

Red Hat Openshift Developer EX288 Practice Exam (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	18

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 role of a Deployment in OpenShift, and how does it relate to ReplicaSets?**
 - A. A Deployment stores image streams for versions.**
 - B. A Deployment serves as a storage class for volumes.**
 - C. A Deployment exposes services to external clients.**
 - D. A Deployment manages a set of Pods through an associated ReplicaSet and handles rolling updates.**

- 2. In a Helm chart, how would you define a dependency on mysql version 5.6 from the repository https://mysql.helm.com/charts?**
 - A. - name: dependencies
version: 5.6
repository: https://mysql.helm.com/charts**
 - B. - name: mysql
version: 5.6
repository: https://mysql.helm.com/charts**
 - C. dependencies:
- name: mysql
version: 5.6
repository: https://mysql.helm.com/charts**
 - D. dependencies:
- name: mysql
version: 5.6
repository: https://charts.helm.sh/stable**

- 3. Which command verifies the current OpenShift project you are in after switching?**
 - A. oc project show**
 - B. oc project current**
 - C. oc project status**
 - D. oc project**

4. What is the purpose of a Service in OpenShift, and how does it enable pod communication?
- A. A Service stores secrets for containers.
 - B. A Service defines build steps for a container image.
 - C. A Service provides a stable endpoint and load-balancing to a set of Pods.
 - D. A Service is used to configure Routes.
5. Which command links the secret quayio to the builder service account so S2I can use it?
- A. `oc create rolebinding image-puller --group system:image-puller`
 - B. `oc policy add-role-to-group system:image-puller system:serviceaccount:myphp -n otherphp`
 - C. `oc policy add-role-to-group system:image-puller system:serviceaccount:myphp -n shared`
 - D. `oc policy remove-role-to-group system:image-puller`
6. In a Kubernetes Deployment, which field controls the number of desired pod replicas?
- A. replicas
 - B. replicasCount
 - C. desiredReplicas
 - D. podCount
7. How do you pause and resume a Deployment rollout in OpenShift?
- A. Pause with `oc rollout pause deployment/<name>` and resume with `oc rollout resume deployment/<name>`.
 - B. Pause with `oc rollout stop deployment/<name>` and resume with `oc rollout continue`.
 - C. Pause with `oc rollout pause deployment/<name>` and resume with `oc rollout resume deployment/<name>`.
 - D. Pause via `oc scale deployment/<name> to 0` and resume to original.

- 8. In an overlay kustomization.yaml, how do you refer to an overlay file called replica_count.yaml?**
- A. patches: - replica_count.yaml**
 - B. overlays: replica_count.yaml**
 - C. resources: replica_count.yaml**
 - D. patchesStrategicMerge: - replica_count.yaml**
- 9. What is required to persist a database's data across restarts in OpenShift?**
- A. A PersistentVolumeClaim bound to a PersistentVolume; mount the PVC to the database container.**
 - B. Use an emptyDir to persist data.**
 - C. Store data in the container's writable layer.**
 - D. Use a HostPath volume.**
- 10. In the image change trigger example, which namespace is specified?**
- A. default**
 - B. phpns**
 - C. openshift**
 - D. mysqlns**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is the role of a Deployment in OpenShift, and how does it relate to ReplicaSets?

- A. A Deployment stores image streams for versions.**
- B. A Deployment serves as a storage class for volumes.**
- C. A Deployment exposes services to external clients.**
- D. A Deployment manages a set of Pods through an associated ReplicaSet and handles rolling updates.**

A Deployment provides a declarative way to manage a set of Pods by coordinating a ReplicaSet and handling rolling updates. You declare how many replicas you want, the Pod template, and how updates should be performed. The Deployment controller ensures there is a corresponding ReplicaSet for those pods. When you change the Deployment (for example, update the image), it creates a new ReplicaSet with the updated Pod template and starts a rolling update, gradually replacing old Pods with new ones while keeping the system available. If something goes wrong, you can roll back to a previous revision. In short, the Deployment manages the lifecycle and updates of Pods indirectly through a ReplicaSet, while the ReplicaSet ensures the desired number of Pod replicas are running. Deployments aren't for storing image streams, selecting storage classes, or exposing services—those are handled by other resources.

