


GOOGLE CLOUD CERTIFIED ASSOCIATE CLOUD ENGINEER PRACTICE (SAMPLE)

STUDY GUIDE



**SAMPLE STUDY GUIDE
WITH LIMITED QUESTIONS**

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

<https://certifiedassociatecloudengineer.examzify.com>



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. What is the command used to create a Cloud Function in Google Cloud?**
 - A. ``gcloud compute networks subnets list``
 - B. Enable binary logging
 - C. Custom monitoring metric
 - D. ``gcloud functions deploy``
- 2. What are mounted directories that are accessible from inside containers in Google Cloud?**
 - A. BigQuery
 - B. Custom machine type
 - C. Volumes
 - D. ``gsutil``
- 3. Which GCP service allows for managed container orchestration?**
 - A. Google Kubernetes Engine (GKE)
 - B. Cloud Run
 - C. App Engine
 - D. Cloud Functions
- 4. The format used for Kubernetes resource files.**
 - A. CSV
 - B. JSON
 - C. XML
 - D. YAML
- 5. Which command is used to enable an API in Google Cloud?**
 - A. `gcloud services activate [ID]`
 - B. `gcloud services enable [ID]`
 - C. `gcloud services start [ID]`
 - D. `gcloud services deploy [ID]`
- 6. Where SSH keys for connecting to Compute Engine instances are stored.**
 - A. User's home directory
 - B. Cloud Storage
 - C. Instance Disk
 - D. Project or Instance metadata

7. Which service in Google Cloud provides infrastructure as code functionality?
- A. Billing account
 - B. Deployment manager
 - C. Pricing Calculator
 - D. Spanner
8. Your coworker has helped you set up several configurations for gcloud. You've noticed that you're running commands against the wrong project. Being new to the company, you haven't yet memorized any of the projects. With the fewest steps possible, what's the fastest way to switch to the correct configuration?
- A. Run gcloud configurations list followed by gcloud configurations activate .
 - B. Run gcloud config list followed by gcloud config activate.
 - C. Run gcloud config configurations list followed by gcloud config configurations activate.
 - D. Re-authenticate with the gcloud auth login command and select the correct configurations on login.
9. Which flag is used with `kubectl` to specify the JSON path of properties in the JSON output?
- A. -j
 - B. -p
 - C. -json
 - D. -o jsonpath

10. Your developers have been thoroughly logging everything that happens in the API. The API allows end users to request the data as JSON, XML, CSV, and XLS. Supporting all of these formats is taking a lot of developer effort. Management would like to start tracking which options are used over the next month. Without modifying the code, what's the fastest way to be able to report on this data at the end of the month?
- A. Create a custom counter logging metric that uses a regex to extract the data format into a label. At the end of the month, use the metric viewer to see the group by the label.
 - B. Create a log sink that filters for rows that mention the data format. Export that to BigQuery, and run a query at the end of the month.
 - C. Create a custom monitoring metric in code and edit the API code to set the metric each time the API is called.
 - D. Export the logs to excel, and search for the different fields.

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is the command used to create a Cloud Function in Google Cloud?

- A. ``gcloud compute networks subnets list``
- B. Enable binary logging
- C. Custom monitoring metric
- D. ``gcloud functions deploy``

The command used to create a Cloud Function in Google Cloud is ``gcloud functions deploy``. This command is specifically designed for deploying serverless functions that run in response to events or HTTP requests without the need to manage the underlying infrastructure. By using this command, you can specify various parameters, such as the function name, runtime, entry point, and event trigger. The other options do not relate to creating a Cloud Function. For instance, listing network subnets is relevant to Google Compute Engine, and enabling binary logging pertains to Cloud SQL configurations. Custom monitoring metrics involve setting up monitoring features but do not directly relate to deploying functions. Thus, the context around ``gcloud functions deploy`` highlights its specific role in the creation and deployment of serverless functions on Google Cloud.

