cloud Natural Language API**
 - C. Video Intelligence API**
 - D. None of the above**
- 8. In Google Kubernetes Engine (GKE), which object should be used when applications require stable network identities and persistent disks?**
- A. Deployment**
 - B. StatefulSet**
 - C. ReplicaSet**
 - D. ReplicaController**
- 9. What deployment strategy allows for maintaining a previous version while switching all traffic to a new application version?**
- A. Utilize a Blue/green deployment strategy for minimal downtime**
 - B. Employ a canary deployment to test the new version gradually**
 - C. Adopt a rolling deployment to update with minimal disruption**
 - D. Implement a recreate deployment for complete replacement**

10. How can pre-trained machine learning APIs be invoked from an application?

- A. Use TensorFlow.**
- B. Use the Google Cloud console.**
- C. Use the gcloud command.**
- D. Use the REST API.**

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. When migrating from an on-premises Hadoop environment, which approach minimizes changes to existing jobs and architecture?

- A. Migrate your data stored in Hadoop to BigQuery.**
- B. Create Compute Engine instances with HDD instead of SSD.**
- C. Create a Cloud Dataproc cluster on Google Cloud Platform.**
- D. Create a Cloud Dataproc cluster and move HDFS data to larger disks.**

Creating a Cloud Dataproc cluster on Google Cloud Platform is the approach that minimizes changes to existing jobs and architecture when migrating from an on-premises Hadoop environment. Cloud Dataproc is a fully managed service that allows you to run Apache Hadoop and Apache Spark jobs in the cloud. It is designed to be compatible with existing Hadoop environments, which means that your existing jobs, code, and workflows can be deployed with minimal modifications. By using Cloud Dataproc, you can instantiate a cluster that closely mirrors your on-premises setup, utilizing familiar tools and configurations. This helps maintain operational continuity, allowing your development teams to leverage their existing knowledge and skills without needing to learn new systems or architectures. The option of moving HDFS data to larger disks might imply some additional steps or changes in how data is managed, as you're shifting not just the computation environment but also the storage architecture. Although this could be beneficial in some contexts, it can complicate the migration process if significant changes are needed. The primary goal is to keep the architecture as unchanged as possible while still utilizing cloud capabilities effectively. In summary, creating a Cloud Dataproc cluster enables a more straightforward transition and allows you to leverage your current investment in Hadoop tools and processes while benefiting from the

2. When testing Cloud Functions, what is a crucial action to perform with your edited code?

- A. Deploy it directly to production**
- B. Run it on a separate server for load testing**
- C. Send mock requests to evaluate functionality**
- D. Only review it in code comments**

Sending mock requests to evaluate functionality is essential when testing Cloud Functions. This process allows developers to simulate real-world scenarios and ensure that the function behaves as expected. By crafting various test cases that mimic actual triggers, developers can verify that their code handles inputs correctly, processes them as intended, and produces the right outputs. Mock requests enable thorough testing of all aspects, including error handling and performance under different conditions. This allows for identifying and fixing issues before the function is deployed live, ensuring a smoother execution when real users interact with the function. In contrast, deploying directly to production without testing can lead to unforeseen issues that impact users. Running the function on a separate server might be useful for load testing but doesn't necessarily verify the functionality of the code changes. Simply reviewing code comments lacks practical verification of how the function operates in response to actual events. Thus, sending mock requests is the most effective way to validate the changes made to the code.

3. What is the best approach to enhance the resilience of a MySQL deployment?

- A. Use single instance MySQL on Compute Engine with read-only servers.
- B. Replicate data from single instance MySQL to Cloud SQL.
- C. Replace with Cloud SQL and configure for high availability.**
- D. Switch to Cloud SQL for redundancy without extra configuration.