SAMPLE

2. In a Helm chart, how would you define a dependency on mysql version 5.6 from the repository <https://mysql.helm.com/charts/>?

A. - name: dependencies

version: 5.6

repository: <https://mysql.helm.com/charts>

B. - name: mysql

version: 5.6

repository: <https://mysql.helm.com/charts>

C. dependencies:

- name: mysql

version: 5.6

repository: <https://mysql.helm.com/charts>

D. dependencies:

- name: mysql

version: 5.6

repository: <https://charts.helm.sh/stable>

Defining dependencies in a Helm chart is about listing other charts your chart relies on, inside Chart.yaml under a dependencies section. Each dependency is a YAML item that specifies the chart name, the version you want, and the repository where that chart is hosted. This structure lets Helm locate and fetch the exact dependency during installation or when you run helm dependency update. The reason this option is the best fit is that it shows the correct dependencies block at the top level, then a single dependency entry with three fields: name set to mysql, version set to 5.6, and repository pointing to the chart repository that hosts that mysql chart. This is the proper way to declare a dependent chart: you specify which chart you want, which version of that chart, and where to find it. After adding this, you would typically run helm repo update and then helm dependency update to fetch the dependency into your chart's charts/ directory. Context: in Helm 3, dependencies are declared in Chart.yaml (not requirements.yaml), and the repository must be a reachable chart repository URL. The version is the version of the mysql chart you want, not the version of MySQL itself. If you need to pull from a different repository, you would replace the repository URL with the correct host and then update the repo so Helm can locate the chart.

3. Which command verifies the current OpenShift project you are in after switching?

- A. oc project show**
- B. oc project current**
- C. oc project status**
- D. oc project**

To verify the active OpenShift project, use the project command with no arguments. After you switch projects, typing `oc project` will display the currently selected project, typically showing a message like “Using project 'my-project' on server ...”. This simple check confirms your current context so subsequent commands run against the right project. The other variants aren't standard ways to check the active project in OpenShift, so they don't reliably indicate which project you're in.

4. What is the purpose of a Service in OpenShift, and how does it enable pod communication?

- A. A Service stores secrets for containers.**
- B. A Service defines build steps for a container image.**
- C. A Service provides a stable endpoint and load-balancing to a set of Pods.**
- D. A Service is used to configure Routes.**

In OpenShift, a Service provides a stable endpoint for a group of pods and handles distributing traffic among them. It uses a selector to identify the pods that belong to the service and creates a stable IP and DNS name that clients can use, so they don't need to track individual pod IPs. The platform forwards requests to one of the healthy pods, effectively load-balancing across the set. This decouples clients from the ephemeral pod IPs, so communication remains reliable as pods are created, scaled, or replaced. The service is primarily for internal communication (ClusterIP by default), though it can be exposed externally via other mechanisms. Other options aren't the purpose of a Service: secrets are managed by Secrets, build steps are defined by BuildConfig, and Routes expose services externally rather than define the service's role.

5. Which command links the secret quayio to the builder service account so S2I can use it?

- A. `oc create rolebinding image-puller --group system:image-puller`
- B. `oc policy add-role-to-group system:image-puller system:serviceaccount:myphp -n otherphp`**
- C. `oc policy add-role-to-group system:image-puller system:serviceaccount:myphp -n shared`
- D. `oc policy remove-role-to-group system:image-puller`

OpenShift needs the builder's service account to have permission to pull the base image from a registry for S2I to run using the credentials in your quay.io secret. The way to do this is by binding the image-puller role to that specific service account in the namespace where the builder operates. The command that attaches the image-puller role to the service account named myphp in the otherphp namespace does exactly that: it grants the necessary image pull permission to the builder's service account in the correct scope, enabling S2I to access the private quay.io image using the secret. The other options either try to create or remove bindings without tying them to the specific service account in the right namespace, or refer to removing access, which would not enable the build to proceed. This targeted role binding is the appropriate step to authorize image pulls needed by the builder during the S2I process.

