



Red Hat Training and Certification

Student Workbook (ROLE)

OCP 4.10 DO280

Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster

Edition 2



Red Hat
Learning Community

Join a community dedicated to learning open source

The Red Hat® Learning Community is a collaborative platform for users to accelerate open source skill adoption while working with Red Hat products and experts.



Network with tens of thousands of community members




Engage in thousands of active conversations and posts



Join and interact with hundreds of certified training instructors



Unlock badges as you participate and accomplish new goals



This knowledge-sharing platform creates a space where learners can connect, ask questions, and collaborate with other open source practitioners.

Access free Red Hat training videos

Discover the latest Red Hat Training and Certification news

Connect with your instructor – and your classmates – before, after, and during your training course.

Join peers as you explore Red Hat products

Join the conversation learn.redhat.com



Copyright © 2020 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Red Hat logo, and Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

The background of the slide is a long-exposure photograph of a city street at night. The image shows light trails from cars and streetlights, creating a sense of motion. In the foreground, there is a faint, light-colored grid pattern that resembles a circuit board or a digital grid. The overall color palette is dominated by warm yellows and oranges from the streetlights, with cooler blues and purples from the distant city lights and the sky.

Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster

OCP 4.10 DO280
Red Hat OpenShift Administration II: Operating a Production
Kubernetes Cluster
Edition 2 20220811
Publication date 20220811

Authors: Zach Gutterman, Dan Kolepp, Eduardo Ramirez Ronco,
Jordi Sola Alaball, Richard Allred, Michael Jarrett, Harpal Singh,
Federico Fapitalle, Maria Fernanda Ordonez Casado,
Andres Hernandez, Ivan Chavero
Editors: Seth Kenlon, Dave Sacco, Connie Petlitzer, Nicole Muller, Sam
Ffrench

Copyright © 2021 Red Hat, Inc.

The contents of this course and all its modules and related materials, including handouts to audience members, are
Copyright © 2021 Red Hat, Inc.

No part of this publication may be stored in a retrieval system, transmitted or reproduced in any way, including, but
not limited to, photocopy, photograph, magnetic, electronic or other record, without the prior written permission of
Red Hat, Inc.

This instructional program, including all material provided herein, is supplied without any guarantees from Red Hat,
Inc. Red Hat, Inc. assumes no liability for damages or legal action arising from the use or misuse of contents or details
contained herein.

If you believe Red Hat training materials are being used, copied, or otherwise improperly distributed, please send
email to training@redhat.com or phone toll-free (USA) +1 (866) 626-2994 or +1 (919) 754-3700.

Red Hat, Red Hat Enterprise Linux, the Red Hat logo, JBoss, OpenShift, Fedora, Hibernate, Ansible, CloudForms,
RHCA, RHCE, RHCSA, Ceph, and Gluster are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries
in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

Java® is a registered trademark of Oracle American, Inc. and/or its affiliates.

XFS® is a registered trademark of Hewlett Packard Enterprise Development LP or its subsidiaries in the United
States and/or other countries.

MySQL® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js® is a trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open
source or commercial project.

The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks
of OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's
permission. Red Hat, Inc. is not affiliated with, endorsed by, or sponsored by the OpenStack Foundation or the
OpenStack community.

All other trademarks are the property of their respective owners.

Contributors: Forrest Taylor, Manuel Aude Morales, James Mighion, Michael Phillips, and Fiona
Allen

Document Conventions	ix
.....	ix
Introduction	xi
DO280 Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster	xi
Orientation to the Classroom Environment	xii
Performing Lab Exercises	xx
1. Describing the Red Hat OpenShift Container Platform	1
Describing OpenShift Container Platform Features	2
Quiz: Describing OpenShift Container Platform Features	7
Describing the Architecture of OpenShift	11
Quiz: Describing the Architecture of OpenShift	14
Describing Cluster Operators	18
Quiz: Describing Cluster Operators	21
Summary	23
2. Verifying the Health of a Cluster	25
Describing Installation Methods	26
Quiz: Describing Installation Methods	28
Troubleshooting OpenShift Clusters and Applications	30
Guided Exercise: Troubleshooting OpenShift Clusters and Applications	38
Introducing OpenShift Dynamic Storage	45
Guided Exercise: Introducing OpenShift Dynamic Storage	49
Summary	54
3. Configuring Authentication and Authorization	55
Configuring Identity Providers	56
Guided Exercise: Configuring Identity Providers	63
Defining and Applying Permissions using RBAC	73
Guided Exercise: Defining and Applying Permissions using RBAC	77
Lab: Configuring Authentication and Authorization	83
Summary	92
4. Configuring Application Security	93
Managing Sensitive Information with Secrets	94
Guided Exercise: Managing Sensitive Information with Secrets	99
Controlling Application Permissions with Security Context Constraints	105
Guided Exercise: Controlling Application Permissions with Security Context Constraints ...	108
Lab: Configuring Application Security	112
Summary	120
5. Configuring OpenShift Networking for Applications	121
Troubleshooting OpenShift Software-defined Networking	122
Guided Exercise: Troubleshooting OpenShift Software-defined Networking	128
Exposing Applications for External Access	137
Guided Exercise: Exposing Applications for External Access	143
Configuring Network Policies	154
Guided Exercise: Configuring Network Policies	158
Lab: Configuring OpenShift Networking for Applications	167
Summary	180
6. Controlling Pod Scheduling	181
Controlling Pod Scheduling Behavior	182
Guided Exercise: Controlling Pod Scheduling Behavior	189
Limiting Resource Usage by an Application	195
Guided Exercise: Limiting Resource Usage by an Application	206

Scaling an Application	216
Guided Exercise: Scaling an Application	219
Lab: Controlling Pod Scheduling	225
Summary	233
7. Describing Cluster Updates	235
Describing the Cluster Update Process	236
Quiz: Describing the Cluster Update Process	247
Summary	251
8. Managing a Cluster with the Web Console	253
Performing Cluster Administration	254
Guided Exercise: Performing Cluster Administration	257
Managing Workloads and Operators	264
Guided Exercise: Managing Workloads and Operators	268
Examining Cluster Metrics	272
Guided Exercise: Examining Cluster Metrics	276
Lab: Managing a Cluster with the Web Console	281
Summary	290
9. Comprehensive Review	291
Comprehensive Review	292
Lab: Troubleshoot an OpenShift Cluster and Applications	294
Lab: Configure a Project Template with Resource and Network Restrictions	310

Document Conventions

This section describes various conventions and practices used throughout all Red Hat Training courses.

Admonitions

Red Hat Training courses use the following admonitions:



References

These describe where to find external documentation relevant to a subject.



Note

These are tips, shortcuts, or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on something that makes your life easier.



Important

These provide details of information that is easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring these admonitions will not cause data loss, but may cause irritation and frustration.



Warning

These should not be ignored. Ignoring these admonitions will most likely cause data loss.

Inclusive Language

Red Hat Training is currently reviewing its use of language in various areas to help remove any potentially offensive terms. This is an ongoing process and requires alignment with the products and services covered in Red Hat Training courses. Red Hat appreciates your patience during this process.

Introduction

DO280 Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster

Red Hat® OpenShift® Container Platform is a containerized application platform that allows enterprises to manage and scale applications utilizing container deployments. OpenShift provides predefined application environments, based upon Kubernetes, to support DevOps principles such as reduced time to market, infrastructure-as-code, continuous integration (CI), and continuous delivery (CD).

Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster (DO280) teaches students how to configure, troubleshoot, and manage the Red Hat® OpenShift® Container Platform. This hands-on, lab-based course shows students how to review the installation of a cluster, configure it, and manage it day-to-day.

Course Objectives

Install, configure, manage, and troubleshoot OpenShift clusters. This course, together with Red Hat OpenShift I: Containers & Kubernetes (DO180), prepares the student to take the Red Hat Certified Specialist in OpenShift Administration exam (EX280).

Audience

System and Software Architects, System Administrators, Cluster Operators, and Site Reliability Engineers

Prerequisites

Either complete the Red Hat OpenShift I: Containers & Kubernetes (DO180) course, or have equivalent knowledge. Either attain the Red Hat Certified System Administrator certification (RHCSA), or have equivalent knowledge.

Orientation to the Classroom Environment

The Workstation Machine

In this course, the main computer system used for hands-on learning activities (exercises) is `workstation`.

The `workstation` machine has a standard user account, `student` with the password `student`. No exercise in this course requires that you log in as `root`, but if you must, the `root` password on the `workstation` machine is `redhat`.

It is from the `workstation` machine that you type `oc` commands to manage the OpenShift cluster that comes preinstalled as part of your classroom environment.

It is also from the `workstation` machine that you run shell scripts and Ansible Playbooks required to complete exercises for this course.

If exercises require that you open a web browser to access any application or website, then you are required to use the graphical console of the `workstation` machine and use the Firefox web browser from there.



Important

The first time you start your classroom environment, OpenShift clusters take a little longer to become fully available. The `lab` command at the beginning of each exercise can fail if the clusters are not fully available.

If you try to access your cluster using either the `oc` command or the web console without first running a `lab` command, then you might find that your cluster is not yet available. If that happens, then wait a few minutes and try again.

You can run the `wait.sh` script in the `utility` machine that waits until your OpenShift cluster is ready to accept authentication and API requests from remote clients.


```
[student@workstation ~]$ ssh lab@utility
[lab@utility ~]$ ./wait.sh
```

The Classroom Environment

Every student gets a complete remote classroom environment. As part of that environment, every student gets a dedicated OpenShift cluster to perform administration tasks.

The classroom environment runs entirely as virtual machines in a large Red Hat OpenStack Platform cluster that is shared among many students.

Red Hat Training maintains many OpenStack clusters, in different data centers across the globe, to provide lower latency to students from many countries.



All machines on the Student, Classroom, and Cluster Networks run Red Hat Enterprise Linux 8 (RHEL 8), except those machines that are nodes of the OpenShift cluster. These run RHEL CoreOS.

The systems called `bastion`, `utility`, and `classroom` must always be running. They provide infrastructure services required by the classroom environment and its OpenShift cluster. You are not expected to interact with any of these systems directly.

Usually, the `lab` commands from exercises access these machines when there is a requirement to set up your environment for the exercise, and will require no further action from you.

All systems in the *Student Network* are in the `lab.example.com` DNS domain, and all systems in the *Classroom Network* are in the `example.com` DNS domain.

The systems called `masterXX` are nodes of the OpenShift 4 cluster that is part of your classroom environment.

All systems in the *Cluster Network* are in the `ocp4.example.com` DNS domain.

Classroom Machines

Machine name	IP addresses	ROLE
workstation.lab.example.com	172.25.250.9	Graphical workstation used for system administration.

Machine name	IP addresses	ROLE
classroom.example.com	172.25.254.254	Router linking the Classroom Network to the Internet.
bastion.lab.example.com	172.25.250.254	Router linking the Student Network to the Classroom Network.
utility.lab.example.com	172.25.250.253	Router linking the Student Network to the Cluster Network and also storage server.
master01.ocp4.example.com	192.168.50.10	Control plane and compute node.
master02.ocp4.example.com	192.168.50.11	Control plane and compute node.
master03.ocp4.example.com	192.168.50.12	Control plane and compute node.

Dependencies on Internet Services

Red Hat OpenShift Container Platform 4 requires access to two container registries to download container images for operators, S2I builders, and other cluster services. These registries are:

- registry.redhat.io
- quay.io

If either registry is unavailable when starting the classroom environment, then the OpenShift cluster might not start or could enter a degraded state. If either of these container registries experiences an outage while the classroom environment is up and running, then it might not be possible to complete exercises until the outage is resolved.

The Dedicated OpenShift Cluster

The Red Hat OpenShift Container Platform 4 cluster inside the classroom environment is preinstalled using the pre-existing infrastructure installation method; all nodes are treated as bare metal servers, even though they are actually virtual machines in an OpenStack cluster.

OpenShift cloud-provider integration capabilities are not enabled and a few features that depend on that integration, such as machine sets and autoscaling of cluster nodes, are not available.

Restoring Access to your OpenShift Cluster

If you suspect that you cannot log in to your OpenShift cluster as the `admin` user anymore because you incorrectly changed your cluster authentication settings, then run the `lab finish` command from your current exercise and restart the exercise by running its `lab start` command.

For labs that expect the `admin` and `developer` users, the `lab` command resets cluster authentication settings and restores passwords so that the `admin` user has a password of `redhat` and the `developer` user has a password of `developer`.

If running a `lab` command is not sufficient, then you can follow the instructions in the next section to use the `utility` machine to access your OpenShift cluster.

Troubleshooting Access to your OpenShift Cluster

The `utility` machine was used to run the OpenShift installer inside your classroom environment, and it is a useful resource to troubleshoot cluster issues. You can view the installer manifests in the `/home/lab/ocp4` folder of the `utility` machine.

Logging in to the `utility` server is not required to perform exercises. If it looks like your OpenShift cluster is taking too long to start, or is in a degraded state, then you can log in on the `utility` machine as the `lab` user to troubleshoot your classroom environment.

The `student` user on the `workstation` machine is already configured with SSH keys that enable logging in to the `utility` machine without a password.

```
[student@workstation ~]$ ssh lab@utility
```

In the `utility` machine, the `lab` user is preconfigured with a `.kube/config` file that grants access as `system:admin` without first requiring `oc login`.

This allows you to run troubleshooting commands, such as `oc get node`, if they fail from the `workstation` machine.

You should not require SSH access to your OpenShift cluster nodes for regular administration tasks because OpenShift 4 provides the `oc debug` command; if necessary, the `lab` user on the `utility` server is preconfigured with SSH keys to access all cluster nodes. For example:

```
[lab@utility ~]$ ssh -i ~/.ssh/lab_rsa core@master01.ocp4.example.com
```

In the preceding example, replace `master01` with the name of the desired cluster node.

Approving Node Certificates on your OpenShift Cluster

Red Hat OpenShift Container Platform clusters are designed to run continuously, 24x7, until they are decommissioned. Unlike a production cluster, the classroom environment contains a cluster that was stopped after installation, and will be stopped and restarted a few times before you finish this course. This presents a scenario that requires special handling that would not be required by a production cluster.

The control plane and compute nodes in an OpenShift cluster communicate frequently with each other. All communication between cluster nodes is protected by mutual authentication based on per-node TLS certificates.

The OpenShift installer handles creating and approving TLS certificate signing requests (CSRs) for the full-stack automation installation method. The system administrator is expected to manually approve these CSRs for the pre-existing infrastructure installation method.

All per-node TLS certificates have a short expiration life of 24 hours (the first time) and 30 days (after renewal). When they are about to expire, the affected cluster nodes create new CSRs and the control plane automatically approves them. If the control plane is offline when the TLS certificate of a node expires, then a cluster administrator is required to approve the pending CSR.

The `utility` machine includes a system service that approves CSRs from the cluster when you start your classroom, to ensure that your cluster is ready when you begin the exercises. If you create or start your classroom and begin an exercise too quickly, then you might find that your cluster is not yet ready. If so, wait a few minutes while the `utility` machine handles CSRs, and then try again.

Sometimes, the `utility` machine fails to approve all required CSRs, for example, because the cluster took too long to generate all required CSRs requests and the system service did not wait long enough. It's also possible that some OpenShift cluster nodes did not wait long enough for their CSRs to be approved, issuing new CSRs that superseded previous ones.

If these issues arise, then you will notice that your cluster is taking too long to come up, and your `oc login` or `lab` commands keep failing. To resolve the problem, you can log in to the `utility` machine, as explained previously, and run the `sign.sh` script to approve any additional and pending CSRs.

```
[lab@utility ~]$ ./sign.sh
```

The `sign.sh` script loops a few times just in case your cluster nodes issue new CSRs that supersede the ones it approved.

After either you approve, or the system service in the `utility` machine approves all CSRs, then OpenShift must restart a few cluster operators; it takes a few moments before your OpenShift cluster is ready to answer requests from clients. To help you handle this scenario, the `utility` machine provides the `wait.sh` script that waits until your OpenShift cluster is ready to accept authentication and API requests from remote clients.

```
[lab@utility ~]$ ./wait.sh
```

Although unlikely, if neither the service on the `utility` machine nor running the `sign.sh` and `wait.sh` scripts make your OpenShift cluster available to begin exercises, then open a customer support ticket.



Note

You can run troubleshooting commands from the `utility` machine at any time, even if you have control plane nodes that are not ready. Some useful commands include:

```
oc get nodes
```

To verify if all of your cluster nodes are ready.

```
oc get csr
```

To verify if your cluster still has any pending, unapproved CSRs.

```
oc get clusteroperators
```

To verify if any of your cluster operators are unavailable, in a degraded state, or progressing through configuration and rolling out pods.

If these fail, you can try destroying and recreating your classroom as a last resort before creating a customer support ticket.

Controlling Your Systems

Students are assigned remote computers in a Red Hat Online Learning classroom. They are accessed through a web application hosted at <https://ol.redhat.com/>. Students should log in to this site using their Red Hat Customer Portal user credentials.

Controlling the Virtual Machines

The virtual machines in your classroom environment are controlled through a web page. The state of each virtual machine in the classroom is displayed on the page under the Lab Environment tab.

Machine States

Virtual Machine State	Description
Active	The virtual machine is running and available (or, when booting, soon will be).
Stopped	The virtual machine is completely shut down.
Building	The initial creation of the virtual machine is being performed.

Depending on the state of a machine, a selection of the following actions is available.

Classroom/Machine Actions

Button or Action	Description
CREATE	Create the ROL classroom. Creates all of the virtual machines needed for the classroom and starts them. Can take several minutes to complete.
DELETE	Delete the ROL classroom. Destroys all virtual machines in the classroom. Caution: Any work generated on the disks is lost.
START	Start all virtual machines in the classroom.
STOP	Stop all virtual machines in the classroom.
OPEN CONSOLE	Open a new tab in the browser and connect to the console of the virtual machine. Students can log in directly to the virtual machine and run commands. In most cases, students should log in to the workstation virtual machine and use ssh to connect to the other virtual machines.
ACTION > Start	Start (power on) the virtual machine.
ACTION > Shutdown	Gracefully shut down the virtual machine, preserving the contents of its disk.
ACTION > Power Off	Forcefully shut down the virtual machine, preserving the contents of its disk. This is equivalent to removing the power from a physical machine.
ACTION > Reset	Forcefully shut down the virtual machine and reset the disk to its initial state. Caution: Any work generated on the disk is lost.

At the start of an exercise, if instructed to reset a single virtual machine node, click **ACTION > Reset** for only the specific virtual machine.

At the start of an exercise, if instructed to reset all virtual machines, click **ACTION** > **Reset**.

If you want to return the classroom environment to its original state at the start of the course, you can click **DELETE** to remove the entire classroom environment. After the lab has been deleted, you can click **CREATE** to provision a new set of classroom systems.



Warning

The **DELETE** operation cannot be undone. Any work you have completed in the classroom environment up to that point will be lost.

The Autostop Timer

The Red Hat Online Learning enrollment entitles students to a certain amount of computer time. To help conserve allotted computer time, the ROL classroom has an associated countdown timer, which shuts down the classroom environment when the timer expires.

To adjust the timer, click + to add one hour to the timer. Note that there is a maximum time of twelve hours.

Controlling Your Systems

You are assigned remote computers in a Red Hat Online Learning classroom. They are accessed through a web application hosted at rol.redhat.com [<http://rol.redhat.com>]. You should log in to this site using your Red Hat Customer Portal user credentials.

Controlling the Virtual Machines

The virtual machines in your classroom environment are controlled through a web page. The state of each virtual machine in the classroom is displayed on the page under the **Online Lab** tab.

Machine States

Virtual Machine State	Description
STARTING	The virtual machine is in the process of booting.
STARTED	The virtual machine is running and available (or, when booting, soon will be).
STOPPING	The virtual machine is in the process of shutting down.
STOPPED	The virtual machine is completely shut down. Upon starting, the virtual machine boots into the same state as when it was shut down (the disk will have been preserved).
PUBLISHING	The initial creation of the virtual machine is being performed.
WAITING_TO_START	The virtual machine is waiting for other virtual machines to start.

Depending on the state of a machine, a selection of the following actions is available.

Classroom/Machine Actions

Button or Action	Description
PROVISION LAB	Create the ROL classroom. Creates all of the virtual machines needed for the classroom and starts them. Can take several minutes to complete.
DELETE LAB	Delete the ROL classroom. Destroys all virtual machines in the classroom. Caution: Any work generated on the disks is lost.
START LAB	Start all virtual machines in the classroom.
SHUTDOWN LAB	Stop all virtual machines in the classroom.
OPEN CONSOLE	Open a new tab in the browser and connect to the console of the virtual machine. You can log in directly to the virtual machine and run commands. In most cases, you should log in to the workstation virtual machine and use ssh to connect to the other virtual machines.
ACTION > Start	Start (power on) the virtual machine.
ACTION > Shutdown	Gracefully shut down the virtual machine, preserving the contents of its disk.
ACTION > Power Off	Forcefully shut down the virtual machine, preserving the contents of its disk. This is equivalent to removing the power from a physical machine.
ACTION > Reset	Forcefully shut down the virtual machine and reset the disk to its initial state. Caution: Any work generated on the disk is lost.

At the start of an exercise, if instructed to reset a single virtual machine node, click **ACTION > Reset** for only the specific virtual machine.

At the start of an exercise, if instructed to reset all virtual machines, click **ACTION > Reset**

If you want to return the classroom environment to its original state at the start of the course, you can click **DELETE LAB** to remove the entire classroom environment. After the lab has been deleted, you can click **PROVISION LAB** to provision a new set of classroom systems.

**Warning**

The **DELETE LAB** operation cannot be undone. Any work you have completed in the classroom environment up to that point will be lost.

The Autostop Timer

The Red Hat Online Learning enrollment entitles you to a certain amount of computer time. To help conserve allotted computer time, the ROL classroom has an associated countdown timer, which shuts down the classroom environment when the timer expires.

To adjust the timer, click **MODIFY** to display the **New Autostop Time** dialog box. Set the number of hours until the classroom should automatically stop. Note that there is a maximum time of ten hours. Click **ADJUST TIME** to apply this change to the timer settings.

Performing Lab Exercises

Performing Lab Exercises

Run the `lab` command from the `workstation` machine to prepare your environment before each hands-on exercise, and again to clean up after an exercise. Each hands-on exercise has a unique name within a course. The exercise is prepended with `lab-` as its file name in `/usr/local/lib`. For example, the `instances-cli` exercise has the file name `/usr/local/lib/lab-instances-cli`. To list the available exercises, use tab completion in the `lab` command. Note that the word `Tab` in the following command refers to pressing the `Tab` key on your keyboard:

```
[student@workstation ~]$ lab Tab Tab
administer-users  deploy-overcloud-lab  prep-deploy-ips      stacks-autoscale
analyze-metrics  instances-cli          prep-deploy-router    stacks-deploy
assign-roles     manage-interfaces     public-instance-deploy verify-overcloud
```

There are two types of exercises. The first type, a guided exercise, is a practice exercise that follows a course narrative. If a narrative is followed by a quiz, this usually indicates that the topic did not have an achievable practice exercise. The second type, an end-of-chapter lab, is a gradable exercise to help verify your learning. When a course includes a comprehensive review, the review exercises are structured as gradable labs. The syntax for running an exercise script is:

```
[student@workstation ~]$ lab exercise action
```

The *action* is a choice of `start`, `grade`, or `finish`. All exercises support `start` and `finish`. Only end-of-chapter labs and comprehensive review labs support `grade`. Older courses might still use `setup` and `cleanup` instead of the current `start` and `finish` actions.

start

Formerly `setup`. The start logic of the script verifies the resources required to begin an exercise. This might include configuring settings, creating resources, checking prerequisite services, and verifying necessary outcomes from previous exercises.

grade

End-of-chapter labs help verify what you have learned, after practicing with earlier guided exercises. The grade action directs the `lab` command to display a list of grading criteria, with a `PASS` or `FAIL` status for each. To achieve a `PASS` status for all criteria, fix the failures and re-run the grade action.

finish

Formerly `cleanup`. The finish logic of the script deletes exercise resources that are no longer necessary.

Exercise scripts do not exist on the `workstation` machine until each is first run. When you run the `lab` command with a valid exercise and action, the script named `lab-exercise` is downloaded from the `classroom` server content share to `/usr/local/lib` on the `workstation` machine. The `lab` command creates two log files in `/var/tmp/labs`, plus the directory if it does not exist. One file, named `exercise`, captures standard output messages that normally display on your terminal. The other file, named `exercise.err`, captures error messages.

```
[student@workstation ~]$ ls -l /usr/local/lib
-rwxr-xr-x. 1 root root 4131 May 9 23:38 lab-instances-cli
-rwxr-xr-x. 1 root root 93461 May 9 23:38 labtool.cl110.shlib
-rwxr-xr-x. 1 root root 10372 May 9 23:38 labtool.shlib

[student@workstation ~]$ ls -l /var/tmp/labs
-rw-r--r--. 1 root root 113 May 9 23:38 instances-cli
-rw-r--r--. 1 root root 113 May 9 23:38 instances-cli.err
```

**Note**

Scripts download from the <http://content.example.com/courses/COURSE/RELEASE/grading-scripts> share, but only if the script does not yet exist on the workstation machine.

If you must download a script again, such as when a script on the share is modified, manually delete the current exercise script from `/usr/local/lib` on the workstation machine, then run the `lab` command for the exercise again. The newer exercise script then downloads from the `grading-scripts` share.

To delete all current exercise scripts on the workstation machine, use the `lab --refresh` command. A refresh deletes all scripts in `/usr/local/lib` but does not delete the log files.

```
[student@workstation ~]$ lab --refresh
[student@workstation ~]$ ls -l /usr/local/lib

[student@workstation ~]$ ls -l /var/tmp/labs
-rw-r--r--. 1 root root 113 May 9 23:38 instances-cli
-rw-r--r--. 1 root root 113 May 9 23:38 instances-cli.err
```

Exercise scripts send output to log files, even when the scripts are successful. Step header text is added between steps and additional date and time headers are added at the start of each script run. The exercise log normally contains messages that indicate the successful completion of command steps. Therefore, the exercise output log is useful for observing messages that are expected if no problems occur, but offers no additional help when failures occur.

Instead, the exercise error log is more useful for troubleshooting. Even when the scripts succeed, messages are still sent to the exercise error log. For example, a script that verifies that an object already exists before attempting to create it should cause an *object not found* message when the object does not exist yet. In this scenario, that message is expected and does not indicate a failure. Actual failure messages are typically more verbose, and experienced system administrators should recognize common log message entries.

Although exercise scripts are always run from the workstation machine, they perform tasks on other systems in the course environment. Many course environments, including OpenStack and OpenShift, use a command-line interface (CLI) invoked from the workstation machine to communicate with server systems using API calls. Because script actions typically distribute tasks to multiple systems, additional troubleshooting is necessary to determine where a failed task occurred. Log in to those other systems and use Linux diagnostic skills to read local system log files and determine the root cause of the lab script failure.

Chapter 1

Describing the Red Hat OpenShift Container Platform

Goal

Describe the architecture of OpenShift Container Platform.

Objectives

- Describe the typical usage of the product and its features.
- Describe the architecture of Red Hat OpenShift Container Platform.
- Describe what a cluster operator is, how it works, and name the major cluster operators.

Sections

- Describing OpenShift Container Platform Features (and Quiz)
- Describing the Architecture of OpenShift (and Quiz)
- Describing Cluster Operators (and Quiz)

Describing OpenShift Container Platform Features

Objectives

After completing this section, you should be able to describe the typical usage of the product and its features.

Introducing OpenShift Container Platform

Container orchestration is a fundamental enabler of digital transformation initiatives. However, as monolithic applications transition to containerized services, it can be tedious to manage these applications with legacy infrastructure. Red Hat OpenShift Container Platform (RHOCP) helps developers and IT organizations to better manage application lifecycles.

RHOCP is based on the Kubernetes open source project and extends the platform with features that bring a robust, flexible, and scalable container platform to customer data centers, enabling developers to run workloads in a high availability environment.

A container orchestrator, such as OpenShift Container Platform, manages a cluster of servers that runs multiple containerized applications. The Red Hat OpenShift product family includes a set of solutions to improve the delivery of business applications in a variety of environments.

Red Hat OpenShift Container Platform

Provides an enterprise-ready Kubernetes environment for building, deploying, and managing container-based applications in any public or private data center, including bare metal servers. RHOCP is compatible with multiple cloud and virtualization providers, isolating application developers and administrators from differences between these providers. You decide when to update to newer releases and which additional components to enable.

Red Hat OpenShift Dedicated

Provides a managed OpenShift environment in a public cloud, such as Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, or IBM Cloud. This product provides all the features that RHOCP provides, however, Red Hat manages the cluster for you. You retain some control of decisions such as when to update to a newer release of OpenShift or to install add-on services.

Red Hat OpenShift Online

Provides a hosted, public container orchestration platform that offers an application development, build, deployment, and hosting solution in a cloud environment. The solution is shared across multiple customers, and Red Hat manages the cluster lifecycle, which includes applying updates or integrating new features.

Red Hat OpenShift Kubernetes Engine

Provides a subset of the features present in OpenShift Container Platform, such as the Red Hat Enterprise Linux CoreOS lightweight transactional operating system, the CRI-O engine, the Kubernetes container orchestration platform, and the core cluster services (web console, Over-the-air updates, internal registry, and Operator Lifecycle Manager, among others).

Red Hat Code Ready Containers

Provides a minimal installation of OpenShift that you can run on a laptop for local development and experimentation.

Some cloud providers also have offerings based on RHOCP that add tight integration with other services from their platforms. Those offerings are supported by the cloud provider in partnership with Red Hat. One example is Microsoft Azure Red Hat OpenShift.

The following figure describes the services and features of OpenShift:




Figure 1.1: OpenShift services and features

The Red Hat OpenShift product family integrates many components:

- The Red Hat Enterprise Linux CoreOS container-optimized, immutable operating system.
- The CRI-O engine, a small footprint, Open Container Initiative (OCI)-compliant container runtime engine with a reduced attack surface.
- Kubernetes, an open source container orchestration platform.
- A self-service web console.
- A number of preinstalled application services, such as an internal container image registry and monitoring framework.
- Certified container images for multiple programming language runtimes, databases, and other software packages.

Introducing OpenShift Features

OpenShift offers many features to automate, scale, and maintain your applications. All of these features are enabled by Kubernetes and most of them require additional components that you need to add and configure on a build-your-own (BYO) Kubernetes setup.




Figure 1.2: Feature comparison between OpenShift Container Platform and OpenShift Kubernetes Engine

High Availability

Kubernetes has been designed with high availability in mind, for both internal components and user applications. A highly available etcd cluster stores the state of the OpenShift cluster and its applications. Resources stored in etcd, such as deployment configurations, provide automatic restarting of containers to ensure that your application is always running and that faulty containers are terminated. This applies not only to your applications, but also to containerized services that make up the cluster, such as the web console and the internal image registry.

Lightweight Operating System

RHOC runs on Red Hat Enterprise Linux CoreOS, Red Hat's lightweight operating system that focuses on agility, portability, and security.

Red Hat Enterprise Linux CoreOS (RHEL CoreOS) is an immutable operating system that is optimized for running containerized applications. The entire operating system is updated as a single image, instead of on a package-by-package basis, and both user applications and system components such as network services run as containers.

RHOC controls updates to RHEL CoreOS and its configurations, and so managing an OpenShift cluster includes managing the operating system on cluster nodes, freeing system administrators from these tasks and reducing the risk of human error.

Load Balancing

Clusters provide three types of load balancers: an external load balancer, which manages access to the OpenShift API; the HAProxy load balancer, for external access to applications; and the internal load balancer, which uses Netfilter rules for internal access to applications and services.

Route resources use HAProxy to manage external access to the cluster. Service resources use Netfilter rules to manage traffic from inside the cluster. The technology that external load balancers use is dependent on the cloud provider that runs your cluster.

Automating Scaling

OpenShift clusters can adapt to increased application traffic in real time by automatically starting new containers, and terminating containers when the load decreases. This feature ensures that your application's access time remains optimal regardless of the number of concurrent connections or activity.

OpenShift clusters can also add or remove more compute nodes from the cluster according to the aggregated load from many applications, keeping responsiveness and costs down on public and private clouds.

Logging and Monitoring

RHOC ships with an advanced monitoring solution, based on Prometheus, which gathers hundreds of metrics about your cluster. This solution interacts with an alerting system that allows you to obtain detailed information about your cluster activity and health.

RHOC ships with an advanced aggregated logging solution, based on Elasticsearch, which allows long-term retention of logs from cluster nodes and containers.

Services Discovery

RHOC runs an internal DNS service on the cluster, and configures all containers to use that internal DNS for name resolution. This means that applications can rely on friendly names to find other applications and services, without the overhead of an external services catalog.

Storage

Kubernetes adds an abstraction layer between the storage back end and the storage consumption. As such, applications can consume long-lived, short-lived, block, and file-based storage using unified storage definitions that are independent of the storage back end. This way your applications are not dependent on particular cloud provider storage APIs.

RHOC embeds a number of storage providers that allow for automatic provisioning of storage on popular cloud providers and virtualization platforms, and so cluster administrators do not need to manage the fine details of proprietary storage arrays.

Application Management

RHOC empowers developers to automate the development and deployment of their applications. Use the OpenShift Source-to-Image (S2I) feature to automatically build containers based on your source code and run them in OpenShift. The internal registry stores application container images, which can be reused. This decreases the time it takes to publish your applications.

The developer catalog, accessible from the web console, is a place for publishing and accessing application templates. It supports many runtime languages, such as Python, Ruby, Java, and Node.js, and also database and messaging servers. You can expand the catalog by installing new operators, which are prepackaged applications and services that embed operational intelligence for deploying, updating, and monitoring their applications.

Cluster Extensibility

RHOC relies on standard extension mechanisms from Kubernetes, such as extension APIs and custom resource definitions, to add features that are otherwise not provided by upstream Kubernetes. OpenShift packages these extensions as operators for ease of installation, update, and management.

OpenShift also includes the Operator Lifecycle Manager (OLM), which facilitates the discovery, installation, and update of applications and infrastructure components packaged as operators.

Red Hat, in collaboration with AWS, Google Cloud, and Microsoft, launched the OperatorHub, accessible at <https://operatorhub.io/>. The platform is a public repository and marketplace for operators compatible with OpenShift and other distributions of Kubernetes that include the OLM.

Red Hat Marketplace is a platform that allows access to certified software packaged as Kubernetes operators that can be deployed in an OpenShift cluster. The certified software includes automatic deployments and seamless upgrades for an integrated experience.



References

Further information is available in the Red Hat OpenShift Container Platform 4.10 product documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/

Red Hat OpenShift Kubernetes Engine

<https://www.openshift.com/products/kubernetes-engine>

► Quiz

Describing OpenShift Container Platform Features

Choose the correct answers to the following questions:

- 1. **Which of the following definitions best describes container orchestration platforms?**
 - a. They extend your application's operational knowledge and provide a way to package and distribute them.
 - b. They allow you to manage a cluster of servers that run containerized applications. They add features such as self-service, high availability, monitoring, and automation.
 - c. They allow you to provision Infrastructure-as-a-Service clusters on a variety of cloud providers, including AWS, GCP, and Microsoft Azure.
 - d. They enable developers to write, package, and publish their applications as operators to the operator catalog.

- 2. **Which three of the following key features enable high availability for your applications? (Choose three.)**
 - a. An OpenShift etcd cluster keeps the cluster state available for all nodes.
 - b. OpenShift HAProxy load balancers allow external access to applications.
 - c. OpenShift services load balance access to applications from inside the cluster.
 - d. OpenShift deployment configurations ensure application containers are restarted in scenarios such as loss of a node.

- 3. **Which two of the following statements about OpenShift are true? (Choose two.)**
 - a. Developers can create and start cloud applications directly from a source code repository.
 - b. OpenShift patches Kubernetes to add features that would not be available to other distributions of Kubernetes.
 - c. OpenShift Dedicated gives you access to an exclusive set of operators that Red Hat curates and maintains. This helps to ensure that the operators are secure and safe to run in your environment.
 - d. OpenShift cluster administrators can discover and install new operators from the operator catalog.

- 4. Which two of the following services do OpenShift components use for load balancing their traffic? (Choose two.)
- a. The OpenShift API, which is accessible over the external load balancer.
 - b. Services, which use Netfilter for load balancing.
 - c. Services, which use HAProxy for load balancing.
 - d. Routes, which use Netfilter for load balancing.
 - e. Routes, which use the HAProxy for load balancing.
- 5. Which two of the following statements about OpenShift high availability and scaling are true? (Choose two.)
- a. OpenShift does not provide high availability by default. You need to use third-party high availability products.
 - b. OpenShift uses metrics from Prometheus to dynamically scale application pods.
 - c. High availability and scaling are restricted to applications that expose a REST API.
 - d. OpenShift can scale applications up and down based on demand.

