

Certified Kubernetes Application Developer (CKAD) Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

This is a sample study guide. To access the full version with hundreds of questions,

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 from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Table of Contents

| | |
|------------------------------------|-----------|
| Copyright | 1 |
| Table of Contents | 2 |
| Introduction | 3 |
| How to Use This Guide | 4 |
| Questions | 6 |
| 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. Don't worry about getting everything right, 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, and take breaks to retain information better.

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.

7. Use Other Tools

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!

SAMPLE

Questions

SAMPLE

- 1. In the context of Kubernetes networking, what does the default network policy allow?**
 - A. Nothing**
 - B. All Deny**
 - C. All Allow**
 - D. Specific Allow rules defined by the user**
- 2. How do you set the maximum number of failures for a probe in Kubernetes?**
 - A. By adjusting the successThreshold property**
 - B. By modifying the failThreshold property**
 - C. By using the maxFailures property under the probe**
 - D. By changing the failureThreshold property**
- 3. In the context of Kubernetes, what does a headless service NOT provide?**
 - A. Load balancing**
 - B. Direct connection to individual pods**
 - C. Stable network identity for pods**
 - D. Internal DNS for pod communication**
- 4. Which command is used to create a deployment in Kubernetes?**
 - A. kubectl start deployment [deployment-name]**
 - B. kubectl create deploy [deployment-name]
--image=[image-name]**
 - C. kubectl create deployment [deployment-name]
--image=[image-name]**
 - D. kubectl deploy [deployment-name] --image=[image-name]**
- 5. To associate a stateful set with a headless service, which property should be set in the stateful set spec?**
 - A. hostname property**
 - B. serviceName property**
 - C. volumeClaimTemplate property**
 - D. subdomain property**

6. What defines a ReplicaSet in Kubernetes?

- A. It ensures a specified number of pod replicas are running.**
- B. It provides a mechanism for automatic scaling of pods.**
- C. It routes traffic to multiple pods serving the same application.**
- D. It manages the storage of container images.**

7. What is the role of a Service Account in Kubernetes?

- A. It provides storage for persistent data**
- B. It serves as an identity for processes in pods**
- C. It manages network policies for applications**
- D. It oversees logging and monitoring of pods**

8. What command would you use to retrieve the cluster's metadata information?

- A. kubectl cluster-info**
- B. describe cluster**
- C. fetch cluster-info**
- D. cluster-info**

9. What command would you use to see the logs of a specific pod?

- A. kubectl show logs <pod>**
- B. kubectl get logs <pod>**
- C. kubectl logs <pod>**
- D. kubectl watch logs <pod>**

10. Which type of probe can run commands directly inside the container?

- A. httpGet**
- B. tcpSocket**
- C. exec**
- D. none**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. In the context of Kubernetes networking, what does the default network policy allow?

- A. Nothing
- B. All Deny
- C. All Allow**
- D. Specific Allow rules defined by the user

In Kubernetes networking, when no network policies are applied to a namespace, the default behavior allows all traffic to and from all pods within that namespace and to/from pods in other namespaces. This means that, by default, any pod can communicate with any other pod without restrictions, implying that all ingress (incoming) and egress (outgoing) traffic is permitted. The option stating "All Allow" correctly reflects this default setting, as it means that there are no constraints on pod communication unless explicitly defined by a custom network policy that enforces specific rules. Custom network policies can limit traffic by specifying which pods are allowed to communicate with one another or with external services, but unless such policies are created, the default state is one of openness, allowing all traffic.

2. How do you set the maximum number of failures for a probe in Kubernetes?

- A. By adjusting the successThreshold property
- B. By modifying the failThreshold property
- C. By using the maxFailures property under the probe
- D. By changing the failureThreshold property**

The correct approach to set the maximum number of failures for a probe in Kubernetes is by modifying the failureThreshold property. This property specifically defines the minimum consecutive failures for the probe to be considered failed. It allows you to specify how many times a pod must fail the probe before Kubernetes takes action, such as restarting the container or marking the pod as unhealthy. Understanding the function of failureThreshold is crucial because it directly impacts the responsiveness of the system in terms of detecting and responding to issues with pods. If this value is set too low, it could lead to unnecessary restarts, while setting it too high might delay the detection of a real problem, potentially affecting application reliability. In contrast, adjusting the successThreshold property pertains to readiness probes and defines how many consecutive successful probes must occur before a pod is considered ready. The maxFailures property does not exist in the Kubernetes probe specifications, making it an invalid option in this context. The failureThreshold property, as mentioned, is the actual attribute used for configuring how many times a failure must occur before acknowledging that there is a problem with the pod.

