

Kubernetes Cloud Native Associate (KCNA) Certification 1 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	8
Explanations	10
Next Steps	15

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 purpose of labeling in Kubernetes?**
 - A. To enforce network policies**
 - B. To set CPU limits**
 - C. To designate attributes for grouping and selection**
 - D. To set memory quotas**

- 2. What best describes a Headless service's behavior?**
 - A. A ClusterIP service that has no IP**
 - B. A NodePort that exposes a port externally**
 - C. An ExternalName pointing to a DNS alias**
 - D. A LoadBalancer without an external IP**

- 3. If you want to prevent pod creation unless it passes a policy check, which Kubernetes feature would you rely on?**
 - A. HorizontalPodAutoscaler**
 - B. Admission Controllers**
 - C. ConfigMaps**
 - D. DaemonSets**

- 4. How do users typically interact with the Kubernetes API?**
 - A. Through a heavy GUI only**
 - B. Using command line tools like kubectl and Helm**
 - C. By editing YAML in a text editor without API**
 - D. Via FTP**

- 5. Which OCI specification defines how content is distributed between registries?**
 - A. Runtime Specification**
 - B. Distribution Specification**
 - C. Image Specification**
 - D. OCI Image Specification**

- 6. Vertical Pod Autoscalers (VPA) adjust which pod characteristics?**
- A. The number of replicas for a deployment**
 - B. The scheduling constraints of pods**
 - C. The resource requests and limits of a pod**
 - D. The namespace quotas**
- 7. Which RBAC component binds a user or group to a ClusterRole for cluster-wide permissions?**
- A. RoleBinding binds a user to a Role within a namespace**
 - B. A Namespace is required to bind a user to a ClusterRole**
 - C. A ServiceAccount binds a user to a Role in a namespace**
 - D. A ClusterRoleBinding binds a user or group to a ClusterRole for cluster-wide permissions**
- 8. Which role is commonly considered to have crossover with DevOps due to overlapping skillsets?**
- A. Security Engineer**
 - B. Cloud Architect**
 - C. Data Scientist**
 - D. Site Reliability Engineer**
- 9. In Kubernetes, the data inside a Secret is typically stored in what form?**
- A. Base64-encoded**
 - B. Encrypted with AES-256**
 - C. Plain text**
 - D. Compressed with gzip**
- 10. In kubectl, what does the -A flag do when listing resources?**
- A. All namespaces**
 - B. All pods**
 - C. All nodes**
 - D. All resources**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is the purpose of labeling in Kubernetes?

- A. To enforce network policies
- B. To set CPU limits
- C. To designate attributes for grouping and selection**
- D. To set memory quotas

Labels in Kubernetes are a way to attach attributes to objects as key-value pairs, enabling grouping and selection across the cluster. You can tag resources (like pods, services, or deployments) with labels such as `app=frontend` or `env=prod`. This lets you treat all resources with a given label as a group, and lets controllers and tools pick that group with label selectors. For example, a Service can route traffic to pods whose labels match its selector, and you can quickly filter resources with commands like `kubectl get pods -l app=frontend`. These labels are not about setting resource limits or enforcing network rules; CPU/memory settings pertain to container resources, and network policies are separate objects that control traffic. By design, labels provide flexible organization and precise selection to manage even large clusters.

2. What best describes a Headless service's behavior?

- A. A ClusterIP service that has no IP**
- B. A NodePort that exposes a port externally
- C. An ExternalName pointing to a DNS alias
- D. A LoadBalancer without an external IP

Headless service means no cluster IP. You create it by setting `clusterIP` to `None`, so Kubernetes does not allocate a virtual IP for the service. Instead, the DNS entry for the service returns multiple records pointing to each backing pod, allowing direct access to individual pods or custom load balancing. This is handy with StatefulSets where each pod has a stable identity. Because there is no cluster IP, the service isn't reached through a single IP and is not typically exposed via a NodePort or LoadBalancer; ExternalName and other patterns describe different mechanisms for DNS-based or external exposure, not headless behavior. So the description that fits is a service without a cluster IP.

3. If you want to prevent pod creation unless it passes a policy check, which Kubernetes feature would you rely on?

- A. HorizontalPodAutoscaler
- B. Admission Controllers**
- C. ConfigMaps
- D. DaemonSets

Admission controllers are the mechanism that intercepts requests to the API server after authentication and authorization and before the object is stored. They can validate inputs and reject the request if policies aren't met, so pod creation is prevented unless the policy check passes. You can use built-in validating admission controllers or set up validating webhooks to enforce custom policies (for example around security context, allowed fields, or image sources). The other options don't serve this purpose: HorizontalPodAutoscaler scales based on metrics, ConfigMaps store configuration data, and DaemonSets manage pod placement on nodes. So admission controllers are the way to enforce policy checks at creation time.