► Solution

Describing OpenShift Container Platform Features

Choose the correct answers to the following questions:

- 1. **Which of the following definitions best describes container orchestration platforms?**
 - a. They extend your application's operational knowledge and provide a way to package and distribute them.
 - b. They allow you to manage a cluster of servers that run containerized applications. They add features such as self-service, high availability, monitoring, and automation.
 - c. They allow you to provision Infrastructure-as-a-Service clusters on a variety of cloud providers, including AWS, GCP, and Microsoft Azure.
 - d. They enable developers to write, package, and publish their applications as operators to the operator catalog.

- 2. **Which three of the following key features enable high availability for your applications? (Choose three.)**
 - a. An OpenShift etcd cluster keeps the cluster state available for all nodes.
 - b. OpenShift HAProxy load balancers allow external access to applications.
 - c. OpenShift services load balance access to applications from inside the cluster.
 - d. OpenShift deployment configurations ensure application containers are restarted in scenarios such as loss of a node.

- 3. **Which two of the following statements about OpenShift are true? (Choose two.)**
 - a. Developers can create and start cloud applications directly from a source code repository.
 - b. OpenShift patches Kubernetes to add features that would not be available to other distributions of Kubernetes.
 - c. OpenShift Dedicated gives you access to an exclusive set of operators that Red Hat curates and maintains. This helps to ensure that the operators are secure and safe to run in your environment.
 - d. OpenShift cluster administrators can discover and install new operators from the operator catalog.

- 4. Which two of the following services do OpenShift components use for load balancing their traffic? (Choose two.)
- a. The OpenShift API, which is accessible over the external load balancer.
 - b. Services, which use Netfilter for load balancing.
 - c. Services, which use HAProxy for load balancing.
 - d. Routes, which use Netfilter for load balancing.
 - e. Routes, which use the HAProxy for load balancing.
- 5. Which two of the following statements about OpenShift high availability and scaling are true? (Choose two.)
- a. OpenShift does not provide high availability by default. You need to use third-party high availability products.
 - b. OpenShift uses metrics from Prometheus to dynamically scale application pods.
 - c. High availability and scaling are restricted to applications that expose a REST API.
 - d. OpenShift can scale applications up and down based on demand.

Describing the Architecture of OpenShift

Objectives

After completing this section, you should be able to describe the architecture of Red Hat OpenShift Container Platform.

Introducing the Declarative Architecture of Kubernetes

The architecture of OpenShift is based on the declarative nature of Kubernetes. Most system administrators are used to imperative architectures, where you perform actions that indirectly change the state of the system, such as starting and stopping containers on a given server. In a declarative architecture, you change the state of the system and the system updates itself to comply with the new state. For example, with Kubernetes, you define a pod resource that specifies that a certain container should run under specific conditions. Then Kubernetes finds a server (a node) that can run that container under these specific conditions.

Declarative architectures allow for self-optimizing and self-healing systems that are easier to manage than imperative architectures.

Kubernetes defines the state of its cluster, including the set of deployed applications, as a set of resources stored in the etcd database. Kubernetes also runs controllers that monitor these resources and compares them to the current state of the cluster. These controllers take any action necessary to reconcile the state of the cluster with the state of the resources, for example by finding a node with sufficient CPU capacity to start a new container from a new pod resource.

Kubernetes provides a REST API to manage these resources. All actions that an OpenShift user takes, either using the command-line interface or the web console, are performed by invoking this REST API.

Introducing the OpenShift Control Plane

A Kubernetes cluster consists of a set of nodes that run the kubelet system service and a container engine. OpenShift runs exclusively the CRI-O container engine. Some nodes are control plane nodes that run the REST API, the etcd database, and the platform controllers. OpenShift configures its control plane nodes so that they are not schedulable to run end-user application pods and are dedicated to running the control plane services. OpenShift schedules end-user application pods to be executed on the compute nodes.

The following graphic provides an overview of an OpenShift control plane node, illustrating the main processes that run in a regular node and in a control plane node, as either system services or containers.




Figure 1.3: Architecture of an OpenShift control plane node

Depending on the node settings, the `kubelet` agent starts different sets of static pods. Static pods are pods that do not require connection to the API server to start. The `kubelet` agent manages the pod's lifecycle. Static pods can provide either control plane services, such as the scheduler, or node services, such as software-defined networking (SDN). OpenShift provides operators that create pod resources for these static pods so that they are monitored like regular pods.

Describing OpenShift Extensions

A lot of functionality from Kubernetes depends on external components, such as ingress controllers, storage plug-ins, network plug-ins, and authentication plug-ins. Similar to Linux distributions, there are many ways to build a Kubernetes distribution by picking and choosing different components.

A lot of functionality from Kubernetes also depends on extension APIs, such as access control and network isolation.

OpenShift is a Kubernetes distribution that provides many of these components already integrated and configured, and managed by operators. OpenShift also provides preinstalled applications, such as a container image registry and a web console, managed by operators.

OpenShift also adds to Kubernetes a series of extension APIs and custom resources. For example, build configurations for the Source-to-Image process, and route resources to manage external access to the cluster.

Red Hat develops all extensions as open source projects and works with the larger Kubernetes community not only to make these official components of Kubernetes but also to evolve the Kubernetes platform to allow easier maintainability and customization.

With OpenShift 3 these extensions were sometimes patches (or forks) of upstream Kubernetes. With OpenShift 4 and operators, these extensions are standard Kubernetes extensions that could be added to any distribution of Kubernetes.

Introducing the OpenShift Default Storage Class

Unlike many container platforms that focus on cloud-native, stateless applications, OpenShift also supports stateful applications that do not follow the standard *Twelve-Factor App* methodology. OpenShift supports stateful applications by offering a comprehensive set of storage capabilities and supporting operators. OpenShift ships with integrated storage plug-ins and storage classes

that rely on the underlying cloud or virtualization platform to provide dynamically provisioned storage.

For example, if you install OpenShift on Amazon Web Services (AWS), your OpenShift cluster comes preconfigured with a default storage class that uses Amazon Elastic Block Store (EBS) service automatically to provision storage volumes on-demand. Users can deploy an application that requires persistent storage, such as a database, and OpenShift automatically creates an EBS volume to host the application data.

OpenShift cluster administrators can later define additional storage classes that use different EBS service tiers. For example, you could have one storage class for high-performance storage that sustains a high input-output operations per second (IOPS) rate, and another storage class for low-performance, low-cost storage. Cluster administrators can then allow only certain applications to use the high-performance storage class, and configure data archiving applications to use the low-performance storage class.

**References**

Further information is available in the Red Hat OpenShift Container Platform 4.10 product documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/

The Twelve-Factor App

<https://12factor.net/>

► Quiz

Describing the Architecture of OpenShift

Choose the correct answers to the following questions:

- 1. **OpenShift is based on which of the following container orchestration technologies?**
 - a. Docker Swarm
 - b. Rancher
 - c. Kubernetes
 - d. Mesosphere Marathon
 - e. CoreOS Fleet

- 2. **Which two of the following statements are true of OpenShift Container Platform? (Choose two.)**
 - a. OpenShift provides an OAuth server that authenticates calls to its REST API.
 - b. OpenShift requires the CRI-O container engine.
 - c. Kubernetes follows a declarative architecture, but OpenShift follows a more traditional imperative architecture.
 - d. OpenShift extension APIs run as system services.

- 3. **Which of the following servers runs Kubernetes API components?**
 - a. Compute nodes
 - b. Nodes
 - c. Control plane nodes

- 4. **Which of the following components does OpenShift add to upstream Kubernetes?**
 - a. The etcd database
 - b. A container engine
 - c. A registry server
 - d. A scheduler
 - e. The Kubelet

- 5. Which of the following sentences is true regarding support for storage with OpenShift?
- a. Users can only store persistent data in the etcd database.
 - b. Users can only deploy on OpenShift cloud-native applications that conform to the Twelve-Factor App methodology.
 - c. Administrators must configure storage plug-ins appropriate for their cloud providers.
 - d. Administrators must define persistent volumes before any user can deploy applications that require persistent storage.
 - e. Users can deploy applications that require persistent storage by relying on the default storage class.

► Solution

Describing the Architecture of OpenShift

Choose the correct answers to the following questions:

- 1. **OpenShift is based on which of the following container orchestration technologies?**
 - a. Docker Swarm
 - b. Rancher
 - c. Kubernetes
 - d. Mesosphere Marathon
 - e. CoreOS Fleet

- 2. **Which two of the following statements are true of OpenShift Container Platform? (Choose two.)**
 - a. OpenShift provides an OAuth server that authenticates calls to its REST API.
 - b. OpenShift requires the CRI-O container engine.
 - c. Kubernetes follows a declarative architecture, but OpenShift follows a more traditional imperative architecture.
 - d. OpenShift extension APIs run as system services.

- 3. **Which of the following servers runs Kubernetes API components?**
 - a. Compute nodes
 - b. Nodes
 - c. Control plane nodes

- 4. **Which of the following components does OpenShift add to upstream Kubernetes?**
 - a. The etcd database
 - b. A container engine
 - c. A registry server
 - d. A scheduler
 - e. The Kubelet

► 5. Which of the following sentences is true regarding support for storage with OpenShift?

- a. Users can only store persistent data in the etcd database.
- b. Users can only deploy on OpenShift cloud-native applications that conform to the Twelve-Factor App methodology.
- c. Administrators must configure storage plug-ins appropriate for their cloud providers.
- d. Administrators must define persistent volumes before any user can deploy applications that require persistent storage.
- e. Users can deploy applications that require persistent storage by relying on the default storage class.

Describing Cluster Operators

Objectives

After completing this section, you should be able to describe what a cluster operator is, how it works, and name the major cluster operators.

Introducing Kubernetes Operators

Kubernetes operators are applications that invoke the Kubernetes API to manage Kubernetes resources. As for any Kubernetes application, you deploy an operator by defining Kubernetes resources such as services and deployments that reference the operator's container image. Because operators, unlike common applications, require direct access to the Kubernetes resources, they usually require custom security settings.

Operators usually define custom resources (CR) that store their settings and configurations. An OpenShift administrator manages an operator by editing its custom resources. The syntax of a custom resource is defined by a custom resource definition (CRD).

Most operators manage another application; for example, an operator that manages a database server. In that case, the operator creates the resources that describe that other application using the information from its custom resource.

The purpose of an operator is usually to automate tasks that a human administrator (or human operator) would perform to deploy, update, and manage an application.

Introducing the Operator Framework

You can develop operators using your preferred programming language. Technically you do not need a special-purpose SDK to develop an operator. All you need is the ability to invoke REST APIs and consume secrets that contain access credentials to the Kubernetes APIs.

The Operator Framework is an open source toolkit for building, testing, and packaging operators. The Operator Framework makes these tasks easier than coding directly to low-level Kubernetes APIs by providing the following components:

Operator Software Development Kit (Operator SDK)

Provides a set of Golang libraries and source code examples that implement common patterns in operator applications. It also provides a container image and playbook examples that allow you to develop operators using Ansible.

Operator Lifecycle Manager (OLM)

Provides an application that manages the deployment, resource utilization, updates, and deletion of operators that have been deployed through an operator catalog. The OLM itself is an operator that comes preinstalled with OpenShift.

The Operator Framework also defines a set of recommended practices for implementing operators and CRDs and a standard way of packaging an operator manifest, as a container image, that allows an operator to be distributed using an operator catalog. The most common form of an operator catalog is an image registry server.

An operator container image that follows the Operator Framework standards contains all resource definitions required to deploy the operator application. This way the OLM can install an operator automatically. If an operator is not built and packaged by following the Operator Framework standards, the OLM will not be able to install nor manage that operator.

Introducing OperatorHub

OperatorHub provides a web interface to discover and publish operators that follow the Operator Framework standards. Both open source operators and commercial operators can be published to the Operator hub. Operator container images can be hosted at different image registries, for example `quay.io`.

Introducing Red Hat Marketplace

Red Hat Marketplace is a platform that allows access to a curated set of enterprise grade operators that can be deployed on an OpenShift or a Kubernetes cluster. Operators available in the Red Hat Marketplace have gone through a certification process to ensure the software follows best practices and also the containers are scanned for vulnerabilities.

The Red Hat Marketplace Operator allows a seamless integration between an OpenShift cluster and the Red Hat Marketplace. This integration manages updates and consolidates billing and reporting, simplifying the deployment of certified operators. Vendors provide several pricing options for their operators, such as free trials, different editions, and discounts for large customers.

Introducing OpenShift Cluster Operators

Cluster operators are regular operators except that they are not managed by the OLM. They are managed by the OpenShift Cluster Version Operator, which is sometimes called a first-level operator. All cluster operators are also called second-level operators.

OpenShift cluster operators provide OpenShift extension APIs and infrastructure services such as:

- The OAuth server, which authenticates access to the control plane and extensions APIs.
- The core DNS server, which manages service discovery inside the cluster.
- The web console, which allows graphical management of the cluster.
- The internal image registry, which allow developers to host container images inside the cluster, using either S2I or another mechanism.
- The monitoring stack, which generates metrics and alerts about the cluster health.

Some cluster operators manage node or control plane settings. For example, with upstream Kubernetes you edit a node configuration file to add storage and network plug-ins, and these plug-ins may require additional configuration files. OpenShift supports operators that manage configuration files in all nodes and reload the node services that are affected by changes to these files.



Important

OpenShift 4 deprecates the usage of SSH sessions to manage nodes configuration and system services. This ensures that you do not customize the nodes, and that they can be safely added or removed from a cluster. You are expected to perform all administrative actions indirectly by editing custom resources and then wait for their respective operators to apply your changes.

Exploring OpenShift Cluster Operators

Usually an operator and its managed application share the same project. In the case of cluster operators, these are in the `openshift-*` projects. Every cluster operator defines a custom resource of type `ClusterOperator`. Cluster operators manage the cluster itself, including the API server, the web console, or the network stack. Each cluster operator defines a set of custom resources, to further control its components. The `ClusterOperator` API resource exposes information such as the health of the update, or the version of the component.

Operators are apparent from their name, for example, the `console` cluster operator provides the web console, the `ingress` cluster operator enables ingresses and routes. The following lists some of the cluster operators:

- `authentication`
- `cloud-credential`
- `cluster-autoscaler`
- `console`
- `dns`
- `image-registry`
- `ingress`
- `monitoring`
- `network`
- `openshift-apiserver`
- `openshift-controller-manager`
- `storage`



References

Further information is available in the Red Hat OpenShift Container Platform 4.10 product documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/

Introducing the Operator Framework

<https://blog.openshift.com/introducing-the-operator-framework/>

Getting started with Red Hat Marketplace

<https://marketplace.redhat.com/en-us/documentation/getting-started>

► Quiz

Describing Cluster Operators

Match the items below to their counterparts in the table.

Custom Resource Definition

Operator

Operator Catalog

Operator Image

Operator Lifecycle Manager (OLM)

Operator SDK

OperatorHub

Red Hat Marketplace

Operator Terminology	Name
An open source toolkit for building, testing, and packaging operators.	
A repository for discovering and installing operators.	
An extension of the Kubernetes API that defines the syntax of a custom resource.	
The artifact defined by the Operator Framework that you can publish for consumption by an OLM instance.	
An application that manages Kubernetes resources.	
An application that manages Kubernetes operators.	
A public web service where you can publish operators that are compatible with the OLM.	
Platform that allows access to certified software packaged as Kubernetes operators that can be deployed in an OpenShift cluster.	

► Solution

Describing Cluster Operators

Match the items below to their counterparts in the table.

Operator Terminology	Name
An open source toolkit for building, testing, and packaging operators.	Operator SDK
A repository for discovering and installing operators.	Operator Catalog
An extension of the Kubernetes API that defines the syntax of a custom resource.	Custom Resource Definition
The artifact defined by the Operator Framework that you can publish for consumption by an OLM instance.	Operator Image
An application that manages Kubernetes resources.	Operator
An application that manages Kubernetes operators.	Operator Lifecycle Manager (OLM)
A public web service where you can publish operators that are compatible with the OLM.	OperatorHub
Platform that allows access to certified software packaged as Kubernetes operators that can be deployed in an OpenShift cluster.	Red Hat Marketplace

Summary

In this chapter, you learned:

- Red Hat OpenShift Container Platform executes the Kubernetes pods in the CRI-O container engine.
- RHOCP 4 provides services on top of Kubernetes, such as an internal container image registry, storage, and networking providers.
- Red Hat OpenShift Container Platform nodes are based on Red Hat Enterprise Linux CoreOS.
- RHOCP 4 implements centralized logging and monitoring.
- Operators package applications that manage Kubernetes resources.
- The Operator Lifecycle Manager (OLM) handles installation and management of operators.
- Operator Hub is an online catalog for discovering operators.

Chapter 2

Verifying the Health of a Cluster

Goal

Describe OpenShift installation methods and verify the health of a newly installed cluster.

Objectives

- Describe the OpenShift installation process, full-stack automation, and pre-existing infrastructure installation methods.
- Execute commands that assist in troubleshooting, verify that the OpenShift nodes are healthy, and troubleshoot common issues with OpenShift and Kubernetes deployments.
- Identify the components and resources of persistent storage and deploy an application that uses a persistent volume claim.

Sections

- Describing Installation Methods (and Quiz)
- Troubleshooting OpenShift Clusters and Applications (and Guided Exercise)
- Introducing OpenShift Dynamic Storage (and Guided Exercise)

Describing Installation Methods

Objectives

After completing this section, you should be able to describe the OpenShift installation process, full-stack automation, and pre-existing infrastructure installation methods.

Introducing OpenShift Installation Methods

Red Hat OpenShift Container Platform provides two main installation methods:

Full-stack Automation

With this method, the OpenShift installer provisions all compute, storage, and network resources from a cloud or virtualization provider. You provide the installer with minimum data, such as credentials to a cloud provider and the size of the initial cluster, and then the installer deploys a fully functional OpenShift cluster.

Pre-existing Infrastructure

With this method, you configure a set of compute, storage, and network resources and the OpenShift installer configures an initial cluster using these resources. You can use this method to set up an OpenShift cluster using bare-metal servers and cloud or virtualization providers that are not supported by the full-stack automation method.

When using a pre-existing infrastructure, you must provide all of the cluster infrastructure and resources, including the bootstrap node. You must run the installation program to generate the required configuration files, and then run the installation program again to deploy an OpenShift cluster on your infrastructure.

At the time of the Red Hat OpenShift Container Platform 4.10 release, the set of cloud providers supported for the full-stack automation method includes Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform using the standard Intel architecture (x86). Supported virtualization providers and architectures for full-stack automation include VMware, Red Hat Virtualization, IBM Power, and IBM System Z.

Every minor release of the 4.x stream adds more capabilities and more support for customizations, such as reusing precreated cloud resources.

Comparing OpenShift Installation Methods

Certain features of OpenShift require using the full-stack automation method, for example, cluster automatic scaling. However, it is expected that future releases might relax such requirements.

Using the full-stack automation method, all nodes of the new cluster run Red Hat Enterprise Linux CoreOS (RHEL CoreOS). Using the pre-existing infrastructure method, compute nodes can be set up using Red Hat Enterprise Linux (RHEL), but the control plane still requires RHEL CoreOS.

Describing the Deployment Process

The installation takes place in several stages, starting with the creation of a bootstrap machine that runs Red Hat Enterprise Linux CoreOS using the assets that the installer generates.

The bootstrapping process for the cluster is as follows:

1. The bootstrap machine boots, and then starts hosting the remote resources required for booting the control plane machines.
2. The control plane machines fetch the remote resources from the bootstrap machine and finish booting.
3. The control plane machines form an Etcd cluster.
4. The bootstrap machine starts a temporary Kubernetes control plane using the newly-created Etcd cluster.
5. The temporary control plane schedules the control plane to the control plane machines.
6. The temporary control plane shuts down and yields to the control plane.
7. The bootstrap node injects components specific to OpenShift into the control plane.
8. Finally, the installer tears down the bootstrap machine.

The result of this bootstrapping process is a fully running OpenShift control plane, which includes the API server, the controllers (such as the SDN), and the Etcd cluster. The cluster then downloads and configures the remaining components needed for day-to-day operation via the Cluster Version operator, including the automated creation of compute machines on supported platforms.

Customizing an OpenShift Installation

The OpenShift installer allows very little customization of the initial cluster that it provisions. Most customization is performed after installation, including:

- Defining custom storage classes for dynamic storage provisioning.
- Changing the custom resources of cluster operators.
- Adding new operators to a cluster.
- Defining new machine sets.



References

For more information on the various installation methods, refer to the Red Hat OpenShift Container Platform 4.10 *Installing* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/installing/index

For more information on Installer Provisioned Infrastructure, refer to the Red Hat OpenShift Container Platform 4.10 *OpenShift 4.x Installation - Quick Overview (IPI Installation)* video at <https://youtu.be/uBsilb4cual>

For more information on User Provisioned Infrastructure, refer to the Red Hat OpenShift Container Platform 4.10 *OpenShift 4 User Provisioned Infrastructure with VMware vSphere* video at <https://youtu.be/TsAJEEDv-gg>

► Quiz

Describing Installation Methods

Choose the correct answers to the following questions:

- 1. **Which of the following installation methods requires using the OpenShift installer to configure control plane and compute nodes?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 2. **Which of the following installation methods allows setting up nodes using Red Hat Enterprise Linux?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 3. **Which of the following installation methods allows using an unsupported virtualization provider at the expense of some OpenShift features?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 4. **Which installation method allows using several supported cloud providers with minimum effort?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 5. **Which of the following installation methods allows extensive customization of the cluster settings by providing input to the OpenShift installer?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.

► Solution

Describing Installation Methods

Choose the correct answers to the following questions:

- 1. **Which of the following installation methods requires using the OpenShift installer to configure control plane and compute nodes?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 2. **Which of the following installation methods allows setting up nodes using Red Hat Enterprise Linux?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 3. **Which of the following installation methods allows using an unsupported virtualization provider at the expense of some OpenShift features?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 4. **Which installation method allows using several supported cloud providers with minimum effort?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.
- 5. **Which of the following installation methods allows extensive customization of the cluster settings by providing input to the OpenShift installer?**
 - a. Full-stack automation.
 - b. Pre-existing infrastructure.
 - c. Both full-stack automation and pre-existing infrastructure.
 - d. Neither full-stack automation nor pre-existing infrastructure.

Troubleshooting OpenShift Clusters and Applications

Objectives

After completing this section, you should be able to execute commands that assist in troubleshooting, verify that the OpenShift nodes are healthy, and troubleshoot common issues with OpenShift and Kubernetes deployments.

Troubleshooting Common Issues with an OpenShift Cluster

Most troubleshooting of the OpenShift cluster is very similar to troubleshooting application deployments, because most components of Red Hat OpenShift 4 are operators, and operators are Kubernetes applications. For each operator, you can identify the project where it resides, the deployment that manages the operator application, and its pods. If that operator has configuration settings that you need to change, then you can identify the custom resource (CR), or sometimes the configuration map or secret resource that stores these settings.

Most OpenShift operators manage applications that are also deployed from standard Kubernetes Workload API resources, such as daemon sets and deployments. The role of the operator is usually to create these resources and keep them in sync with the CR.

This section begins by focusing on cluster issues that are not directly related to operators or application deployments; later in this section, you learn how to troubleshoot application deployments.

Verifying the Health of OpenShift Nodes

The following commands display information about the status and health of nodes in an OpenShift cluster:

`oc get nodes`

Displays a column with the status of each node. If a node is not **Ready**, then it cannot communicate with the OpenShift control plane, and is effectively dead to the cluster.

`oc adm top nodes`

Displays the current CPU and memory usage of each node. These are actual usage numbers, not the resource requests that the OpenShift scheduler considers as the available and used capacity of the node.

`oc describe node my-node-name`

Displays the resources available and used from the scheduler point of view, and other information. Look for the headings "Capacity", "Allocatable", and "Allocated resources" in the output. The heading "Conditions" indicates whether the node is under memory pressure, disk pressure, or some other condition that would prevent the node from starting new containers.

Reviewing the Cluster Version Resource

The OpenShift installer creates an `auth` directory containing the `kubeconfig` and `kubeadmin-password` files. Run the `oc login` command to connect to the cluster with the `kubeadmin` user. The password of the `kubeadmin` user is in the `kubeadmin-password` file.

```
[user@host ~]$ oc login -u kubeadmin -p MMTUc-TnXjo-NFyh3-aeWmC \
> https://api.ocp4.example.com:6443
Login successful.
...output omitted...
```

`ClusterVersion` is a custom resource that holds high-level information about the cluster, such as the update channels, the status of the cluster operators, and the cluster version (for example, 4.10.3). Use this resource to declare the version of the cluster you want to run. Defining a new version for the cluster instructs the `cluster-version` operator to upgrade the cluster to that version.

You can retrieve the cluster version to verify that it is running the desired version, and also to ensure that the cluster uses the right subscription channel.

- Run `oc get clusterversion` to retrieve the cluster version. The output lists the version, including minor releases, the cluster uptime for a given version, and the overall status of the cluster.

```
[user@host ~]$ oc get clusterversion
```

NAME	VERSION	AVAILABLE	PROGRESSING	SINCE	STATUS
version	4.10.3	True	False	35d	Cluster version is 4.10.3

- Run `oc describe clusterversion` to obtain more detailed information about the cluster status.

```
[user@host ~]$ oc describe clusterversion
Name:          version
Namespace:
Labels:        <none>
Annotations:   <none>
API Version:   config.openshift.io/v1
Kind:          ClusterVersion
...output omitted...
Spec:
  Channel:      stable-4.10 1
  Cluster ID:   f33267f8-260b-40c1-9cf3-ecc406ce035e 2
Status:
  Available Updates: <nil> 3
  Channels:
    candidate-4.10
    candidate-4.11
    eus-4.10
    fast-4.10
    stable-4.10
  Image:        quay.io/openshift-release-dev/ocp-release...
  URL:          https://access.redhat.com/errata/RHBA-2022:0811
  Version:      4.10.4
Conditions:
  Last Transition Time: 2022-03-15T12:39:02Z
  Message:             Done applying 4.10.3 4
  Status:              True
  Type:               Available
...output omitted...
Desired:
```

```

Channels:
  candidate-4.10
  candidate-4.11
  eus-4.10
  fast-4.10
  stable-4.10
...output omitted...
History:
  Completion Time:    2022-03-15T12:39:02Z ❸
...output omitted...
  State:              Completed ❹
  Verified:           false
  Version:            4.10.3
Observed Generation: 2
...output omitted...

```

- ❶ Displays the version of the cluster and its channel. Depending on your subscription, the channel might be different.
- ❷ Displays the unique identifier for the cluster. Red Hat uses this identifier to identify clusters and cluster entitlements.
- ❸ This entry lists the available images for updating the cluster.
- ❹ This entry lists the history. The output indicates that an update completed.
- ❺ This entry shows when the cluster deployed the version indicated in the `Version` entry
- ❻ This entry indicates that the version successfully deployed. Use this entry to determine if the cluster is healthy.

Reviewing Cluster Operators

OpenShift Container Platform *cluster operators* are top level operators that manage the cluster. They are responsible for the main components, such as the API server, the web console, storage, or the SDN. Their information is accessible through the `ClusterOperator` resource, which allows you to access the overview of all cluster operators, or detailed information on a given operator.

Run `oc get clusteroperators` to retrieve the list of all cluster operators:

```

[user@host ~]$ oc get clusteroperators

```

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.10.3	True	False	False	35d ❶
baremetal	4.10.3	True	False	False	35d
cloud-controller-manager	4.10.3	True	False	False	35d
cloud-credential	4.10.3	True	False	False	35d
cluster-autoscaler	4.10.3	True	False	False	35d
config-operator	4.10.3	True	False	False	35d
console	4.10.3	True	False	False	3h22m
csi-snapshot-controller	4.10.3	True	False	False	35d
dns	4.10.3	True	False	False	35d
etcd	4.10.3	True	False	False	35d
image-registry	4.10.3	True	False	False	35d
...output omitted...					

- 1 Each row describes a cluster operator.

The **NAME** field indicates the name of the operator. This operator is responsible for managing authentication.

The **AVAILABLE** field indicates that the authentication operator deployed successfully and is available for use in the cluster. Notice that a cluster operator might return a status of available even if its degraded. An operator reports *degraded* when its current state does not match its desired state over a period of time. For example, if the operator requires three running pods, but one pod is crashing, the operator is available but in a degraded state.

The **PROGRESSING** field indicates whether an operator is being updated to a newer version by the top level operator. If new resources are being deployed by the `cluster version` operator, then the columns read **True**.

The **DEGRADED** field returns the health of the operator. The entry reads **True** if the operator encounters an error that prevents it from working properly. The operator services might still be available, however, all the requirements might not be satisfied. This can indicate that the operator will fail and require user intervention.

Displaying the Logs of OpenShift Nodes

Most of the infrastructure components of OpenShift are containers inside pods; you can view their logs the same way you view logs for any end-user application. Some of these containers are created by the Kubelet, and thus invisible to most distributions of Kubernetes, but OpenShift cluster operators create pod resources for them.

An OpenShift node based on Red Hat Enterprise Linux CoreOS runs very few local services that would require direct access to a node to inspect their status. Most of the system services in Red Hat Enterprise Linux CoreOS run as containers. The main exceptions are the CRI-O container engine and the Kubelet, which are Systemd units. To view these logs, use the `oc adm node-logs` command as shown in the following examples:

```
[user@host ~]$ oc adm node-logs -u cri-o my-node-name
```

```
[user@host ~]$ oc adm node-logs -u kubelet my-node-name
```

You can also display all journal logs of a node:

```
[user@host ~]$ oc adm node-logs my-node-name
```

Opening a Shell Prompt on an OpenShift Node

Administrators who manage Red Hat OpenShift Cluster Platform 3 and other distributions of Kubernetes frequently open SSH sessions to their nodes to inspect the state of the control plane and the container engine, or to make changes to configuration files. Although this can still be done, it is no longer recommended with Red Hat OpenShift Cluster Platform 4.

If you install your cluster using the full-stack automation method, then your cluster nodes are not directly accessible from the internet because they are on a virtual private network, which AWS calls Virtual Private Cloud (VPC). To open SSH sessions, a bastion server on the same VPC of your cluster that is also assigned a public IP address is required. Creating a bastion server depends on your cloud provider and is out of scope for this course.

The `oc debug node` command provides a way to open a shell prompt in any node of your cluster. That prompt comes from a special-purpose tools container that mounts the node root file system at the `/host` folder, and allows you to inspect any files from the node.

To run local commands directly from the node, while in a `oc debug node` session, you must start a chroot shell in the `/host` folder. Then you can inspect the local file systems of the node, the status of its systemd services, and perform other tasks that would otherwise require a SSH session. The following is an example `oc debug node` session:

```
[user@host ~]$ oc debug node/my-node-name
...output omitted...
sh-4.4# chroot /host
sh-4.4# systemctl is-active kubelet
active
```

A shell session started from the `oc debug node` command depends on the OpenShift control plane to work. It uses the same tunneling technology that allows opening a shell prompt inside a running pod (see the `oc rsh` command later in this section). The `oc debug node` command is not based on the SSH or RSH protocols.

If your control plane is not working, your node is not ready, or for some reason your node is not able to communicate with the control plane, then you cannot rely on the `oc debug node` command and will require a bastion host.



Warning

Exercise care when using the `oc debug node` command. Some actions can render your node unusable, such as stopping the Kubelet, and you cannot recover using only `oc` commands.

Troubleshooting The Container Engine

From an `oc debug node` session, use the `crictl` command to get low-level information about all local containers running on the node. You cannot use the `podman` command for this task because it does not have visibility on containers created by CRI-O. The following example lists all containers running on a node. The `oc describe node` command provides the same information but organized by pod instead of by container.

```
[user@host ~]$ oc debug node/my-node-name
...output omitted...
sh-4.4# chroot /host
sh-4.4# crictl ps
...output omitted...
```

Troubleshooting Application Deployments

You can usually ignore the differences between Kubernetes deployments and OpenShift deployment configurations when troubleshooting applications. The common failure scenarios and the ways to troubleshoot them are essentially the same.

There are many scenarios that will be described in later chapters of this course, such as pods that cannot be scheduled. This section focuses on common scenarios that apply to generic applications, and the same scenarios usually apply to operators also.

Troubleshooting Pods That Fail to Start

A common scenario is that OpenShift creates a pod and that pod never establishes a `Running` state. This means that OpenShift could not start the containers inside that pod. Start troubleshooting using the `oc get pod` and `oc status` commands to verify whether your pods and containers are running. At some point, the pods are in an error state, such as `ErrImagePull` or `ImagePullBackOff`.

When this happens, the first step is listing events from the current project using the `oc get events` command. If your project contains many pods, then you can get a list of events filtered by pod using the `oc describe pod` command. You can also run similar `oc describe` commands to filter events by deployments and deployment configurations.

Troubleshooting Running and Terminated Pods

Another common scenario is that OpenShift creates a pod, and for a short time no problem is encountered. The pod enters the `Running` state, which means at least one of its containers started running. Later, an application running inside one of the pod containers stops working. It might either terminate or return error messages to user requests.

If the application is managed by a properly designed deployment, then it should include health probes that will eventually terminate the application and stop its container. If that happens, then OpenShift tries to restart the container several times. If the application continues terminating, due to health probes or other reasons, then the pod will be left in the `CrashLoopBackOff` state.

A container that is running, even for a very short time, generates logs. These logs are not discarded when the container terminates. The `oc logs` command displays the logs from any container inside a pod. If the pod contains a single container, then the `oc logs` command only requires the name of the pod.

```
[user@host ~]$ oc logs my-pod-name
```

If the pod contains multiple containers, then the `oc logs` command requires the `-c` option.

```
[user@host ~]$ oc logs my-pod-name -c my-container-name
```

Interpreting application logs requires specific knowledge of that particular application. If all goes well, the application provides clear error messages that can help you find the problem.

Introducing OpenShift Aggregated Logging

Red Hat OpenShift Container Platform 4 provides the Cluster Logging subsystem, based on Elasticsearch, Fluentd or Rsyslog, and Kibana, which aggregates logs from the cluster and its containers.

Deploying and configuring the OpenShift Cluster Logging subsystem through its operator is beyond the scope of this course. Refer to the references section at the end of this section for more information.

Creating Troubleshooting Pods

If you are not sure whether your issues relate to the application container image, or to the settings it gets from its OpenShift resources, then the `oc debug` command is very useful. This command creates a pod based on an existing pod, a deployment configuration, a deployment, or any other resource from the Workloads API.

The new pod runs an interactive shell instead of the default entry point of the container image. It also runs with health probes disabled. This way, you can easily verify environment variables, network access to other services, and permissions inside the pod.

The command-line options of the `oc debug` command allow you to specify settings that you do not want to clone. For example, you could change the container image, or specify a fixed user id. Some settings might require cluster administrator privileges.

A common scenario is creating a pod from a deployment, but running as the root user and thus proving that the deployment references a container image that was not designed to run under the default security policies of OpenShift:

```
[user@host ~]$ oc debug deployment/my-deployment-name --as-root
```

Changing a Running Container

Because container images are immutable, and containers are supposed to be ephemeral, it is not recommended that you make changes to running containers. However, sometimes making these changes can help with troubleshooting application issues. After you try changing a running container, do not forget to apply the same changes back to the container image and its application resources, and then verify that the now permanent fixes work as expected.

The following commands help with making changes to running containers. They all assume that pods contain a single container. If not, you must add the `-c my-container-name` option.

`oc rsh my-pod-name`

Opens a shell inside a pod to run shell commands interactively and non-interactively.

`oc cp /local/path my-pod-name:/container/path`

Copies local files to a location inside a pod. You can also reverse arguments and copy files from inside a pod to your local file system. See also the `oc rsync` command for copying multiple files at once.

`oc port-forward my-pod-name local-port:remote-port`

Creates a TCP tunnel from *local-port* on your workstation to *local-port* on the pod. The tunnel is alive as long as you keep the `oc port-forward` running. This allows you to get network access to the pod without exposing it through a route. Because the tunnel starts at your localhost, it cannot be accessed by other machines.

Troubleshooting OpenShift CLI Commands

Sometimes, you cannot understand why an `oc` command fails and you need to troubleshoot its low-level actions to find the cause. Maybe you need to know what a particular invocation of the `oc` command does behind the scenes, so you can replicate the behavior with an automation tool that makes OpenShift and Kubernetes API requests, such as Ansible Playbooks using the `k8s` module.

