

KUBERNETES CERTIFIED NETWORK ADMINISTRATOR (KCNA) PART 2 PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

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

<https://kcnapt2.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	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. Which is a use case for combining Init Containers with Persistent Volume Claims (PVCs)?**
 - A. To run sidecar logging
 - B. To initialize a database schema before the app starts
 - C. To perform health checks
 - D. To manage secrets
- 2. In Kubernetes, how do you request and attach persistent storage to a Pod?**
 - A. By defining a PersistentVolumeClaim in the Pod specification
 - B. By creating a ConfigMap in the Pod spec
 - C. By using a StorageClass inside the Pod spec
 - D. By mounting a host path directly
- 3. Which tool exposes metrics describing the state of Kubernetes objects, such as pods and deployments, for monitoring?**
 - A. Prometheus
 - B. CoreDNS
 - C. Kube-state-metrics
 - D. Kubelet
- 4. In Kubernetes, which security-focused tool is commonly used for runtime security monitoring and detection of anomalous activities within containers and pods?**
 - A. Falco
 - B. AquaSec
 - C. Twistlock
 - D. Sysdig
- 5. What is a sidecar container?**
 - A. A container that stores persistent data for the application
 - B. A container that runs alongside the main container to handle administrative tasks
 - C. A separate pod that hosts the main application
 - D. A security container that blocks unauthorized access

- 6. The ServiceMeshInterface (SMI) specification provides what kind of abstraction for service meshes?**
- A. A deployment model for pods
 - B. A data plane component for mesh networking
 - C. A standardised interface for traffic policy across meshes
 - D. A monitoring dashboard for service meshes
- 7. Which of the following best describes a headless service in Kubernetes?**
- A. It uses a load balancer to distribute traffic
 - B. It has no cluster IP, allowing direct pod DNS resolution
 - C. It requires an external DNS
 - D. It is only accessible within the cluster
- 8. Which of the following is a component of the Kubernetes Control Plane?**
- A. kube-scheduler
 - B. cloud-controller-manager
 - C. kubelet
 - D. kube-proxy
- 9. Which Kubernetes object is designed for stateful workloads requiring stable network identities and persistent storage for each pod?**
- A. Deployment
 - B. DaemonSet
 - C. StatefulSet
 - D. Job
- 10. In the Open Container Initiative (OCI) Specification, which areas are standardised?**
- A. Container image format and distribution only
 - B. Runtime environment only
 - C. Distribution and runtime environment
 - D. Container image format, runtime environment, and distribution

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which is a use case for combining Init Containers with Persistent Volume Claims (PVCs)?

- A. To run sidecar logging
- B. To initialize a database schema before the app starts**
- C. To perform health checks
- D. To manage secrets

Init containers run to completion before the main application containers, and when they mount the same persistent volume, they can set up the data that the app will use. This makes it ideal to pre-create and populate a database's schema or seed data in a volume before the application starts. By initializing the data directory on the PVC, the app container starts with a ready-to-use database, avoiding migrations or setup at first run and reducing startup failures. For example, a PVC backs the database data directory; an Init Container runs the schema creation and seeds, writes the initial state to the mounted path, and only then does the main app container start using that pre-initialized data. Other patterns like sidecar logging, health checks, or secret management serve different purposes and don't specifically involve preparing a persistent data store before startup.

2. In Kubernetes, how do you request and attach persistent storage to a Pod?

- A. By defining a PersistentVolumeClaim in the Pod specification**
- B. By creating a ConfigMap in the Pod spec
- C. By using a StorageClass inside the Pod spec
- D. By mounting a host path directly

In Kubernetes, to request and attach persistent storage to a Pod you use a PersistentVolumeClaim and connect it to the Pod through a volume. The claim specifies how much storage you need and with what access mode, and Kubernetes will bind it to an existing PersistentVolume or dynamically provision one if a StorageClass is configured. In the Pod spec you then declare a volume that uses that claim (persistentVolumeClaim: claimName: your-pvc) and mount it into the container via a volumeMounts entry. This approach decouples the Pod from the storage lifecycle and enables portability and dynamic provisioning. ConfigMaps are for configuration data, not storage. A StorageClass defines how storage is provisioned but isn't directly attached to a Pod; it's used by the PVC to provision PVs. Mounting a hostPath attaches a directory from the node's filesystem, which isn't portable or durable across nodes.