4. How do users typically interact with the Kubernetes API?

- A. Through a heavy GUI only
- B. Using command line tools like kubectl and Helm**
- C. By editing YAML in a text editor without API
- D. Via FTP

Users interact with Kubernetes by talking to the API server through clients that send RESTful requests. The standard workflow uses command-line tools like kubectl to apply manifests or run commands that translate into API calls to the API server. Helm is another common client; it acts as a package manager that uses the Kubernetes API to install, upgrade, and manage charts, which creates and updates resources in the cluster. These tools handle authentication and RBAC, giving you a practical, scriptable way to manage resources. Relying on a heavy GUI alone isn't the typical workflow, and editing YAML in a text editor without applying it to the cluster doesn't affect the cluster state since nothing is sent to the API. FTP isn't applicable to Kubernetes management.

5. Which OCI specification defines how content is distributed between registries?

- A. Runtime Specification
- B. Distribution Specification**
- C. Image Specification
- D. OCI Image Specification

OCI Distribution Specification defines how content is exchanged between registries. It specifies the protocol and semantics for pushing and pulling image artifacts—such as manifests, blobs, and indices—across registries, enabling interoperable distribution and synchronization of content. The OCI Image Specification, in contrast, describes the layout and metadata of an image itself (how the image is structured on disk and in transfers), while the Runtime Specification outlines how containers are executed. The OCI Image Specification is essentially the image format, not the distribution mechanism.

6. Vertical Pod Autoscalers (VPA) adjust which pod characteristics?

- A. The number of replicas for a deployment
- B. The scheduling constraints of pods
- C. The resource requests and limits of a pod**
- D. The namespace quotas

Vertical Pod Autoscaler focuses on giving a pod the right amount of resources by adjusting its container resource requests and limits (CPU and memory). This helps the scheduler reserve appropriate resources and prevents under- or over-provisioning of a pod. It does not change how many pods exist—that's handled by horizontal scaling. It also isn't about scheduling constraints like affinity rules, tolerations, or quotas at the namespace level. Those are separate concerns: constraints determine where a pod can run, and quotas limit total resources in a namespace. When VPA updates a PodSpec, you typically apply new resource requests/limits to trigger a rollout so the running pods use the updated resources. In short, VPA adjusts resource requests and limits, not the number of replicas or scheduling constraints.

7. Which RBAC component binds a user or group to a ClusterRole for cluster-wide permissions?
- A. RoleBinding binds a user to a Role within a namespace
 - B. A Namespace is required to bind a user to a ClusterRole
 - C. A ServiceAccount binds a user to a Role in a namespace
 - D. A ClusterRoleBinding binds a user or group to a ClusterRole for cluster-wide permissions**

Cluster-wide access in Kubernetes RBAC is granted by binding subjects to a ClusterRole using a ClusterRoleBinding. This binding connects users, groups, or service accounts to a ClusterRole, so the resulting permissions apply across the entire cluster. In contrast, binding a subject to a Role within a specific namespace (RoleBinding) limits those permissions to that namespace. A Namespace is just a scope, not a binding mechanism, and while a ServiceAccount can be a subject, binding it to a Role in a namespace yields namespace-scoped access. So the mechanism that provides cluster-wide permissions is ClusterRoleBinding.

8. Which role is commonly considered to have crossover with DevOps due to overlapping skillsets?
- A. Security Engineer
 - B. Cloud Architect
 - C. Data Scientist
 - D. Site Reliability Engineer**

Site Reliability Engineers are a natural fit with DevOps because both roles treat operations as software and focus on building automated, reliable systems. They share practices like infrastructure as code, continuous integration and deployment, extensive monitoring and alerting, and rigorous incident response and postmortems. SREs explicitly codify reliability through concepts like SLIs/SLOs and error budgets, which aligns closely with the DevOps goal of continual improvement of software delivery and operations. Other roles focus more on distinct domains—Security Engineers on protecting systems, Cloud Architects on designing cloud solutions, and Data Scientists on extracting insights from data. While they collaborate with DevOps, the day-to-day overlap isn't as direct as with Site Reliability Engineers.

9. In Kubernetes, the data inside a Secret is typically stored in what form?
- A. Base64-encoded**
 - B. Encrypted with AES-256
 - C. Plain text
 - D. Compressed with gzip

Base64-encoded data. Kubernetes stores Secret values as base64-encoded strings in etcd. The Secret's data field is a map where each value is encoded, and you can provide plaintext with stringData—the API server encodes it to base64 when persisting. Remember, base64 is encoding, not encryption, so if you need secrecy you must enable encryption at rest or use a KMS provider.

10. In kubectl, what does the -A flag do when listing resources?

A. All namespaces

B. All pods

C. All nodes

D. All resources

The flag expands the scope to include every namespace. It's shorthand for --all-namespaces, so when you list resources with it, kubectl queries all namespaces in the cluster instead of just the current one. For example, `kubectl get pods -A` will show pods from every namespace. Keep in mind some resources aren't namespaced (like nodes), so -A doesn't change their scope in the same way, but for namespaced resources it clearly broadens the view to all namespaces.

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://kubernetescloudnativeassockcna1.examzify.com>

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

SAMPLE