Choosing to replace a single instance MySQL deployment with Cloud SQL configured for high availability is the best approach to enhance resilience. Cloud SQL is a fully managed database service that offers built-in high availability configurations which automatically manage failover processes, backups, and patching, thus reducing the risk of downtime due to hardware or software failures. High availability in Cloud SQL is typically achieved through a primary and a standby instance setup. This ensures that if the primary instance becomes unavailable for any reason, the system can seamlessly failover to the standby instance, maintaining availability of the database. This level of redundancy and automatic management significantly enhances the resilience of the database deployment compared to traditional single instance setups, where downtime can lead to service interruptions while recovery efforts are made. Other options do not provide the same level of resilience. Using a single instance MySQL on Compute Engine, even with read-only replicas, does not offer automatic failover, thus presenting single points of failure. Replicating data from a single instance MySQL to Cloud SQL may improve some aspects of resilience but still relies on the single MySQL instance, which can cause downtimes if that instance fails. The option that suggests switching to Cloud SQL for redundancy without extra configuration lacks the assurance that the setup will be efficiently

4. Does Migrate for Anthos simply put a VM into a container?

- A. True
- B. False**

Migrate for Anthos does not merely package a virtual machine (VM) into a container; it facilitates the transformation of VMs into containerized applications that run in Kubernetes clusters. This process goes beyond simply placing a VM within a container; it includes re-architecting the application to take full advantage of container orchestration and microservices principles. When migrating applications, Migrate for Anthos analyzes the VM and its workloads to generate the necessary Kubernetes resources to effectively manage that application in the new environment. This allows for improved scalability, resilience, and operational efficiency that containers and Kubernetes clusters inherently provide. Therefore, the correct understanding is that Migrate for Anthos enables a more sophisticated application migration process that involves reworking the application architecture rather than simply encapsulating a VM within a container.

5. What is the appropriate action to secure a multi-tenant platform on Google Kubernetes Engine?

- A. Enable Application-layer Secrets on the GKE cluster
- B. Deploy a namespace per tenant and use Network Policies in each blog deployment**
- C. Use GKE Audit Logging to identify malicious containers
- D. Build a custom image of the blogging software and use Binary Authorization

Using a namespace per tenant and employing Network Policies is a well-established approach for securing multi-tenant applications deployed on Google Kubernetes Engine (GKE). Namespaces in Kubernetes provide a mechanism for isolating resources, allowing you to maintain a clear boundary between different tenants. This helps prevent tenants from accessing each other's resources, providing an essential layer of security in a multi-tenant environment. Network Policies further enhance this security by defining rules for how pods communicate with each other and with other network endpoints. By applying specific network policies to the namespace, you can restrict communication between different tenants' pods. This way, even if multiple tenants are deployed in the same cluster, the risk of data leakage or unauthorized access is minimized, as network traffic can be controlled effectively. Although other options may provide additional security measures, they do not directly address the fundamental need for isolation and controlled interaction between tenants in a multi-tenant architecture. Proper namespace management complemented by Network Policies establishes a foundational security posture that is crucial for protecting sensitive tenant data within a shared environment.

6. How should you secure a Cloud Function that accesses other Google Cloud resources with the least privilege?

- A. Create a service account with Editor authority for resource access.
- B. Create a new service account with a custom IAM role for resource access.
- C. Create a service account with Editor authority and allow the deployer to act on its behalf.
- D. Create a custom IAM service account and allow the deployer to act on its behalf.**

To secure a Cloud Function that needs to access other Google Cloud resources with the principle of least privilege, creating a custom IAM service account and allowing the deployer to act on its behalf is an effective strategy. This approach ensures that the service account has specifically tailored permissions that only allow the necessary access to resources required for the Cloud Function's operation. By defining a custom IAM role, you can limit the permissions strictly to what the Cloud Function needs, thereby reducing the risk of unnecessary exposure or unauthorized access to other resources within your Google Cloud environment. Moreover, by allowing the deployment service account to act on behalf of the custom IAM service account, you ensure that the operations performed by the Cloud Function are done securely and under a controlled identity. This adds an additional layer of security since actions are logged and can be traced back to a specific entity, reducing the surface area for potential vulnerabilities or misconfigurations. The other choices involve broader permissions than necessary. Using an Editor role grants excessive access that goes against the principle of least privilege, which is crucial for maintaining secure and manageable cloud environments.

7. Which API is designed to transcribe audio into text?

- A. Speech-to-Text API**
- B. Cloud Natural Language API**
- C. Video Intelligence API**
- D. None of the above**

The Speech-to-Text API is specifically designed to convert spoken language into text format. It utilizes advanced machine learning models to accurately transcribe audio input, making it a robust solution for applications requiring voice recognition capabilities. The API supports multiple languages and various audio formats, allowing for real-time processing and transcription of both live and pre-recorded audio. This functionality is essential for creating applications such as voice commands, transcription services, and more, where turning audio speech into text is critical for analysis and integration into other systems. The other options, while useful in their contexts, do not focus on audio transcription. The Cloud Natural Language API is intended for analyzing and understanding text data, while the Video Intelligence API analyzes video content rather than audio transcription directly.