2. What are mounted directories that are accessible from inside containers in Google Cloud?

- A. BigQuery
- B. Custom machine type
- C. Volumes
- D. ``gsutil``

The correct choice here is volumes, which serve as an important concept in container orchestration and management systems such as Google Kubernetes Engine (GKE) and Docker. Volumes are mounted directories that provide a way to persist data generated by and used by Docker containers. They enable sharing data between containers and provide a mechanism to store data outside the lifecycle of the individual containers. When a volume is mounted into a container, it can be accessed as part of the filesystem within that container, making it possible to read from and write to that directory. This allows for the persistence of data beyond the lifespan of a single container instance, which is crucial for applications that require stateful data storage (e.g., databases or user-generated content). In contrast, the other options do not represent mounted directories accessible from inside containers. BigQuery is a serverless data warehouse designed for big data analysis, which does not directly relate to container file management. Custom machine types refer to configuring virtual machine types with specific resource allocations in Google Cloud, and while ``gsutil`` is a command-line tool for interacting with Google Cloud Storage, it doesn't represent a mounted directory but rather a method of managing cloud storage resources. Thus, volumes stand out as the correct answer as they directly relate to the

3. Which GCP service allows for managed container orchestration?

A. Google Kubernetes Engine (GKE)

B. Cloud Run

C. App Engine

D. Cloud Functions

Google Kubernetes Engine (GKE) is the service that provides managed container orchestration on Google Cloud Platform. It simplifies the deployment, management, and scaling of containerized applications using Kubernetes, which is a powerful platform for orchestrating containers. GKE automates many of the operational tasks associated with managing Kubernetes clusters, such as upgrades, monitoring, and scaling, allowing developers to focus on building and deploying their applications without getting bogged down by infrastructure concerns. GKE integrates with other GCP services and utilizes Google's infrastructure to offer high availability, security, and performance. This makes it suitable for running complex applications that require container orchestration, leveraging features like auto-scaling, load balancing, and seamless integration with CI/CD pipelines. The other services mentioned, while also useful for deploying applications, do not provide the same level of orchestration capabilities as GKE. Cloud Run focuses on managing serverless applications, App Engine handles platform-as-a-service deployment for web applications, and Cloud Functions is designed for event-driven microservices. Each of these services operates at different levels of abstraction compared to GKE, making GKE the ideal choice for those specifically seeking managed container orchestration.

4. The format used for Kubernetes resource files.

A. CSV

B. JSON

C. XML

D. YAML

YAML is indeed the correct format used for Kubernetes resource files. This format is favored in Kubernetes due to its readability and ease of use for defining configuration data. YAML allows for hierarchical data representation which makes it simpler to organize complex configurations, such as those found in Kubernetes manifests. Each resource in Kubernetes, including Deployments, Services, and Pods, can be defined in a YAML file, making it a standard format for deployment and configuration management in a Kubernetes environment. While JSON is another format that can be used for Kubernetes resource files, it is not as commonly utilized primarily due to its more verbose syntax, which can make it less user-friendly for configuration purposes. CSV and XML are not typically used for defining Kubernetes resources, as they do not provide the necessary structure and readability that YAML or JSON offer for such configurations.

5. Which command is used to enable an API in Google Cloud?

A. gcloud services activate [ID]

B. gcloud services enable [ID]

C. gcloud services start [ID]

D. gcloud services deploy [ID]

The correct command to enable an API in Google Cloud is 'gcloud services enable [ID]'. This is because the 'activate', 'start', and 'deploy' commands are used for different purposes. For example, the 'activate' command is used to activate GCP services on the current project, the 'start' command is used to start a stopped service, and the 'deploy' command is used to deploy an application. Therefore, these commands are not relevant for enabling an API in Google Cloud. It is important to read and understand the documentation to ensure the correct command is used for a specific task.