The `--loglevel level` option displays OpenShift API requests, starting with level 6. As you increase the level, up to 10, more information about those requests is added, such as their HTTP request headers and response bodies. Level 10 also includes a `curl` command to replicate each request.

You can try these two commands, from any project, and compare their outputs.

```
[user@host ~]$ oc get pods --loglevel 6
```

```
[user@host ~]$ oc get pods --loglevel 10
```

Sometimes, you only need the authentication token that the `oc` command uses to authenticate OpenShift API requests. With this token, an automation tool can make OpenShift API requests as if it was logged in as your user. To get your token, use the `-t` option of the `oc whoami` command:

```
[user@host ~]$ oc whoami -t
```



References

For more information about OpenShift events, refer to the *Viewing system event information in an OpenShift Container Platform cluster* section in the *Working with clusters* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-containers-events

For more information about how to copy files to running containers, refer to the *Copying files to or from an OpenShift Container Platform container* section in the *Working with containers* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-containers-copying-files

For more information about how to execute commands on running containers, refer to the *Executing remote commands in an OpenShift Container Platform container* section in the *Working with containers* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-containers-remote-commands

For more information about how to forward local ports to running containers, refer to the *Using port forwarding to access applications in a container* section in the *Working with containers* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-containers-port-forwarding

For more information about aggregated logging, refer to the Red Hat OpenShift Container Platform 4.10 *Logging* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/logging/index

ClusterOperator Custom Resource

<https://github.com/openshift/cluster-version-operator>

► Guided Exercise

Troubleshooting OpenShift Clusters and Applications

In this exercise, you will execute commands that assist in troubleshooting common problems with the OpenShift control plane and with application deployments.

Outcomes

You should be able to:

- Inspect the general state of an OpenShift cluster.
- Inspect local services and pods running in an OpenShift compute node.
- Diagnose and fix issues with the deployment of an application.

Before You Begin



Important

The first time you start your classroom environment, OpenShift clusters take a little longer to become fully available. The `lab` command at the beginning of each exercise can fail if the clusters are not fully available.

If you try to access your cluster using either the `oc` command or the web console without first running a `lab` command, then you might find that your cluster is not yet available. If that happens, then wait a few minutes and try again.

You can run the `wait.sh` script in the `utility` machine that waits until your OpenShift cluster is ready to accept authentication and API requests from remote clients.

```
[student@workstation ~]$ ssh lab@utility
[lab@utility ~]$ ./wait.sh
```

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable, and creates the resource files that you will be using in the activity. It also creates the `install-troubleshoot` project with an application that you will diagnose and fix during this exercise.

```
[student@workstation ~]$ lab install-troubleshoot start
```

Instructions

- 1. Log in to the OpenShift cluster and inspect the status of your cluster nodes.

- 1.1. Source the classroom configuration file that is accessible at `/usr/local/etc/ocp4.config`.

```
[student@workstation ~]$ source /usr/local/etc/ocp4.config
```

- 1.2. Log in to the cluster as the kubeadmin user. When prompted, accept the insecure certificate.

```
[student@workstation ~]$ oc login -u kubeadmin -p ${RHT_OCP4_KUBEADM_PASSWD} \
> https://api.ocp4.example.com:6443
The server uses a certificate signed by an unknown authority.
You can bypass the certificate check, but any data you send to the server could be
intercepted by others.
Use insecure connections? (y/n): y

Login successful.
...output omitted...
```

- 1.3. Verify that all nodes on your cluster are ready.

```
[student@workstation ~]$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
master01	Ready	master,worker	36d	v1.23.3+e419edf
master02	Ready	master,worker	36d	v1.23.3+e419edf
master03	Ready	master,worker	36d	v1.23.3+e419edf

- 1.4. Verify whether any of your nodes are close to using all of the CPU and memory available to them.
Repeat the following command a few times to prove that you see actual usage of CPU and memory from your nodes. The numbers you see should change slightly each time you repeat the command.

```
[student@workstation ~]$ oc adm top node
```

NAME	CPU(cores)	CPU%	MEMORY(bytes)	MEMORY%
master01	743m	21%	9002Mi	60%
master02	747m	21%	6112Mi	41%
master03	1047m	29%	10150Mi	68%

- 1.5. Use the `oc describe` command to verify that all of the conditions that might indicate problems are false.

```
[student@workstation ~]$ oc describe node master01
...output omitted...
Conditions:
```

Type	Status	...	Message
MemoryPressure	False	...	kubelet has sufficient memory available
DiskPressure	False	...	kubelet has no disk pressure
PIDPressure	False	...	kubelet has sufficient PID available

```

Ready      True      ... kubelet is posting ready status
Addresses:
...output omitted...

```

- ▶ 2. Review the logs of the internal registry operator, the internal registry server, and the Kubelet of a node.
 - 2.1. List all pods inside the `openshift-image-registry` project, and then identify the pod that runs the operator and the pod that runs the internal registry server.

```

[student@workstation ~]$ oc get pod -n openshift-image-registry
NAME                                     READY   STATUS    ...
cluster-image-registry-operator-6f7784dc86-qq258  1/1     Running   ...
image-pruner-27506880-dcn8n                0/1     Completed ...
image-pruner-27508320-2vfph                 0/1     Completed ...
image-registry-867979b75b-xk4ns             1/1     Running   ...
...output omitted...

```

- 2.2. Follow the logs of the operator pod (`cluster-image-registry-operator-xxx`). Your output might be different than the following example.

```

[student@workstation ~]$ oc logs --tail 3 -n openshift-image-registry \
> cluster-image-registry-operator-564bd5dd8f-s46bz
I0421 01:30:11.357218      1 generator.go:60] object *v1.DaemonSet,
  Namespace=openshift-image-registry, ...
I0421 01:30:18.290064      1 recorder_logging.go:44] &Event{ObjectMeta:
{dummy.16e7c534a15072a6 dummy ...
I0421 01:30:18.290674      1 generator.go:60] object *v1.DaemonSet,
  Namespace=openshift-image-registry, ...
...output omitted...

```

- 2.3. Follow the logs of the image registry server pod (`image-registry-xxx` from the output of the `oc get pod` command run previously). Your output might be different than the following example.

```

[student@workstation ~]$ oc logs --tail 1 -n openshift-image-registry \
> image-registry-794dfc7978-w7w69
I0421 01:30:11.357218      1 generator.go:60] object *v1.DaemonSet,
  Namespace=openshift-image-registry, ...
I0421 01:30:18.290064      1 recorder_logging.go:44] &Event{ObjectMeta:
{dummy.16e7c534a15072a6 dummy ...
...output omitted...

```

- 2.4. Follow the logs of the Kubelet from the same node that you inspected for CPU and memory usage in the previous step. Your output might be different than the following example.

```
[student@workstation ~]$ oc adm node-logs --tail 1 -u kubelet master01
-- Logs begin at Tue 2022-03-15 12:21:41 UTC, end at Thu 2022-04-21 02:17:38 UTC.
--
Apr 20 17:30:28.172354 master01 systemd[1]: kubelet.service: Consumed 2h 12min
42.989s CPU time
-- Logs begin at Tue 2022-03-15 12:21:41 UTC, end at Thu 2022-04-21 02:17:38 UTC.
--
Apr 21 02:17:05.808392 master01 hyperkube[1844]: I0421 02:17:05.808374      1844
kubelet_getters.go:176] ...
```

- ▶ 3. Start a shell session to the same node that you previously used to inspect its OpenShift services and pods. Do not make any change to the node, such as stopping services or editing configuration files.

- 3.1. Start a shell session on the node, and then use the `chroot` command to enter the local file system of the host.

```
[student@workstation ~]$ oc debug node/master01
Creating debug namespace/openshift-debug-node-5zsch ...
Starting pod/master01-debug ...
To use host binaries, run `chroot /host`
Pod IP: 192.168.50.10
If you do not see a command prompt, try pressing enter.
sh-4.4# chroot /host
sh-4.4#
```

- 3.2. Still using the same shell session, verify that the Kubelet and the CRI-O container engine are running. Type `q` to exit the command.

```
sh-4.4# systemctl status kubelet
● kubelet.service - Kubernetes Kubelet
   Loaded: loaded (/etc/systemd/system/kubelet.service; enabled; vendor preset: disabled)
   Drop-In: /etc/systemd/system/kubelet.service.d
            └─10-mco-default-madv.conf, 20-logging.conf, 20-nodenet.conf
   Active: active (running) since Wed 2022-04-20 20:36:05 UTC; 5h 47min ago
   ...output omitted...
q
```

Rerun the same command against the `crio` service. Type `q` to exit from the command.

```
sh-4.4# systemctl status crio
● crio.service - Container Runtime Interface for OCI (CRI-O)
   Loaded: loaded (/usr/lib/systemd/system/crio.service; disabled; vendor preset: disabled)
   Drop-In: /etc/systemd/system/crio.service.d
            └─10-mco-default-madv.conf, 10-mco-profile-unix-socket.conf, 20-nodenet.conf
   Active: active (running) since Wed 2022-04-20 20:35:59 UTC; 5h 59min ago
   ...output omitted...
q
```

3.3. Still using the same shell session, verify that the `etcd` pod is running.

```
sh-4.4# crictl ps --name etcd
CONTAINER    ... STATE    NAME           ATTEMPT  POD ID
5e18a8220a9d5 ... Running   etcd-operator  2        90b6b1150e3fe8d9adceeb5549
19bc3ed5e8643 ... Running   etcd-metrics   3        c94baa4f55618
...output omitted...
```

3.4. Terminate the `chroot` session and shell session to the node. This also terminates the `oc debug node` command.

```
sh-4.4# exit
exit

sh-4.4# exit
exit

Removing debug pod ...
[student@workstation ~]$
```

► 4. Enter the `install-troubleshoot` project to diagnose a pod that is in an error state.

4.1. Use the `install-troubleshoot` project.

```
[student@workstation ~]$ oc project install-troubleshoot
Now using project "install-troubleshoot" on server ...
```

4.2. Verify that the project has a single pod in either the `ErrImagePull` or `ImagePullBackOff` status.

```
[student@workstation ~]$ oc get pods
NAME                                READY   STATUS    ...
psql-7d4cc9d6d-m5r59               0/1     ImagePullBackOff ...
```

4.3. Verify that the project includes a Kubernetes deployment that manages the pod.

```
[student@workstation ~]$ oc status
...output omitted...
svc/psql - 172.30.13.140:5432

deployment/psql deploys registry.redhat.io/rhel8/postgresql-13:1
  deployment #1 running for 7 hours - 0/1 pods
...output omitted...
```

4.4. List all events from the current project and look for error messages related to the pod.

```
[student@workstation ~]$ oc get events
LAST SEEN   TYPE      REASON          OBJECT                                MESSAGE
7h17m       Normal    Scheduled        pod/psql-768d797946-d2gj4            Successfully assigned install-troubleshoot/psql-768d797946-d2gj4 to master0
2
```

```

7h17m      Normal    AddedInterface    pod/psql-768d797946-d2gj4    Add eth0
[10.9.0.21/23] from openshift-sdn
3h22m      Normal    Pulling           pod/psql-768d797946-d2gj4    Pulling
image "registry.redhat.io/rhel8/postgresq-13:1"
7h16m      Warning   Failed            pod/psql-768d797946-d2gj4    Failed to
pull image "registry.redhat.io/rhel8/postgresq-13:1": rpc error: code = Unknown
desc = reading manifest 1 in registry.redhat.io/rhel8/postgresq-13: unknown: Not
Found
7h16m      Warning   Failed            pod/psql-768d797946-d2gj4    Error:
ErrImagePull
2m32s      Normal    BackOff           pod/psql-768d797946-d2gj4    Back-off
pulling image "registry.redhat.io/rhel8/postgresq-13:1"
7h16m      Warning   Failed            pod/psql-768d797946-d2gj4    Error:
ImagePullBackOff
7h17m      Normal    SuccessfulCreate   replicaset/psql-768d797946    Created
pod: psql-768d797946-d2gj4
7h17m      Normal    ScalingReplicaSet  deployment/psql                Scaled up
replica set psql-768d797946 to 1

```

This output also indicates a problem getting the image for deploying the pod.

- 4.5. Log in to the Red Hat Container Catalog with your Red Hat account.

```

[student@workstation ~]$ podman login registry.redhat.io
Username: your_username
Password: your_password
Login Succeeded!

```

- 4.6. Use Skopeo to find information about the container image from the events.

```

[student@workstation ~]$ skopeo inspect \
> docker://registry.redhat.io/rhel8/postgresq-13:1
FATA[0000] Error parsing image name "docker://registry.redhat.io/rhel8/
postgresq-13:1": Error reading manifest 1 in registry.redhat.io/rhel8/
postgresq-13: unknown: Not Found

```

- 4.7. It looks like the container image is misspelled. Verify that it works if you replace `postgresq-13` with `postgresql-13`.

```

[student@workstation ~]$ skopeo inspect \
> docker://registry.redhat.io/rhel8/postgresql-13:1
{
  "Name": "registry.redhat.io/rhel8/postgresql-13",
  ...output omitted...
}

```

- 4.8. To verify that the image name is the root cause of the error, edit the `psql` deployment to correct the name of the container image. The `oc edit` commands uses `vi` as the default editor.



Warning

In a real-world scenario, you would ask whoever deployed the PostgreSQL database to fix their YAML and redeploy their application.

```
[student@workstation ~]$ oc edit deployment/psql
...output omitted...
spec:
  containers:
  - env:
    - name: POSTGRESQL_DATABASE
      value: db
    - name: POSTGRESQL_PASSWORD
      value: pass
    - name: POSTGRESQL_USER
      value: user
    image: registry.redhat.io/rhel8/postgresql-13:1-7
...output omitted...
```

4.9. Verify that a new deployment is active.

```
[student@workstation ~]$ oc status
...output omitted...
deployment/psql deploys registry.redhat.io/rhel8/postgresql-13:1
deployment #2 running for 30 seconds - 1 pod
deployment #1 deployed 7 hours ago
```

4.10. List all pods in the current project. You might see both the old failing pod and the new pod for a few moments. Repeat the following command until you see that the new pod is ready and running, and you no longer see the old pod.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
psql-544c9c666f-btlw8	1/1	Running	0	55s

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab install-troubleshoot finish
```

This concludes the guided exercise.

Introducing OpenShift Dynamic Storage

Objectives

After completing this section, you should be able to identify the components and resources of persistent storage and deploy an application that uses a persistent volume claim.

Persistent Storage Overview

Containers have ephemeral storage by default. For example, when a container is deleted, all the files and data inside it are deleted also. To preserve the files, containers offer two main ways of maintaining persistent storage: volumes and bind mounts. Volumes are the preferred OpenShift way of managing persistent storage. Volumes are managed manually by the administrator or dynamically through a storage class. Developers working with containers on a local system can mount a local directory into a container using a bind mount.

OpenShift cluster administrators use the Kubernetes persistent volume framework to manage persistent storage for the users of a cluster. There are two ways of provisioning storage for the cluster: static and dynamic. Static provisioning requires the cluster administrator to create persistent volumes manually. Dynamic provisioning uses storage classes to create the persistent volumes on demand.

The OpenShift Container Platform uses storage classes to allow administrators to provide persistent storage. Storage classes are a way to describe types of storage for the cluster and provision dynamic storage on demand.

Developers use persistent volume claims to add persistent volumes dynamically to their applications; it is not necessary for the developer to know details of the storage infrastructure. With static provisioning, developers use precreated PVs or ask a cluster administrator to manually create persistent volumes for their applications.

A persistent volume claim (PVC) belongs to a specific project. To create a PVC, you must specify the access mode and size, among other options. Once created, a PVC cannot be shared between projects. Developers use a PVC to access a persistent volume (PV). Persistent volumes are not exclusive to projects and are accessible across the entire OpenShift cluster. When a persistent volume binds to a persistent volume claim, the persistent volume cannot be bound to another persistent volume claim.

Persistent Volume and Persistent Volume Claim Lifecycle

Persistent volume claims request persistent volume resources. To be eligible, a PV must not be bound to another PVC. Additionally, the PV must provide the access mode specified in the PVC and it must be at least as large as the size requested in the PVC. A PVC can specify additional criteria, such as the name of a storage class. If a PVC cannot find a PV that matches all criteria, the PVC enters a pending state and waits until an appropriate PV becomes available. A cluster administrator can manually create the PV or a storage class can dynamically create the PV. A bound persistent volume can be mounted as a volume to a specific mount point in the pod (for example, `/var/lib/pgsql` for a PostgreSQL database).

Verifying the Dynamic Provisioned Storage

Use the `oc get storageclass` command to view available storage classes. The output identifies the default storage class. If a storage class exists, then persistent volumes are created dynamically to match persistent volume claims. A persistent volume claim that does not specify a storage class uses the default storage class.

```
[user@host ~]$ oc get storageclass
NAME                                PROVISIONER    ...
nfs-storage (default)             nfs-storage-provisioner    ...
```



Note

The classroom environment uses an external, open source NFS provisioner. The provisioner dynamically creates NFS persistent volumes from an existing NFS server. Red Hat does not recommend using this provisioner in production environments.

Deploying Dynamically Provisioned Storage

To add a volume to an application create a `PersistentVolumeClaim` resource and add it to the application as a volume. Create the persistent volume claim using either a Kubernetes manifest or the `oc set volumes` command. In addition to either creating a new persistent volume claim or using an existing persistent volume claim, the `oc set volumes` command can modify a deployment to mount the persistent volume claim as a volume within the pod.

To add a volume to an application, use the `oc set volumes` command:

```
[user@host ~]$ oc set volumes deployment/example-application \
> --add --name example-pv-storage --type pvc --claim-class nfs-storage \
> --claim-mode rwo --claim-size 15Gi --mount-path /var/lib/example-app \
> --claim-name example-pv-claim
```

The command creates a persistent volume claim resource and adds it to the application as a volume within the pod.

The following YAML example specifies a persistent volume claim.

To create a `PersistentVolumeClaim` API object:

```
---
apiVersion: v1
kind: PersistentVolumeClaim ❶
metadata:
  name: example-pv-claim ❷
  labels:
    app: example-application
spec:
  accessModes:
    - ReadWriteOnce ❸
```

```
resources:
  requests:
    storage: 156i ④
```

- ① Indicates that it is a persistent volume claim.
- ② The name to use in the `claimName` field of the `persistentVolumeClaim` element in the `volumes` section of a deployment manifest.
- ③ The storage class provisioner must provide this access mode. If persistent volumes are created statically, then an eligible persistent volume must provide this access mode.
- ④ The storage class will create a persistent volume matching this size request. If persistent volumes are created statically, then an eligible persistent volume must be at least the requested size.

OpenShift defines three access modes that are summarized in the following table.

Access Mode	CLI Abbreviation	Description
ReadWriteMany	RWX	Kubernetes can mount the volume as read-write on many nodes.
ReadOnlyMany	ROX	Kubernetes can mount the volume as read-only on many nodes.
ReadWriteOnce	RWO	Kubernetes can mount the volume as read-write on only a single node.

It is important to mention that claims are matched with the best available PV, usually with a similar access mode, but the supported modes depends of the capabilities of the provider. For example, you can have a PVC with RWO requesting a NFS PV, and it can be matched because NFS supports RWO, but this can not happen in reverse, and the request will remain in the `pending` status.

To add the PVC to the application:

```
...output omitted...
spec:
  volumes:
    - name: example-pv-storage
      persistentVolumeClaim:
        claimName: example-pv-claim
  containers:
    - name: example-application
      image: registry.redhat.io/rhel8/example-app
      ports:
        - containerPort: 1234
      volumeMounts:
        - mountPath: "/var/lib/example-app"
          name: example-pv-storage
...output omitted...
```

Deleting Persistent Volume Claims

To delete a volume, use the `oc delete` command to delete the persistent volume claim. The storage class will reclaim the volume after the PVC is removed.

```
[user@host ~]$ oc delete pvc/example-pvc-storage
```



References

For more information on persistent storage, refer to the *Understanding persistent storage* chapter in the Red Hat OpenShift Container Platform 4.10 *Storage* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/storage/index#understanding-persistent-storage

For more information on ephemeral storage, refer to the *Understanding ephemeral storage* chapter of the Red Hat OpenShift Container Platform 4.10 *Storage* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/storage/index#understanding-ephemeral-storage

► Guided Exercise

Introducing OpenShift Dynamic Storage

In this exercise, you will deploy a PostgreSQL database using a persistent volume claim and identify its dynamically allocated volume.

Outcomes

You should be able to:

- Identify the default storage settings of an OpenShift cluster.
- Create persistent volume claims.
- Manage persistent volumes.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and downloads files required for this exercise.

```
[student@workstation ~]$ lab install-storage start
```

Instructions

- 1. Log in to the OpenShift cluster.
 - 1.1. Source the classroom configuration file that is accessible at `/usr/local/etc/ocp4.config`.

```
[student@workstation ~]$ source /usr/local/etc/ocp4.config
```

- 1.2. Log in to the cluster as the `kubeadmin` user. If prompted, accept the insecure certificate.

```
[student@workstation ~]$ oc login -u kubeadmin -p ${RHT_OCP4_KUBEADM_PASSWD} \
> https://api.ocp4.example.com:6443
...output omitted...
```

- 2. Create a new project named `install-storage`.

```
[student@workstation ~]$ oc new-project install-storage
Now using project "install-storage" on server "https://api.ocp4.example.com:6443".
...output omitted...
```

- 3. Verify the default storage class.

```
[student@workstation ~]$ oc get storageclass
NAME                                     PROVISIONER                               ...
nfs-storage (default)                  k8s-sigs.io/nfs-subdir-external-provisioner ...
```

- 4. Create a new database deployment using the container image located at `registry.redhat.io/rhel8/postgresql-13:1-7`.

```
[student@workstation ~]$ oc new-app --name postgresql-persistent \
> --image registry.redhat.io/rhel8/postgresql-13:1-7 \
> -e POSTGRES_USER=redhat \
> -e POSTGRES_PASSWORD=redhat123 \
> -e POSTGRES_DATABASE=persistentdb
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "postgresql-persistent" created
    deployment.apps "postgresql-persistent" created
    service "postgresql-persistent" created
--> Success
...output omitted...
```

**Note**

For convenience, the `~/D0280/labs/install-storage/commands.txt` file contains some commands that you can copy and paste.

- 5. Add a persistent volume for the PostgreSQL database.

- 5.1. Create a new persistent volume claim to add a new volume to the `postgresql-persistent` deployment.

```
[student@workstation ~]$ oc set volumes deployment/postgresql-persistent \
> --add --name postgresql-storage --type pvc --claim-class nfs-storage \
> --claim-mode rwo --claim-size 10Gi --mount-path /var/lib/pgsql \
> --claim-name postgresql-storage
deployment.apps/postgresql-persistent volume updated
```

- 5.2. Verify that you successfully created the new PVC.

```
[student@workstation ~]$ oc get pvc
NAME                                STATUS ... CAPACITY  ACCESS MODES  STORAGECLASS  AGE
postgresql-storage                 Bound    ...  10Gi         RWO           nfs-storage   25s
```

- 5.3. Verify that you successfully created the new PV.

```
[student@workstation ~]$ oc get pv \
> -o custom-columns=NAME:.metadata.name,CLAIM:.spec.claimRef.name
NAME                                CLAIM
pvc-26cc804a-4ec2-4f52-b6e5-84404b4b9def  image-registry-storage
pvc-65c3cce7-45eb-482d-badf-a6469640bd75  postgresql-storage
```

- ▶ 6. Populate the database using the `~/D0280/labs/install-storage/init_data.sh` script.

6.1. Execute the `init_data.sh` script.

```
[student@workstation ~]$ cd ~/D0280/labs/install-storage
[student@workstation install-storage]$ ./init_data.sh
Populating characters table
CREATE TABLE
INSERT 0 5
```

6.2. Use the `~/D0280/labs/install-storage/check_data.sh` script to verify that the database was populated successfully.

```
[student@workstation install-storage]$ ./check_data.sh
Checking characters table
  id |          name          |          nationality
-----+-----+-----
  1 | Wolfgang Amadeus Mozart | Prince-Archbishopric of Salzburg
  2 | Ludwig van Beethoven   | Bonn, Germany
  3 | Johann Sebastian Bach  | Eisenach, Germany
  4 | José Pablo Moncayo      | Guadalajara, México
  5 | Niccolò Paganini        | Genoa, Italy
(5 rows)
```

- ▶ 7. Remove the `postgresql-persistent` deployment and create another deployment named `postgresql-deployment2` that uses the same persistent volume; verify that the data persisted.

7.1. Delete all resources that contain the `app=postgresql-persistent` label.

```
[student@workstation install-storage]$ oc delete all -l app=postgresql-persistent
service "postgresql-persistent" deleted
deployment.apps "postgresql-persistent" deleted
imagestream.image.openshift.io "postgresql-persistent" deleted
```

7.2. Create the `postgresql-persistent2` deployment with the same initialization data as the `postgresql-persistent` deployment.

```
[student@workstation install-storage]$ oc new-app --name postgresql-persistent2 \
> --image registry.redhat.io/rhel8/postgresql-13:1-7 \
> -e POSTGRESQL_USER=redhat \
> -e POSTGRESQL_PASSWORD=redhat123 \
> -e POSTGRESQL_DATABASE=persistentdb
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "postgresql-persistent2" created
    deployment.apps "postgresql-persistent2" created
    service "postgresql-persistent2" created
--> Success
...output omitted...
```

- 7.3. Use the `~/D0280/labs/install-storage/check_data.sh` script to verify that the database does not have the characters table.

```
[student@workstation install-storage]$ ./check_data.sh
Checking characters table
ERROR: 'characters' table does not exist
```

- 7.4. Add the existing `postgresql-storage` persistent volume claim to the `postgresql-persistent2` deployment.

```
[student@workstation install-storage]$ oc set volumes \
> deployment/postgresql-persistent2 \
> --add --name postgresql-storage --type pvc \
> --claim-name postgresql-storage --mount-path /var/lib/pgsql
deployment.apps/postgresql-persistent2 volume updated
```

- 7.5. Use the `~/D0280/labs/install-storage/check_data.sh` script to verify that the persistent volume was successfully added and that the `postgresql-persistent2` pod can access the previously created data.

```
[student@workstation install-storage]$ ./check_data.sh
Checking characters table
  id |          name          |          nationality
-----+-----+-----
   1 | Wolfgang Amadeus Mozart | Prince-Archbishopric of Salzburg
   2 | Ludwig van Beethoven    | Bonn, Germany
...output omitted...
```

- 8. Remove the `postgresql-persistent2` deployment and the persistent volume claim.

- 8.1. Delete all resources that contain the `app=postgresql-persistent2` label.

```
[student@workstation install-storage]$ oc delete all -l app=postgresql-persistent2
service "postgresql-persistent2" deleted
deployment.apps "postgresql-persistent2" deleted
imagestream.image.openshift.io "postgresql-persistent2" deleted
```

- 8.2. Delete the persistent volume by removing the `postgresql-storage` persistent volume claim, and then return to the home directory.

```
[student@workstation install-storage]$ oc delete pvc/postgresql-storage
persistentvolumeclaim "postgresql-storage" deleted
```

- 8.3. Return to the `/home/student` directory.

```
[student@workstation install-storage]$ cd ~
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab install-storage finish
```

This concludes the guided exercise.

Summary

In this chapter, you learned:

- Red Hat OpenShift Container Platform provides two main installation methods: full-stack automation, and pre-existing infrastructure.
- Future releases are expected to add more cloud and virtualization providers, such as VMware, Red Hat Virtualization, and IBM System Z.
- An OpenShift node based on Red Hat Enterprise Linux CoreOS runs very few local services that would require direct access to a node to inspect their status. Most of the system services run as containers, the main exceptions are the CRI-O container engine and the Kubelet.
- The `oc get node`, `oc adm top`, `oc adm node-logs`, and `oc debug` commands provide troubleshooting information about OpenShift nodes.

Chapter 3

Configuring Authentication and Authorization

Goal

Configure authentication with the HTPasswd identity provider and assign roles to users and groups.

Objectives

- Configure the HTPasswd identity provider for OpenShift authentication.
- Define role-based access controls and apply permissions to users.

Sections

- Configuring Identity Providers (and Guided Exercise)
- Defining and Applying Permissions using RBAC (and Guided Exercise)

Lab

Verifying the Health of a Cluster

Configuring Identity Providers

Objectives

After completing this section, you should be able to configure the HTPasswd identity provider for OpenShift authentication.

Describing OpenShift Users and Groups

There are several OpenShift resources related to authentication and authorization. The following is a list of the primary resource types and their definitions:

User

In the OpenShift Container Platform architecture, users are entities that interact with the API server. The user resource represents an actor within the system. Assign permissions by adding roles to the user directly or to the groups of which the user is a member.

Identity

The identity resource keeps a record of successful authentication attempts from a specific user and identity provider. Any data concerning the source of the authentication is stored on the identity. Only a single user resource is associated with an identity resource.

Service Account

In OpenShift, applications can communicate with the API independently when user credentials cannot be acquired. To preserve the integrity of a regular user's credentials, credentials are not shared and service accounts are used instead. Service accounts enable you to control API access without the need to borrow a regular user's credentials.

Group

Groups represent a specific set of users. Users are assigned to one or to multiple groups. Groups are leveraged when implementing authorization policies to assign permissions to multiple users at the same time. For example, if you want to allow twenty users access to objects within a project, then it is advantageous to use a group instead of granting access to each of the users individually. OpenShift Container Platform also provides system groups or virtual groups that are provisioned automatically by the cluster.

Role

A role defines a set of permissions that enables a user to perform API operations over one or more resource types. You grant permissions to users, groups, and service accounts by assigning roles to them.

User and identity resources are usually not created in advance. They are usually created automatically by OpenShift after a successful interactive log in using OAuth.

Authenticating API Requests

Authentication and authorization are the two security layers responsible for enabling user interaction with the cluster. When a user makes a request to the API, the API associates the user with the request. The authentication layer authenticates the user. Upon successful authentication, the authorization layer decides to either honor or reject the API request. The authorization layer uses role-based access control (RBAC) policies to determine user privileges.

The OpenShift API has two methods for authenticating requests:

- OAuth Access Tokens
- X.509 Client Certificates

If the request does not present an access token or certificate, then the authentication layer assigns it the `system:anonymous` virtual user, and the `system:unauthenticated` virtual group.

Introducing the Authentication Operator

The OpenShift Container Platform provides the Authentication operator, which runs an OAuth server. The OAuth server provides OAuth access tokens to users when they attempt to authenticate to the API. An identity provider must be configured and available to the OAuth server. The OAuth server uses an identity provider to validate the identity of the requester. The server reconciles the user with the identity and creates the OAuth access token for the user. OpenShift automatically creates identity and user resources after a successful login.

Introducing Identity Providers

OpenShift OAuth server can be configured to use many identity providers. The following lists includes the more common ones:

HTPasswd

Validates user names and passwords against a secret that stores credentials generated using the `htpasswd` command.

Keystone

Enables shared authentication with an OpenStack Keystone v3 server.

LDAP

Configures the LDAP identity provider to validate user names and passwords against an LDAPv3 server, using simple bind authentication.

GitHub or GitHub Enterprise

Configures a GitHub identity provider to validate user names and passwords against GitHub or the GitHub Enterprise OAuth authentication server.

OpenID Connect

Integrates with an OpenID Connect identity provider using an Authorization Code Flow.

The OAuth custom resource must be updated with your desired identity provider. You can define multiple identity providers, of the same or different kinds, on the same OAuth custom resource.

Authenticating as a Cluster Administrator

Before you can configure an identity provider and manage users, you must access your OpenShift cluster as a cluster administrator. A newly-installed OpenShift cluster provides two ways to authenticate API requests with cluster administrator privileges:

- Use the `kubeconfig` file, which embeds an X.509 client certificate that never expires.
- Authenticate as the `kubeadmin` virtual user. Successful authentication grants an OAuth access token.

To create additional users and grant them different access levels, you must configure an identity provider and assign roles to your users.

Authenticating Using the X.509 Certificate

During installation, the OpenShift installer creates a unique `kubeconfig` file in the `auth` directory. The `kubeconfig` file contains specific details and parameters used by the CLI to connect a client to the correct API server, including an X.509 certificate.

The installation logs provide the location of the `kubeconfig` file:

```
INFO Run 'export KUBECONFIG=root/auth/kubeconfig' to manage the cluster with 'oc'.
```



Note

In the classroom environment, the utility machine stores the `kubeconfig` file at `/home/lab/ocp4/auth/kubeconfig`.

To use the `kubeconfig` file to authenticate `oc` commands, you must copy the file to your workstation and set the absolute or relative path to the `KUBECONFIG` environment variable. Then, you can run any `oc` that requires cluster administrator privileges without logging in to OpenShift.

```
[user@host ~]$ export KUBECONFIG=/home/user/auth/kubeconfig
[user@host ~]$ oc get nodes
```

As an alternative, you can use the `--kubeconfig` option of the `oc` command.

```
[user@host ~]$ oc --kubeconfig /home/user/auth/kubeconfig get nodes
```

Authenticating Using the kubeadmin Virtual User

After installation completes, OpenShift creates the `kubeadmin` virtual user. The `kubeadmin` secret in the `kube-system` namespace contains the hashed password for the `kubeadmin` user. The `kubeadmin` user has cluster administrator privileges.

The OpenShift installer dynamically generates a unique `kubeadmin` password for the cluster. The installation logs provide the `kubeadmin` credentials used to log in to the cluster. The cluster installation logs also provide log in, password, and the URL for console access.

```
...output omitted...
INFO The cluster is ready when 'oc login -u kubeadmin -p shdU_trbi_6ucX_edbu_aqop'
...output omitted...
INFO Access the OpenShift web-console here:
      https://console-openshift-console.apps.ocp4.example.com
INFO Login to the console with user: kubeadmin, password: shdU_trbi_6ucX_edbu_aqop
```



Note

In the classroom environment, the utility machine stores the password for the `kubeadmin` user in the `/home/lab/ocp4/auth/kubeadmin-password` file.

Deleting the Virtual User

After you define an identity provider, create a new user, and assign that user the `cluster-admin` role, you can remove the `kubeadmin` user credentials to improve cluster security.

```
[user@host ~]$ oc delete secret kubeadmin -n kube-system
```



Warning

If you delete the `kubeadmin` secret before you configure another user with cluster admin privileges, then the only way you can administer your cluster is using the `kubeconfig` file. If you do not have a copy of this file in a safe location, then you cannot recover administrative access to your cluster. The only alternative is destroying and reinstalling your cluster.



Warning

Do not delete the `kubeadmin` user at any time during this course. The `kubeadmin` user is essential to the course lab architecture. Deleting the `kubeadmin` user damages the lab environment, requiring that you create a new lab environment.

Configuring the HTPasswd Identity Provider

The HTPasswd identity provider validates users against a secret that contains user names and passwords generated with the `htpasswd` command from the Apache HTTP Server project. Only a cluster administrator can change the data inside the HTPasswd secret. Regular users cannot change their own passwords.

Managing users using the HTPasswd identity provider might suffice for a proof-of-concept environment with a small set of users. However, most production environments require a more powerful identity provider that integrates with the organization's identity management system.

Configuring the OAuth Custom Resource

To use the HTPasswd identity provider, the OAuth custom resource must be edited to add an entry to the `.spec.identityProviders` array:

```
apiVersion: config.openshift.io/v1
kind: OAuth
metadata:
  name: cluster
spec:
  identityProviders:
  - name: my_htpasswd_provider ❶
    mappingMethod: claim ❷
    type: HTPasswd
    htpasswd:
      fileData:
        name: htpasswd-secret ❸
```

❶ This provider name is prefixed to provider user names to form an identity name.

- 2 Controls how mappings are established between provider identities and user objects.
- 3 An existing secret containing data generated using the `htpasswd` command.

Updating the OAuth Custom Resource

To update the OAuth custom resource, use the `oc get` command to export the existing OAuth cluster resource to a file in YAML format.

```
[user@host ~]$ oc get oauth cluster -o yaml > oauth.yaml
```

Then, open the resulting file in a text editor and make the needed changes to the embedded identity provider settings.

After completing modifications and saving the file, you must apply the new custom resource using the `oc replace` command.

```
[user@host ~]$ oc replace -f oauth.yaml
```

Managing Users with the HTPasswd Identity Provider

Managing user credentials with the HTPasswd Identity Provider requires creating a temporary `htpasswd` file, making changes to the file, and applying these changes to the secret.