3. Which tool exposes metrics describing the state of Kubernetes objects, such as pods and deployments, for monitoring?

- A. Prometheus
- B. CoreDNS
- C. Kube-state-metrics**
- D. Kubelet

The thing being tested is how you get metrics that describe the real-time state of Kubernetes objects. The tool that does this is kube-state-metrics. It runs in the cluster, queries the Kubernetes API for resources like pods and deployments, and exposes metrics in Prometheus format. This lets Prometheus scrape data such as pod status, deployment replicas, and ready counts to monitor how resources are actually behaving. Prometheus is the scraping and storage system, not the source of the object-state metrics. CoreDNS handles DNS inside the cluster, not object-state monitoring. Kubelet runs on each node and exposes node-level and container metrics, not a broad view of Kubernetes object states.

4. In Kubernetes, which security-focused tool is commonly used for runtime security monitoring and detection of anomalous activities within containers and pods?

- A. Falco
- B. AquaSec
- C. Twistlock
- D. Sysdig

Focusing on runtime behavior inside containers and pods, this question centers on real-time security monitoring that detects anomalous activities as workloads run. Falco is designed for exactly that: a runtime security tool that watches what the system is actually doing, rather than just scanning configurations or images. Falco observes kernel and system call events from the host (and can be deployed in Kubernetes as a DaemonSet) and uses a flexible rules engine to flag suspicious activity. This lets you detect things like a shell being spawned inside a running container, a container accessing sensitive host paths, unusual process trees, or unexpected network connections from a pod. Because it analyzes behavior in real time, it provides immediate alerts about active threats or policy violations, which is essential for protecting workloads in a dynamic Kubernetes environment. Other options offer valuable security capabilities, like image scanning, secrets protection, or broader runtime protection suites, but Falco's strength lies in its focused, open-source, runtime, rule-based detection tailored to Kubernetes workloads.

5. What is a sidecar container?

- A. A container that stores persistent data for the application
- B. A container that runs alongside the main container to handle administrative tasks
- C. A separate pod that hosts the main application
- D. A security container that blocks unauthorized access

A sidecar container is a companion container that runs in the same Kubernetes Pod as the main application container and performs supporting tasks. It handles auxiliary work like log collection, monitoring, data synchronization, or proxying requests, so the primary process can stay focused on its core function. Because it shares the Pod's network space and can mount the same volumes, the sidecar and the main container can communicate directly (often via localhost) and coordinate through the shared filesystem. This pattern keeps operational concerns separate from the main application logic without needing a separate Pod. For example, a logging agent or a token refresher running as a sidecar can handle those tasks transparently while the main app processes data.

6. The ServiceMeshInterface (SMI) specification provides what kind of abstraction for service meshes?

- A. A deployment model for pods
- B. A data plane component for mesh networking
- C. A standardised interface for traffic policy across meshes**
- D. A monitoring dashboard for service meshes

SMI provides a standardised interface for traffic policy across service meshes. It defines a vendor-agnostic set of abstractions that describe how traffic should be routed and governed between services, and are implemented by different meshes. This lets you declare policies once and apply them across multiple meshes, enabling portability and interoperability. It decouples policy from the underlying data plane, so changing the mesh doesn't require rewriting routing rules. It's not about pod deployment models, its own data plane components, or a monitoring dashboard—the policy interface is the common layer that sits above those mesh-specific implementations.

