

# RED HAT OPENSIFT DEVELOPER II DO288 PRACTICE EXAM (SAMPLE)

**STUDY GUIDE**



**SAMPLE STUDY GUIDE  
WITH LIMITED QUESTIONS**

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

<https://redhatopenshiftdev2do288.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 command requests a new build for an application?**
  - A. `oc start-build <application-name>`
  - B. `oc start-build application-name`
  - C. `oc new-build <application-name>`
  - D. `oc trigger-build <application-name>`
- 2. Which command removes a local image?**
  - A. `sudo podman tag mysql-custom devops/mysql:snapshot`
  - B. `sudo podman rmi devops/mysql:snapshot`
  - C. `oc rmi devops/mysql:snapshot`
  - D. `docker rmi devops/mysql:snapshot`
- 3. Which command creates a service of type ExternalName that points to an external host name?**
  - A. `oc create service externalname <servicename> --external-name <mysql_db>`
  - B. `oc create service externalName <servicename> --external-name <mysql_db>`
  - C. `oc create service clusterip <servicename> --external-name <mysql_db>`
  - D. `oc create service externalname <servicename> --external-name <db>`
- 4. Which command creates a deployment config using a shared image stream named shared/myis?**
  - A. `oc new-app --as-deployment-config -i shared/myis`
  - B. `oc new-app -i shared/myis --as-deployment-config`
  - C. `oc new-app --name myis --image shared/myis`
  - D. `oc create deploymentconfig myis --image shared/myis`
- 5. In a workflow that shares a private image stream, which action is performed first to enable cross-project access to the image?**
  - A. Create secret regtoken in the shared project
  - B. Grant system:image-puller to the target service accounts in the target project
  - C. Switch to the target project and then grant pull access
  - D. Import the image into the target project's image stream



6. Which tool provides a Kubernetes-native approach to building and running CI/CD pipelines in OpenShift?
- A. OpenShift Pipelines (Tekton).
  - B. Jenkins.
  - C. GitLab CI.
  - D. Argo CD.
7. Which command would patch a deployment config to set the service account name to myserviceaccount?
- A. `oc patch dc/demo-app --patch '{"spec":{"template":{"spec":{"serviceAccountName": "myserviceaccount"}}}}'`
  - B. `oc patch dc/demo-app --patch '{"spec":{"template":{"serviceAccountName": "other"}}}'`
  - C. `oc patch dc/demo-app --patch '{"spec":{"template":{"serviceAccount": "myserviceaccount"}}}'`
  - D. `oc patch dc/demo-app --patch '{"template":{"spec":{"serviceAccountName": "myserviceaccount"}}}'`
8. How do you provide persistent storage to a pod in OpenShift?
- A. PersistentVolume
  - B. StorageClass
  - C. ConfigMap
  - D. PersistentVolumeClaim
9. Which OpenShift object represents a non-human identity used by pods to access resources?
- A. ServiceAccount
  - B. SecurityContext
  - C. PodSecurityPolicy
  - D. Namespace
10. How do you revert a failed deployment to the previous revision?
- A. `oc rollback deployment/<name>`
  - B. `oc revert deployment/<name> --to-revision=previous`
  - C. `oc patch deployment/<name> --to-revision=previous`
  - D. `oc rollout undo deployment/<name>`



## **Answers**

SAMPLE

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

SAMPLE

## **Explanations**

SAMPLE

### 1. Which command requests a new build for an application?

- A. oc start-build <application-name>
- B. oc start-build application-name
- C. oc new-build <application-name>
- D. oc trigger-build <application-name>

To request a new build for an application, you trigger the BuildConfig with the start-build command and specify the BuildConfig (application) name. BuildConfigs define how to build the app; starting a build tells OpenShift to run that process again using the current source and configuration. The other options serve different purposes: creating a new BuildConfig with new-build, or using a non-existent trigger-build command. So the correct approach is to run the build trigger command followed by the application's BuildConfig name.