Creating an HTPasswd File

The `httpd-tools` package provides the `htpasswd` utility, it must be installed and available on your system.

Create the `htpasswd` file.

```
[user@host ~]$ htpasswd -c -B -b /tmp/htpasswd student redhat123
```



Important

Use the `-c` option only when creating a new file. The `-c` option replaces all file content if the file already exists.

Add or update credentials.

```
[user@host ~]$ htpasswd -b /tmp/htpasswd student redhat1234
```

Delete credentials.

```
[user@host ~]$ htpasswd -D /tmp/htpasswd student
```

Creating the HTPasswd Secret

To use the HTPasswd provider, you must create a secret that contains the `htpasswd` file data. The following example uses a secret named `htpasswd-secret`.

```
[user@host ~]$ oc create secret generic htpasswd-secret \
> --from-file htpasswd=/tmp/htpasswd -n openshift-config
```

**Important**

A secret used by the HTPasswd identity provider requires adding the `htpasswd=` prefix before specifying the path to the file.

Extracting Secret Data

When adding or removing users, an administrator cannot assume the validity of a local `htpasswd` file. Moreover, the administrator might not be on a system that has the `htpasswd` file. In a real world scenario, it would behoove the administrator to use the `oc extract` command.

By default, the `oc extract` command saves each key within a configuration map or secret as a separate file. Alternatively, all data can then be redirected to a file or displayed as standard output. To extract data from the `htpasswd-secret` secret to the `/tmp/` directory, use the following command. The `--confirm` option replaces the file if it already exists.

```
[user@host ~]$ oc extract secret/htpasswd-secret -n openshift-config \
> --to /tmp/ --confirm /tmp/htpasswd
```

Updating the HTPasswd Secret

The secret must be updated after adding, changing, or deleting users. Use the `oc set data secret` command to update a secret. Unless the file name is `htpasswd`, you must specify `htpasswd=` to update the `htpasswd` key within the secret.

The following command updates the `htpasswd-secret` secret in the `openshift-config` namespace using the content of the `/tmp/htpasswd` file.

```
[user@host ~]$ oc set data secret/htpasswd-secret \
> --from-file htpasswd=/tmp/htpasswd -n openshift-config
```

After updating the secret, the OAuth operator redeploys pods in the `openshift-authentication` namespace. Monitor the redeployment of the new OAuth pods by running:

```
[user@host ~]$ watch oc get pods -n openshift-authentication
```

Test additions, changes, or deletions to the secret after the new pods finish deploying.

Deleting Users and Identities

When a scenario occurs that requires you to delete a user, it is not sufficient to delete the user from the identity provider. The user and identity resources must also be deleted.

You must remove the password from the `htpasswd` secret, remove the user from the local `htpasswd` file, and then update the secret.

To delete the user from `htpasswd`, run the following command:


```
[user@host ~]$ htpasswd -D /tmp/htpasswd manager
```

Update the secret to remove all remnants of the user's password.

```
[user@host ~]$ oc set data secret/htpasswd-secret \
> --from-file htpasswd=/tmp/htpasswd -n openshift-config
```

Remove the user resource with the following command:

```
[user@host ~]$ oc delete user manager
user.user.openshift.io "manager" deleted
```

Identity resources include the name of the identity provider. To delete the identity resource for the `manager` user, find the resource and then delete it.

```
[user@host ~]$ oc get identities | grep manager
my_htpasswd_provider:manager    my_htpasswd_provider    manager    manager    ...

[user@host ~]$ oc delete identity my_htpasswd_provider:manager
identity.user.openshift.io "my_htpasswd_provider:manager" deleted
```

Assigning Administrative Privileges

The cluster-wide `cluster-admin` role grants cluster administration privileges to users and groups. This role enables the user to perform any action on any resources within the cluster. The following example assigns the `cluster-admin` role to the `student` user.

```
[user@host ~]$ oc adm policy add-cluster-role-to-user cluster-admin student
```



References

For more information on identity providers, refer to the Understanding identity provider configuration chapter in the Red Hat OpenShift Container Platform 4.10 Authentication and authorization documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/authentication_and_authorization/index#understanding-identity-provider

► Guided Exercise

Configuring Identity Providers

In this exercise, you will configure the HTPasswd identity provider and create users for cluster administrators.

Outcomes

You should be able to:

- Create users and passwords for HTPasswd authentication.
- Configure the Identity Provider for HTPasswd authentication.
- Assign cluster administration rights to users.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

```
[student@workstation ~]$ lab auth-provider start
```

The command ensures that the cluster API is reachable, the `httpd-utils` package is installed, and that the authentication settings are configured to the installation defaults.

Instructions

- 1. Add an entry for two `htpasswd` users, `admin` and `developer`. Assign `admin` a password of `redhat` and `developer` a password of `developer`.

- 1.1. Source the classroom configuration file that is accessible at `/usr/local/etc/ocp4.config`.

```
[student@workstation ~]$ source /usr/local/etc/ocp4.config
```



Warning

Do **not** delete the `kubeadmin` user at any time during this course. The `kubeadmin` user is essential to the course lab architecture. Deleting the `kubeadmin` user damages the lab environment, requiring that you create a new lab environment.

- 1.2. Create an HTPasswd authentication file named `htpasswd` in the `~/D0280/labs/auth-provider/` directory. Add the `admin` user with the password of `redhat`. The name of the file is arbitrary, but this exercise use the `~/D0280/labs/auth-provider/htpasswd` file.

Use the `htpasswd` command to populate the HTPasswd authentication file with the user names and encrypted passwords. The `-B` option uses `bcrypt` encryption. By default, the `htpasswd` command uses the MD5 hashing algorithm when you do not specify other algorithm.

```
[student@workstation ~]$ htpasswd -c -B -b ~/D0280/labs/auth-provider/htpasswd \
> admin redhat
Adding password for user admin
```

- 1.3. Add the `developer` user with a password of `developer` to the `~/D0280/labs/auth-provider/htpasswd` file. The password for the `developer` user is stored using MD5, because no hashing algorithm was specified when invoking the `htpasswd` command.

```
[student@workstation ~]$ htpasswd -b ~/D0280/labs/auth-provider/htpasswd \
> developer developer
Adding password for user developer
```

- 1.4. Review the contents of the `~/D0280/labs/auth-provider/htpasswd` file and verify that it includes two entries with hashed passwords: one for the `admin` user and another for the `developer` user.

```
[student@workstation ~]$ cat ~/D0280/labs/auth-provider/htpasswd
admin:$2y$05$QPuzHdl06IDkJsST.tdkZuSmgjUHV1XeYU4FjxhQrFqKL7hs2ZU16
developer:$apr1$0Nzmc1rh$yGtne1k.JX6L5s5wNa2ye.
```

- ▶ 2. Log in to OpenShift and create a secret that contains the HTPasswd users file.

- 2.1. Log in to the cluster as the `kubeadmin` user.

```
[student@workstation ~]$ oc login -u kubeadmin -p ${RHT_OCP4_KUBEADM_PASSWD} \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 2.2. Create a secret from the `~/D0280/labs/auth-provider/htpasswd` file. To use the HTPasswd identity provider, you must define a secret with a key named `htpasswd` that contains the HTPasswd user file `~/D0280/labs/auth-provider/htpasswd`.



Important

A secret that is used by the HTPasswd identity provider requires adding the `htpasswd=` prefix before specifying the path to the file.

```
[student@workstation ~]$ oc create secret generic localusers \
> --from-file htpasswd=~/D0280/labs/auth-provider/htpasswd \
> -n openshift-config
secret/localusers created
```

- 2.3. Assign the `admin` user the `cluster-admin` role.

```
[student@workstation ~]$ oc adm policy add-cluster-role-to-user \
> cluster-admin admin
Warning: User 'admin' not found
clusterrole.rbac.authorization.k8s.io/cluster-admin added: "admin"
```

**Note**

The output indicates that the `admin` user is not found and can be safely ignored.

- ▶ 3. Update the HTPasswd identity provider for the cluster so that your users can authenticate. Configure the custom resource file and update the cluster.

- 3.1. Export the existing OAuth resource to a file named `oauth.yaml` in the `~/D0280/labs/auth-provider` directory.

```
[student@workstation ~]$ oc get oauth cluster \
> -o yaml > ~/D0280/labs/auth-provider/oauth.yaml
```

**Note**

An `oauth.yaml` file containing the completed custom resource file is downloaded to `~/D0280/solutions/auth-provider` for your convenience.

- 3.2. Edit the `~/D0280/labs/auth-provider/oauth.yaml` file with your preferred text editor. You can choose the names of the `identityProviders` and `fileData` structures. For this exercise, use the `myusers` and `localusers` values respectively. The completed custom resource should match the following. Note that `htpasswd`, `mappingMethod`, `name` and `type` are at the same indentation level.

```
apiVersion: config.openshift.io/v1
kind: OAuth
...output omitted...
spec:
  identityProviders:
    - htpasswd:
        fileData:
          name: localusers
        mappingMethod: claim
        name: myusers
        type: HTPasswd
```

- 3.3. Apply the custom resource defined in the previous step.

```
[student@workstation ~]$ oc replace -f ~/D0280/labs/auth-provider/oauth.yaml
oauth.config.openshift.io/cluster replaced
```

**Note**

Pods in the `openshift-authentication` namespace will redeploy if the `oc replace` command succeeds. Provided the previously created secret was created correctly, you can log in using the HTPasswd identity provider.

- ▶ 4. Log in as `admin` and as `developer` to verify the HTPasswd user configuration.
 - 4.1. Log in to the cluster as the `admin` user to verify the HTPasswd authentication is configured correctly. The authentication operator takes some time to load the configuration changes from the previous step.

**Note**

If the authentication fails, wait a few moments and try again.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.
```

...output omitted...

- 4.2. Use the `oc get nodes` command to verify that the `admin` user has the `cluster-admin` role.

```
[student@workstation ~]$ oc get nodes
NAME      STATUS    ROLES          AGE   VERSION
master01  Ready     master,worker  2d    v1.23.5+9ce5071
master02  Ready     master,worker  2d    v1.23.5+9ce5071
master03  Ready     master,worker  2d    v1.23.5+9ce5071
```

- 4.3. Log in to the cluster as the `developer` user to verify the HTPasswd authentication is configured correctly.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.
```

...output omitted...

- 4.4. Use the `oc get nodes` command to verify that the `developer` and `admin` users do not share the same level of access.

```
[student@workstation ~]$ oc get nodes
Error from server (Forbidden): nodes is forbidden:
User "developer" cannot list resource "nodes" in API group "" at the cluster scope
```

- 4.5. Log in as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

4.6. List the current users

```
[student@workstation ~]$ oc get users
```

NAME	UID	FULL NAME	IDENTITIES
admin	31f6ccd2-6c58-47ee-978d-5e5e3c30d617		myusers:admin
developer	d4e77b0d-9740-4f05-9af5-ecfc08a85101		myusers:developer

4.7. Display the list of current identities.

```
[student@workstation ~]$ oc get identity
```

NAME	IDP NAME	IDP USER NAME	USER NAME	USER UID
myusers:admin	myusers	admin	admin	31f6ccd2-6c58-47...
myusers:developer	myusers	developer	developer	d4e77b0d-9740-4f...

- 5. As the `admin` user, create a new HTPasswd user named `manager` with a password of `redhat`.

5.1. Extract the file data from the secret to the `~/D0280/labs/auth-provider/htpasswd` file.

```
[student@workstation ~]$ oc extract secret/localusers -n openshift-config \
> --to ~/D0280/labs/auth-provider/ --confirm
/home/student/D0280/labs/auth-provider/htpasswd
```

5.2. Add an entry to your `~/D0280/labs/auth-provider/htpasswd` file for the additional user `manager` with a password of `redhat`.

```
[student@workstation ~]$ htpasswd -b ~/D0280/labs/auth-provider/htpasswd \
> manager redhat
Adding password for user manager
```

5.3. Review the contents of your `~/D0280/labs/auth-provider/htpasswd` file and verify that it includes three entries with hashed passwords: one each for the `admin`, `developer` and `manager` users.

```
[student@workstation ~]$ cat ~/D0280/labs/auth-provider/htpasswd
admin:$2y$05$QPuzHdl06IDkJssT.tdkZuSmgjUHV1XeYU4FjxhQrFqKL7hs2ZU16
developer:$apr1$0Nzmc1rh$yGtne1k.JX6L5s5wNa2ye.
manager:$apr1$CJ/tpa6a$SLhjPkIIAy755ZArTT5EH/
```

5.4. You must update the secret after adding additional users. Use the `oc set data secret` command to update the secret. If you receive a failure, then rerun the command again after a few moments as the `oauth operator` might still be reloading.

```
[student@workstation ~]$ oc set data secret/localusers \
> --from-file htpasswd=~/.DO280/labs/auth-provider/htpasswd \
> -n openshift-config
secret/localusers data updated
```

- 5.5. Wait a few moments for the authentication operator to reload, and then log in to the cluster as the `manager` user.

**Note**

If the authentication fails, wait a few moments and try again.

```
[student@workstation ~]$ oc login -u manager -p redhat
Login successful.
```

...output omitted...

- ▶ 6. Create a new project named `auth-provider`, and then verify that the `developer` user cannot access the project.

- 6.1. As the `manager` user, create a new `auth-provider` project.

```
[student@workstation ~]$ oc new-project auth-provider
Now using project "auth-provider" on server https://api.ocp4.example.com:6443".
...output omitted...
```

- 6.2. Log in as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.
```

...output omitted...

- 6.3. Attempt to delete the `auth-provider` project.

```
[student@workstation ~]$ oc delete project auth-provider
Error from server (Forbidden): projects.project.openshift.io "auth-provider"
is forbidden: User "developer" cannot delete resource "projects"
in API group "project.openshift.io" in the namespace "auth-provider"
```

- ▶ 7. Change the password for the `manager` user.

- 7.1. Log in as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.
```

...output omitted...

- 7.2. Extract the file data from the secret to the ~/D0280/labs/auth-provider/htpasswd file.

```
[student@workstation ~]$ oc extract secret/localusers -n openshift-config \
> --to ~/D0280/labs/auth-provider/ --confirm
/home/student/D0280/labs/auth-provider/htpasswd
```

- 7.3. Generate a random user password and assign it to the MANAGER_PASSWD variable.

```
[student@workstation ~]$ MANAGER_PASSWD="$(openssl rand -hex 15)"
```

- 7.4. Update the manager user to use the password stored in the MANAGER_PASSWD variable.

```
[student@workstation ~]$ htpasswd -b ~/D0280/labs/auth-provider/htpasswd \
> manager ${MANAGER_PASSWD}
Updating password for user manager
```

- 7.5. Update the secret.

```
[student@workstation ~]$ oc set data secret/localusers \
> --from-file htpasswd=~/D0280/labs/auth-provider/htpasswd \
> -n openshift-config
secret/localusers data updated
```

- 7.6. Log in as the manager user to verify the updated password.

```
[student@workstation ~]$ oc login -u manager -p ${MANAGER_PASSWD}
Login successful.

...output omitted...
```



Note

If the authentication fails, wait a few moments and try again.

- 8. Remove the manager user.

- 8.1. Log in as the admin user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

- 8.2. Extract the file data from the secret to the ~/D0280/labs/auth-provider/htpasswd file.


```
[student@workstation ~]$ oc extract secret/localusers -n openshift-config \
> --to ~/D0280/labs/auth-provider/ --confirm
/home/student/D0280/labs/auth-provider/htpasswd
```

- 8.3. Delete the manager user from the ~/D0280/labs/auth-provider/htpasswd file.

```
[student@workstation ~]$ htpasswd -D ~/D0280/labs/auth-provider/htpasswd manager
Deleting password for user manager
```

- 8.4. Update the secret.

```
[student@workstation ~]$ oc set data secret/localusers \
> --from-file htpasswd=~/D0280/labs/auth-provider/htpasswd \
> -n openshift-config
secret/localusers data updated
```

- 8.5. Delete the identity resource for the manager user.

```
[student@workstation ~]$ oc delete identity "myusers:manager"
identity.user.openshift.io "myusers:manager" deleted
```

- 8.6. Delete the user resource for the manager user.

```
[student@workstation ~]$ oc delete user manager
user.user.openshift.io manager deleted
```

- 8.7. Now, attempts to log in as the manager user fail.

```
[student@workstation ~]$ oc login -u manager -p ${MANAGER_PASSWD}
Login failed (401 Unauthorized)
Verify you have provided correct credentials.
```

- 8.8. List the current users to verify that the manager user is deleted.

```
[student@workstation ~]$ oc get users
```

NAME	UID	FULL NAME	IDENTITIES
admin	31f6ccd2-6c58-47ee-978d-5e5e3c30d617		myusers:admin
developer	d4e77b0d-9740-4f05-9af5-ecfc08a85101		myusers:developer

- 8.9. Display the list of current identities to verify that the manager identity is deleted.

```
[student@workstation ~]$ oc get identity
```

NAME	IDP NAME	IDP USER NAME	USER NAME	
myusers:admin	myusers	admin	admin	...
myusers:developer	myusers	developer	developer	...

- 8.10. Extract the secret and verify that only the users `admin` and `developer` are displayed. Using `--to -` sends the secret to STDOUT rather than saving it to a file.

```
[student@workstation ~]$ oc extract secret/localusers -n openshift-config --to -
# httpasswd
admin:$2y$05$TizWp/2ct4Edn08gmeMBIO9IXujpLqkKAJ0Nldxc/V2XYMBf6WBy
developer:$apr1$8Bc6txgb$bwHke4cGRGk9C8tQLg.hi1
```

- 9. Remove the identity provider and clean up all users.

9.1. Log in as the kubeadmin user.

```
[student@workstation ~]$ oc login -u kubeadmin -p ${RHT_OCP4_KUBEADM_PASSWD}
Login successful.

...output omitted...
```

9.2. Delete the auth-provider project.

```
[student@workstation ~]$ oc delete project auth-provider
project.project.openshift.io "auth-provider" deleted
```

9.3. Edit the resource in place to remove the identity provider from OAuth:

```
[student@workstation ~]$ oc edit oauth
```

Delete all the lines under `spec:`, and then append `{}` after `spec:`. Leave all the other information in the file unchanged. Your `spec:` line should match the following:

```
...output omitted...
spec: {}
```

Save your changes, and then verify that the `oc edit` command applied your changes:

```
oauth.config.openshift.io/cluster edited
```

9.4. Delete the `localusers` secret from the `openshift-config` namespace.

```
[student@workstation ~]$ oc delete secret localusers -n openshift-config
secret "localusers" deleted
```

9.5. Delete all user resources.

```
[student@workstation ~]$ oc delete user --all
user.user.openshift.io "admin" deleted
user.user.openshift.io "developer" deleted
```

9.6. Delete all identity resources.

```
[student@workstation ~]$ oc delete identity --all
identity.user.openshift.io "myusers:admin" deleted
identity.user.openshift.io "myusers:developer" deleted
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab auth-provider finish
```

This concludes the guided exercise.

Defining and Applying Permissions using RBAC

Objectives

After completing this section, you should be able to define role-based access controls and apply permissions to users.

Role-based Access Control (RBAC)

Role-based access control (RBAC) is a technique for managing access to resources in a computer system. In Red Hat OpenShift, RBAC determines if a user can perform certain actions within the cluster or project. There are two types of roles that can be used depending on the user's level of responsibility: cluster and local.



Note

Authorization is a separate step from authentication.

Authorization Process

The authorization process is managed by rules, roles, and bindings.

RBAC Object	Description
Rule	Allowed actions for objects or groups of objects.
Role	Sets of rules. Users and groups can be associated with multiple roles.
Binding	Assignment of users or groups to a role.

RBAC Scope

Red Hat OpenShift Container Platform (RHOCP) defines two groups of roles and bindings depending on the user's scope and responsibility: cluster roles and local roles.

Role Level	Description
Cluster Role	Users or groups with this role level can manage the OpenShift cluster.
Local Role	Users or groups with this role level can only manage elements at a project level.



Note

Cluster role bindings take precedence over local role bindings.

Managing RBAC Using the CLI

Cluster administrators can use the `oc adm policy` command to both add and remove cluster roles and namespace roles.

To add a cluster role to a user, use the `add-cluster-role-to-user` subcommand:

```
[user@host ~]$ oc adm policy add-cluster-role-to-user cluster-role username
```

For example, to change a regular user to a cluster administrator, use the following command:

```
[user@host ~]$ oc adm policy add-cluster-role-to-user cluster-admin username
```

To remove a cluster role from a user, use the `remove-cluster-role-from-user` subcommand:

```
[user@host ~]$ oc adm policy remove-cluster-role-from-user cluster-role username
```

For example, to change a cluster administrator to a regular user, use the following command:

```
[user@host ~]$ oc adm policy remove-cluster-role-from-user cluster-admin username
```

Rules are defined by an action and a resource. For example, the `create user` rule is part of the `cluster-admin` role.

You can use the `oc adm policy who-can` command to determine if a user can execute an action on a resource. For example:

```
[user@host ~]$ oc adm policy who-can delete user
```

Default Roles

OpenShift ships with a set of default cluster roles that can be assigned locally or to the entire cluster. You can modify these roles for fine-grained access control to OpenShift resources, but additional steps are required that are outside the scope of this course.

Default roles	Description
<code>admin</code>	Users with this role can manage all project resources, including granting access to other users to access the project.
<code>basic-user</code>	Users with this role have read access to the project.
<code>cluster-admin</code>	Users with this role have superuser access to the cluster resources. These users can perform any action on the cluster, and have full control of all projects.
<code>cluster-status</code>	Users with this role can get cluster status information.

Default roles	Description
<code>edit</code>	Users with this role can create, change, and delete common application resources from the project, such as services and deployments. These users cannot act on management resources such as limit ranges and quotas, and cannot manage access permissions to the project.
<code>self-provisioner</code>	Users with this role can create new projects. This is a cluster role, not a project role.
<code>view</code>	Users with this role can view project resources, but cannot modify project resources.

The `admin` role gives a user access to project resources such as quotas and limit ranges, and also the ability to create new applications. The `edit` role gives a user sufficient access to act as a developer inside the project, but working under the constraints configured by a project administrator.

Project administrators can use the `oc policy` command to add and remove namespace roles.

Add a specified role to a user with the `add-role-to-user` subcommand. For example:

```
[user@host ~]$ oc policy add-role-to-user role-name username -n project
```

For example, to add the user `dev` to the role `basic-user` in the `wordpress` project:

```
[user@host ~]$ oc policy add-role-to-user basic-user dev -n wordpress
```

User Types

Interaction with OpenShift Container Platform is associated with a user. An OpenShift Container Platform user object represents a user who can be granted permissions in the system by adding roles to that user or to a user's group via `rolebindings`.

Regular users

Most interactive OpenShift Container Platform users are regular users, represented with the `User` object. This type of user represents a person that has been allowed access to the platform.

Examples of regular users include `user1` and `user2`.

System users

Many system users are created automatically when the infrastructure is defined, mainly for the purpose of enabling the infrastructure to securely interact with the API. System users include a cluster administrator (with access to everything), a per-node user, users for routers and registries, and various others. An anonymous system user is used by default for unauthenticated requests.

Examples of system users include: `system:admin`, `system:openshift-registry`, and `system:node:node1.example.com`.

Service accounts

These are special system users associated with projects. Some service account users are created automatically when the project is first created. Project administrators can create more

for the purpose of defining access to the contents of each project. Service accounts are often used to give extra privileges to pods or deployments. Service accounts are represented with the `ServiceAccount` object.

Examples of service account users include `system:serviceaccount:default:deployer` and `system:serviceaccount:foo:builder`.

Every user must authenticate before they can access OpenShift Container Platform. API requests with no authentication or invalid authentication are authenticated as requests by the anonymous system user. After successful authentication, policy determines what the user is authorized to do.



References

For more information about Kubernetes namespaces refer to the **Kubernetes Documentation**

<https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/>

For more information about RBAC, refer to the *Using RBAC to define and apply permissions* chapter in the Red Hat OpenShift Container Platform 4.10 *Authentication and authorization* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/authentication_and_authorization/index#using-rbac

For more information about groups, refer to the *Understanding authentication* chapter in the Red Hat OpenShift Container Platform 4.10 *Authentication and authorization* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/authentication_and_authorization/index#understanding-authentication

► Guided Exercise

Defining and Applying Permissions using RBAC

In this exercise, you will define role-based access controls and apply permissions to users.

Outcomes

You should be able to:

- Remove project creation privileges from users who are not OpenShift cluster administrators.
- Create OpenShift groups and add members to these groups.
- Create a project and assign project administration privileges to the project.
- As a project administrator, assign read and write privileges to different groups of users.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates some HTPasswd users for the exercise.

```
[student@workstation ~]$ lab auth-rbac start
```

Instructions

- 1. Log in to the OpenShift cluster and determine which cluster role bindings assign the `self-provisioner` cluster role.

- 1.1. Log in to the cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. List all cluster role bindings that reference the `self-provisioner` cluster role.

```
[student@workstation ~]$ oc get clusterrolebinding -o wide | \
> grep -E 'NAME|self-provisioner'
NAME                                ROLE                                ...
self-provisioners                  ... ClusterRole/self-provisioner ...
```


- ▶ 2. Remove the privilege to create new projects from all users who are not cluster administrators by deleting the `self-provisioner` cluster role from the `system:authenticated:oauth` virtual group.
 - 2.1. Confirm that the `self-provisioners` cluster role binding that you found in the previous step assigns the `self-provisioner` cluster role to the `system:authenticated:oauth` group.

```
[student@workstation ~]$ oc describe clusterrolebindings self-provisioners
Name:          self-provisioners
Labels:        <none>
Annotations:   rbac.authorization.kubernetes.io/autoupdate: true
Role:
  Kind:  ClusterRole
  Name:  self-provisioner
Subjects:
  Kind  Name                               Namespace
  ----  ---                               -
  Group system:authenticated:oauth
```

- 2.2. Remove the `self-provisioner` cluster role from the `system:authenticated:oauth` virtual group, which deletes the `self-provisioners` role binding.

```
[student@workstation ~]$ oc adm policy remove-cluster-role-from-group \
> self-provisioner system:authenticated:oauth
Warning: Your changes may get lost whenever a master is restarted, unless you
prevent reconciliation of this rolebinding using the following command:
oc annotate clusterrolebinding.rbac self-provisioner
'rbac.authorization.kubernetes.io/autoupdate=false' --overwrite
clusterrole.rbac.authorization.k8s.io/self-provisioner removed:
"system:authenticated:oauth"
```

**Note**

You can safely ignore the warning about your changes being lost.

- 2.3. Verify that the role has been removed from the group. The cluster role binding `self-provisioners` should not exist.

```
[student@workstation ~]$ oc describe clusterrolebindings self-provisioners
Error from server (NotFound): clusterrolebindings.rbac.authorization.k8s.io "self-provisioners" not found
```

- 2.4. Determine if any other cluster role bindings reference the `self-provisioner` cluster role.

```
[student@workstation ~]$ oc get clusterrolebinding -o wide | \
> grep -E 'NAME|self-provisioner'
NAME                               ROLE    ...
```

- 2.5. Log in as the `leader` user with a password of `redhat`.

```
[student@workstation ~]$ oc login -u leader -p redhat
Login successful.

...output omitted...
```

2.6. Try to create a project, the operation should fail.

```
[student@workstation ~]$ oc new-project test
Error from server (Forbidden): You may not request a new project via this API.
```

► **3.** Create a project and add project administration privileges to the leader user.

3.1. Log in as the admin user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

3.2. Create the auth-rbac project.

```
[student@workstation ~]$ oc new-project auth-rbac
Now using project "auth-rbac" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

3.3. Grant project administration privileges to the leader user on the auth-rbac project.

```
[student@workstation ~]$ oc policy add-role-to-user admin leader
clusterrole.rbac.authorization.k8s.io/admin added: "leader"
```

► **4.** Create the dev-group and qa-group groups and add their respective members.

4.1. Create a group called dev-group.

```
[student@workstation ~]$ oc adm groups new dev-group
group.user.openshift.io/dev-group created
```

4.2. Add the developer user to dev-group.

```
[student@workstation ~]$ oc adm groups add-users dev-group developer
group.user.openshift.io/dev-group added: "developer"
```

4.3. Create a second group called qa-group.

```
[student@workstation ~]$ oc adm groups new qa-group
group.user.openshift.io/qa-group created
```

4.4. Add the qa-engineer user to qa-group.

```
[student@workstation ~]$ oc adm groups add-users qa-group qa-engineer
group.user.openshift.io/qa-group added: "qa-engineer"
```

4.5. Review all existing OpenShift groups to verify that they have the correct members.

```
[student@workstation ~]$ oc get groups
NAME      USERS
dev-group developer
qa-group  qa-engineer
```

- 5. As the leader user, assign write privileges for dev-group and read privileges for qa-group to the auth-rbac project.

5.1. Log in as the leader user.

```
[student@workstation ~]$ oc login -u leader -p redhat
Login successful.

...output omitted...

Using project "auth-rbac".
```

5.2. Add write privileges to dev-group on the auth-rbac project.

```
[student@workstation ~]$ oc policy add-role-to-group edit dev-group
clusterrole.rbac.authorization.k8s.io/edit added: "dev-group"
```

5.3. Add read privileges to qa-group on the auth-rbac project.

```
[student@workstation ~]$ oc policy add-role-to-group view qa-group
clusterrole.rbac.authorization.k8s.io/view added: "qa-group"
```

5.4. Review all role bindings on the auth-rbac project to verify that they assign roles to the correct groups and users. The following output omits default role bindings assigned by OpenShift to service accounts.

```
[student@workstation ~]$ oc get rolebindings -o wide
NAME      ROLE              AGE   USERS   GROUPS   ...
admin     ClusterRole/admin  58s   admin
admin-0    ClusterRole/admin  51s   leader
edit       ClusterRole/edit   12s
...output omitted...
view       ClusterRole/view   8s    qa-group
```

- 6. As the developer user, deploy an Apache HTTP Server to prove that the developer user has write privileges in the project. Also try to grant write privileges to the qa-engineer user to prove that the developer user has no project administration privileges.

6.1. Log in as the developer user.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.

...output omitted...

Using project "auth-rbac".
```

6.2. Deploy an Apache HTTP Server using the standard image stream from OpenShift.

```
[student@workstation ~]$ oc new-app --name httpd httpd:2.4
...output omitted...
--> Creating resources ...
    imagestreamtag.image.openshift.io "httpd:2.4" created
    deployment.apps "httpd" created
    service "httpd" created
--> Success
...output omitted...
```

6.3. Try to grant write privileges to the qa-engineer user, the operation should fail.

```
[student@workstation ~]$ oc policy add-role-to-user edit qa-engineer
Error from server (Forbidden): rolebindings.rbac.authorization.k8s.io is
  forbidden: User "developer" cannot list resource "rolebindings" in API group
  "rbac.authorization.k8s.io" in the namespace "auth-rbac"
```

► 7. Verify that the qa-engineer user only has read privileges on the httpd application.

7.1. Log in as the qa-engineer user.

```
[student@workstation ~]$ oc login -u qa-engineer -p redhat
Login successful.

...output omitted...

Using project "auth-rbac".
```

7.2. Attempt to scale the httpd application, the operation should fail.

```
[student@workstation ~]$ oc scale deployment httpd --replicas 3
Error from server (Forbidden): deployments.apps "httpd" is forbidden: User "qa-
engineer" cannot patch resource "deployments/scale" in API group "apps" in the
  namespace "auth-rbac"
```

► 8. Restore project creation privileges to all users.

8.1. Log in as the admin user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

- 8.2. Restore project creation privileges for all users by recreating the `self-provisioners` cluster role binding created by the OpenShift installer.

```
[student@workstation ~]$ oc adm policy add-cluster-role-to-group \
> --rolebinding-name self-provisioners \
> self-provisioner system:authenticated:oauth
Warning: Group 'system:authenticated:oauth' not found
clusterrole.rbac.authorization.k8s.io/self-provisioner added:
"system:authenticated:oauth"
```

**Note**

You can safely ignore the warning that the group was not found.

Finish

On the workstation machine, run the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab auth-rbac finish
```

This concludes the guided exercise.

► Lab

Configuring Authentication and Authorization

Performance Checklist

In this lab, you will configure the HTPasswd identity provider, create groups, and assign roles to users and groups.

Outcomes

You should be able to:

- Create users and passwords for HTPasswd authentication.
- Configure the Identity Provider for HTPasswd authentication.
- Assign cluster administration rights to users.
- Remove the ability to create projects at the cluster level.
- Create groups and add users to groups.
- Manage user privileges in projects by granting privileges to groups.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

```
[student@workstation ~]$ lab auth-review start
```

The command ensures that the cluster API is reachable, the `httpd-util` package is installed, and the authentication settings are configured to installation defaults.

Instructions

1. Update the existing `~/DO280/labs/auth-review/tmp_users` HTPasswd authentication file to remove the `analyst` user. Ensure that the `tester` and `leader` users in the file use a password of `L@bR3v!ew`. Add two new entries to the file for the users `admin` and `developer`. Use `L@bR3v!ew` as the password for each new user.
2. Log in to your OpenShift cluster as the `kubeadmin` user using the `RHT_OCP4_KUBEADM_PASSWD` variable defined in the `/usr/local/etc/ocp4.config` file as the password. Configure your cluster to use the HTPasswd identity provider using the user names and passwords defined in the `~/DO280/labs/auth-review/tmp_users` file.
3. Make the `admin` user a cluster administrator. Log in as both `admin` and as `developer` to verify HTPasswd user configuration and cluster privileges.
4. As the `admin` user, remove the ability to create projects cluster wide.

5. Create a group named `managers`, and add the `leader` user to the group. Grant project creation privileges to the `managers` group. As the `leader` user, create the `auth-review` project.
6. Create a group named `developers` and grant edit privileges on the `auth-review` project. Add the `developer` user to the group.
7. Create a group named `qa` and grant view privileges on the `auth-review` project. Add the `tester` user to the group.

Evaluation

On the `workstation` machine, run the `lab` command to grade your work. Correct any reported failures and rerun the script until successful.

```
[student@workstation ~]$ lab auth-review grade
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab auth-review finish
```

This concludes the lab.

► Solution

Configuring Authentication and Authorization

Performance Checklist

In this lab, you will configure the HTPasswd identity provider, create groups, and assign roles to users and groups.

Outcomes

You should be able to:

- Create users and passwords for HTPasswd authentication.
- Configure the Identity Provider for HTPasswd authentication.
- Assign cluster administration rights to users.
- Remove the ability to create projects at the cluster level.
- Create groups and add users to groups.
- Manage user privileges in projects by granting privileges to groups.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

```
[student@workstation ~]$ lab auth-review start
```

The command ensures that the cluster API is reachable, the `httpd-util` package is installed, and the authentication settings are configured to installation defaults.

Instructions

1. Update the existing `~/D0280/labs/auth-review/tmp_users` HTPasswd authentication file to remove the `analyst` user. Ensure that the `tester` and `leader` users in the file use a password of `L@bR3v!ew`. Add two new entries to the file for the users `admin` and `developer`. Use `L@bR3v!ew` as the password for each new user.
 - 1.1. Remove the `analyst` user from the `~/D0280/labs/auth-review/tmp_users` HTPasswd authentication file.

```
[student@workstation ~]$ htpasswd -D ~/D0280/labs/auth-review/tmp_users analyst
Deleting password for user analyst
```

- 1.2. Update the entries for the `tester` and `leader` users so that they use a password of `L@bR3v!ew`. Add entries for the `admin` and `developer` users using a password of `L@bR3v!ew`.


```
[student@workstation ~]$ for NAME in tester leader admin developer
> do
> htpasswd -b ~/D0280/labs/auth-review/tmp_users ${NAME} 'L@bR3v!ew'
> done
Updating password for user tester
Updating password for user leader
Adding password for user admin
Adding password for user developer
```

- 1.3. Review the contents of the ~/D0280/labs/auth-review/tmp_users file. It does not contain a line for the analyst user. It includes two new entries with hashed passwords for the admin and developer users.

```
[student@workstation ~]$ cat ~/D0280/labs/auth-review/tmp_users
tester:$apr1$0eqhKgbU$Dwd0CB4IumhasaRuEr6hp0
leader:$apr1$.EB5IXlu$FDV.Av16nj10CMzgolScr/
admin:$apr1$ItcCncDS$xFCUjQGTsXAup00KQfmw0
developer:$apr1$D8F1Hren$izDhAwq5DRjUHPv0i7FHN.
```

2. Log in to your OpenShift cluster as the kubeadmin user using the RHT_OCP4_KUBEADM_PASSWD variable defined in the /usr/local/etc/ocp4.config file as the password. Configure your cluster to use the HTPasswd identity provider using the user names and passwords defined in the ~/D0280/labs/auth-review/tmp_users file.