6. In a Kubernetes Deployment, which field controls the number of desired pod replicas?

- A. `replicas`**
- B. `replicasCount`
- C. `desiredReplicas`
- D. `podCount`

The number of desired pod replicas in a Deployment is controlled by the replicas field in the Deployment spec. This field tells the Deployment how many pod instances it aims to have running at any given time. The Deployment controller reconciles the actual state to match this desired state by creating or terminating pods (via the underlying ReplicaSet) as needed. If you set replicas to a certain number, the controller works to maintain that many pods; if the field is omitted, the default is 1. To change the number, you can modify this field in the manifest or use commands like `kubectl scale`. The other options aren't valid fields in a Deployment spec—there isn't a replicasCount, desiredReplicas, or podCount field, so they don't define the desired scale.

7. How do you pause and resume a Deployment rollout in OpenShift?

- A. Pause with `oc rollout pause deployment/<name>` and resume with `oc rollout resume deployment/<name>`.**
- B. Pause with `oc rollout stop deployment/<name>` and resume with `oc rollout continue`.**
- C. Pause with `oc rollout pause deployment/<name>` and resume with `oc rollout resume deployment/<name>`.**
- D. Pause via `oc scale deployment/<name>` to 0 and resume to original.**

Pausing and resuming a Deployment rollout is controlled with the rollout commands on the oc CLI. To pause a rollout, use `oc rollout pause deployment/<name>`. To continue after pausing, use `oc rollout resume deployment/<name>`. Pausing stops the rollout at its current step so no further updates are applied, giving you a window to inspect or fix issues before continuing. Resuming picks up from where it left off and proceeds with the rollout. The other options aren't correct because one uses non-existent rollout subcommands (`stop/continue`) and the other scales the deployment to zero, which changes the number of pods but doesn't pause or resume the rollout process. You can also check progress with `oc rollout status deployment/<name>`.

8. In an overlay `kustomization.yaml`, how do you refer to an overlay file called `replica_count.yaml`?

- A. `patches: - replica_count.yaml`**
- B. `overlays: replica_count.yaml`**
- C. `resources: replica_count.yaml`**
- D. `patchesStrategicMerge: - replica_count.yaml`**

In an overlay, changes to existing resources are applied through patch files listed under `patches`. Each entry points to a patch file that contains the modification you want to apply. So to reference a patch file named `replica_count.yaml`, you place it in the `patches` list like this: `patches: - replica_count.yaml` This tells Kustomize to apply that patch when building the overlay, modifying the target resource (for example, adjusting the `replicas` field in a Deployment) according to the patch contents. The other options don't reflect how patches are referenced in an overlay: `overlays` is not a standard field for listing patch files; `resources` is used to bring in base resources, not to apply patches; and `patchesStrategicMerge` is used only when you intend to apply a strategic merge patch, which is a different patch mechanism.

9. What is required to persist a database's data across restarts in OpenShift?

A. A PersistentVolumeClaim bound to a PersistentVolume; mount the PVC to the database container.

B. Use an emptyDir to persist data.

C. Store data in the container's writable layer.

D. Use a HostPath volume.

Persisting a database's data across restarts requires durable storage attached to the database pod. In OpenShift this is done by creating a PersistentVolume and a PersistentVolumeClaim, then mounting the PVC into the database container. The database writes go to that mounted volume, which lives beyond the lifecycle of any single pod. If the pod restarts or is rescheduled to another node, the same PVC is reattached and the data remains available. Why the other options aren't suitable: an emptyDir is ephemeral and disappears when the pod stops, so no data persistence across restarts. Writing data to the container's writable layer is also ephemeral and tied to the container's lifecycle, not a durable external store. A HostPath volume binds data to a specific node's filesystem, which breaks portability and reliability if the pod moves to a different node or if the cluster scales. Using a PVC backed by a PV (ideally provisioned by a StorageClass) is the correct approach for persistent database storage in OpenShift.

10. In the image change trigger example, which namespace is specified?

A. default

B. phpns

C. openshift

D. mysqlns

In OpenShift, an image change trigger watches a specific ImageStreamTag within a chosen namespace, so it knows where to look for new images to trigger deployments. The example explicitly sets the namespace to phpns, meaning the trigger will monitor the ImageStream in the phpns project. That's why phpns is the best choice here. The other names would point to different projects or rely on a default namespace, which isn't what the example demonstrates.

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://redhatopenshiftdevex288.examzify.com>

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

SAMPLE