### 2. Which command removes a local image?

- A. sudo podman tag mysql-custom devops/mysql:snapshot
- B. sudo podman rmi devops/mysql:snapshot
- C. oc rmi devops/mysql:snapshot
- D. docker rmi devops/mysql:snapshot

Removing a local image is done with the rmi command, which stands for remove image. The tag command would just create or update a tag for an existing image, not delete it. OpenShift's oc rmi isn't the standard way to remove a local image, and docker rmi would apply only if you're using Docker instead of Podman. Using podman rmi devops/mysql:snapshot removes the specified image from the local image store (assuming no containers are using it, or you use a force option). So the command shown is the correct way to remove a local image.

### 3. Which command creates a service of type ExternalName that points to an external host name?

- A. oc create service externalname <servicename> --external-name <mysql\_db>
- B. oc create service externalName <servicename> --external-name <mysql\_db>
- C. oc create service clusterip <servicename> --external-name <mysql\_db>
- D. oc create service externalname <servicename> --external-name <db>

ExternalName services route to an external DNS name instead of a cluster IP. You create one by specifying the service type as externalname and providing the external host with the --external-name flag. The command that fits this pattern is oc create service externalname <servicename> --external-name <mysql\_db>. It uses the correct resource type and the correct flag to point to the external host. The other options fail because they either use a different service type (like clusterip), or use incorrect capitalization for the resource type or the flag, or supply a less appropriate external name.

4. Which command creates a deployment config using a shared image stream named shared/myis?

- A. oc new-app --as-deployment-config -i shared/myis
- B. oc new-app -i shared/myis --as-deployment-config
- C. oc new-app --name myis --image shared/myis
- D. oc create deploymentconfig myis --image shared/myis

OpenShift uses the new-app flow to generate resources from an image source, and the flag that forces the result to be a DeploymentConfig ensures that the created object runs as a DeploymentConfig rather than a Deployment. By pointing to a shared image stream with the image spec shared/myis and explicitly requesting a DeploymentConfig, the command builds a DeploymentConfig whose pod template uses the image from that shared image stream. This combination is the direct, unambiguous way to create a DeploymentConfig from a shared image stream in one step.

5. In a workflow that shares a private image stream, which action is performed first to enable cross-project access to the image?

- A. Create secret regtoken in the shared project
- B. Grant system:image-puller to the target service accounts in the target project
- C. Switch to the target project and then grant pull access
- D. Import the image into the target project's image stream

Cross-project image sharing in OpenShift relies on RBAC to grant pull rights. To let workloads in the consuming project pull a private image from the image stream in the source project, you grant the system:image-puller role to the consuming project's service accounts for the source image stream's namespace. This permission is what enables the actual pull of the image across projects, and it must be in place before any pod can start using the image. Creating a secret for a regtoken isn't needed for internal OpenShift image streams, switching projects isn't required to grant the access, and importing the image into the target project would duplicate the image rather than enable ongoing cross-project access. So, granting pull access to the consuming project's service accounts is the correct first step.

6. Which tool provides a Kubernetes-native approach to building and running CI/CD pipelines in OpenShift?

- A. OpenShift Pipelines (Tekton).
- B. Jenkins.
- C. GitLab CI.
- D. Argo CD.

OpenShift Pipelines, built on Tekton, provides a Kubernetes-native way to define and run CI/CD pipelines. It exposes pipeline components as Kubernetes resources (Tasks, Pipelines, PipelineRuns) defined in YAML and runs inside the cluster, leveraging Kubernetes RBAC, service accounts, and the cluster's scheduling and isolation features. This native integration means pipelines can be triggered by git events, run as pods, and be portable across Kubernetes environments, all without resorting to external agents or separate CI servers. Other tools can operate in OpenShift, but they aren't built around Kubernetes-native pipeline primitives. Jenkins runs as a traditional application with its own agents; GitLab CI uses runners and external orchestration, not Kubernetes-native CRDs; Argo CD focuses on GitOps for application delivery rather than building and running CI pipelines.