- 2.1. Source the classroom configuration file that is accessible at /usr/local/etc/ocp4.config.

```
[student@workstation ~]$ source /usr/local/etc/ocp4.config
```

- 2.2. Log in to the cluster as the kubeadmin user.

```
[student@workstation ~]$ oc login -u kubeadmin -p ${RHT_OCP4_KUBEADM_PASSWD} \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 2.3. Create a secret named auth-review using the ~/D0280/labs/auth-review/tmp_users file.

```
[student@workstation ~]$ oc create secret generic auth-review \
> --from-file htpasswd=/home/student/D0280/labs/auth-review/tmp_users \
> -n openshift-config
secret/auth-review created
```

- 2.4. Export the existing OAuth resource to ~/D0280/labs/auth-review/oauth.yaml.

```
[student@workstation ~]$ oc get oauth cluster \
> -o yaml > ~/D0280/labs/auth-review/oauth.yaml
```

- 2.5. Edit the `~/D0280/labs/auth-review/oauth.yaml` file to replace the `spec: {}` line with the following bold lines. Note that `htpasswd`, `mappingMethod`, `name` and `type` are at the same indentation level.

```
apiVersion: config.openshift.io/v1
kind: OAuth
...output omitted...
spec:
  identityProviders:
  - htpasswd:
      fileName:
        name: auth-review
      mappingMethod: claim
      name: htpasswd
      type: HTPasswd
```

**Note**

For convenience, the `~/D0280/solutions/auth-review/oauth.yaml` file contains a minimal version of the OAuth configuration with the specified customizations.

- 2.6. Apply the customized resource defined in the previous step.

```
[student@workstation ~]$ oc replace -f ~/D0280/labs/auth-review/oauth.yaml
oauth.config.openshift.io/cluster replaced
```

- 2.7. A successful update to the `oauth/cluster` resource recreates the `oauth-openshift` pods in the `openshift-authentication` namespace.

```
[student@workstation ~]$ watch oc get pods -n openshift-authentication
```

Wait until the new `oauth-openshift` pods are ready and running and the previous pods have terminated.

```
Every 2.0s: oc get pods -n openshift-authentication    ...

NAME                                READY   STATUS    RESTARTS   AGE
oauth-openshift-6755d8795-h8bgv     1/1     Running   0           34s
oauth-openshift-6755d8795-rk4m6     1/1     Running   0           38s
oauth-openshift-6755d8795-2859w     1/1     Running   0           53s
```

Press `Ctrl+C` to exit the `watch` command.

3. Make the `admin` user a cluster administrator. Log in as both `admin` and as `developer` to verify HTPasswd user configuration and cluster privileges.

- 3.1. Assign the `admin` user the `cluster-admin` role.

```
[student@workstation ~]$ oc adm policy add-cluster-role-to-user \
> cluster-admin admin
Warning: User 'admin' not found
clusterrole.rbac.authorization.k8s.io/cluster-admin added: "admin"
```

**Note**

You can safely ignore the warning that the `admin` user is not found.

- 3.2. Log in to the cluster as the `admin` user to verify that `HTPasswd` authentication was configured correctly.

```
[student@workstation ~]$ oc login -u admin -p 'L@bR3v!ew'
Login successful.

...output omitted...
```

- 3.3. Use `oc get nodes` command to verify the `admin` user has the `cluster-admin` role. The names of the nodes from your cluster might be different.

```
[student@workstation ~]$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
master01	Ready	master,worker	4d	v1.23.5+9ce5071
master02	Ready	master,worker	4d	v1.23.5+9ce5071
master03	Ready	master,worker	4d	v1.23.5+9ce5071

- 3.4. Log in to the cluster as the `developer` user to verify the `HTPasswd` authentication is configured correctly.

```
[student@workstation ~]$ oc login -u developer -p 'L@bR3v!ew'
Login successful.

...output omitted...
```

- 3.5. Use the `oc get nodes` command to verify that the `developer` user does not have cluster administration privileges.

```
[student@workstation ~]$ oc get nodes
Error from server (Forbidden): nodes is forbidden: User "developer" cannot list resource "nodes" in API group "" at the cluster scope
```

4. As the `admin` user, remove the ability to create projects cluster wide.

- 4.1. Log in to the cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p 'L@bR3v!ew'
Login successful.

...output omitted...
```

- 4.2. Remove the `self-provisioner` cluster role from the `system:authenticated:oauth` virtual group.

```
[student@workstation ~]$ oc adm policy remove-cluster-role-from-group \
> self-provisioner system:authenticated:oauth
clusterrole.rbac.authorization.k8s.io/self-provisioner removed:
"system:authenticated:oauth"
```

**Note**

You can safely ignore the warning about your changes being lost.

5. Create a group named `managers`, and add the `leader` user to the group. Grant project creation privileges to the `managers` group. As the `leader` user, create the `auth-review` project.

- 5.1. Create a group named `managers`.

```
[student@workstation ~]$ oc adm groups new managers
group.user.openshift.io/managers created
```

- 5.2. Add the `leader` user to the `managers` group.

```
[student@workstation ~]$ oc adm groups add-users managers leader
group.user.openshift.io/managers added: "leader"
```

- 5.3. Assign the `self-provisioner` cluster role to the `managers` group.

```
[student@workstation ~]$ oc adm policy add-cluster-role-to-group \
> self-provisioner managers
clusterrole.rbac.authorization.k8s.io/self-provisioner added: "managers"
```

- 5.4. As the `leader` user, create the `auth-review` project.

```
[student@workstation ~]$ oc login -u leader -p 'L@bR3v!ew'
Login successful.
```

...output omitted...

The user who creates a project is automatically assigned the `admin` role on the project.

```
[student@workstation ~]$ oc new-project auth-review
Now using project "auth-review" on server "https://api.ocp4.example.com:6443".
```

...output omitted...

6. Create a group named `developers` and grant edit privileges on the `auth-review` project. Add the `developer` user to the group.

- 6.1. Log in to the cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p 'L@bR3v!ew'
Login successful.

...output omitted...
```

6.2. Create a group named `developers`.

```
[student@workstation ~]$ oc adm groups new developers
group.user.openshift.io/developers created
```

6.3. Add the `developer` user to the `developers` group.

```
[student@workstation ~]$ oc adm groups add-users developers developer
group.user.openshift.io/developers added: "developer"
```

6.4. Grant edit privileges to the `developers` group on the `auth-review` project.

```
[student@workstation ~]$ oc policy add-role-to-group edit developers
clusterrole.rbac.authorization.k8s.io/edit added: "developers"
```

7. Create a group named `qa` and grant view privileges on the `auth-review` project. Add the `tester` user to the group.

7.1. Create a group named `qa`.

```
[student@workstation ~]$ oc adm groups new qa
group.user.openshift.io/qa created
```

7.2. Add the `tester` user to the `qa` group.

```
[student@workstation ~]$ oc adm groups add-users qa tester
group.user.openshift.io/qa added: "tester"
```

7.3. Grant view privileges to the `qa` group on the `auth-review` project.

```
[student@workstation ~]$ oc policy add-role-to-group view qa
clusterrole.rbac.authorization.k8s.io/view added: "qa"
```

Evaluation

On the `workstation` machine, run the `lab` command to grade your work. Correct any reported failures and rerun the script until successful.

```
[student@workstation ~]$ lab auth-review grade
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab auth-review finish
```

This concludes the lab.

Summary

In this chapter, you learned:

- A newly-installed OpenShift cluster provides two authentication methods that grant administrative access: the `kubeconfig` file and the `kubeadmin` virtual user.
- The `HTPasswd` identity provider authenticates users against credentials stored in a secret. The name of the secret, and other settings for the identity provider, are stored inside the `OAuth` custom resource.
- To manage user credentials using the `HTPasswd` identity provider, you must extract data from the secret, change that data using the `htpasswd` command, and then apply the data back to the secret.
- Creating OpenShift users requires valid credentials, managed by an identity provider, plus user and identity resources.
- Deleting OpenShift users requires deleting their credentials from the identity provider, and also deleting their user and identity resources.
- OpenShift uses role-based access control (RBAC) to control user actions. A role is a collection of rules that govern interaction with OpenShift resources. Default roles exist for cluster administrators, developers, and auditors.
- To control user interaction, assign a user to one or more roles. A role binding contains all of the associations of a role to users and groups.
- To grant a user cluster administrator privileges, assign the `cluster-admin` role to that user.

Chapter 4

Configuring Application Security

Goal

Restrict permissions of applications using security context constraints and protect access credentials using secrets.

Objectives

- Create and apply secrets to manage sensitive information and share secrets between applications.
- Create service accounts and apply permissions, and manage security context constraints.

Sections

- Managing Sensitive Information with Secrets (and Guided Exercise)
- Controlling Application Permissions with Security Context Constraints (and Guided Exercise)

Lab

Configuring Application Security

Managing Sensitive Information with Secrets

Objectives

After completing this section, you should be able to create and apply secrets to manage sensitive information and share secrets between applications.

Secrets Overview

Modern applications are designed to loosely couple code, configuration, and data. Configuration files and data are not hard-coded as part of the software. Instead, the software loads configuration and data from an external source. This enables deploying an application to different environments without requiring a change to the application source code.

Applications often require access to sensitive information. For example, a back-end web application requires access to database credentials to perform a database query. Kubernetes and OpenShift use secret resources to hold sensitive information, such as:

- Passwords.
- Sensitive configuration files.
- Credentials to an external resource, such as an SSH key or OAuth token.

A secret can store any type of data. Data in a secret is Base64-encoded, not stored in plain text. Secret data is not encrypted; you can decode the secret from Base64 format to access the original data.

Although secrets can store any type of data, Kubernetes and OpenShift support different types of secrets. Different types of secret resources exist, including service account tokens, SSH keys, and TLS certificates. When you store information in a specific secret resource type, Kubernetes validates that the data conforms to the type of secret.



Note

You can encrypt the Etcd database, although this is not the default setting. When enabled, Etcd encrypts the following resources: secrets, configuration maps, routes, OAuth access tokens, and OAuth authorize tokens. Enabling Etcd encryption is outside the scope of this class.

Features of Secrets

The main features of secrets include:

- Secret data can be shared within a project namespace.
- Secret data is referenced independently of secret definition. Administrators can create and manage a secret resource that other team members can reference in their deployment configurations.
- Secret data is injected into pods when OpenShift creates a pod. You can expose a secret as an environment variable or as a mounted file in the pod.

- If the value of a secret changes during pod execution, the secret data in the pod does not update. After a secret value changes, you must create new pods to inject the new secret data.
- Any secret data that OpenShift injects into a pod is ephemeral. If OpenShift exposes sensitive data to a pod as environment variables, then those variables are destroyed when the pod is destroyed.

Secret data volumes are backed by temporary file storage. If a secret is mounted as a file in the pod, then the file is also destroyed when the pod is destroyed. A stopped pod does not contain secret data.

Use Cases for Secrets

Two primary use cases for secrets are storing credentials and securing communication between services.

Credentials

Store sensitive information, such as passwords and user names, in a secret.

If an application expects to read sensitive information from a file, then you mount the secret as a data volume to the pod. The application can read the secret as an ordinary file to access the sensitive information. Some databases, for example, read credentials from a file to authenticate users.

Some applications use environment variables to read configuration and sensitive data. You can link secret variables to pod environment variables in a deployment configuration.

Transport Layer Security (TLS) and Key Pairs

Use a TLS certificate and key to secure communication to a pod. A TLS secret stores the certificate as `tls.crt` and the certificate key as `tls.key`. Developers can mount the secret as a volume and create a pass through route to the application.

Creating a Secret

If a pod requires access to sensitive information, then create a secret for the information before you deploy the pod. Use one of the following commands to create a secret:

- Create a generic secret containing key-value pairs from literal values typed on the command line:

```
[user@host ~]$ oc create secret generic secret_name \
> --from-literal key1=secret1 \
> --from-literal key2=secret2
```

- Create a generic secret using key names specified on the command line and values from files:

```
[user@host ~]$ oc create secret generic ssh-keys \
> --from-file id_rsa=/path-to/id_rsa \
> --from-file id_rsa.pub=/path-to/id_rsa.pub
```

- Create a TLS secret specifying a certificate and the associated key:

```
[user@host ~]$ oc create secret tls secret-tls \
> --cert /path-to-certificate --key /path-to-key
```

Exposing Secrets to Pods

To expose a secret to a pod, first create the secret. Assign each piece of sensitive data to a key. After creation, the secret contains key-value pairs. The following command creates a generic secret named `demo-secret` with two keys: `user` with the `demo-user` value and `root_password` with the `zT1KTgk` value.

```
[user@host ~]$ oc create secret generic demo-secret \
> --from-literal user=demo-user
> --from-literal root_password=zT1KTgk
```

Secrets as Pod Environment Variables

Consider a database application that reads the database administrator password from the `MYSQL_ROOT_PASSWORD` environment variable. Modify the environment variables section of the deployment configuration to use values from the secret:

```
env:
- name: MYSQL_ROOT_PASSWORD ❶
  valueFrom:
    secretKeyRef: ❷
      name: demo-secret ❸
      key: root_password ❹
```

- ❶ The environment variable name in the pod that contains data from a secret.
- ❷ The `secretKeyRef` key expects a secret. Use the `configMapKeyRef` key for configuration maps.
- ❸ The name of the secret that contains the desired sensitive information.
- ❹ The name of the key that contains the sensitive information in the secret.

You can also use the `oc set env deployment/demo --from secret/demo-secret \` command to set application environment variables from either secrets or configuration maps. In some cases, the names of the keys can be modified to match the names of environment variables by using the `--prefix` option. In the following example, the `user` key sets the `MYSQL_USER` environment variable, and the `root_password` key sets the `MYSQL_ROOT_PASSWORD` environment variable. If the key name is lowercase, then the corresponding environment variable is converted to uppercase.

```
[user@host ~]$ oc set env deployment/demo --from secret/demo-secret \
> --prefix MYSQL_
```

Secrets as Files in a Pod

A secret can be mounted to a directory within a pod. A file is created for each key in the secret using the name of the key. The content of each file is the decoded value of the secret. The following command shows how to mount secrets in a pod:

```
[user@host ~]$ oc set volume deployment/demo \ ❶
> --add --type secret \ ❷
> --secret-name demo-secret \ ❸
> --mount-path /app-secrets ❹
```

- ❶ Modify the volume configuration in the `demo` deployment.
- ❷ Add a new volume from a secret. Configuration maps can also be mounted as volumes.
- ❸ Use the `demo-secret` secret.
- ❹ Make the secret data available in the `/app-secrets` directory in the pod. The content of the `/app-secrets/user` file is `demo-user`. The content of the `/app-secrets/root_password` file is `zT1KTgk`.

The container image can dictate the location of the mount point and the expected file names. For example, a container image running NGINX can specify the SSL certificate location and the SSL certificate key location in the `/etc/nginx/nginx.conf` configuration file. If the expected files are not found, then the container might fail to run.



Important

If the mount point already exists in the pod, then any existing files at the mount point are obscured by the mounted secret. The existing files are not visible and are not accessible.

Configuration Map Overview

Similar to secrets, configuration maps decouple configuration information from container images. Unlike secrets, the information contained in configuration maps does not require protection. You can use the data in a configuration map to set environment variables in the container image, or mount the configuration map as a volume within the container image.

Container images do not need to be rebuilt when a secret or a configuration map changes. New pods use the updated secrets and configuration maps. You can delete pods using the older secrets and configuration maps.

The syntax for creating a configuration map closely matches the syntax for creating a secret. Key-value pairs can be entered on the command line or the content of a file can be used as the value of a specified key. The following command shows how to create a configuration map:

```
[user@host ~]$ oc create configmap my-config \
> --from-literal key1=config1 --from-literal key2=config2
```

Updating Secrets and Configuration Maps

Secrets and configuration maps occasionally require updates. Use the `oc extract` command to ensure you have the latest data. Save the data to a specific directory using the `--to` option. Each key in the secret or configuration map creates a file with the same name as the key. The content of each file is the value of the associated key. If you run the `oc extract` command more than once, then use the `--confirm` option to overwrite the existing files.

```
[user@host ~]$ oc extract secret/htpasswd-ppklq -n openshift-config \
> --to /tmp/ --confirm
```

After updating the locally saved files, use the `oc set data` command to update the secret or configuration map. For each key that requires an update, specify the name of a key and the associated value. If a file contains the value, use the `--from-file` option.

In the previous `oc extract` example, the `htpasswd-ppklq` secret contained only one key named `htpasswd`. Using the `oc set data` command, you can explicitly specify the `htpasswd` key name using `--from-file htpasswd=/tmp/htpasswd`. If the key name is not specified, the file name is used as the key name.

```
[user@host ~]$ oc set data secret/htpasswd-ppklq -n openshift-config \
> --from-file /tmp/htpasswd
```



References

For more information on secrets, refer to the *Understanding secrets* section in the *Working with pods* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-pods-secrets-about_nodes-pods-secrets

For more information on Etcd encryption, refer to the *Encrypting Etcd data* chapter in the Red Hat OpenShift Container Platform 4.10 *Security* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/security_and_compliance/index#encrypting-etcd

► Guided Exercise

Managing Sensitive Information with Secrets

In this exercise, you will manage information using secrets.

Outcomes

You should be able to:

- Manage secrets and use them to initialize environment variables in applications.
- Use secrets for a MySQL database application.
- Assign secrets to an application that connects to a MySQL database.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and downloads the resource files necessary for this exercise.

```
[student@workstation ~]$ lab authorization-secrets start
```

Instructions

- 1. Log in to the OpenShift cluster and create the `authorization-secrets` project.

- 1.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.
...output omitted...
```

- 1.2. Create the `authorization-secrets` project.

```
[student@workstation ~]$ oc new-project authorization-secrets
Now using project "authorization-secrets" on server
"https://api.ocp4.example.com:6443".
...output omitted...
```

- 2. Create a secret with the credentials and connection information to access a MySQL database.

```
[student@workstation ~]$ oc create secret generic mysql \
> --from-literal user=myuser --from-literal password=redhat123 \
> --from-literal database=test_secrets --from-literal hostname=mysql
secret/mysql created
```

► 3. Deploy a database and add the secret for user and database configuration.

- 3.1. Try to deploy an ephemeral database server. This should fail because the MySQL image needs environment variables for its initial configuration. The values for these variables cannot be assigned from a secret using the `oc new-app` command.

```
[student@workstation ~]$ oc new-app --name mysql \
> --image registry.redhat.io/rhel8/mysql-80:1
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "mysql" created
    deployment.apps "mysql" created
    service "mysql" created
--> Success
...output omitted...
```

- 3.2. Run the `oc get pods` command with the `-w` option to retrieve the status of the deployment in real time. Notice how the database pod is in a failed state. Press `Ctrl+C` to exit the command.

```
[student@workstation ~]$ oc get pods -w
```

NAME	READY	STATUS	RESTARTS	AGE
mysql-7b58b9b68b-ftp6j	0/1	CrashLoopBackOff	3 (46s ago)	106s
mysql-7b58b9b68b-ftp6j	1/1	Running	4 (53s ago)	113s
mysql-7b58b9b68b-ftp6j	0/1	Error	4 (54s ago)	114s
mysql-7b58b9b68b-ftp6j	0/1	CrashLoopBackOff	4 (16s ago)	2m9s



Note

It might take a while for the pod to reach the error state.

- 3.3. Use the `mysql` secret to initialize environment variables on the `mysql` deployment. The deployment needs the `MYSQL_USER`, `MYSQL_PASSWORD`, and `MYSQL_DATABASE` environment variables for a successful initialization. The secret has the `user`, `password`, and `database` keys that can be assigned to the deployment as environment variables, adding the prefix `MYSQL_`.

```
[student@workstation ~]$ oc set env deployment/mysql --from secret/mysql \
> --prefix MYSQL_
deployment.apps/mysql updated
```

- 3.4. To demonstrate how a secret can be mounted as a volume, mount the `mysql` secret to the `/run/secrets/mysql` directory within the pod. This step only shows how to mount a secret as a volume, it is not required to fix the deployment.

```
[student@workstation ~]$ oc set volume deployment/mysql --add --type secret \
> --mount-path /run/secrets/mysql --secret-name mysql
info: Generated volume name: volume-nrh7r
deployment.apps/mysql volume updated
```

- 3.5. Modifying the deployment using the `oc set env` command or the `oc set volume` command triggers a new application deployment. Verify that the `mysql` application deploys successfully after adding the environment variables.

```
[student@workstation ~]$ oc get pods
NAME                                READY   STATUS    RESTARTS   AGE
mysql-7cd7499d66-gm2rh             1/1     Running   0           21s
```

Take note of the pod name in the `Running` state; you will need it in upcoming steps.

- 4. Verify that the database now authenticates using the environment variables initialized from the `mysql` secret.

- 4.1. Open a remote shell session to the `mysql` pod in the `Running` state.

```
[student@workstation ~]$ oc rsh mysql-7cd7499d66-gm2rh
sh-4.4$
```

- 4.2. Start a MySQL session to verify that the environment variables initialized by the `mysql` secret were used to configure the `mysql` application.

```
sh-4.4$ mysql -u myuser --password=redhat123 test_secrets -e 'show databases;'
mysql: [Warning] Using a password on the command line interface can be insecure.
+-----+
| Database                |
+-----+
| information_schema      |
| test_secrets            |
+-----+
```

- 4.3. List mount points in the pod containing the `mysql` pattern. Notice that the mount point is backed by a temporary file system (`tmpfs`). This is true for all secrets that are mounted as volumes.

```
sh-4.4$ df -h | grep mysql
tmpfs           15G  16K   15G   1% /run/secrets/mysql
```

- 4.4. Examine the files mounted at the `/run/secrets/mysql` mount point. Each file matches a key name in the secret, and the content of each file is the value of the associated key.


```
sh-4.4$ for FILE in $(ls /run/secrets/mysql)
> do
> echo "${FILE}: $(cat /run/secrets/mysql/${FILE})"
> done
database: test_secrets
hostname: mysql
password: redhat123
user: myuser
```

4.5. Close the remote shell session to continue from your workstation machine.

```
sh-4.4$ exit
exit
[student@workstation ~]$
```

► 5. Create a new application that uses the MySQL database.

5.1. Create a new application using the `redhattraining/famous-quotes` image from Quay.io.

```
[student@workstation ~]$ oc new-app --name quotes \
> --image quay.io/redhattraining/famous-quotes:2.1
--> Found container image 7ff1a7b (19 months old) from quay.io for "quay.io/
redhattraining/famous-quotes:2.1"
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "quotes" created
    deployment.apps "quotes" created
    service "quotes" created
--> Success
...output omitted...
```

5.2. Verify the status of the `quotes` application pod. The pod displays an error because it cannot connect to the database. This might take a while to display in the output. Press `Ctrl+C` to exit the command.

```
[student@workstation ~]$ oc get pods -l deployment=quotes -w
```

NAME	READY	STATUS	RESTARTS	AGE
quotes-66b987df5d-59flf	0/1	CrashLoopBackOff	2 (20s ago)	50s
quotes-66b987df5d-59flf	0/1	Error	3	59s
quotes-66b987df5d-59flf	0/1	CrashLoopBackOff	3 (16s ago)	74s

► 6. The `quotes` application requires several environment variables. The `mysql` secret can initialize environment variables for the `quotes` application by adding the `QUOTES_` prefix.

6.1. Use the `mysql` secret to initialize the following environment variables that the `quotes` application needs to connect to the database: `QUOTES_USER`, `QUOTES_PASSWORD`, `QUOTES_DATABASE`, and `QUOTES_HOSTNAME`, which correspond to the `user`, `password`, `database`, and `hostname` keys of the `mysql` secret.

```
[student@workstation ~]$ oc set env deployment/quotes --from secret/mysql \
> --prefix QUOTES_
deployment.apps/quotes updated
```

- 6.2. Wait until the `quotes` application pod redeploys. The older pods terminate automatically.

```
[student@workstation ~]$ oc get pods -l deployment=quotes
NAME                                READY   STATUS    RESTARTS   AGE
quotes-77df54758b-mqdtf            1/1     Running   3           7m17s
```

**Note**

It might take a while for the pod to finish the deployment.

- 7. Verify that the `quotes` pod successfully connects to the database and that the application displays a list of quotes.

- 7.1. Examine the pod logs using the `oc logs` command. The logs indicate a successful database connection.

```
[student@workstation ~]$ oc logs quotes-77df54758b-mqdtf | head -n2
... Connecting to the database: myuser:redhat123@tcp(mysql:3306)/test_secrets
... Database connection OK
```

- 7.2. Expose the `quotes` service so that it can be accessed from outside the cluster.

```
[student@workstation ~]$ oc expose service quotes \
> --hostname quotes.apps.ocp4.example.com
route.route.openshift.io/quotes exposed
```

- 7.3. Verify the application host name.

```
[student@workstation ~]$ oc get route quotes
NAME      HOST/PORT                                PATH   SERVICES   PORT      ...
quotes    quotes.apps.ocp4.example.com             quotes 8000-tcp   ...
```

- 7.4. Verify that the variables are properly set in the application by accessing the `env` REST API entry point.

```
[student@workstation ~]$ curl -s \
> http://quotes.apps.ocp4.example.com/env | grep QUOTES_
<li>QUOTES_USER: myuser </li>
<li>QUOTES_PASSWORD: redhat123 </li>
<li>QUOTES_DATABASE: test_secrets</li>
<li>QUOTES_HOST: mysql</li>
```

- 7.5. Access the application status REST API entry point to test the database connection.

```
[student@workstation ~]$ curl -s http://quotes.apps.ocp4.example.com/status
Database connection OK
```

7.6. Test application functionality by accessing the random REST API entry point.

```
[student@workstation ~]$ curl -s http://quotes.apps.ocp4.example.com/random
8: Those who can imagine anything, can create the impossible.
- Alan Turing
```

► 8. Remove the authorization-secrets project.

```
[student@workstation ~]$ oc delete project authorization-secrets
project.project.openshift.io "authorization-secrets" deleted
```

Finish

On the workstation machine, run the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab authorization-secrets finish
```

This concludes the guided exercise.

Controlling Application Permissions with Security Context Constraints

Objectives

After completing this section, you should be able to create service accounts and apply permissions, and manage security context constraints.

Security Context Constraints (SCCs)

Red Hat OpenShift provides *security context constraints* (SCCs), a security mechanism that restricts access to resources, but not to operations in OpenShift.

SCCs limit the access from a running pod in OpenShift to the host environment. SCCs control:

- Running privileged containers.
- Requesting extra capabilities for a container
- Using host directories as volumes.
- Changing the SELinux context of a container.
- Changing the user ID.

Some containers developed by the community might require relaxed security context constraints to access resources that are forbidden by default, such as file systems, sockets, or to access a SELinux context.

You can run the following command as a cluster administrator to list the SCCs defined by OpenShift:

```
[user@host ~]$ oc get scc
```

OpenShift provides eight default SCCs:

- anyuid
- hostaccess
- hostmount-anyuid
- hostnetwork
- node-exporter
- nonroot
- privileged
- restricted

To get additional information about an SCC, use the `oc describe` command:

```
[user@host ~]$ oc describe scc anyuid
Name:          anyuid
Priority:       10
Access:
  Users:        <none>
  Groups:       system:cluster-admins
Settings:
...output omitted...
```

Most pods created by OpenShift use the SCC named `restricted`, which provides limited access to resources external to OpenShift. Use the `oc describe` command to view the security context constraint that a pod uses.

```
[user@host ~]$ oc describe pod console-5df4fcbb47-67c52 \
> -n openshift-console | grep scc
openshift.io/scc: restricted
```

Container images downloaded from public container registries, such as Docker Hub, might fail to run using the `restricted` SCC. For example, a container image that requires running as a specific user ID can fail because the `restricted` SCC runs the container using a random user ID. A container image that listens on port 80 or port 443 can fail for a related reason. The random user ID used by the `restricted` SCC cannot start a service that listens on a privileged network port (port numbers less than 1024). Use the `scc-subject-review` subcommand to list all the security context constraints that can overcome the limitations of a container:

```
[user@host ~]$ oc get pod podname -o yaml | \
> oc adm policy scc-subject-review -f -
```

For the `anyuid` SCC, the `run as user` strategy is defined as `RunAsAny`, which means that the pod can run as any user ID available in the container. This strategy allows containers that require a specific user to run the commands using a specific user ID.

To change the container to run using a different SCC, you must create a service account bound to a pod. Use the `oc create serviceaccount` command to create the service account, and use the `-n` option if the service account must be created in a namespace different than the current one:

```
[user@host ~]$ oc create serviceaccount service-account-name
```

To associate the service account with an SCC, use the `oc adm policy` command. Use the `-z` option to identify a service account, and use the `-n` option if the service account exists in a namespace different than the current one:

```
[user@host ~]$ oc adm policy add-scc-to-user SCC -z service-account
```



Important

Assigning an SCC to a service account or removing an SCC from a service account must be performed by a cluster administrator. Allowing pods to run with a less restrictive SCC can make your cluster less secure. Use with caution.

Modify an existing deployment or deployment configuration to use the service account using the `oc set serviceaccount` command:

```
[user@host ~]$ oc set serviceaccount deployment/deployment-name \
> service-account-name
```

If the command succeeds, then the pods associated with the deployment or deployment configuration redeploy.

Privileged Containers

Some containers might need to access the runtime environment of the host. For example, the S2I builder containers are a class of privileged containers that require access beyond the limits of their own containers. These containers can pose security risks because they can use any resources on an OpenShift node. Use SCCs to enable access for privileged containers by creating service accounts with privileged access.



References

For more information, refer to the *Managing Security Context Constraints* chapter in the Red Hat OpenShift Container Platform 4.10 *Authentication and Authorization* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/security_and_compliance/index#encrypting-etcd

► Guided Exercise

Controlling Application Permissions with Security Context Constraints

In this exercise, you will deploy applications that require pods with extended permissions.

Outcomes

You should be able to:

- Create service accounts and assign security context constraints (SCCs) to them.
- Assign a service account to a deployment configuration.
- Run applications that need root privileges.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates some HTPasswd users for the exercise.

```
[student@workstation ~]$ lab authorization-scc start
```

Instructions

- 1. Log in to the OpenShift cluster and create the `authorization-scc` project.

- 1.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.
...output omitted...
```

- 1.2. Create the `authorization-scc` project.

```
[student@workstation ~]$ oc new-project authorization-scc
Now using project "authorization-scc" on server ...
...output omitted...
```

- 2. Deploy an application named `gitlab` using the container image located at `quay.io/redhattraining/gitlab-ce:8.4.3-ce.0`. This image is a copy of the container image available at `docker.io/gitlab/gitlab-ce:8.4.3-ce.0`. Verify that the pod fails because the container image needs root privileges.

- 2.1. Deploy the `gitlab` application.

```
[student@workstation ~]$ oc new-app --name gitlab \
> --image quay.io/redhattraining/gitlab-ce:8.4.3-ce.0
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "gitlab" created
    deployment.apps "gitlab" created
    service "gitlab" created
--> Success
...output omitted...
```

- 2.2. Determine if the application is successfully deployed. It should give an error because this image needs `root` privileges to properly deploy.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
gitlab-d89cd88f8-jwqbp	0/1	ContainerCreating	0	19s
gitlab-d89cd88f8-jwqbp	1/1	Running	0	30s
gitlab-d89cd88f8-jwqbp	0/1	Error	0	36s



Note

It might take some time for the image to reach the `Error` state. You might also see the `CrashLoopBackOff` status while checking the pod's health.

- 2.3. Review the application logs to confirm that the failure is caused by insufficient privileges.

```
[student@workstation ~]$ oc logs pod/gitlab-7d67db7875-gcsjl
...output omitted...

=====
Recipe Compile Error in /opt/gitlab/embedded/cookbooks/cache/cookbooks/gitlab/
recipes/default.rb
=====

Chef::Exceptions::InsufficientPermissions
-----
directory[/etc/gitlab] (gitlab::default line 26) had an error:
  Chef::Exceptions::InsufficientPermissions: Cannot create directory[/etc/gitlab]
  at /etc/gitlab due to insufficient permissions
...output omitted...
```

- 2.4. Log in as the admin user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.
...output omitted...
```

- 2.5. Check if using a different SCC can resolve the permissions problem.


```
[student@workstation ~]$ oc get pod/gitlab-7d67db7875-gcsjl -o yaml | \
>   oc adm policy scc-subject-review -f -
RESOURCE                                ALLOWED BY
Pod/gitlab-7d67db7875-gcsjl             anyuid
```

- ▶ 3. Create a new service account and assign the anyuid SCC to it.

3.1. Create a service account named `gitlab-sa`.

```
[student@workstation ~]$ oc create sa gitlab-sa
serviceaccount/gitlab-sa created
```

3.2. Assign the anyuid SCC to the `gitlab-sa` service account.

```
[student@workstation ~]$ oc adm policy add-scc-to-user anyuid -z gitlab-sa
clusterrole.rbac.authorization.k8s.io/system:openshift:scc:anyuid added: "gitlab-sa"
```

- ▶ 4. Modify the `gitlab` application so that it uses the newly created service account. Verify that the new deployment succeeds.

4.1. Log in as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.
...output omitted...
```

4.2. Assign the `gitlab-sa` service account to the `gitlab` deployment.

```
[student@workstation ~]$ oc set serviceaccount deployment/gitlab gitlab-sa
deployment.apps/gitlab serviceaccount updated
```

4.3. Verify that the `gitlab` redeployment succeeds. You might need to run the `oc get pods` command multiple times until you see a running application pod.

```
[student@workstation ~]$ oc get pods
NAME                                READY   STATUS    RESTARTS   AGE
gitlab-86d6d65-zm2fd               1/1     Running   0           55s
```

- ▶ 5. Verify that the `gitlab` application works properly.

5.1. Expose the `gitlab` application. Because the `gitlab` service listens on ports 22, 80, and 443, you must use the `--port` option.

```
[student@workstation ~]$ oc expose service/gitlab --port 80 \
>   --hostname gitlab.apps.ocp4.example.com
route.route.openshift.io/gitlab exposed
```

5.2. Get the exposed route.

```
[student@workstation ~]$ oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT	...
gitlab	gitlab.apps.ocp4.example.com		gitlab	80	...

5.3. Verify that the `gitlab` application is answering HTTP queries.

```
[student@workstation ~]$ curl -sL http://gitlab.apps.ocp4.example.com/ | \
> grep '<title>'
<title>Sign in · GitLab</title>
```

► 6. Delete the `authorization-scc` project.

```
[student@workstation ~]$ oc delete project authorization-scc
project.project.openshift.io "authorization-scc" deleted
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab authorization-scc finish
```

This concludes the guided exercise.

► Lab

Configuring Application Security

In this lab, you will create a secret to share sensitive configuration information and modify a deployment so that it runs with less restrictive settings.

Outcomes

You should be able to:

- Provide sensitive information to deployments using secrets.
- Allow applications to run in a less restrictive environment using security context constraints.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and that you can log in to the cluster.

```
[student@workstation ~]$ lab authorization-review start
```

Instructions

1. As the `developer` user, create the `authorization-review` project. All additional tasks in this exercise use the `authorization-review` project.
2. Create a secret named `review-secret` that you will use with the MySQL database and WordPress applications. The secret must include three key-value pairs: `user=wpuser`, `password=redhat123`, and `database=wordpress`.
3. Deploy a MySQL database application named `mysql` using the container image located at `registry.redhat.io/rhel8/mysql-80:1`. Modify the `mysql` deployment to use the `review-secret` secret as environment variables. The environment variables must use the `MYSQL_` prefix.
4. Deploy a WordPress application named `wordpress` using the container image located at `quay.io/redhattraining/wordpress:5.7-php7.4-apache`. Once deployed, modify the `wordpress` deployment to use the `review-secret` secret as additional environment variables.

The Wordpress application requires that you set several environment variables when executing the `oc new-app` command:

- `WORDPRESS_DB_HOST` with a value of `mysql`.
- `WORDPRESS_DB_NAME` with a value of `wordpress`.
- `WORDPRESS_TITLE` with a value of `auth-review`.
- `WORDPRESS_USER` with a value of `wpuser`.
- `WORDPRESS_PASSWORD` with a value of `redhat123`.
- `WORDPRESS_EMAIL` with a value of `student@redhat.com`.
- `WORDPRESS_URL` with a value of `wordpress-review.apps.ocp4.example.com`.

The additional environment variables must use the `WORDPRESS_DB_` prefix.