3. In the context of Kubernetes, what does a headless service NOT provide?

- A. Load balancing**
- B. Direct connection to individual pods**
- C. Stable network identity for pods**
- D. Internal DNS for pod communication**

A headless service in Kubernetes is designed to provide a way for clients to connect directly to individual pods rather than utilizing a load balancer. This type of service is essential when applications require direct access to the underlying pods, enabling clients to connect specifically to one pod's IP rather than via a stable service IP. Choosing "A" as the correct answer highlights that headless services do not provide load balancing capabilities. When a typical service is created in Kubernetes, it might distribute incoming traffic across multiple pods, serving as a load balancer. However, with a headless service (prompted by setting the ClusterIP field to "None"), incoming traffic is not balanced; instead, it goes directly to the individual pods based on their unique IPs. The other options reflect the characteristics of headless services: - Direct connection to individual pods is a primary feature, making it suitable for certain applications that benefit from reduced latency or specific routing. - While headless services do not offer a stable service IP, they enable pods to maintain stable network identities through their individual pod IPs, which can be crucial for certain workloads needing specific addressing. - Internal DNS for pod communication is still available in the context of headless services, where DNS entries are created directly for the

4. Which command is used to create a deployment in Kubernetes?

- A. kubectl start deployment [deployment-name]**
- B. kubectl create deploy [deployment-name]
--image=[image-name]**
- C. kubectl create deployment [deployment-name]
--image=[image-name]**
- D. kubectl deploy [deployment-name] --image=[image-name]**

The command to create a deployment in Kubernetes is structured as follows: `kubectl create deployment [deployment-name] --image=[image-name]`. This command uses the `kubectl` command-line tool, which is the primary way to interact with the Kubernetes cluster. In this command: - `create deployment` specifies that you want to create a new deployment resource. This part of the command precisely tells Kubernetes what type of resource you are creating. - `[deployment-name]` is a placeholder for the name you wish to give to the deployment, and it must be unique within the namespace. - `--image=[image-name]` indicates the container image to be used for the deployment. This part is crucial as it defines what application will run inside the pods created by the deployment. This command is complete and properly formatted, allowing Kubernetes to understand the intent and parameters provided. The other command options do not correctly reflect the standard syntax or capabilities provided by `kubectl` for deployment creation, which is why they are not valid. The correct structure and use of deployment commands are essential for managing applications effectively in Kubernetes.

5. To associate a stateful set with a headless service, which property should be set in the stateful set spec?

- A. hostname property**
- B. serviceName property**
- C. volumeClaimTemplate property**
- D. subdomain property**

The stateful set specification requires the `serviceName` property to be set when associating it with a headless service. This property indicates which service will be used for the stateful set's associated pods, enabling the stateful set to seamlessly integrate with the service. In Kubernetes, a headless service is defined by setting the `ClusterIP` field to "None", which allows the DNS system to create individual A records for each pod associated with the stateful set. By specifying the `serviceName` in the stateful set, you ensure that the pods can be reached using their DNS names, which follow the pattern `'<pod-name>.<service-name>'`. This is fundamental for stateful applications that rely on stable network identities. Although other properties such as `hostname` or `volumeClaimTemplate` are important in their contexts, they do not directly contribute to the association with a headless service in the way that `serviceName` does. The `subdomain` property could be involved in filtering DNS records, but it is secondary to the explicit connection provided by setting the `serviceName`. Therefore, the correct focus on `serviceName` marks the definitive link between the stateful set and its service.

6. What defines a ReplicaSet in Kubernetes?

- A. It ensures a specified number of pod replicas are running.**
- B. It provides a mechanism for automatic scaling of pods.**
- C. It routes traffic to multiple pods serving the same application.**
- D. It manages the storage of container images.**

A ReplicaSet in Kubernetes is fundamentally defined by its role in ensuring that a specified number of pod replicas are running at any given time. This means that if some pods fail or are deleted, the ReplicaSet controller will automatically create new pods to maintain the desired count. This self-healing capability is a key feature of ReplicaSets, making them essential for maintaining the availability and reliability of applications deployed in a Kubernetes cluster. While other options describe important components of Kubernetes, they do not accurately reflect the primary function of a ReplicaSet. Automatic scaling of pods is typically the responsibility of a Horizontal Pod Autoscaler, which adjusts the number of pods based on resource metrics. Traffic routing to multiple pods is managed by Services in Kubernetes, which provide stable endpoints to direct client requests. Managing the storage of container images falls under container registries, which are separate from ReplicaSet functionality. Thus, the specific role of maintaining a desired state of pod replicas distinctly characterizes a ReplicaSet.