7. Which command would patch a deployment config to set the service account name to myserviceaccount?

- A. oc patch dc/demo-app --patch '{"spec":{"template":{"spec":{"serviceAccountName": "myserviceaccount"}}}}'
- B. oc patch dc/demo-app --patch '{"spec":{"template":{"serviceAccountName": "other"}}}'
- C. oc patch dc/demo-app --patch '{"spec":{"template":{"serviceAccount": "myserviceaccount"}}}'
- D. oc patch dc/demo-app --patch '{"template":{"spec":{"serviceAccountName": "myserviceaccount"}}}'

The key idea is updating the service account for pods created by the DeploymentConfig, which lives in the pod template inside the DeploymentConfig. To change it, you patch the field at `spec.template.spec.serviceAccountName` with the desired value. The correct patch does exactly that: it targets `spec.template.spec.serviceAccountName` and sets it to "myserviceaccount". This aligns with where OpenShift stores the pod's service account in a DeploymentConfig. The other options fail because they either place the field in the wrong level (for example, setting `serviceAccountName` at the template level rather than inside `template.spec`), use a different field name (`serviceAccountName` vs `serviceAccount`), or have invalid JSON structure (such as an extra brace or incorrect nesting).

8. How do you provide persistent storage to a pod in OpenShift?

- A. PersistentVolume
- B. StorageClass
- C. ConfigMap
- D. PersistentVolumeClaim

Providing persistent storage to a pod centers on requesting storage through a PersistentVolumeClaim. A pod never uses raw storage directly; it mounts a volume that is backed by a PersistentVolume, and the binding between them is done via a PVC. The PVC expresses how much storage is needed and with what access mode, and the cluster binds it to a suitable PV. If no matching PV exists, a StorageClass can dynamically provision one, creating a PV behind the scenes. A ConfigMap, by contrast, stores configuration data and is not used for persistent storage. So, the mechanism a pod uses to access persistent storage is the PersistentVolumeClaim, which the pod references to mount the storage once the claim is bound to a PV.

9. Which OpenShift object represents a non-human identity used by pods to access resources?

- A. ServiceAccount
- B. SecurityContext
- C. PodSecurityPolicy
- D. Namespace

*Pods access resources in the cluster using a non-human identity called a ServiceAccount. A ServiceAccount is an object that represents an application's identity for pods, enabling them to authenticate to the API server and perform actions based on RBAC permissions. When a pod runs, it receives credentials for its ServiceAccount (a token) that is mounted into the pod, and the pod's processes use this token to access other resources. You can grant specific permissions by binding the ServiceAccount to roles with RoleBinding or ClusterRoleBinding, and every project/namespace has a default ServiceAccount if none is specified. SecurityContext controls container-level privileges, not identity. PodSecurityPolicy (now deprecated in favor of alternatives) governs security constraints on pods. Namespace defines a scope for resources, not an identity.*

10. How do you revert a failed deployment to the previous revision?

- A. `oc rollback deployment/<name>`
- B. `oc revert deployment/<name> --to-revision=previous`
- C. `oc patch deployment/<name> --to-revision=previous`
- D. `oc rollout undo deployment/<name>`

*Reverting a deployment to a prior revision is done with the rollout undo operation. OpenShift tracks rollout revisions for a Deployment, so you can undo the current rollout to the previous state. Running the command to revert targets the Deployment and, by default, restores the last revision (you can use a specific revision with --to-revision if needed). This approach is the proper way to roll back a failed deployment; other commands like patching or non-existent rollback options don't perform a proper rollout rollback.*



## 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](mailto:hello@examzify.com).

Or visit your dedicated course page for more study tools and resources:

<https://redhatopenshiftdev2do288.examzify.com>

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

SAMPLE