Note

The `wordpress` pod does not run successfully until you modify the deployment to use a less restrictive security context constraint.

5. As the `admin` user, identify a less restrictive SCC that allows the `wordpress` deployment to run successfully. Create a service account named `wordpress-sa` and grant the `anyuid` SCC to it. Modify the `wordpress` deployment to use the `wordpress-sa` service account.
6. As the `developer` user, make the `wordpress` service accessible to external requests using the `wordpress-review.apps.ocp4.example.com` host name. Access the route using a web browser and verify the WordPress application displays the setup wizard.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab authorization-review grade
```

Finish

As the `student` user on the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab authorization-review finish
```

This concludes the lab.

► Solution

Configuring Application Security

In this lab, you will create a secret to share sensitive configuration information and modify a deployment so that it runs with less restrictive settings.

Outcomes

You should be able to:

- Provide sensitive information to deployments using secrets.
- Allow applications to run in a less restrictive environment using security context constraints.

Before You Begin

As the `student` user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and that you can log in to the cluster.

```
[student@workstation ~]$ lab authorization-review start
```

Instructions

1. As the `developer` user, create the `authorization-review` project. All additional tasks in this exercise use the `authorization-review` project.

- 1.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.
...output omitted...
```

- 1.2. Create the `authorization-review` project.

```
[student@workstation ~]$ oc new-project authorization-review
Now using project "authorization-review" on server
"https://api.ocp4.example.com:6443".
...output omitted...
```

2. Create a secret named `review-secret` that you will use with the MySQL database and WordPress applications. The secret must include three key-value pairs: `user=wpuser`, `password=redhat123`, and `database=wordpress`.

- 2.1. Create a secret named `review-secret`.

```
[student@workstation ~]$ oc create secret generic review-secret \
> --from-literal user=wpuser --from-literal password=redhat123 \
> --from-literal database=wordpress
secret/review-secret created
```

3. Deploy a MySQL database application named `mysql` using the container image located at `registry.redhat.io/rhel8/mysql-80:1`. Modify the `mysql` deployment to use the `review-secret` secret as environment variables. The environment variables must use the `MYSQL_` prefix.

3.1. Create a new application to deploy a `mysql` database server.

```
[student@workstation ~]$ oc new-app --name mysql \
> --image registry.redhat.io/rhel8/mysql-80:1
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "mysql" created
    deployment.apps "mysql" created
    service "mysql" created
--> Success
...output omitted...
```

- 3.2. Use the `review-secret` secret to initialize the environment variables on the `mysql` deployment. The `--prefix` option ensures that all the variables injected from the secret into the pod start with `MYSQL_`.

```
[student@workstation ~]$ oc set env deployment/mysql --prefix MYSQL_ \
> --from secret/review-secret
deployment.apps/mysql updated
```

- 3.3. Verify that the `mysql` pod redeploys successfully.

```
[student@workstation ~]$ watch oc get pods
```

Press `Ctrl+C` to exit the `watch` command after the `mysql` pod displays both `1/1` and `Running`.

```
Every 2.0s: oc get pods          ...
NAME                                READY  STATUS   RESTARTS  AGE
mysql-f675b96f8-vspb9              1/1    Running   0          20s
```



Note

It may take a few minutes for the deployment to roll out successfully after setting the secret.

4. Deploy a WordPress application named `wordpress` using the container image located at `quay.io/redhattraining/wordpress:5.7-php7.4-apache`. Once deployed, modify the `wordpress` deployment to use the `review-secret` secret as additional environment variables.

The Wordpress application requires that you set several environment variables when executing the `oc new-app` command:

- `WORDPRESS_DB_HOST` with a value of `mysql`.
- `WORDPRESS_DB_NAME` with a value of `wordpress`.
- `WORDPRESS_TITLE` with a value of `auth-review`.
- `WORDPRESS_USER` with a value of `wpuser`.
- `WORDPRESS_PASSWORD` with a value of `redhat123`.
- `WORDPRESS_EMAIL` with a value of `student@redhat.com`.
- `WORDPRESS_URL` with a value of `wordpress-review.apps.ocp4.example.com`.

The additional environment variables must use the `WORDPRESS_DB_` prefix.



Note

The wordpress pod does not run successfully until you modify the deployment to use a less restrictive security context constraint.

4.1. Deploy a wordpress application.

```
[student@workstation ~]$ oc new-app --name wordpress \
> --image quay.io/redhattraining/wordpress:5.7-php7.4-apache \
> -e WORDPRESS_DB_HOST=mysql \
> -e WORDPRESS_DB_NAME=wordpress \
> -e WORDPRESS_TITLE=auth-review \
> -e WORDPRESS_USER=wpuser \
> -e WORDPRESS_PASSWORD=redhat123 \
> -e WORDPRESS_EMAIL=student@redhat.com \
> -e WORDPRESS_URL=wordpress-review.apps.ocp4.example.com
...output omitted...
-> Creating resources ...
    imagestream.image.openshift.io "wordpress" created
    deployment.apps "wordpress" created
    service "wordpress" created
--> Success
...output omitted...
```

4.2. Use the `review-secret` secret to initialize the environment variables on the wordpress deployment. The `--prefix` option ensures that the variables injected from the secret into the pod all start with `WORDPRESS_DB_`.

```
[student@workstation ~]$ oc set env deployment/wordpress \
> --prefix WORDPRESS_DB_ --from secret/review-secret
deployment.apps/wordpress updated
```

4.3. Verify that the `wordpress` pod does not successfully redeploy, even after injecting variables from the `review-secret` secret.

```
[student@workstation ~]$ watch oc get pods -l deployment=wordpress
```

Wait up to one minute, and then press `Ctrl+C` to exit the `watch` command. The `wordpress` pod continually restarts. Each time the pod transitions to a status of `Error` and then `CrashLoopBackOff`.

```
Every 2.0s: oc get pods -l deployment=wordpress ...
```

NAME	READY	STATUS	RESTARTS	AGE
wordpress-68c49c9d4-wq46g	0/1	CrashLoopBackOff	5	4m30s

4.4. Check the pod logs for error messages.

```
[student@workstation ~]$ oc logs wordpress-68c49c9d4-wq46g
...output omitted...
(13)Permission denied: AH00072: make_sock: could not bind to address [::]:80
(13)Permission denied: AH00072: make_sock: could not bind to address 0.0.0.0:80
no listening sockets available, shutting down
AH00015: Unable to open logs
```

By default, OpenShift prevents pods from starting services that listen on ports lower than 1024.

- As the `admin` user, identify a less restrictive SCC that allows the `wordpress` deployment to run successfully. Create a service account named `wordpress-sa` and grant the `anyuid` SCC to it. Modify the `wordpress` deployment to use the `wordpress-sa` service account.

5.1. Log in as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.
...output omitted...
```

5.2. Check whether using a different SCC resolves the permissions problem.

```
[student@workstation ~]$ oc get pod/wordpress-68c49c9d4-wq46g -o yaml \
> | oc adm policy scc-subject-review -f -
```

RESOURCE	ALLOWED BY
Pod/wordpress-68c49c9d4-wq46g	anyuid



Important

The `oc adm policy` command must be run as the `admin` user.

5.3. Create a service account named `wordpress-sa`.

```
[student@workstation ~]$ oc create serviceaccount wordpress-sa
serviceaccount/wordpress-sa created
```

5.4. Grant the `anyuid` SCC to the `wordpress-sa` service account. If the `wordpress` pod runs as the `root` user, then OpenShift allows the pod to start a service on port 80.

```
[student@workstation ~]$ oc adm policy add-scc-to-user anyuid -z wordpress-sa
clusterrole.rbac.authorization.k8s.io/system:openshift:scc:anyuid added:
"wordpress-sa"
```

5.5. Configure the `wordpress` deployment to use the `wordpress-sa` service account.


```
[student@workstation ~]$ oc set serviceaccount deployment/wordpress \
>   wordpress-sa
deployment.apps/wordpress serviceaccount updated
```

- 5.6. Verify that the `wordpress` pod successfully redeploys after configuring the service account.

```
[student@workstation ~]$ watch oc get pods -l deployment=wordpress
```

Press `Ctrl+C` to exit the `watch` command after the `wordpress` pod displays both `1/1` and `Running`.

```
Every 2.0s: oc get pods -l deployment=wordpress ...
```

NAME	READY	STATUS	RESTARTS	AGE
wordpress-bcb5d97f6-mwljs	1/1	Running	0	21s

6. As the `developer` user, make the `wordpress` service accessible to external requests using the `wordpress-review.apps.ocp4.example.com` host name. Access the route using a web browser and verify the WordPress application displays the setup wizard.

- 6.1. Use the `oc expose` command to create a route to the `wordpress` application.

```
[student@workstation ~]$ oc expose service/wordpress \
>   --hostname wordpress-review.apps.ocp4.example.com
route.route.openshift.io/wordpress exposed
```

- 6.2. Use a web browser to verify access to the URL `http://wordpress-review.apps.ocp4.example.com`. When you correctly deploy the application, a setup wizard displays in the browser.

Alternatively, use the `curl` command to access the installation URL directly.

```
[student@workstation ~]$ curl -s \
>   http://wordpress-review.apps.ocp4.example.com/wp-admin/install.php \
>   | grep Installation
<title>WordPress &rsquo; Installation</title>
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab authorization-review grade
```

Finish

As the `student` user on the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab authorization-review finish
```

This concludes the lab.

Summary

In this chapter, you learned:

- Secret resources allow you to separate sensitive information from application pods. You expose secrets to an application pod either as environment variables or as ordinary files.
- OpenShift uses security context constraints (SCCs) to define allowed pod interactions with system resources. By default, pods operate under the `restricted` context which limits access to node resources.

Chapter 5

Configuring OpenShift Networking for Applications

Goal

Troubleshoot OpenShift software-defined networking (SDN) and configure network policies.

Objectives

- Troubleshoot OpenShift software-defined networking using the command-line interface.
- Allow and protect network connections to applications inside an OpenShift cluster.
- Restrict network traffic between projects and pods.

Sections

- Troubleshooting OpenShift Software-defined Networking (and Guided Exercise)
- Exposing Applications for External Access (and Guided Exercise)
- Configuring Network Policies (and Guided Exercise)

Lab

Configuring OpenShift Networking for Applications

Troubleshooting OpenShift Software-defined Networking

Objectives

After completing this section, you should be able to troubleshoot OpenShift software-defined networking using the command-line interface.

Introducing OpenShift Software-defined Networking

OpenShift implements a software-defined network (SDN) to manage the network infrastructure of the cluster and user applications. Software-defined networking is a networking model that allows you to manage network services through the abstraction of several networking layers. It decouples the software that handles the traffic, called the *control plane*, and the underlying mechanisms that route the traffic, called the *data plane*. Among the many features of SDN, open standards enable vendors to propose their solutions, centralized management, dynamic routing, and tenant isolation.

In OpenShift Container Platform, the SDN satisfies the following five requirements:

- Managing the network traffic and network resources programmatically, so that the organization teams can decide how to expose their applications.
- Managing communication between containers that run in the same project.
- Managing communication between pods, whether they belong to a same project or run in separate projects.
- Managing network communication from a pod to a service.
- Managing network communication from an external network to a service, or from containers to external networks.

The SDN implementation creates a backward-compatible model, in which pods are akin to virtual machines in terms of port allocation, IP address leasing, and reservation.

Discussing OpenShift Networking Model

The Container Network Interface (CNI) is a common interface between the network provider and the container execution and is implemented as network plug-ins. The CNI provides the specification for plug-ins to configure network interfaces inside containers. Plug-ins written to the specification allow different network providers to control the OpenShift cluster network.

The SDN uses CNI plug-ins to create Linux namespaces to partition the usage of resources and processes on physical and virtual hosts. This implementation allows containers inside pods to share network resources, such as devices, IP stacks, firewall rules, and routing tables. The SDN allocates a unique routable IP to each pod so that you can access the pod from any other service in the same network.

Some common CNI plug-ins used in OpenShift are:

- OpenShift SDN
- OVN-Kubernetes
- Kuryr

In OpenShift 4.10, both OpenShift SDN and OVN-Kubernetes are the default network providers.

The OpenShift SDN network provider uses Open vSwitch (OVS) to connect pods on the same node and Virtual Extensible LAN (VXLAN) tunneling to connect nodes. OVN-Kubernetes uses Open Virtual Network (OVN) to manage the cluster network. OVN extends OVS with virtual network abstractions. Kuryr provides networking through the Neutron and Octavia Red Hat OpenStack Platform services.

Migrating Legacy Applications

The SDN design makes it easy to containerize your legacy applications because you do not need to change the way the application components communicate with each other. If your application is comprised of many services that communicate over the TCP/UDP stack, this approach still works as containers in a pod share the same network stack.

The following diagram shows how all pods are connected to a shared network:




Figure 5.1: Kubernetes basic networking



Note

The OpenShift *Cluster Network Operator* manages the network, as discussed later.

Using Services for Accessing Pods

Kubernetes provides the concept of a *service*, which is an essential resource in any OpenShift application. Services allow for the logical grouping of pods under a common access route. A service acts as a load balancer in front of one or more pods, thus decoupling the application specifications (such as the number of running replicas) from the access to the application. It load balances client requests across member pods, and provides a stable interface that enables communication with pods without tracking individual pod IP addresses.

Most real-world applications do not run as a single pod. They need to scale horizontally, so an application could run on many pods to meet growing user demand. In an OpenShift cluster, pods are constantly created and destroyed across the nodes in the cluster, such as during the deployment of a new application version or when draining a node for maintenance. Pods are assigned a different IP address each time they are created; thus, pods are not easily addressable. Instead of having a pod discover the IP address of another pod, you can use services to provide a single, unique IP address for other pods to use, independent of where the pods are running.

Services rely on selectors (labels) that indicate which pods receive the traffic through the service. Each pod matching these selectors is added to the service resource as an endpoint. As pods are created and killed, the service automatically updates the endpoints. Using selectors brings flexibility to the way you design the architecture and routing of your applications. For example, you can divide the application into tiers and decide to create a service for each tier. Selectors allow a design that is flexible and highly resilient.

OpenShift uses two subnets: one subnet for pods, and one subnet for services. The traffic is forwarded in a transparent way to the pods; an agent (depending on the network mode that you use) manages routing rules to route traffic to the pods that match the selectors.

The following diagram shows how three API pods are running on separate nodes. The `service1` service balances the load between these three pods.




Figure 5.2: Using services for accessing applications

The following YAML definition shows you how to create a service. This defines the `application-frontend` service, which creates a virtual IP that exposes the TCP port 443. The front-end application listens on the unprivileged port 8843.

```
kind: Service
apiVersion: v1
metadata:
  name: application-frontend ❶
  labels:
    app: frontend-svc ❷
spec:
  ports: ❸
  - name: HTTP ❹
    protocol: TCP
    port: 443 ❺
    targetPort: 8843 ❻
  selector: ❼
```

```

app: shopping-cart
name: frontend
type: ClusterIP 8

```

- 1 The name of the service. This identifier allows you to manage the service after its creation.
- 2 A label that you can use as a selector. This allow you to logically group your services.
- 3 An array of objects that describes network ports to expose.
- 4 Each entry defines the name for the port mapping. This value is generic and is used for identification purposes only.
- 5 This is the port that the service exposes. You use this port to connect to the application that the service exposes.
- 6 This is the port on which the application listens. The service creates a forwarding rule from the service port to the service target port.
- 7 The selector defines which pods are in the service pool. Services use this selector to determine where to route the traffic. In this example, the service targets all pods whose labels match `app: shopping-cart` and `name: frontend`.
- 8 This is how the service is exposed. `ClusterIP` exposes the service using an IP address internal to the cluster and is the default value. Other service types will be described in elsewhere this course.

Discussing the DNS Operator

The DNS operator deploys and runs a DNS server managed by CoreDNS, a lightweight DNS server written in GoLang. The DNS operator provides DNS name resolution between pods, which enables services to discover their endpoints.

Every time you create a new application, OpenShift configures the pods so that they contact the CoreDNS service IP for DNS resolution.

Run the following command to review the configuration of the DNS operator:

```

[user@host ~]$ oc describe dns.operator/default
Name:          default
...output omitted...
API Version:   operator.openshift.io/v1
Kind:          DNS
...output omitted...
Spec:
Status:
  Cluster Domain: cluster.local
  Cluster IP:     172.30.0.10
...output omitted...

```

The DNS operator is responsible for the following:

- Creating a default cluster DNS name (`cluster.local`)
- Assigning DNS names to services that you define (for example, `db.backend.svc.cluster.local`)

For example, from a pod named `example-6c4984d949-7m26r`, the following command demonstrates you can reach the `hello` pod through the `hello` service in the `test` project via the FQDN for the service:

```
[user@host ~]$ oc rsh example-6c4984d949-7m26r \
> curl hello.test.svc.cluster.local:8080
```

Managing DNS Records for Services

This DNS implementation allows pods to seamlessly resolve DNS names for resources in a project or the cluster. Pods can use a predictable naming scheme for accessing a service. For example, querying the `db.backend.svc.cluster.local` host name from a container returns the IP address of the service. In this example, `db` is the name of the service, `backend` is the project name, and `cluster.local` is the cluster DNS name.

CoreDNS creates two kind of records for services: A records that resolve to services, and SRV records that match the following format:

```
_port-name._port-protocol.svc-name.namespace.svc.cluster-domain.cluster-domain.
```

For example, if you use a service that exposes the TCP port 443 through the HTTPS service, the SRV record is created as follows:

```
_443-tcp._tcp.https.frontend.svc.cluster.local.
```



Note

When services do not have a cluster IP, the DNS operator assigns them a DNS A record that resolves to the set of IPs of the pods behind the service.

Similarly, the newly created SRV record resolves to all the pods that are behind the service.

Introducing the Cluster Network Operator

OpenShift Container Platform uses the Cluster Network Operator for managing the SDN. This includes the network CIDR to use, the network provider, and the IP address pools. Configuration of the Cluster Network Operator is done before installation, although it is possible to migrate from the OpenShift SDN default CNI network provider to the OVN-Kubernetes network provider.

Run the following `oc get` command as an administrative user to consult the SDN configuration, which is managed by the `Network.config.openshift.io` custom resource definition:

```
[user@host ~]$ oc get network/cluster -o yaml
apiVersion: config.openshift.io/v1
kind: Network
...output omitted...
spec:
  clusterNetwork:
    - cidr: 10.128.0.0/14 ❶
      hostPrefix: 23 ❷
  externalIP:
```

```

policy: {}
networkType: OpenShiftSDN ❸
serviceNetwork:
- 172.30.0.0/16
...output omitted...

```

- ❶ Defines the CIDR for all pods in the cluster. In this example, the SDN has a netmask of 255.252.0.0 and can allocate 262144 IP addresses.
- ❷ Defines the host prefix. A value of 23 indicates a netmask of 255.255.254.0, which translates to 512 allocatable IPs.
- ❸ Shows the current SDN provider. You can choose between `OpenShiftSDN`, `OVNKubernetes`, and `Kuryr`.

Introducing Multus CNI

Multus is an open source project to support multiple network cards in OpenShift. One of the challenges that Multus solves is the migration of network function virtualization to containers. Multus acts as a broker and arbiter of other CNI plug-ins for managing the implementation and lifecycle of supplementary network devices in containers. Multus supports plug-ins, such as SR-IOV, vHost CNI, Flannel, and Calico. Special edge cases often seen in telecommunication services, edge computing, and virtualization are handled by Multus, allowing multiple network interfaces to pods.



Note

Be aware that all Kubernetes and OpenShift networking features, such as services, ingress, and routes, ignore the extra network devices provided by Multus.



References

For more information, refer to the *Cluster Network Operator in OpenShift Container Platform* chapter in the Red Hat OpenShift Container Platform 4.10 *Networking* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/networking/index#cluster-network-operator

Cluster Networking

<https://kubernetes.io/docs/concepts/cluster-administration/networking/>

CoreDNS: DNS and Service Discovery

<https://coredns.io/>

Multus-CNI

<https://github.com/intel/multus-cni>

► Guided Exercise

Troubleshooting OpenShift Software-defined Networking

In this exercise, you will diagnose and fix connectivity issues with a Kubernetes-style application deployment.

Outcomes

You should be able to:

- Deploy the To Do Node.js application.
- Create a route to expose an application service.
- Troubleshoot communication between pods in your application using `oc debug`.
- Update an OpenShift service.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable.

```
[student@workstation ~]$ lab network-sdn start
```

Instructions

As an OpenShift developer, you just completed the migration of a To Do Node.js application to OpenShift. The application is comprised of two deployments, one for the database, and one for the front end. It also contains two services for communication between pods.

Although the application seems to initialize, you cannot access it via a web browser. In this activity, you will troubleshoot your application and correct the issue.

- 1. Log in to the OpenShift cluster and create the `network-sdn` project.

- 1.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create the `network-sdn` project.

```
[student@workstation ~]$ oc new-project network-sdn
Now using project "network-sdn" on server "https://api.ocp4.example.com:6443".
...output omitted...
```

- 2. Deploy the database and restore its data. The `~/D0280/labs/network-sdn/todo-db.yaml` file defines the following resources:

- A deployment that creates a container based on a MySQL image.
- A service that points to the `mysql` application.

2.1. Change to the `~/D0280/labs/network-sdn` directory.

```
[student@workstation ~]$ cd ~/D0280/labs/network-sdn
```

2.2. Go to the `network-sdn` directory and list the files. In a later step, you will use `db-data.sql` to initialize the database for the application.

```
[student@workstation network-sdn]$ ls
db-data.sql  todo-db.yaml  todo-frontend.yaml
```

2.3. Use `oc create` with the `-f` option against `todo-db.yaml` to deploy the database server pod.

```
[student@workstation network-sdn]$ oc create -f todo-db.yaml
deployment.apps/mysql created
service/mysql created
```

2.4. Run the `oc status` command to review the resources that are present in the project. The `mysql` service points to the database pod.

```
[student@workstation network-sdn]$ oc status
In project network-sdn on server https://api.ocp4.example.com:6443

svc/mysql - 172.30.223.41:3306
  deployment/mysql deploys registry.redhat.io/rhel8/mysql-80:1
  deployment #1 running for 4 seconds - 0/1 pods
...output omitted...
```

2.5. Wait a few moments to ensure that the database pod is running. Retrieve the name of the database pod to restore the tables of the `items` database.

```
[student@workstation network-sdn]$ oc get pods
NAME                                READY   STATUS    RESTARTS   AGE
mysql-94dc6645b-hjjqb              1/1     Running   0           33m
```

2.6. Use the `oc cp` command to transfer the database dump to the pod. Make sure to replace the pod name with the one you obtained in the previous step.

```
[student@workstation network-sdn]$ oc cp db-data.sql mysql-94dc6645b-hjjqb:/tmp/
```

- 2.7. Use the `oc rsh` command to connect to the pod and restore the database.

```
[student@workstation network-sdn]$ oc rsh mysql-94dc6645b-hjjqb /bin/bash
bash-4.4$ mysql -u root items < /tmp/db-data.sql
```

- 2.8. Ensure that the `Item` table is present in the database.

```
bash-4.4$ mysql -u root items -e "SHOW TABLES;"
+-----+
| Tables_in_items |
+-----+
| Item             |
+-----+
```

- 2.9. Exit the container.

```
bash-4.4$ exit
exit
```

- ▶ 3. Deploy the front end application. The `~/DO280/labs/network-sdn/todo-frontend.yaml` file defines the following resources:

- A deployment that creates the `Todo` Node.js application.
- A service that points to the `frontend` application.

- 3.1. Use the `oc create` command to create the front-end application.

```
[student@workstation network-sdn]$ oc create -f todo-frontend.yaml
deployment.apps/frontend created
service/frontend created
```

- 3.2. Wait a few moments for the front end container to start, and then run the `oc get pods` command.

```
[student@workstation network-sdn]$ oc get pods
NAME                                READY   STATUS    RESTARTS   AGE
frontend-57b8b445df-f56qh           1/1     Running   0           34s
...output omitted...
```

- ▶ 4. Create a route to access the `frontend` service and access the application.

- 4.1. You must create a route to access the application from an external network. To create this route, use the `oc expose` command against the `frontend` service. Use the `--hostname` option to override the default FQDN that OpenShift creates.

```
[student@workstation network-sdn]$ oc expose service frontend \
> --hostname todo.apps.ocp4.example.com
route.route.openshift.io/frontend exposed
```

- 4.2. List the routes in the project.

```
[student@workstation network-sdn]$ oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT	...
frontend	todo.apps.ocp4.example.com		frontend	8080	...

As you can see in the example, OpenShift detects the port on which the application listens and creates a forwarding rule from port 80 to port 8080, which is the *target* port.

- 4.3. From *workstation*, open Firefox and access the URL, make sure to add the trailing slash at the end.

- `http://todo.apps.ocp4.example.com/todo/`

The application is not reachable, as shown in the following screen capture.

Application is not available

The application is currently not serving requests at this endpoint. It may not have been started or is still starting.

i Possible reasons you are seeing this page:

- **The host doesn't exist.** Make sure the hostname was typed correctly and that a route matching this hostname exists.
- **The host exists, but doesn't have a matching path.** Check if the URL path was typed correctly and that the route was created using the desired path.
- **Route and path matches, but all pods are down.** Make sure that the resources exposed by this route (pods, services, deployment configs, etc) have at least one pod running.

- 4.4. Inspect the pod logs for errors. The output does not indicate any errors.

```
[student@workstation network-sdn]$ oc logs frontend-57b8b445df-f56qh
App is ready at : 8080
```

- ▶ 5. Run `oc debug` to create a carbon copy of an existing pod in the *frontend* deployment. You use this pod to check connectivity to the database.

- 5.1. Before creating a debug pod, retrieve the database service IP. In a later step, you use the `curl` command to access the database endpoint.

The JSONPath expression allows you to retrieve the service IP.

```
[student@workstation network-sdn]$ oc get service/mysql \
> -o jsonpath="{.spec.clusterIP}{'\n'}"
172.30.103.29
```

- 5.2. Run the `oc debug` command against the *frontend* deployment, which runs the web application pod.

```
[student@workstation network-sdn]$ oc debug -t deployment/frontend
Starting pod/frontend-debug ...
Pod IP: 10.131.0.144
If you don't see a command prompt, try pressing enter.
sh-4.4$
```

- 5.3. One way to test the connectivity between the `frontend` deployment and the database is using `curl`, which supports a variety of protocols.

Use `curl` to connect to the database over port 3306, which is the MySQL default port. Make sure to replace the IP address with the one that you obtained previously for the `mysql` service. When finished, type `Ctrl+C` to exit the session, and then type `exit` to exit the debug pod.

```
sh-4.4$ curl -v telnet://172.30.103.29:3306
* About to connect() to 172.30.103.29 port 3306 (#0)
* Trying 172.30.103.29
* Connected to 172.30.103.29 (172.30.103.29) port 3306 (#0)
J
8.0.21
* RCVD IAC 2
* RCVD IAC 199
^C
sh-4.4$ exit
exit

Removing debug pod ...
```

The output indicates that the database is up and running, and that it is accessible from the `frontend` deployment.

- 6. In the following steps, you ensure that the network connectivity inside the cluster is operational by connecting to the front end container from the database container. Obtain some information about the `frontend` pod, and then use the `oc debug` command to diagnose the issue from the `mysql` deployment.

- 6.1. Before creating a debug pod, retrieve IP address of the `frontend` service.

```
[student@workstation network-sdn]$ oc get service/frontend \
> -o jsonpath="{.spec.clusterIP}{'\n'}"
172.30.23.147
```

- 6.2. Run the `oc debug` command to create a container for troubleshooting based on the `mysql` deployment. You must override the container image because the MySQL Server image does not provide the `curl` command.

```
[student@workstation network-sdn]$ oc debug -t deployment/mysql \
> --image registry.access.redhat.com/ubi8/ubi:8.4
Starting pod/mysql-debug ...
Pod IP: 10.131.0.146
If you don't see a command prompt, try pressing enter.
sh-4.4$
```

- 6.3. Use the `curl` command to connect to the `frontend` application over port 8080. Make sure to replace the IP address with the one that you obtained previously for the `frontend` service.

The following output indicates that the `curl` command times out. This could indicate either that the application is not running or that the service is not able to access the application.

```
sh-4.4$ curl -m 10 -v http://172.30.23.147:8080
* Rebuilt URL to: http://172.30.23.147:8080/
* Trying 172.30.23.147...
* TCP_NODELAY set
* Connection timed out after 10000 milliseconds
* Closing connection 0
curl: (28) Connection timed out after 10000 milliseconds
```

- 6.4. Exit the debug pod.

```
sh-4.4$ exit
exit

Removing debug pod ...
```

- 7. In the following steps, you connect to the `frontend` pod through its private IP. This allows testing, whether or not the issue is related to the service.

- 7.1. Retrieve the IP of the `frontend` pod.

```
[student@workstation network-sdn]$ oc get pods -o wide -l name=frontend
NAME                                READY   STATUS    RESTARTS   AGE   IP              ...
frontend-57b8b445df-f56qh          1/1     Running   0           39m   10.128.2.61    ...
```

- 7.2. Create a debug pod from the `mysql` deployment.

```
[student@workstation network-sdn]$ oc debug -t deployment/mysql \
> --image registry.access.redhat.com/ubi8/ubi:8.4
Starting pod/mysql-debug ...
Pod IP: 10.131.1.27
If you don't see a command prompt, try pressing enter.
sh-4.4$
```

- 7.3. Run the `curl` command in verbose mode against the `frontend` pod on port 8080. Replace the IP address with the one that you obtained previously for the `frontend` pod.

```
sh-4.4$ curl -v http://10.128.2.61:8080/todo/
* Trying 10.128.2.61...
* TCP_NODELAY set
* Connected to 10.128.2.61 (10.128.2.61) port 8080 (#0)
> GET /todo/ HTTP/1.1
> Host: 10.128.2.61:8080
> User-Agent: curl/7.61.1
```



```
> Accept: /
>
< HTTP/1.1 200 OK
...output omitted...
```

The `curl` command can access the application through the pod's private IP.

7.4. Exit the debug pod.

```
sh-4.2$ exit
exit

Removing debug pod ...
```

8. Review the configuration of the `frontend` service.

8.1. List the services in the project and ensure that the `frontend` service exists.

```
[student@workstation network-sdn]$ oc get svc
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
frontend	ClusterIP	172.30.23.147	<none>	8080/TCP	93m
mysql	ClusterIP	172.30.103.29	<none>	3306/TCP	93m

8.2. Review the configuration and status of the `frontend` service. Notice the value of the service selector that indicates to which pod the service should forward packets.

```
[student@workstation network-sdn]$ oc describe svc/frontend
```

```
Name:                frontend
Namespace:           network-sdn
Labels:              app=todonodejs
                    name=frontend
Annotations:         <none>
Selector:            name=api
Type:                ClusterIP
IP:                  172.30.23.147
Port:                <unset> 8080/TCP
TargetPort:          8080/TCP
Endpoints:           <none>
Session Affinity:    None
Events:              <none>
```

This output also indicates that the service has no endpoint, so it is not able to forward incoming traffic to the application.

8.3. Retrieve the labels of the `frontend` deployment. The output shows that pods are created with a `name` label that has a value of `frontend`, whereas the service in the previous step uses `api` as the value.

```
[student@workstation network-sdn]$ oc describe deployment/frontend | \
> grep -A1 Labels
Labels:                app=todonodejs
                   name=frontend
--
Labels: app=todonodejs
       name=frontend
```

► 9. Update the frontend service and access the application.

- 9.1. Run the `oc edit` command to edit the frontend service. Update the selector to match the correct label.

```
[student@workstation network-sdn]$ oc edit svc/frontend
```

Locate the section that defines the selector, and then update the `name: frontend` label inside the selector. After making the changes, exit the editor.

```
...output omitted...
selector:
  name: frontend
...output omitted...
```

Save your changes and verify that the `oc edit` command applied them.

```
service/frontend edited
```

- 9.2. Review the service configuration to ensure that the service has an endpoint.

```
[student@workstation network-sdn]$ oc describe svc/frontend
Name:                frontend
Namespace:           network-sdn
Labels:              app=todonodejs
                   name=frontend
Annotations:         <none>
Selector:            name=frontend
Type:                ClusterIP
IP:                  172.30.169.113
Port:                <unset> 8080/TCP
TargetPort:          8080/TCP
Endpoints:           10.128.2.61:8080
Session Affinity:    None
Events:              <none>
```

- 9.3. From the `workstation` machine, open Firefox and access the To Do application at the following URL.

- `http://todo.apps.ocp4.example.com/todo/`

You should see the To Do application.

To Do List Application

To Do List

Id	Description	Done	
1	Pick up newsp...	false	✖
2	Buy groceries	true	✖

First Previous **1** Next Last

Add Task

Description:

Add Description.

Completed:

☐

Clear Save

10. Perform cleanup.

10.1. Change to the user home directory.

```
[student@workstation network-sdn]$ cd
```

10.2. Delete the `network-sdn` project.

```
[student@workstation ~]$ oc delete project network-sdn
project.project.openshift.io "network-sdn" deleted
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab network-sdn finish
```

This concludes the guided exercise.

Exposing Applications for External Access

Objectives

After completing this section, you should be able to allow and protect network connections to applications inside an OpenShift cluster.

Accessing Applications from External Networks

OpenShift Container Platform offers many ways to expose your applications to external networks. You can expose HTTP and HTTPS traffic, TCP applications, and also non-TCP traffic. Some of these methods are service types, such as NodePort or load balancer, while others use their own API resource, such as Ingress and Route.

OpenShift *routes* allow you to expose your applications to external networks. With routes, you can access your application with a unique host name that is publicly accessible. Routes rely on a router *plug-in* to redirect the traffic from the public IP to pods.

The following diagram shows how a route exposes an application running as pods in your cluster:




Figure 5.5: Using routes to expose applications

**Note**

For performance reasons, routers send requests directly to pods based on service configuration.

The dotted line indicates this implementation. That is, the router accesses the pods through the services network.

Describing Methods for Managing Ingress Traffic

The most common way to manage ingress traffic is with the Ingress Controller. OpenShift implements the Ingress Controller with a shared router service that runs as a pod inside the cluster. You can scale and replicate this pod like any other regular pod. This router service is based on the open source software HAProxy.

Routes and ingress are the main resources for handling ingress traffic:

Route

Routes provide ingress traffic to services in the cluster. Routes were created before Kubernetes ingress objects and provide more features. Routes provide advanced features that may not be supported by Kubernetes ingress controllers through a standard interface, such as TLS re-encryption, TLS passthrough, and split traffic for blue-green deployments.

Ingress

An ingress is a Kubernetes resource that provides some of the same features as routes (which are an OpenShift resource). Ingresses accept external requests and proxy them based on the route. You can only allow certain types of traffic: HTTP, HTTPS and *server name identification* (SNI), and TLS with SNI. In OpenShift, routes are generated to meet the conditions specified by the ingress object.

There are alternatives to ingress and routes, but they are for special use cases. The following service types provide external access to services:

External load balancer

This resource instructs OpenShift to spin up a load balancer in a cloud environment. A load balancer instructs OpenShift to interact with the cloud provider in which the cluster is running to provision a load balancer.

Service external IP

This method instructs OpenShift to set NAT rules to redirect traffic from one of the cluster IPs to the container.

NodePort

With this method, OpenShift exposes a service on a static port on the node IP address. You must ensure that the external IP addresses are properly routed to the nodes.

Creating Routes

The easiest and preferred way to create a route (secure or insecure) is to use the `oc expose service service` command, where `service` corresponds to a service. Use the `--hostname` option to provide a custom host name for the route.