8. In Google Kubernetes Engine (GKE), which object should be used when applications require stable network identities and persistent disks?

- A. Deployment**
- B. StatefulSet**
- C. ReplicaSet**
- D. ReplicaController**

When applications require stable network identities and persistent disks in Google Kubernetes Engine (GKE), a StatefulSet is the appropriate choice. StatefulSets are specifically designed for managing stateful applications, allowing them to maintain a unique identity and stable network identity. Each pod in a StatefulSet is assigned a unique ordinal index, which helps in distinguishing them, and this is essential for scenarios where the identity of the application instance matters, such as databases or applications that require consistent storage. Moreover, StatefulSets enable the use of persistent storage by automatically provisioning a PersistentVolume for each pod. This is critical for applications that need data to persist even when pods are restarted, scaled, or rescheduled. The configuration of StatefulSets ensures that not only is the storage persistent, but the data remains tied to the lifecycle of the respective pods, providing both stability and consistency across the application instances. In contrast, other options like Deployments and ReplicaSets are more suited for stateless applications, where the instances can be easily recreated and do not require stable identities or persistent storage. A ReplicaController is an earlier iteration of managing replication in Kubernetes with similar capabilities to ReplicaSets, but does not specifically cater to the requirements of network identities or persistent storage in stateful scenarios.

9. What deployment strategy allows for maintaining a previous version while switching all traffic to a new application version?

- A. Utilize a Blue/green deployment strategy for minimal downtime**
- B. Employ a canary deployment to test the new version gradually**
- C. Adopt a rolling deployment to update with minimal disruption**
- D. Implement a recreate deployment for complete replacement**

A blue/green deployment strategy is designed to minimize downtime and ensure a smooth transition when deploying new application versions. This strategy involves maintaining two identical environments, referred to as "blue" (the current version) and "green" (the new version). Initially, the blue environment serves all the production traffic. When you are ready to deploy a new version of the application, you provision it in the green environment. Once the new version is tested and confirmed to be functioning correctly in the green environment, you can switch all traffic to the green environment with a simple change in the routing configuration. This switch allows for a seamless transition since the blue environment is still available and can be rolled back to in case any issues arise with the new version. This capability to maintain the previous version while simultaneously deploying and testing the new version is a key advantage of the blue/green deployment strategy, promoting higher application availability and reducing deployment-related risks. The other deployment strategies, such as canary, rolling, and recreate deployments, each serve different purposes and have their own strengths. However, they either introduce risks during the transition (like in rolling or canary deployments) or require downtime for a complete replacement (as with recreate deployments), which is not the case with blue/green

10. How can pre-trained machine learning APIs be invoked from an application?

- A. Use TensorFlow.**
- B. Use the Google Cloud console.**
- C. Use the gcloud command.**
- D. Use the REST API.**

Pre-trained machine learning APIs can be invoked from an application using the REST API. This method allows developers to leverage the power of pre-built models trained by Google and integrate advanced machine learning capabilities into their applications without needing to build or train models from scratch. By using the REST API, developers can send requests to these machine learning services over HTTP. This provides flexibility across different programming languages and platforms, as long as the application can make web requests. The REST API also supports standard data formats like JSON, which makes it easy to handle inputs and outputs between applications and the machine learning service. The other methods listed, such as using TensorFlow or the gcloud command, aren't as direct for invoking pre-trained APIs within an application context. TensorFlow is primarily a library for training and deploying machine learning models rather than consuming pre-trained models. The Google Cloud console is primarily a web interface for managing services and configurations, and while it allows you to interact with the APIs, it does not enable direct invocation from an application. Similarly, the gcloud command is a command-line tool that is useful for managing Google Cloud resources but is not designed for embedding or calling machine learning APIs directly from application code. Therefore, using the REST API is the most straightforward and effective method for

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://googlecloudprofessionalclouddev.examzify.com>

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

SAMPLE