7. What is the role of a Service Account in Kubernetes?

- A. It provides storage for persistent data
- B. It serves as an identity for processes in pods**
- C. It manages network policies for applications
- D. It oversees logging and monitoring of pods

The role of a Service Account in Kubernetes is to serve as an identity for processes running in pods. Each pod can be associated with a Service Account, allowing it to access the Kubernetes API and other resources in a secure way. This is essential for enabling pods to authenticate themselves and perform actions according to the permissions granted to the Service Account. By using Service Accounts, Kubernetes can enforce access controls and manage permissions based on roles. This means that specific pods can be restricted to only the permissions they need to function, enhancing security within the cluster. The other options focus on different aspects of Kubernetes functionality. For instance, persistent storage is managed through Persistent Volume Claims and Persistent Volumes, network policies are established using NetworkPolicy resources, and logging and monitoring are typically handled by specialized tools integrated into the Kubernetes ecosystem, not Service Accounts. Thus, the correct answer adeptly highlights the primary purpose of a Service Account within Kubernetes.

8. What command would you use to retrieve the cluster's metadata information?

- A. kubectl cluster-info**
- B. describe cluster
- C. fetch cluster-info
- D. cluster-info

The command "get cluster-info" is valid and retrieves the essential metadata about the Kubernetes cluster, including the master node's details and the services that are running on the cluster. This command is pivotal for anyone needing a quick overview of cluster state and configuration, facilitating operational awareness and troubleshooting insights in a Kubernetes environment. In contrast, the other command options provided do not align with Kubernetes' command syntax or do not exist as valid commands in `kubectl`, which is the command-line tool for interacting with Kubernetes clusters. For example, "describe cluster" is not a recognized command in the `kubectl` CLI; the `describe` command is typically used with specific resources like pods, nodes, or services rather than at the cluster level. Similarly, "fetch cluster-info" and "cluster-info" are not valid commands, as there is no `fetch` command available in `kubectl`, and commands must follow the `get` or `describe` formats with appropriate resource types. Therefore, using "get cluster-info" is indeed the correct and valid way to access the cluster metadata.

9. What command would you use to see the logs of a specific pod?

- A. `kubectl show logs <pod>`
- B. `kubectl get logs <pod>`
- C. `kubectl logs <pod>`**
- D. `kubectl watch logs <pod>`

To view the logs of a specific pod in Kubernetes, the command you would use is '`kubectl logs <pod>`'. This command is specifically designed for retrieving the logs that a container in a specified pod produces. It fetches the log output from the container's standard output (stdout) and standard error (stderr), allowing developers and operators to diagnose issues or monitor the behavior of applications running within the Kubernetes cluster. The '`kubectl logs <pod>`' command is versatile; it can be used to view logs from a pod with a single container or specify which container to view in a pod that has multiple containers. This is particularly useful in a microservices architecture where multiple components are deployed as separate containers within the same pod. In contrast, the other commands do not perform this specific action. For instance, '`kubectl show logs <pod>`' does not exist in Kubernetes, which means it would not execute successfully. Similarly, '`kubectl get logs <pod>`' does not align with the Kubernetes command structure; 'get' is used for listing resources rather than fetching logs. Lastly, '`kubectl watch logs <pod>`' is not a valid command either since Kubernetes doesn't support the 'watch' aspect directly applied to logs retrieval in that manner. Thus,

10. Which type of probe can run commands directly inside the container?

- A. `httpGet`
- B. `tcpSocket`
- C. `exec`**
- D. `none`

The correct answer is "exec," as this type of probe is designed to run commands directly within the container's environment. The exec probe allows Kubernetes to execute a specified command inside the container, enabling the system to assess the container's health based on the command's exit status. If the command returns a successful exit code (typically 0), the container is considered healthy. In contrast, the other options do not allow for direct execution of commands inside the container. The `httpGet` probe checks the health of a container by making an HTTP GET request to a specified endpoint; it evaluates the response status code to determine health. The `tcpSocket` probe tests connectivity by initiating a TCP connection to a specified port on the container, but it does not execute commands. The option "`none`" is not applicable since there is indeed a probe type that executes commands. This understanding of probes is crucial in Kubernetes, as they help ensure that applications are running correctly and can recover from failures. Knowing the specific capabilities of each probe type will help in making informed decisions for monitoring and maintaining application health in a Kubernetes environment.

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://ckad.examzify.com>

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

SAMPLE