6. Where SSH keys for connecting to Compute Engine instances are stored.

- A. User's home directory
- B. Cloud Storage
- C. Instance Disk
- D. Project or Instance metadata**

SSH keys for connecting to Compute Engine instances are stored in the project or instance metadata. Other options such as user's home directory, Cloud Storage, and instance disk are incorrect because they do not serve the purpose of storing SSH keys for connecting to Compute Engine instances. The user's home directory contains personal files and settings, Cloud Storage is a file storage service, and the instance disk is used for persistent storage of data and needs to be manually configured by the user. The project or instance metadata is the designated location for storing SSH keys, which allows for automatic configuration of keys during instance creation.

7. Which service in Google Cloud provides infrastructure as code functionality?

- A. Billing account
- B. Deployment manager**
- C. Pricing Calculator
- D. Spanner

Deployment Manager is a service in Google Cloud that allows users to express their desired configuration for a deployment in a template, which can then be used to create and manage resources. The other options listed, Billing account, Pricing Calculator, and Spanner, do not offer infrastructure as code functionality. A billing account is used for managing payment for Google Cloud services, the Pricing Calculator is used for estimating costs of Google Cloud services, and Spanner is a database service. While all of these services can be useful in managing and using Google Cloud, they do not offer the same functionality as Deployment Manager when it comes to infrastructure as code.

8. Your coworker has helped you set up several configurations for gcloud. You've noticed that you're running commands against the wrong project. Being new to the company, you haven't yet memorized any of the projects. With the fewest steps possible, what's the fastest way to switch to the correct configuration?

- A. Run gcloud configurations list followed by gcloud configurations activate .
- B. Run gcloud config list followed by gcloud config activate.
- C. Run gcloud config configurations list followed by gcloud config configurations activate.**
- D. Re-authenticate with the gcloud auth login command and select the correct configurations on login.

The most efficient way to switch to the correct configuration involves using the command to list all available configurations and then activating the desired one. The correct answer involves running "gcloud config configurations list" to see a list of configured contexts, which makes it easy to identify the correct configuration to switch to. After identifying the correct configuration, you execute "gcloud config configurations activate <CONFIG_NAME>" to switch to that configuration. This approach minimizes the steps needed to effectively manage your cloud configurations without needing to re-authenticate or explore options that might not directly lead to switching projects. Additionally, using the full command is direct and provides clarity in ensuring you're selecting the right configuration. In contrast, other choices either use incorrect commands or involve unnecessary steps, such as re-authenticating, which does not directly address the need to simply switch configurations.

9. Which flag is used with `kubectl` to specify the JSON path of properties in the JSON output?

- A. -j
- B. -p
- C. -json
- D. -o jsonpath**

The -j flag stands for "output as json" and is used to get the output as JSON. The -p flag provides the property pattern is used to do partial match with the property name in the output. The -json flag is not a valid flag for `kubectl` and will result in an error. The correct flag to specify the JSON path of properties is -o jsonpath.

10. Your developers have been thoroughly logging everything that happens in the API. The API allows end users to request the data as JSON, XML, CSV, and XLS. Supporting all of these formats is taking a lot of developer effort. Management would like to start tracking which options are used over the next month. Without modifying the code, what's the fastest way to be able to report on this data at the end of the month?

- A. Create a custom counter logging metric that uses a regex to extract the data format into a label. At the end of the month, use the metric viewer to see the group by the label.
- B. Create a log sink that filters for rows that mention the data format. Export that to BigQuery, and run a query at the end of the month.
- C. Create a custom monitoring metric in code and edit the API code to set the metric each time the API is called.
- D. Export the logs to excel, and search for the different fields.

Option B is incorrect because exporting to BigQuery and running a query at the end of the month would require modifying the code. Option C is incorrect because editing the API code to set a custom monitoring metric would also require modifying the code. Option D is incorrect because exporting the logs to excel and searching for different fields would be time consuming and not as efficient as using a custom counter logging metric. A custom counter logging metric allows for easy tracking and analysis of the data formats used without modifying the code.

SAMPLE

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://certifiedassociatecloudengineer.examzify.com>

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

SAMPLE