7. Which of the following best describes a headless service in Kubernetes?

- A. It uses a load balancer to distribute traffic
- B. It has no cluster IP, allowing direct pod DNS resolution**
- C. It requires an external DNS
- D. It is only accessible within the cluster

A headless service is defined by not getting a cluster IP. When you create such a service (clusterIP is None), there isn't a single virtual IP to front the pods. Instead, the service's DNS name resolves to the actual pods backing it, so clients can reach the individual pod IPs directly through DNS, or use SRV records to discover the ports. This is why the description "no cluster IP, allowing direct pod DNS resolution" best captures headless services. That's why the other statements don't fit: there isn't a service-level load balancer for traffic distribution, since there's no cluster IP to balance to; internal DNS is sufficient for DNS resolution—external DNS isn't a requirement; and while headless services are often used within the cluster, they aren't defined by being strictly inaccessible from outside the cluster.

8. Which of the following is a component of the Kubernetes Control Plane?

- A. kube-scheduler
- B. cloud-controller-manager**
- C. kubelet
- D. kube-proxy

The control plane coordinates the whole cluster and includes components that run APIs, manage state, and make global decisions. The cloud-controller-manager is the part of the control plane that talks directly to the cloud provider's API to manage cloud-specific resources, such as load balancers, external IPs, and certain node lifecycle tasks. It's designed to run separately from the core controllers so cloud-specific logic can be decoupled and swapped as needed, which is why it's considered a control-plane component. The other options operate on the nodes rather than in the control plane: the kubelet is the agent that runs on each worker node to start and supervise containers, and kube-proxy runs on each node to implement Kubernetes services networking. Kube-scheduler is indeed a control-plane component that schedules pods, but the question highlights the cloud-provider integration role, which is fulfilled by the cloud-controller-manager.

9. Which Kubernetes object is designed for stateful workloads requiring stable network identities and persistent storage for each pod?

- A. Deployment
- B. DaemonSet
- C. StatefulSet**
- D. Job

Stateful workloads need predictable, enduring identities and dedicated storage for each instance. A StatefulSet is built for exactly that: every pod gets a stable, unique network identity and an ordinal index (like app-0, app-1, etc.) that survives rescheduling. This makes it possible for other components to reliably connect to the same pod address over time, which is crucial for databases, queues, and other stateful services. In addition, StatefulSet provisions per-pod storage, so each instance has its own durable volume that remains attached to the same pod identity even if the pod moves to a different node. This pairing of stable network identity and per-pod storage is what enables stateful applications to recover cleanly and maintain data integrity. Deployments are great for stateless, scalable workloads where pod identities can change and storage is typically shared or ephemeral. DaemonSets ensure one pod per node, but they don't provide stable per-pod identities or per-pod durable storage in the way StatefulSet does. Jobs run tasks to completion and don't support long-running stateful services with persistent storage tied to a fixed identity. So, for stateful workloads needing stable network identities and persistent per-pod storage, StatefulSet is the appropriate Kubernetes object.

10. In the Open Container Initiative (OCI) Specification, which areas are standardised?

- A. Container image format and distribution only
- B. Runtime environment only
- C. Distribution and runtime environment
- D. Container image format, runtime environment, and distribution**

The OCI standardizes three interlocking areas to ensure portability: the container image format, the runtime environment, and the distribution layer. The image format defines how images are packaged—layers, configuration, and manifests—so any OCI-compliant runtime can interpret and extract the image. The runtime environment (the OCI Runtime Specification) dictates how a container is actually executed by a runtime, describing how processes, namespaces, and resources are setup to run the image consistently. The distribution aspect (the OCI Distribution Specification) standardizes how images are stored, pulled, and exchanged with registries, including how manifests and layers are addressed and authenticated. By standardizing all three, images created by one tool can be run by any OCI-compliant runtime against any OCI-compliant registry.

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://kcnapt2.examzify.com>

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

SAMPLE