```
[user@host ~]$ oc expose service api-frontend \
> --hostname api.apps.acme.com
```

When you omit the host name, OpenShift generates a host name for you with the following structure: `<route-name>-<project-name>.<default-domain>`. For example, if you create a `frontend` route in an `api` project, in a cluster that uses `apps.example.com` as the wildcard domain, then the route host name will be:

```
frontend-api.apps.example.com.
```



Important

The DNS server that hosts the wildcard domain is unaware of any route host names; it only resolves any name to the configured IPs. Only the OpenShift router knows about route host names, treating each one as an HTTP virtual host.

Invalid wildcard domain host names, that is, host names that do not correspond to any route, are blocked by the OpenShift router and result in an HTTP 404 error.

Consider the following settings when creating a route:

- The name of a service. The route uses the service to determine the pods to which to route the traffic.
- A host name for the route. A route is always a subdomain of your cluster wildcard domain. For example, if you are using a wildcard domain of `apps.dev-cluster.acme.com`, and need to expose a front-end service through a route, then it will be named:

```
frontend.apps.dev-cluster.acme.com.
```



Note

You can also let OpenShift automatically generate a host name for the route.

- An optional path, for path-based routes.
- A target port on which the application listens. The target port usually corresponds to the port that you define in the `targetPort` key of a service.
- An encryption strategy, depending on whether you need a secure or insecure route.

The following listing shows a minimal definition for a route:

```
kind: Route
apiVersion: route.openshift.io/v1
metadata:
  name: a-simple-route ❶
  labels: ❷
    app: API
    name: api-frontend
spec:
  host: api.apps.acme.com ❸
  to:
    kind: Service
```

```
name: api-frontend ❹
port: ❺
targetPort: 8443
```

- ❶ The name of the route. This name must be unique.
- ❷ A set of labels that you can use as selectors.
- ❸ The host name of the route. This host name must be a subdomain of your wildcard domain because OpenShift routes the wildcard domain to the routers.
- ❹ The service to which to redirect the traffic. Although you use a service name, the route only uses this information to determine the list of pods that receive the traffic.
- ❺ The application port. Because routes bypass services, this must match the application port and not the service port.

Securing Routes

Routes can be either secured or unsecured. Secure routes provide the ability to use several types of TLS termination to serve certificates to the client. Unsecured routes are the simplest to configure because they require no key or certificates, but secured routes encrypt traffic to and from the pods.

A secured route specifies the TLS termination of the route. The available types of termination are presented in the following list:

OpenShift Secure Routes

Edge

With edge termination, TLS termination occurs at the router, before the traffic is routed to the pods. The router serves the TLS certificates, so you must configure them into the route; otherwise, OpenShift assigns its own certificate to the router for TLS termination. Because TLS is terminated at the router, connections from the router to the endpoints over the internal network are not encrypted.

Passthrough

With passthrough termination, encrypted traffic is sent straight to the destination pod without the router providing TLS termination. In this mode, the application is responsible for serving certificates for the traffic. Passthrough is currently the only method that supports mutual authentication between the application and a client that accesses it.

Re-encryption

Re-encryption is a variation on edge termination, whereby the router terminates TLS with a certificate, and then re-encrypts its connection to the endpoint, which might have a different certificate. Therefore, the full path of the connection is encrypted, even over the internal network. The router uses health checks to determine the authenticity of the host.

Securing Applications with Edge Routes

Before creating a secure route, you need to generate a TLS certificate. The following command shows how to create a secure edge route with a TLS certificate:

```
[user@host ~]$ oc create route edge \
>   --service api-frontend --hostname api.apps.acme.com \
>   --key api.key --cert api.crt ❶ ❷
```

- 1 The `--key` option requires the certificate private key.
- 2 The `--cert` option requires the certificate that has been signed with that key.

When using a route in edge mode, the traffic between the client and the router is encrypted, but traffic between the router and the application is not:




Figure 5.6: Securing applications with edge routes



Note

Network policies can help you protect the internal traffic between your applications or between projects.

Securing Applications with Passthrough Routes

The previous example demonstrates how to create an edge route, that is, an OpenShift route that presents a certificate at the edge. Passthrough routes offer a secure alternative because the application exposes its TLS certificate. As such, the traffic is encrypted between the client and the application.

To create a passthrough route, you need a certificate and a way for your application to access it. The best way to accomplish this is by using OpenShift TLS secrets. Secrets are exposed via a mount point into the container.

The following diagram shows how you can mount a `secret` resource in your container. The application is then able to access your certificate. With this mode, there is no encryption between the client and the router.




Figure 5.7: Securing applications with passthrough routes



References

For more information on how to manage routes, refer to *Configuring Routes* chapter in the Red Hat OpenShift Container Platform 4.10 *Networking* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/networking/index#configuring-routes

For more information on how to configure ingress cluster traffic, refer to *Configuring ingress cluster traffic* chapter in the Red Hat OpenShift Container Platform 4.10 *Networking* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/networking/index#configuring-ingress-cluster-traffic

► Guided Exercise

Exposing Applications for External Access

In this exercise, you will expose an application secured by TLS certificates.

Outcomes

You should be able to:

- Deploy an application and create an unencrypted route for it.
- Create an OpenShift edge route with encryption.
- Update an OpenShift deployment to support a new version of the application.
- Create an OpenShift TLS secret and mount it to your application.
- Verify that the communication to the application is encrypted.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable, and creates the `network-ingress` OpenShift project. It also gives the `developer` user edit access on the project.

```
[student@workstation ~]$ lab network-ingress start
```

Instructions

As an application developer, you are ready to deploy your application in OpenShift. In this activity, you will deploy two versions of the application, one that is exposed over unencrypted traffic (HTTP), and one that is exposed over secure traffic.

The container image, accessible at <https://quay.io/redhattraining/todo-angular>, has two tags: `v1.1`, which is the insecure version of the application, and `v1.2`, which is the secure version. Your organization uses its own certificate authority (CA) that can sign certificates for the following domains.

- `*.apps.ocp4.example.com`
- `*.ocp4.example.com`

The CA certificate is accessible at `~/DO280/labs/network-ingress/certs/training-CA.pem`. The `passphrase.txt` contains a unique password that protects the CA key. The `certs` folder also contains the CA key.

- 1. Log in to the OpenShift cluster and create the `network-ingress` project.

- 1.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create the `network-ingress` project.

```
[student@workstation ~]$ oc new-project network-ingress
Now using project "network-ingress" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

- ▶ 2. The OpenShift deployment file for the application is accessible at `~/D0280/labs/network-ingress/todo-app-v1.yaml`. The deployment points to `quay.io/redhattraining/todo-angular:v1.1`, which is the initial and unencrypted version of the application. The file defines the `todo-http` service that points to the application pod. Create the application and expose the service.

- 2.1. Use the `oc create` command to deploy the application in the `network-ingress` OpenShift project.

```
[student@workstation ~]$ oc create -f \
> ~/D0280/labs/network-ingress/todo-app-v1.yaml
deployment.apps/todo-http created
service/todo-http created
```

- 2.2. Wait a couple of minutes, so that the application can start, and then review the resources in the project.

```
[student@workstation ~]$ oc status
In project network-ingress on server https://api.ocp4.example.com:6443

svc/todo-http - 172.30.247.75:80 -> 8080
  deployment/todo-http deploys quay.io/redhattraining/todo-angular:v1.1
  deployment #1 running for 16 seconds - 1 pod
...output omitted...
```

- 2.3. Run the `oc expose` command to create a route for accessing the application. Give the route a host name of `todo-http.apps.ocp4.example.com`.

```
[student@workstation ~]$ oc expose svc todo-http \
> --hostname todo-http.apps.ocp4.example.com
route.route.openshift.io/todo-http exposed
```

- 2.4. Retrieve the name of the route and copy it to the clipboard.

```
[student@workstation ~]$ oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT	...
todo-http	todo-http.apps.ocp4.example.com		todo-http	8080	...

- 2.5. On the **workstation** machine, open Firefox and access the application URL. Confirm that you can see the application.
 - `http://todo-http.apps.ocp4.example.com`
- 2.6. Open a new terminal tab and run the `tcpdump` command with the following options to intercept the traffic on port 80:
 - `-i eth0` intercepts traffic on the main interface.
 - `-A` strips the headers and prints the packets in ASCII format.
 - `-n` disables DNS resolution.
 - `port 80` is the port of the application.

Optionally, the `grep` command allows you to filter on JavaScript resources.

Start by retrieving the name of the main interface whose IP is `172.25.250.9`.

```
[student@workstation ~]$ ip addr | grep 172.25.250.9
inet 172.25.250.9/24 brd 172.25.250.255 scope global noprefixroute eth0

[student@workstation ~]$ sudo tcpdump -i eth0 -A -n port 80 | grep "angular"
```



Note

The full command is available at `~/D0280/labs/network-ingress/tcpdump-command.txt`.

- 2.7. On Firefox, refresh the page and notice the activity in the terminal. Press `Ctrl+C` to stop the capture.

```
...output omitted...
<script type="text/javascript" src="assets/js/libs/angular/angular.min.js">
<script type="text/javascript" src="assets/js/libs/angular/angular-route.min.js">
<script type="text/javascript" src="assets/js/libs/angular/angular-
animate.min.js">
...output omitted...
```

- ▶ 3. Create a secure edge route. Edge certificates encrypt the traffic between the client and the router, but leave the traffic between the router and the service unencrypted. OpenShift generates its own certificate that it signs with its CA.

In later steps, you extract the CA to ensure the route certificate is signed.

- 3.1. Go to `~/D0280/labs/network-ingress` and run the `oc create route` command to define the new route.


Give the route a host name of `todo-https.apps.ocp4.example.com`.

```
[student@workstation ~]$ cd ~/D0280/labs/network-ingress
[student@workstation network-ingress]$ oc create route edge todo-https \
> --service todo-http \
> --hostname todo-https.apps.ocp4.example.com
route.route.openshift.io/todo-https created
```


- 3.2. To test the route and read the certificate, open Firefox and access the application URL.

- `https://todo-https.apps.ocp4.example.com`


Click on the **padlock**, and then click on the arrow next to **Connection secure**.




Click **More Information** to display the page information window.



Click **View Certificate** to display the certificate information.



Locate the CN entry to see that the OpenShift ingress operator created the certificate with its own CA.



- 3.3. From the terminal, use the `curl` command with the `-I` and `-v` options to retrieve the connection headers.

The `Server certificate` section shows some information about the certificate and the alternative name matches the name of the route. The output indicates that the remote certificate is trusted because it matches the CA.

```
[student@workstation network-ingress]$ curl -I -v \
> https://todo-https.apps.ocp4.example.com
...output omitted...
* Server certificate:
*  subject: O=EXAMPLE.COM; CN=.api.ocp4.example.com
*  start date: May 10 11:18:41 2021 GMT
*  expire date: May 10 11:18:41 2026 GMT
*  subjectAltName: host "todo-https.apps.ocp4.example.com" matched cert's
    "*.apps.ocp4.example.com"
*  issuer: O=EXAMPLE.COM; CN=Red Hat Training Certificate Authority
*  SSL certificate verify ok.
...output omitted...
```

- 3.4. Although the traffic is encrypted at the edge with a certificate, you can still access the insecure traffic at the service level because the pod behind the service does not offer an encrypted route.

Retrieve the IP address of the `todo-http` service.

```
[student@workstation network-ingress]$ oc get svc todo-http \
> -o jsonpath="{.spec.clusterIP}"
172.30.102.29
```

- 3.5. Create a debug pod in the `todo-http` deployment. Use the Red Hat Universal Base Image (UBI), which contains some basic tools to interact with containers.

```
[student@workstation network-ingress]$ oc debug -t deployment/todo-http \
> --image registry.access.redhat.com/ubi8/ubi:8.4
Starting pod/todo-http-debug ...
Pod IP: 10.131.0.255
If you don't see a command prompt, try pressing enter.
sh-4.4$
```

- 3.6. From the debug pod, use the `curl` command to access the service over HTTP. Replace the IP address with the one that you obtained in a previous step. The output indicates that the application is available over HTTP.

```
sh-4.4$ curl -v 172.30.102.29
* Rebuilt URL to: 172.30.102.29/
* Trying 172.30.102.29...
* TCP_NODELAY set
* Connected to 172.30.102.29 (172.30.102.29) port 80 (#0)
> GET / HTTP/1.1
> Host: 172.30.102.29
> User-Agent: curl/7.61.1
> Accept: /
>
< HTTP/1.1 200 OK
...output omitted...
```

- 3.7. Exit the debug pod.

```
sh-4.4$ exit
Removing debug pod ...
```

3.8. Delete the edge route. In the next steps, you define the passthrough route.

```
[student@workstation network-ingress]$ oc delete route todo-https
route.route.openshift.io "todo-https" deleted
```

► **4.** Generate TLS certificates for the application.

In the following steps, you generate a CA-signed certificate that you attach as a secret to the pod. You then configure a secure route in passthrough mode and let the application expose that certificate.

4.1. Go to the `~/D0280/labs/network-ingress/certs` directory and list the files.

```
[student@workstation network-ingress]$ cd certs
[student@workstation certs]$ ls -l
total 20
-rw-rw-r--. 1 student student 604 Nov 29 17:35 openssl-commands.txt
-rw-r--r--. 1 student student 33 Nov 29 17:35 passphrase.txt
-rw-r--r--. 1 student student 1743 Nov 29 17:35 training-CA.key
-rw-r--r--. 1 student student 1363 Nov 29 17:35 training-CA.pem
-rw-r--r--. 1 student student 406 Nov 29 17:35 training.ext
```

4.2. Generate the private key for your CA-signed certificate.



Note

The following commands for generating a signed certificate are all available in the `~/D0280/labs/network-ingress/certs/openssl-commands.txt` file.

```
[student@workstation certs]$ openssl genrsa -out training.key 4096
Generating RSA private key, 4096 bit long modulus (2 primes)
...output omitted...
e is 65537 (0x010001)
```

4.3. Generate the certificate signing request (CSR) for `todo-https.apps.ocp4.example.com`.

```
[student@workstation certs]$ openssl req -new \
> -key training.key -out training.csr \
> -subj "/C=US/ST=North Carolina/L=Raleigh/O=Red Hat/\
> CN=todo-https.apps.ocp4.example.com"
```


**Warning**

Make sure to type the **subject** of the request on one line. Alternatively, remove the `-subj` option and its content. Without the `-subj` option, the `openssl` command prompts you for the values; make sure to indicate a common name (CN) of `todo-https.apps.ocp4.example.com`.

- 4.4. Finally, generate the signed certificate. Notice the use of the `-CA` and `-CAkey` options for signing the certificate against the CA. The `-passin` option allows you to reuse the password of the CA. The `extfile` option allows you to define a *Subject Alternative Name (SAN)*.

```
[student@workstation certs]$ openssl x509 -req -in training.csr \
> -passin file:passphrase.txt \
> -CA training-CA.pem -CAkey training-CA.key -CAcreateserial \
> -out training.crt -days 1825 -sha256 -extfile training.ext
Signature ok
subject=C = US, ST = North Carolina, L = Raleigh, O = Red Hat, CN = todo-
https.apps.ocp4.example.com
Getting CA Private Key
```

- 4.5. Ensure that the newly created certificate and key are present in the current directory.

```
[student@workstation certs]$ ls -lrt
total 36
-rw-r--r--. 1 student student 599 Jul 31 09:35 openssl-commands.txt
-rw-r--r--. 1 student student 33 Aug 3 12:38 passphrase.txt
-rw-r--r--. 1 student student 352 Aug 3 12:38 training.ext
-rw-----. 1 student student 1743 Aug 3 12:38 training-CA.key
-rw-r--r--. 1 student student 1334 Aug 3 12:38 training-CA.pem
-rw-----. 1 student student 1675 Aug 3 13:38 training.key
-rw-rw-r--. 1 student student 1017 Aug 3 13:39 training.csr
-rw-rw-r--. 1 student student 41 Aug 3 13:40 training-CA.srl
-rw-rw-r--. 1 student student 1399 Aug 3 13:40 training.crt
```

- 4.6. Return to the `network-ingress` directory. This is important as the next step involves the creation of a route using the self-signed certificate.

```
[student@workstation certs]$ cd ~/D0280/labs/network-ingress
```

► **5.** Deploy a new version of your application.

The new version of the application expects a certificate and a key inside the container at `/usr/local/etc/ssl/certs`. The web server in that version is configured with SSL support. Create a secret to import the certificate from the `workstation` machine. In a later step, the application deployment requests that secret and exposes its content to the container at `/usr/local/etc/ssl/certs`.

- 5.1. Create a `tls` OpenShift secret named `todo-certs`. Use the `--cert` and `--key` options to embed the TLS certificates. Use `training.crt` as the certificate, and `training.key` as the key.

```
[student@workstation network-ingress]$ oc create secret tls todo-certs \
> --cert certs/training.crt --key certs/training.key
secret/todo-certs created
```

- 5.2. The deployment file, accessible at `~/D0280/labs/network-ingress/todo-app-v2.yaml`, points to version 2 of the container image. The new version of the application is configured to support SSL certificates. Run `oc create` to create a new deployment using that image.

```
[student@workstation network-ingress]$ oc create -f todo-app-v2.yaml
deployment.apps/todo-https created
service/todo-https created
```

- 5.3. Wait a couple of minutes to ensure that the application pod is running. Copy the pod name to your clipboard.

```
[student@workstation network-ingress]$ oc get pods
NAME                                READY   STATUS    RESTARTS   AGE
...output omitted...
todo-https-59d8fc9d47-265ds        1/1     Running   0           62s
```

- 5.4. Review the volumes that are mounted inside the pod. The output indicates that the certificates are mounted to `/usr/local/etc/ssl/certs`.

```
[student@workstation network-ingress]$ oc describe pod \
> todo-https-59d8fc9d47-265ds | grep -A2 Mounts
Mounts:
  /usr/local/etc/ssl/certs from tls-certs (ro)
  /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-qrkqj (ro)
```

► 6. Create the secure route.

- 6.1. Run the `oc create route` command to define the new route.
Give the route a host name of `todo-https.apps.ocp4.example.com`.

```
[student@workstation network-ingress]$ oc create route passthrough todo-https \
> --service todo-https --port 8443 \
> --hostname todo-https.apps.ocp4.example.com
route.route.openshift.io/todo-https created
```

- 6.2. Use the `curl` command in verbose mode to test the route and read the certificate. Use the `--cacert` option to pass the CA certificate to the `curl` command.
The output indicates a match between the certificate chain and the application certificate. This match indicates that the OpenShift router only forwards packets that are encrypted by the application web server certificate.

```
[student@workstation network-ingress]$ curl -vv -I \
> --cacert certs/training-CA.pem \
> https://todo-https.apps.ocp4.example.com
...output omitted...
```

```
* Server certificate:
*   subject: C=US; ST=North Carolina; L=Raleigh; O=Red Hat; CN=todo-https.apps.ocp4.example.com
*   start date: Jun 15 01:53:30 2021 GMT
*   expire date: Jun 14 01:53:30 2026 GMT
*   subjectAltName: host "todo-https.apps.ocp4.example.com" matched cert's "*.apps.ocp4.example.com"
*   issuer: C=US; ST=North Carolina; L=Raleigh; O=Red Hat; CN=ocp4.example.com
*   SSL certificate verify ok.
...output omitted...
```

- 7. Create a new debug pod to further confirm proper encryption at the service level.

7.1. Retrieve the IP address of the `todo-https` service.

```
[student@workstation network-ingress]$ oc get svc todo-https \
> -o jsonpath="{.spec.clusterIP}{'\n'}"
172.30.121.154
```

7.2. Create a debug pod in the `todo-https` deployment with the Red Hat UBI.

```
[student@workstation network-ingress]$ oc debug -t deployment/todo-https \
> --image registry.access.redhat.com/ubi8/ubi:8.4
Starting pod/todo-https-debug ...
Pod IP: 10.128.2.129
If you don't see a command prompt, try pressing enter.
sh-4.4$
```

7.3. From the debug pod, use the `curl` command to access the service over HTTP. Replace the IP address with the one that you obtained in a previous step.

The output indicates that the application is not available over HTTP, and the web server redirects you to the secure version.

```
sh-4.4$ curl -I http://172.30.121.154
HTTP/1.1 301 Moved Permanently
Server: nginx/1.14.1
Date: Tue, 15 Jun 2021 02:01:19 GMT
Content-Type: text/html
Connection: keep-alive
Location: https://172.30.121.154:8443/
```

7.4. Finally, access the application over HTTPS. Use the `-k` option because the container does not have access to the CA certificate.

```
sh-4.4$ curl -s -k https://172.30.121.154:8443 | head -n5
<!DOCTYPE html>
<html lang="en" ng-app="todoItemsApp" ng-controller="appCtl">
<head>
  <meta charset="utf-8">
  <title>ToDo app</title>
```

7.5. Exit the debug pod.

```
sh-4.4$ exit  
Removing debug pod ...
```

▶ **8.** Perform cleanup.

8.1. Change to the home directory.

```
[student@workstation network-ingress]$ cd
```

8.2. Delete the network-ingress project.

```
[student@workstation ~]$ oc delete project network-ingress  
project.project.openshift.io "network-ingress" deleted
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab network-ingress finish
```

This concludes the guided exercise.

Configuring Network Policies

Objectives

After completing this section, you should be able to restrict network traffic between projects and pods.

Managing Network Policies in OpenShift

Network policies allow you to configure isolation policies for individual pods. Network policies do not require administrative privileges, giving developers more control over the applications in their projects. You can use network policies to create logical zones in the SDN that map to your organization network zones. The benefit of this approach is that the location of running pods becomes irrelevant because network policies allow you to segregate traffic regardless of where it originates.

To manage network communication between two namespaces, assign a label to the namespace that needs access to another namespace. The following command assigns the `name=network-1` label to the `network-1` namespace:

```
[user@host ~]$ oc label namespace network-1 name=network-1
```

The following examples describe network policies that allow communication between the `network-1` and `network-2` namespaces:

- The following network policy applies to all pods with the label `deployment="product-catalog"` in the `network-1` namespace. The policy allows TCP traffic over port 8080 from pods whose label is `role="qa"` in the `network-2` namespace.

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: network-1-policy
spec:
  podSelector: ❶
    matchLabels:
      deployment: product-catalog
  ingress: ❷
    - from: ❸
      - namespaceSelector:
          matchLabels:
            name: network-2
        podSelector:
          matchLabels:
            role: qa
  ports: ❹
    - port: 8080
      protocol: TCP
```

- ❶ The top-level `podSelector` field is required and defines which pods use the network policy. If the `podSelector` is empty, all pods in the namespace are matched.
 - ❷ The `ingress` field defines a list of ingress traffic rules to apply to the matched pods from the top-level `podSelector`.
 - ❸ The `from` field defines a list of rules to match traffic from all sources. The selectors are not limited to the project in which the network policy is defined.
 - ❹ The `ports` field is a list of destination ports that allow traffic to reach the selected pods.
- The following network policy allows traffic from all the pods and ports in the `network-1` namespace to all pods and ports in the `network-2` namespace. This policy is less restrictive than the `network-1` policy, because it does not restrict traffic from any pods in the `network-1` namespace.

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: network-2-policy
spec:
  podSelector: {}
  ingress:
    - from:
      - namespaceSelector:
          matchLabels:
            name: network-1
```

**Note**

Network policies are Kubernetes resources. As such, you can manage them using `oc` commands.

One benefit of using network policies is the management of security between projects (tenants) that you cannot do with layer 2 technologies, such as VLANs. This approach allows you to create tailored policies between projects to make sure users can only access what they should (which conforms to the least privilege approach).

The fields in the network policy that take a list of objects can either be combined in the same object or listed as multiple objects. If combined, the conditions are combined with a logical *AND*. If separated in a list, the conditions are combined with a logical *OR*. The logic options allow you to create very specific policy rules. The following examples highlight the differences the syntax can make:

- This example combines the selectors into one rule, thereby only allowing access from pods in the `dev` namespace with the `app=mobile` label. This is an example of a logical *AND*.

```
...output omitted...
ingress:
  - from:
    - namespaceSelector:
        matchLabels:
          name: dev
```

```

podSelector:
  matchLabels:
    app: mobile

```

- By changing the `podSelector` field in the previous example to be an item in the `from` list, all pods in the `dev` namespace and all pods from any namespace labeled `app=mobile` can reach the pods that match the top-level `podSelector` field. This is an example of a logical OR.

```

...output omitted...
ingress:
- from:
  - namespaceSelector:
    matchLabels:
      name: dev
  - podSelector:
    matchLabels:
      app: mobile

```

If a pod is matched by selectors in one or more network policies, then the pod will only accept connections that are allowed by at least one of those network policies. A strict example is a policy to deny all ingress traffic to pods in your project, including from other pods inside your project. An empty pod selector means that this policy applies to all pods in this project. The following policy blocks all traffic because no ingress rules are defined. Traffic is blocked unless you also define an explicit policy that overrides this default behavior.

```

kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: default-deny
spec:
  podSelector: {}

```

If you have Cluster Monitoring or exposed routes, then you need to allow ingress from them as well. The following policies allow ingress from OpenShift monitoring and Ingress Controllers:

```

---
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-from-openshift-ingress
spec:
  podSelector: {}
  ingress:
  - from:
    - namespaceSelector:
      matchLabels:
        network.openshift.io/policy-group: ingress
---
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-from-openshift-monitoring
spec:

```

```
podSelector: {}
ingress:
- from:
  - namespaceSelector:
      matchLabels:
        network.openshift.io/policy-group: monitoring
```



Important

If the default Ingress Controller uses the `HostNetwork` endpoint publishing strategy, then the default namespace requires the `network.openshift.io/policy-group=ingress` label.

Check the endpoint publishing strategy using the `oc describe` command to describe the `ingresscontroller/default` resource in the `openshift-ingress-operator` namespace.



References

For more information about network policy, refer to the *Network policy* chapter in the Red Hat OpenShift Container Platform 4.10 *Networking* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/networking/index#network-policy

► Guided Exercise

Configuring Network Policies

In this exercise, you will create network policies and review pod isolation created by these network policies.

Outcomes

You should be able to:

- Create network policies to control communication between pods.
- Verify ingress traffic is limited to pods.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the environment is ready and downloads the resource files necessary for the exercise.

```
[student@workstation ~]$ lab network-policy start
```

Instructions

- 1. Log in to the OpenShift cluster and create the `network-policy` project.

- 1.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create the `network-policy` project.

```
[student@workstation ~]$ oc new-project network-policy
Now using project "network-policy" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

- 2. Create two deployments and create a route for one of them.

- 2.1. Create the `hello` deployment using the `quay.io/redhattraining/hello-world-nginx:v1.0` container image.

```
[student@workstation ~]$ oc new-app --name hello \
> --image quay.io/redhattraining/hello-world-nginx:v1.0
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "hello" created
    deployment.apps "hello" created
    service "hello" created
--> Success
...output omitted...
```

- 2.2. Create the test deployment using the `quay.io/redhattraining/hello-world-nginx:v1.0` container image.

```
[student@workstation ~]$ oc new-app --name test \
> --image quay.io/redhattraining/hello-world-nginx:v1.0
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "test" created
    deployment.apps "test" created
    service "test" created
--> Success
...output omitted...
```

- 2.3. Use the `oc expose` command to create a route to the `hello` service.

```
[student@workstation ~]$ oc expose service hello
route.route.openshift.io/hello exposed
```

- ▶ 3. Verify access to the `hello` pod with the `oc rsh` and `curl` commands.

- 3.1. Open a second terminal and run the script located at `~/D0280/labs/network-policy/display-project-info.sh`. This script provides information about the pods, service, and route used in the rest of this exercise.

```
[student@workstation ~]$ ~/D0280/labs/network-policy/display-project-info.sh
=====
PROJECT: network-policy

POD NAME                IP ADDRESS
hello-6c4984d949-g28c4  10.8.0.13
test-c4d74c9d5-5pq9s    10.8.0.14

SERVICE NAME    CLUSTER-IP
hello            172.30.137.226
test            172.30.159.119

ROUTE NAME    HOSTNAME                                PORT
hello        hello-network-policy.apps.ocp4.example.com  8080-tcp
=====
```

- 3.2. Use the `oc rsh` and `curl` commands to confirm that the `test` pod can access the IP address of the `hello` pod.

```
[student@workstation ~]$ oc rsh test-c4d74c9d5-5pq9s \
> curl 10.8.0.13:8080 | grep Hello
<h1>Hello, world from nginx!</h1>
```

- 3.3. Use the `oc rsh` and `curl` commands to confirm that the test pod can access the IP address of the hello service.

```
[student@workstation ~]$ oc rsh test-c4d74c9d5-5pq9s \
> curl 172.30.137.226:8080 | grep Hello
<h1>Hello, world from nginx!</h1>
```

- 3.4. Verify access to the hello pod using the `curl` command against the URL of the hello route.

```
[student@workstation ~]$ curl -s hello-network-policy.apps.ocp4.example.com | \
> grep Hello
<h1>Hello, world from nginx!</h1>
```

► 4. Create the `network-test` project and a deployment named `sample-app`.

- 4.1. Create the `network-test` project.

```
[student@workstation ~]$ oc new-project network-test
Now using project "network-test" on server "https://api.ocp4.example.com:6443".
...output omitted...
```

- 4.2. Create the `sample-app` deployment with the `quay.io/redhattraining/hello-world-nginx:v1.0` image. The web app listens on port 8080.

```
[student@workstation ~]$ oc new-app --name sample-app \
> --image quay.io/redhattraining/hello-world-nginx:v1.0
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "sample-app" created
    deployment.apps "sample-app" created
    service "sample-app" created
--> Success
...output omitted...
```

► 5. Verify that pods in a different namespace can access the hello and test pods in the `network-policy` namespace.

- 5.1. In the second terminal, run the `display-project-info.sh` script again to view the full name of the `sample-app` pod.

```
[student@workstation ~]$ ~/D0280/labs/network-policy/display-project-info.sh
=====
PROJECT: network-policy

POD NAME                IP ADDRESS
```

```
hello-6c4984d949-g28c4    10.8.0.13
test-c4d74c9d5-5pq9s     10.8.0.14
```

```
SERVICE NAME    CLUSTER-IP
hello           172.30.137.226
test           172.30.159.119
```

```
ROUTE NAME      HOSTNAME                                PORT
hello          hello-network-policy.apps.ocp4.example.com  8080-tcp
=====
```

```
PROJECT: network-test
```

```
POD NAME
sample-app-d5f945-spx9q
=====
```

- 5.2. Returning to the first terminal, use the `oc rsh` and `curl` commands to confirm that the `sample-app` pod can access the IP address of the `hello` pod.

```
[student@workstation ~]$ oc rsh sample-app-d5f945-spx9q \
> curl 10.8.0.13:8080 | grep Hello
<h1>Hello, world from nginx!</h1>
```

- 5.3. Use the `oc rsh` and `curl` commands to confirm access to the `test` pod from the `sample-app` pod. Target the IP address previously retrieved for the `test` pod.

```
[student@workstation ~]$ oc rsh sample-app-d5f945-spx9q \
> curl 10.8.0.14:8080 | grep Hello
<h1>Hello, world from nginx!</h1>
```

- 6. From the `network-policy` project, create the `deny-all` network policy using the resource file available at `~/D0280/labs/network-policy/deny-all.yaml`.

- 6.1. Switch to the `network-policy` project.

```
[student@workstation ~]$ oc project network-policy
Now using project "network-policy" on server "https://api.ocp4.example.com:6443".
```

- 6.2. Go to the `~/D0280/labs/network-policy` directory.

```
[student@workstation ~]$ cd ~/D0280/labs/network-policy
```

- 6.3. Use a text editor to update the `deny-all.yaml` file with an empty `podSelector` to target all pods in the namespace.

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: deny-all
spec:
  podSelector: {}
```

**Note**

A solution is provided at `~/DO280/solutions/network-policy/deny-all.yaml`.

6.4. Create the network policy with the `oc create` command.

```
[student@workstation network-policy]$ oc create -f deny-all.yaml
networkpolicy.networking.k8s.io/deny-all created
```

► 7. Verify there is no longer access to the pods in the `network-policy` namespace.

7.1. Verify there is no longer access to the `hello` pod via the exposed route. Wait a few seconds, and then press `Ctrl+C` to exit the `curl` command that does not reply.

```
[student@workstation network-policy]$ curl -s \
>   hello-network-policy.apps.ocp4.example.com | grep Hello
^C
```

**Important**

If the `hello` pod lands on the same node as a `router-default` pod, then the `curl` command works when traffic goes through that router pod. This is only the case with three-node clusters. In a traditional OpenShift cluster, where the control plane or infrastructure nodes are separated from the compute nodes, the network policy would apply to all the router pods in the cluster.

If the `curl` command succeeds, then run the command again to validate the network policy works as expected. This second attempt should go through the other router pod in the cluster.

7.2. Verify that the `test` pod can no longer access the IP address of the `hello` pod. Wait a few seconds, and then press `Ctrl+C` to exit the `curl` command that does not reply.

```
[student@workstation network-policy]$ oc rsh test-c4d74c9d5-5pq9s \
>   curl 10.8.0.13:8080 | grep Hello
^C
command terminated with exit code 130
```

7.3. Switch to the `network-test` project.

```
[student@workstation network-policy]$ oc project network-test
Now using project "network-test" on server "https://api.ocp4.example.com:6443".
```

7.4. Confirm that the `sample-app` pod can no longer access the IP address of the `test` pod. Wait a few seconds, and then press `Ctrl+C` to exit the `curl` command that does not reply.

```
[student@workstation network-policy]$ oc rsh sample-app-d5f945-spx9q \
> curl 10.8.0.14:8080 | grep Hello
^C
command terminated with exit code 130
```

- ▶ 8. Create a network policy to allow traffic to the `hello` pod in the `network-policy` namespace from the `sample-app` pod in the `network-test` namespace over TCP on port 8080. Use the resource file available at `~/D0280/labs/network-policy/allow-specific.yaml`.

8.1. Use a text editor to replace the `CHANGE_ME` sections in the `allow-specific.yaml` file as follows.

```
...output omitted...
spec:
  podSelector:
    matchLabels:
      deployment: hello
  ingress:
    - from:
      - namespaceSelector:
          matchLabels:
            name: network-test
        podSelector:
          matchLabels:
            deployment: sample-app
      ports:
        - port: 8080
          protocol: TCP
```



Note

A solution has been provided at `~/D0280/solutions/network-policy/allow-specific.yaml`.

8.2. Create the network policy with the `oc create` command.

```
[student@workstation network-policy]$ oc create -n network-policy -f \
> allow-specific.yaml
networkpolicy.networking.k8s.io/allow-specific created
```

8.3. View the network policies in the `network-policy` namespace.

```
[student@workstation network-policy]$ oc get networkpolicies -n network-policy
```

NAME	POD-SELECTOR	AGE
allow-specific	deployment=hello	11s
deny-all	<none>	5m6s

- ▶ 9. As the `admin` user, label the `network-test` namespace with the `name=network-test` label.

9.1. Log in as the `admin` user.

```
[student@workstation network-policy]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

9.2. Use the `oc label` command to apply the `name=network-test` label.

```
[student@workstation network-policy]$ oc label namespace network-test \
>   name=network-test
namespace/network-test labeled
```



Important

The `allow-specific` network policy uses labels to match the name of a namespace. By default, namespaces and projects do not get any labels automatically.

9.3. Confirm the label was applied.

```
[student@workstation network-policy]$ oc describe namespace network-test
Name:         network-test
Labels:       name=network-test
...output omitted...
```

9.4. Log in as the `developer` user

```
[student@workstation network-policy]$ oc login -u developer -p developer
Login successful.

...output omitted...
```

- ▶ 10. Verify that the `sample-app` pod can access the IP address of the `hello` pod, but cannot access the IP address of the `test` pod.

10.1. Switch to the `network-test` project.

```
[student@workstation network-policy]$ oc project network-test
Already on project "network-test" on server "https://api.ocp4.example.com:6443".
```

10.2. Verify access to the `hello` pod in the `network-policy` namespace.

```
[student@workstation network-policy]$ oc rsh sample-app-d5f945-spx9q \
>   curl 10.8.0.13:8080 | grep Hello
<h1>Hello, world from nginx!</h1>
```

10.3. Verify there is no response from the `hello` pod on another port. Because the network policy only allows access to port 8080 on the `hello` pod, requests made to any other port are ignored and eventually time out. Wait a few seconds, and then press `Ctrl+C` to exit the `curl` command that does not reply.

```
[student@workstation network-policy]$ oc rsh sample-app-d5f945-spx9q \
> curl 10.8.0.13:8181 | grep Hello
^C
command terminated with exit code 130
```

- 10.4. Verify there is no access to the test pod. Wait a few seconds, and then press Ctrl+C to exit the curl command that does not reply.

```
[student@workstation network-policy]$ oc rsh sample-app-d5f945-spx9q \
> curl 10.8.0.14:8080 | grep Hello
^C
command terminated with exit code 130
```

11. Create a network policy to allow traffic to the hello pod from the exposed route. Use the resource file available at ~/D0280/labs/network-policy/allow-from-openshift-ingress.yaml.

- 11.1. Use a text editor to replace the CHANGE_ME values in the allow-from-openshift-ingress.yaml file as follows.

```
...output omitted...
spec:
  podSelector: {}
  ingress:
  - from:
    - namespaceSelector:
        matchLabels:
            network.openshift.io/policy-group: ingress
```



Note

A solution is provided at ~/D0280/solutions/network-policy/allow-from-openshift-ingress.yaml.

- 11.2. Create the network policy with the oc create command.

```
[student@workstation network-policy]$ oc create -n network-policy -f \
> allow-from-openshift-ingress.yaml
networkpolicy.networking.k8s.io/allow-from-openshift-ingress created
```

- 11.3. View the network policies in the network-policy namespace.

```
[student@workstation network-policy]$ oc get networkpolicies -n network-policy
```

NAME	POD-SELECTOR	AGE
allow-from-openshift-ingress	<none>	10s
allow-specific	deployment=hello	8m16s
deny-all	<none>	13m

- 11.4. Log in as the admin user.


```
[student@workstation network-policy]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

- 11.5. Apply the `network.openshift.io/policy-group=ingress` label to the default namespace.

```
[student@workstation network-policy]$ oc label namespace default \
>   network.openshift.io/policy-group=ingress
namespace/default labeled
```

**Note**

Applying this label to the `default` namespace is only necessary because the classroom's default Ingress Controller uses the `HostNetwork` endpoint publishing strategy.

- 11.6. Verify access to the `hello` pod via the exposed route.

```
[student@workstation network-policy]$ curl -s \
>   hello-network-policy.apps.ocp4.example.com | grep Hello
<h1>Hello, world from nginx!</h1>
```

- ▶ 12. Close the terminal window that has the output of the `display-project-info.sh` script. Navigate to the home directory.

```
[student@workstation network-policy]$ cd
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab network-policy finish
```

This concludes the guided exercise.

► Lab

Configuring OpenShift Networking for Applications

In this lab, you will configure a TLS passthrough route for your application.

Outcomes

You should be able to:

- Deploy an application and configure an insecure route.
- Restrict traffic to the applications.
- Generate a TLS certificate for an application.
- Configure a passthrough route for an application with a TLS certificate.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates the self-signed certificate authority (CA) that you use in this lab.

```
[student@workstation ~]$ lab network-review start
```

Instructions

In this review, you deploy a PHP application that prints some information about the system. The application is available with two different configurations: one that runs with an unencrypted network that listens on port 8080, and one that uses a TLS certificate to encrypt the network traffic, which listens on port 8443.

The container image for this review is accessible at `quay.io/redhattraining/php-ssl`. It has two tags: `v1.0` for the insecure version of the application, and `v1.1` for the secure version.

1. As the OpenShift `developer` user, create the `network-review` project.
2. As the `developer` user, deploy the insecure version of the PHP application to the `network-review` project. Use the resource file available at `~/D0280/labs/network-review/php-http.yaml`.

Make the necessary changes to the file, use the `quay.io/redhattraining/php-ssl:v1.0` container image and set the port to 8080.

After creating the application, wait a few moments to ensure that one pod is running.

3. Create a route named `php-http`, with a hostname of `php-http.apps.ocp4.example.com`, to access the application.

From the `workstation` machine, use Firefox to access the application URL:

- `http://php-http.apps.ocp4.example.com/`.

Confirm the availability of the application before proceeding to the next step.

4. Create a network policy in the `network-review` namespace to deny all ingress traffic by default. When configured correctly, the network policy also prevents pods within the `network-review` namespace from communicating with each other.
Use the resource file available at `~/D0280/labs/network-review/deny-all.yaml`. Make the necessary changes to target all pods in the namespace.
5. Create a network policy to allow ingress traffic to routes in the `network-review` namespace.
Use the resource file available at `~/D0280/labs/network-review/allow-from-openshift-ingress.yaml`. Make the necessary changes to target all pods in the namespace and allow traffic from the default ingress controller.
Because the classroom environment uses the `HostNetwork` endpoint strategy, label the `default` namespace with the `network.openshift.io/policy-group=ingress` label. This action must be performed as the `admin` user.
6. As the `developer` user, generate and sign a TLS certificate for the encrypted version of the application.
Create a certificate signing request (CSR) for the `php-https.apps.ocp4.example.com` hostname. Save the CSR to `~/D0280/labs/network-review/certs/training.csr`.
Use the CSR to generate a certificate and save it to `~/D0280/labs/network-review/certs/training.crt`. To generate the certificate, pass as arguments the CA certificate accessible at `~/D0280/labs/network-review/certs/training-CA.pem` and the CSR.
You can use the `~/D0280/labs/network-review/certs/openssl-commands.txt` text file for help. This file contains the commands for generating the certificate signing request and the certificate. Make sure to replace the values in the file before copying and running the OpenSSL commands.
7. Create an OpenShift TLS secret named `php-certs` in the `network-review` project. Use the `~/D0280/labs/network-review/certs/training.crt` file for the certificate and the `~/D0280/labs/network-review/certs/training.key` file for the key.
8. Use the resource file, available at `~/D0280/labs/network-review/php-https.yaml`, to deploy the secure version of the PHP application. Deploy the application to the `network-review` project.
Before deploying the application, make the necessary changes to the resources file, specifically:
 - The location of the container.
 - The port the application listens on.
 - The name of the secret to mount as a volume.
9. Create a secure passthrough route named `php-https`, with a hostname of `php-https.apps.ocp4.example.com`, to access the secure version of the application.
From the `workstation` machine, use Firefox to access the application URL:
 - `https://php-https.apps.ocp4.example.com/`.
 Accept the signed certificate and confirm the availability of the application.
10. **Optional step:** from the `workstation` machine, use the `curl` command to access the HTTPS version of the application.
Pass the CA certificate to the `curl` command to validate the secure connection.

11. Return to the home directory.

```
[student@workstation network-review]$ cd
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab network-review grade
```

Finish

As the `student` user on the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab network-review finish
```

This concludes the lab.

► Solution

Configuring OpenShift Networking for Applications

In this lab, you will configure a TLS passthrough route for your application.

Outcomes

You should be able to:

- Deploy an application and configure an insecure route.
- Restrict traffic to the applications.
- Generate a TLS certificate for an application.
- Configure a passthrough route for an application with a TLS certificate.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates the self-signed certificate authority (CA) that you use in this lab.

```
[student@workstation ~]$ lab network-review start
```

Instructions

In this review, you deploy a PHP application that prints some information about the system. The application is available with two different configurations: one that runs with an unencrypted network that listens on port 8080, and one that uses a TLS certificate to encrypt the network traffic, which listens on port 8443.

The container image for this review is accessible at `quay.io/redhattraining/php-ssl`. It has two tags: `v1.0` for the insecure version of the application, and `v1.1` for the secure version.

1. As the OpenShift developer user, create the `network-review` project.

- 1.1. Log in to the cluster as the developer user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create the `network-review` project.

```
[student@workstation ~]$ oc new-project network-review
Now using project "network-review" on server "https://api.ocp4.example.com:6443".
...output omitted...
```

2. As the `developer` user, deploy the insecure version of the PHP application to the `network-review` project. Use the resource file available at `~/D0280/labs/network-review/php-http.yaml`.

Make the necessary changes to the file, use the `quay.io/redhattraining/php-ssl:v1.0` container image and set the port to 8080.

After creating the application, wait a few moments to ensure that one pod is running.

- 2.1. Go to the `~/D0280/labs/network-review` directory.

```
[student@workstation ~]$ cd ~/D0280/labs/network-review
```

- 2.2. Use a text editor to update the `php-http.yaml` file as follows:

- Set the image to use the container image accessible at `quay.io/redhattraining/php-ssl:v1.0`.
- Locate the `containerPort` entry and set the value to `8080`, which corresponds to the **insecure** endpoint.

```
...output omitted...
spec:
  ...output omitted...
  template:
    ...output omitted...
    spec:
      containers:
        ...output omitted...
        image: 'quay.io/redhattraining/php-ssl:v1.0'
        name: php-http
        ports:
          - containerPort: 8080
            name: php-http
        ...output omitted...
```

After making your changes, save and exit the file.

- 2.3. Use the `oc create` command to deploy the application. This creates a deployment and a service.

```
[student@workstation network-review]$ oc create -f php-http.yaml
deployment.apps/php-http created
service/php-http created
```

- 2.4. Wait a few moments, and then run the `oc get pods` command to ensure that there is a pod running.

```
[student@workstation network-review]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
php-http-6cb58c847b-7qsbd	1/1	Running	0	8m11s

3. Create a route named `php-http`, with a hostname of `php-http.apps.ocp4.example.com`, to access the application.

From the `workstation` machine, use Firefox to access the application URL:

- `http://php-http.apps.ocp4.example.com/`.

Confirm the availability of the application before proceeding to the next step.

- 3.1. Run the `oc expose` command to create a route for accessing the application. Give the route a hostname of `php-http.apps.ocp4.example.com`.

```
[student@workstation network-review]$ oc expose svc php-http \
> --hostname php-http.apps.ocp4.example.com
route.route.openshift.io/php-http exposed
```

- 3.2. Retrieve the name of the route and copy it to the clipboard.

```
[student@workstation network-review]$ oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT	...
php-http	php-http.apps.ocp4.example.com		php-http	8080	...

- 3.3. From the `workstation` machine, open Firefox and access the application URL:

- `http://php-http.apps.ocp4.example.com/`.

Confirm that you can see the application.

About this application

⚠ The application is currently served over HTTP

- **Current system load:** 2.5
- **Number of connections:** 1
- **Memory usage:** 8 Mb

4. Create a network policy in the `network-review` namespace to deny all ingress traffic by default. When configured correctly, the network policy also prevents pods within the `network-review` namespace from communicating with each other.

Use the resource file available at `~/D0280/labs/network-review/deny-all.yaml`. Make the necessary changes to target all pods in the namespace.

- 4.1. Use a text editor to update the `deny-all.yaml` file with an empty pod selector to target all pods in the namespace.

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: deny-all
spec:
  podSelector: {}
```

4.2. Use the `oc create` command to create the network policy.

```
[student@workstation network-review]$ oc create -f deny-all.yaml
networkpolicy.networking.k8s.io/deny-all created
```

4.3. Use the `curl` command to verify that there is no access to the `php-http` pod from the route. Wait a few seconds, and then press `Ctrl+C` to exit the `curl` command.

```
[student@workstation network-review]$ curl http://php-http.apps.ocp4.example.com
^C
```



Important

If the `php-http` pod landed on the same node as a `router-default` pod, then the `curl` command works when traffic goes through that router pod. This is only the case with three-node clusters. In a traditional OpenShift cluster, where the control plane or infrastructure nodes are separated from the compute nodes, the network policy would apply to all the router pods in the cluster.

If the `curl` command succeeds, then run the command again to validate the network policy works as expected. This second attempt should go through the other router pod in the cluster.

5. Create a network policy to allow ingress traffic to routes in the `network-review` namespace.

Use the resource file available at `~/DO280/labs/network-review/allow-from-openshift-ingress.yaml`. Make the necessary changes to target all pods in the namespace and allow traffic from the default ingress controller.

Because the classroom environment uses the `HostNetwork` endpoint strategy, label the default namespace with the `network.openshift.io/policy-group=ingress` label. This action must be performed as the `admin` user.

- 5.1. Use a text editor to update the `allow-from-openshift-ingress.yaml` file with an empty pod selector to target all pods in the namespace. Include a namespace selector to match the `network.openshift.io/policy-group=ingress` label.

```
...output omitted...
spec:
  podSelector: {}
  ingress:
  - from:
    - namespaceSelector:
        matchLabels:
          network.openshift.io/policy-group: ingress
```


- 5.2. Use the `oc create` command to create the network policy.

```
[student@workstation network-review]$ oc create -f \
>   allow-from-openshift-ingress.yaml
networkpolicy.networking.k8s.io/allow-from-openshift-ingress created
```

- 5.3. Log in as the admin user.

```
[student@workstation network-review]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

- 5.4. Apply the `network.openshift.io/policy-group=ingress` label to the default namespace.

```
[student@workstation network-policy]$ oc label namespace default \
>   network.openshift.io/policy-group=ingress
namespace/default labeled
```

- 5.5. Use the `curl` command to verify there is access to the `php-http` pod from the route. Because the classroom is running a three-node cluster, run the `curl` command multiple times to validate access through all router pods.

```
[student@workstation network-review]$ for X in {1..4}
>   do
>   curl -s http://php-http.apps.ocp4.example.com | grep "PHP"
>   done
<title>PHP Application</title>
<title>PHP Application</title>
<title>PHP Application</title>
<title>PHP Application</title>
```

6. As the `developer` user, generate and sign a TLS certificate for the encrypted version of the application.

Create a certificate signing request (CSR) for the `php-https.apps.ocp4.example.com` hostname. Save the CSR to `~/D0280/labs/network-review/certs/training.csr`.

Use the CSR to generate a certificate and save it to `~/D0280/labs/network-review/certs/training.crt`. To generate the certificate, pass as arguments the CA certificate accessible at `~/D0280/labs/network-review/certs/training-CA.pem` and the CSR.

You can use the `~/D0280/labs/network-review/certs/openssl-commands.txt` text file for help. This file contains the commands for generating the certificate signing request and the certificate. Make sure to replace the values in the file before copying and running the OpenSSL commands.

- 6.1. Log in as the `developer` user to complete the rest of this lab.

```
[student@workstation network-review]$ oc login -u developer -p developer
Login successful.

...output omitted...
```

- 6.2. Go to the `~/D0280/labs/network-review/certs` directory.

```
[student@workstation network-review]$ cd certs
```

- 6.3. Generate the certificate signing request (CSR) for `php-https.apps.ocp4.example.com`. Make sure to type the subject of the request on one line. Alternatively, remove the `-subj` option and its content. This command prompts you for the values; make sure to indicate a common name (CN) of `php-https.apps.ocp4.example.com`.

**Note**

Make sure there is no space after the trailing slash of the organization (Red Hat) and the common name (CN).

```
[student@workstation certs]$ openssl req -new -key training.key \
> -subj "/C=US/ST=North Carolina/L=Raleigh/O=Red Hat/\
> CN=php-https.apps.ocp4.example.com" \
> -out training.csr
```

Alternatively, open the `openssl-commands.txt` text file. Copy and paste the first `openssl` command to your terminal. Replace the wildcard domain with `apps.ocp4.example.com` and the output file with `training.csr`.

**Note**

The command does not generate any output.

- 6.4. Generate the signed certificate. Notice the usage of the `-CA` and `-CAkey` options for signing the certificate with the CA.

```
[student@workstation certs]$ openssl x509 -req -in training.csr \
> -CA training-CA.pem -CAkey training-CA.key -CAcreateserial \
> -passin file:passphrase.txt \
> -out training.crt -days 3650 -sha256 -extfile training.ext
Signature ok
subject=C = US, ST = North Carolina, L = Raleigh, O = Red Hat, CN = php-
https.apps.ocp4.example.com
Getting CA Private Key
```

Alternatively, copy and paste the second `openssl` command in the `openssl-commands.txt` file to your terminal. Replace the CSR file with `training.csr`, the CA with `training-CA.pem`, and the output certificate with `training.crt`.

- 6.5. Ensure that the newly created certificate and the key are present in the current directory.

```
[student@workstation certs]$ ls -l
total 36
-rw-rw-r--. 1 student student 566 abr 26 07:43 openssl-commands.txt
-rw-rw-r--. 1 student student 33 jun 15 22:20 passphrase.txt
-rw-----. 1 student student 1743 jun 15 22:20 training-CA.key
-rw-r--r--. 1 student student 1334 jun 15 22:20 training-CA.pem
-rw-rw-r--. 1 student student 41 jun 15 22:33 training-CA.srl
-rw-rw-r--. 1 student student 1395 jun 15 22:33 training.crt
-rw-rw-r--. 1 student student 1021 jun 15 22:33 training.csr
-rw-r--r--. 1 student student 352 jun 15 22:20 training.ext
-rw-----. 1 student student 1679 jun 15 22:20 training.key
```

- 6.6. Return to the `network-review` directory. This is important as the next step involves the creation of a route using the signed certificate.

```
[student@workstation certs]$ cd ~/D0280/labs/network-review
```

7. Create an OpenShift TLS secret named `php-certs` in the `network-review` project. Use the `~/D0280/labs/network-review/certs/training.crt` file for the certificate and the `~/D0280/labs/network-review/certs/training.key` file for the key.
 - 7.1. Use the `oc create secret` command to create the `php-certs` TLS secret. Pass the `training.crt` file as the certificate, and `training.key` as the key.

```
[student@workstation network-review]$ oc create secret tls php-certs \
> --cert certs/training.crt --key certs/training.key
secret/php-certs created
```

- 7.2. Retrieve the list of secrets to make sure that it is present.

```
[student@workstation network-review]$ oc get secrets
NAME          TYPE          DATA   AGE
...output omitted...
php-certs     kubernetes.io/tls      2       93s
```

8. Use the resource file, available at `~/D0280/labs/network-review/php-https.yaml`, to deploy the secure version of the PHP application. Deploy the application to the `network-review` project.

Before deploying the application, make the necessary changes to the resources file, specifically:

- The location of the container.
- The port the application listens on.
- The name of the secret to mount as a volume.

- 8.1. Use a text editor to update the `php-https.yaml` file as follows:

- Set the image to use the container image accessible at `quay.io/redhattraining/php-ssl:v1.1`.
- Locate the `containerPort` entry and set the value to `8443`, which corresponds to the **secure** endpoint.

- Locate the `secretName` entry and set the value to `php-certs`, which corresponds to the name of the secret that you created in a previous step.

```
...output omitted...
spec:
...output omitted...
  template:
...output omitted...
    spec:
      containers:
...output omitted...
        image: 'quay.io/redhattraining/php-ssl:v1.1'
        name: php-https
        ports:
        - containerPort: 8443
          name: php-https
...output omitted...
      volumes:
      - name: tls-certs
        secret:
          secretName: php-certs
```

After making your changes, save and exit the file.

- 8.2. Use the `oc create` command to deploy the secure application. This creates a deployment and a service.

```
[student@workstation network-review]$ oc create -f php-https.yaml
deployment.apps/php-https created
service/php-https created
```

- 8.3. Wait a few moments, and then run the `oc get pods` command to ensure that the `php-https` pod is running.

```
[student@workstation network-review]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
php-http-6cb58c847b-7qsbd	1/1	Running	0	8m11s
php-https-84498cd794-hvf7w	1/1	Running	0	26s

9. Create a secure passthrough route named `php-https`, with a hostname of `php-https.apps.ocp4.example.com`, to access the secure version of the application.

From the `workstation` machine, use Firefox to access the application URL:

- `https://php-https.apps.ocp4.example.com/`.

Accept the signed certificate and confirm the availability of the application.

- 9.1. Run the `oc create route` command to create a passthrough route for accessing the application. Give the route a hostname of `php-https.apps.ocp4.example.com`. Use the `port` option to indicate the secure port 8443.

```
[student@workstation network-review]$ oc create route passthrough php-https \
> --service php-https --port 8443 --hostname php-https.apps.ocp4.example.com
route.route.openshift.io/php-https created
```

9.2. Retrieve the name of the route and copy it to the clipboard.

```
[student@workstation network-review]$ oc get routes
```

NAME	HOST/PORT	...	SERVICES	PORT	TERMINATION
php-http	php-http.apps.ocp4.example.com	...	php-http	8080	
php-https	php-https.apps.ocp4.example.com	...	php-https	8443	passthrough

9.3. From the workstation machine, open Firefox and access the application URL:

- <https://php-https.apps.ocp4.example.com/>.

Accept the signed certificate and confirm that you can see secure version of the application.

About this application

 **The application is currently served over TLS**

- **Current system load:** 1
- **Number of connections:** 0
- **Memory usage:** 8 Mb

10. **Optional step:** from the workstation machine, use the `curl` command to access the HTTPS version of the application.

Pass the CA certificate to the `curl` command to validate the secure connection.

Use the `--cacert` option to pass the CA certificate to the `curl` command.

```
[student@workstation network-review]$ curl -v --cacert certs/training-CA.pem \
> https://php-https.apps.ocp4.example.com
...output omitted...
* Server certificate:
* subject: C=US; ST=North Carolina; L=Raleigh; O=Red Hat; \
CN=php-https.apps.ocp4.example.com
...output omitted...
* SSL certificate verify ok.
...output omitted...
```

11. Return to the home directory.

```
[student@workstation network-review]$ cd
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab network-review grade
```

Finish

As the `student` user on the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab network-review finish
```

This concludes the lab.

Summary

In this chapter, you learned:

- OpenShift implements a software-defined networking (SDN) to manage the network infrastructure of the cluster. SDN decouples the software that handles the traffic from the underlying mechanisms that route the traffic.
- Kubernetes provides *services* that allow the logical grouping of pods under a common access route. Services act as load balancers in front of one or more pods.
- Services use selectors (labels) that indicate which pods available to the service.
- There are two kind of routes: secure, and insecure. Secure routes encrypt the traffic using TLS certificates, and insecure routes forward traffic over an unencrypted connection.

Secure routes support three modes: edge, passthrough, and re-encryption.

- Network policies control network traffic to pods. Logical zones can be created in the SDN to segregate traffic among pods in any namespace.

Chapter 6

Controlling Pod Scheduling

Goal

Control the nodes on which a pod runs.

Objectives

- Describe pod scheduling algorithms, the methods used to influence scheduling, and apply these methods.
- Limit the resources consumed by containers, pods, and projects.
- Control the number of replicas of a pod, specify the number of replicas in a deployment, manually scale the number of replicas, and create a horizontal pod autoscaler (HPA) resource.

Sections

- Controlling Pod Scheduling Behavior (and Guided Exercise)
- Limiting Resource Usage by an Application (and Guided Exercise)
- Scaling an Application (and Guided Exercise)

Lab

Controlling Pod Scheduling

Controlling Pod Scheduling Behavior

Objectives

After completing this section, you should be able to describe pod scheduling algorithms, the methods used to influence scheduling, and apply these methods.

Introducing the OpenShift Scheduler Algorithm

The pod scheduler determines placement of new pods onto nodes in the OpenShift cluster. It is designed to be highly configurable and adaptable to different clusters. The default configuration shipped with Red Hat OpenShift Container Platform supports the common data center concepts of zones and regions by using node labels, affinity rules, and anti-affinity rules.

The OpenShift pod scheduler algorithm follows a three step process:

1. Filtering nodes.

The scheduler filters the list of running nodes by evaluating each node against a set of predicates, such as the availability of host ports, or whether a pod can be scheduled to a node experiencing either disk or memory pressure.

Additionally, a pod can define a node selector that matches the labels in the cluster nodes. Nodes whose labels do not match are not eligible.

A pod can also define resource requests for compute resources such as CPU, memory, and storage. Nodes that have insufficient free computer resources are not eligible.

Another filtering check evaluates if the list of nodes have any taints, and if so whether the pod has an associated toleration that can accept the taint. If a pod cannot accept the taint of a node, then the node is not eligible. By default, control plane nodes include the taint `node-role.kubernetes.io/master:NoSchedule`. A pod that does not have a matching toleration for this taint will not be scheduled to a control plane node.



Note

The classroom environment uses a three-node cluster that does not include additional compute nodes. The three-node cluster is available for bare-metal installation in OpenShift Container Platform 4.10. This type of cluster is applicable for resource constrained environments, such as far-edge deployments.

The control plane nodes in a three-node cluster do not have the `node-role.kubernetes.io/master:NoSchedule` taint. Regular application pods can be scheduled to the control plane nodes.

The end result of the filtering step is typically a shorter list of node candidates that are eligible to run the pod. In some cases, none of the nodes are filtered out, which means the pod could run on any of the nodes. In other cases, all of the nodes are filtered out, which means the pod cannot be scheduled until a node with the desired prerequisites becomes available.

A full list of predicates and their descriptions can be found in the references section.

2. Prioritizing the filtered list of nodes.

The list of candidate nodes is evaluated using multiple priority criteria that add up to a weighted score. Nodes with higher values are better candidates to run the pod.

Among the criteria are **affinity** and **anti-affinity** rules. Nodes with higher affinity for the pod have a higher score, and nodes with higher anti-affinity have a lower score.

A common use for **affinity** rules is to schedule related pods to be close to each other, for performance reasons. An example is to use the same network backbone for pods that synchronize with each other.

A common use for **anti-affinity** rules is to schedule related pods that are not too close to each other, for high availability. An example is to avoid scheduling all pods from the same application to the same node.

3. Selecting the best fit node.

The candidate list is sorted based on these scores, and the node with the highest score is selected to host the pod. If multiple nodes have the same high score, then one is selected in a round-robin fashion.

The scheduler is flexible and can be customized for advanced scheduling situations. Additionally, although this course will focus on pod placement using node labels and node selectors, pods can also be placed using pod affinity and anti-affinity rules, as well as node affinity and anti-affinity rules. Customizing the scheduler and covering these alternative pod placement scenarios is outside the scope of this course.

Scheduling and Topology

A common topology for large data centers, such as cloud providers, is to organize hosts into **regions** and **zones**:

- A **region** is a set of hosts in a close geographic area, which guarantees high-speed connectivity between them.
- A **zone**, also called an **availability zone**, is a set of hosts that might fail together because they share common critical infrastructure components, such as a network switch, a storage array, or a uninterruptible power supply (UPS).

As an example of regions and zones, Amazon Web Services (AWS) has a region in northern Virginia (**us-east-1**) with 6 availability zones, and another region in Ohio (**us-east-2**) with 3 availability zones. Each of the AWS availability zones can contain multiple data centers potentially consisting of hundreds of thousands of servers.

The standard configuration of the OpenShift pod scheduler supports this kind of cluster topology by defining predicates based on the **region** and **zone** labels. The predicates are defined in such a way that:

- Replica pods, created from the same deployment, are scheduled to run on nodes that have the same value for the **region** label.
- Replica pods are scheduled to run on nodes that have different values for the **zone** label.

The following figure shows a sample topology that consists of multiple regions, each with multiple zones, and each zone with multiple nodes:




Figure 6.1: Sample cluster topology using regions and zones

Labeling Nodes

As an OpenShift cluster administrator, you can add additional labels to your nodes. For example, you might label nodes with the `env` label using the values of `dev`, `qa`, or `prod` with the intent that development, quality assurance, and production workloads will be deployed to a specific subset of nodes. The labels you choose are arbitrary, but you must publish the labels and their associated values to your developers so that they can configure their applications appropriately.

Use the `oc label` command as a cluster administrator to immediately add, update, or remove a node label. For example, use the following command to label a node with `env=dev`:

```
[user@host ~]$ oc label node master01 env=dev
```

Use the `--overwrite` option to change an existing label:

```
[user@host ~]$ oc label node master01 env=prod --overwrite
```

Remove a label by specifying the label name followed by a hyphen, such as `env-`:

```
[user@host ~]$ oc label node master01 env-
```



Important

Both labels and their values are case-sensitive. An application node selector must match the case of the actual label and the value applied to the node.

Use the `--show-labels` option with the `oc get nodes` command to see the case-sensitive labels assigned to a node:

```
[user@host ~]$ oc get node master02 --show-labels
```

```
NAME      ... ROLES      ... LABELS
master02  ... master,worker ... beta.kubernetes.io/arch=amd64,beta.kubernetes.io/
os=linux,kubernetes.io/arch=amd64,kubernetes.io/hostname=master02,kubernetes.io/
os=linux,node-role.kubernetes.io/master=,node-role.kubernetes.io/
worker=,node.openshift.io/os_id=rhcos
```

Notice that a node might have several default labels assigned by OpenShift. Labels whose keys include the `kubernetes.io` suffix should not be changed by a cluster administrator because they are used internally by the scheduler. The nodes displayed in these sample commands use the AWS full-stack automation setup.

Cluster administrators can also use the `-L` option to determine the value of a single label. For example:

```
[user@host ~]$ oc get node -L failure-domain.beta.kubernetes.io/region
```

```
NAME      STATUS  ROLES      AGE  VERSION  REGION
master01  Ready   master,worker  42d  v1.23.3+e419edf
master02  Ready   master,worker  42d  v1.23.3+e419edf
master03  Ready   master,worker  42d  v1.23.3+e419edf
```

Multiple `-L` options in the same `oc get` command are supported. For example:

```
[user@host ~]$ oc get node -L failure-domain.beta.kubernetes.io/region \
> -L failure-domain.beta.kubernetes.io/zone -L env
```

```
NAME      STATUS  ROLES      AGE  VERSION  ... ENV
master01  Ready   master,worker  42d  v1.23.3+e419edf
master02  Ready   master,worker  42d  v1.23.3+e419edf
master03  Ready   master,worker  42d  v1.23.3+e419edf
```

Labeling Machine Sets

Although node labels are persistent, if your OpenShift cluster contains machine sets, then you should also add labels to the machine set configuration. This ensures that new machines (and the nodes generated from them) will also contain the desired labels. Machine sets are found in clusters using full-stack automation and in some clusters using pre-existing infrastructure that enable cloud provider integration. Bare-metal clusters do not use machine sets.

You can identify the relationship between machines and nodes by listing machines in the `openshift-machine-api` namespace and including the `-o wide` option:

```
[user@host ~]$ oc get machines -n openshift-machine-api -o wide
```

```
NAME      ... NODE
...output omitted...
ocp-qz7hf-worker-us-west-1b-rvx6w ... ip-10-0-139-250.us-west-1.compute.internal
ocp-qz7hf-worker-us-west-1c-v4n4n ... ip-10-0-154-226.us-west-1.compute.internal
```

Machines used for `worker` nodes should come from a machine set. The name of a machine contains the name of the machine set from which it was generated. Use the following command to list machine sets:

```
[user@host ~]$ oc get machineset -n openshift-machine-api
```

NAME	DESIRED	CURRENT	READY	AVAILABLE	...
ocp-qz7hf-worker-us-west-1b	1	1	1	1	...
ocp-qz7hf-worker-us-west-1c	1	1	1	1	...

Edit a machine set so that new machines generated from it will have the desired label or labels. Modifying a machine set will not apply changes to existing machines or nodes. Use the following command to edit a machine set:

```
[user@host ~]$ oc edit machineset ocp-qz7hf-worker-us-west-1b \
> -n openshift-machine-api
```

The highlighted lines below show where to add a label within a machine set:

```
...output omitted...
spec:
  metadata:
    creationTimestamp: null
    labels:
      env: dev
  providerSpec:
...output omitted...
```

Controlling Pod Placement

Many infrastructure-related pods in an OpenShift cluster are configured to run on control plane nodes. Examples include pods for the DNS operator, the OAuth operator, and the OpenShift API server. In some cases, this is accomplished by using the node selector `node-role.kubernetes.io/master: ''` in the configuration of a daemon set or a replica set.

Similarly, some user applications might require running on a specific set of nodes. For example, certain nodes provide hardware acceleration for certain types of workloads, or the cluster administrator does not want to mix production applications with development applications. Use node labels and node selectors to implement these kinds of scenarios.

A node selector is part of an individual pod definition. Define a node selector in a deployment resource, so that any new pod generated from that resource will have the desired node selector. If your deployment resource is under version control, then modify the resource file and apply the changes using the `oc apply -f` command.

Alternatively, a node selector can be added or modified using either the `oc edit` command or the `oc patch` command. For example, to configure the `myapp` deployment so that its pods only run on nodes that have the `env=qa` label, use the `oc edit` command:

```
[user@host ~]$ oc edit deployment/myapp
```

```
...output omitted...
spec:
  ...output omitted...
  template:
    metadata:
      annotations:
```

```

openshift.io/generated-by: OpenShiftNewApp
creationTimestamp: null
labels:
  deployment: myapp
spec:
  nodeSelector:
    env: dev
  containers:
  - image: quay.io/redhattraining/scaling:v1.0
...output omitted...

```

The following `oc patch` command accomplishes the same thing:

```

[user@host ~]$ oc patch deployment/myapp --patch \
> '{"spec":{"template":{"spec":{"nodeSelector":{"env":"dev"}}}}}'

```

Whether using the `oc edit` command or the `oc patch` command, the change triggers a new deployment and the new pods are scheduled according to the node selector.

Configuring a Node Selector for a Project

If the cluster administrator does not want developers controlling the node selector for their pods, then a default node selector should be configured in the project resource. A cluster administrator can either define a node selector when a project is created, or can add or update a node selector after a project is created. Use the `oc adm new-project` command to add the node selector at project creation. For example, the following command creates a new project named `demo`, where all pods will be deployed to nodes that have the label of `tier=1`.

```

[user@host ~]$ oc adm new-project demo --node-selector "tier=1"

```

To configure a default node selector for an existing project, add an annotation to the namespace resource with the `openshift.io/node-selector` key. The `oc annotate` command can add, modify, or remove a node selector annotation:

```

[user@host ~]$ oc annotate namespace demo \
> openshift.io/node-selector="tier=2" --overwrite

```

Scaling the Number of Pod Replicas

Although most deployment resources start with creating a single pod, the number of replicas (or copies) of a pod is frequently increased. This is accomplished by scaling the deployment. Multiple methods for scaling will be covered later, but one method uses the `oc scale` command. For example, the number of pods in the `myapp` deployment can be scaled to three using the following command:

```

[user@host ~]$ oc scale --replicas 3 deployment/myapp

```



References

For more information, refer to the *Controlling pod placement onto nodes (scheduling)* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#controlling-pod-placement-onto-nodes-scheduling

Amazon Web Services Regions and Availability Zones

https://aws.amazon.com/about-aws/global-infrastructure/regions_az/

► Guided Exercise

Controlling Pod Scheduling Behavior

In this exercise, you will configure an application to run on a subset of the cluster compute nodes.

Outcomes

You should be able to use the OpenShift command-line interface to:

- Add a new label to a node.
- Deploy pods to nodes that match a specified label.
- Remove a label from a node.
- Troubleshoot when pods fail to deploy to a node.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates a project that you will be using in the activity.

```
[student@workstation ~]$ lab schedule-pods start
```

Instructions

- 1. As the developer user, create a new project named `schedule-pods`.

- 1.1. Log in to your OpenShift cluster as the developer user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create a new project named `schedule-pods`.

```
[student@workstation ~]$ oc new-project schedule-pods
Now using project "schedule-pods" on server "https://api.ocp4.example.com".

...output omitted...
```

- 2. Deploy and scale a test application.

- 2.1. Create a new application named `hello` using the container located at `quay.io/redhattraining/hello-world-nginx:v1.0`.


```
[student@workstation ~]$ oc new-app --name hello \
> --image quay.io/redhattraining/hello-world-nginx:v1.0
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "hello" created
    deployment.apps "hello" created
    service "hello" created
--> Success
...output omitted...
```

2.2. Create a route to the application.

```
[student@workstation ~]$ oc expose svc/hello
route.route.openshift.io/hello exposed
```

2.3. Manually scale the application so there are four running pods.

```
[student@workstation ~]$ oc scale --replicas 4 deployment/hello
deployment.apps/hello scaled
```

2.4. Verify that the four running pods are distributed between the nodes.

```
[student@workstation ~]$ oc get pods -o wide
```

NAME	READY	STATUS	...	IP	NODE	...
hello-6c4984d949-78qsp	1/1	Running	...	10.9.0.30	master02	...
hello-6c4984d949-cf6tb	1/1	Running	...	10.10.0.20	master01	...
hello-6c4984d949-kwgbg	1/1	Running	...	10.8.0.38	master03	...
hello-6c4984d949-mb8z7	1/1	Running	...	10.10.0.19	master01	...



Note

Depending on the existing load on each node, your output may be different. Although the scheduler will attempt to distribute the pods, the distribution may not be even.

- ▶ 3. Prepare the nodes so that application workloads can be distributed to either development or production nodes by assigning the `env` label. Assign the `env=dev` label to the `master01` node and the `env=prod` label to the `master02` node.

3.1. Log in to your OpenShift cluster as the `admin` user. A regular user does not have permission to view information about nodes and cannot label nodes.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

3.2. Verify that none of the nodes use the `env` label.

```
[student@workstation ~]$ oc get nodes -L env
```

NAME	STATUS	ROLES	AGE	VERSION	ENV
master01	Ready	master,worker	42d	v1.23.3+e419edf	
master02	Ready	master,worker	42d	v1.23.3+e419edf	
master03	Ready	master,worker	42d	v1.23.3+e419edf	

- 3.3. Add the `env=dev` label to the `master01` node to indicate that it is a development node.

```
[student@workstation ~]$ oc label node master01 env=dev
node/master01 labeled
```

- 3.4. Add the `env=prod` label to the `master02` node to indicate that it is a production node.

```
[student@workstation ~]$ oc label node master02 env=prod
node/master02 labeled
```

- 3.5. Verify that the nodes have the correct `env` label set. Make note of the node that has the `env=dev` label, as you will check later to see if the application pods have been deployed to that node.

```
[student@workstation ~]$ oc get nodes -L env
```

NAME	STATUS	ROLES	AGE	VERSION	ENV
master01	Ready	master,worker	42d	v1.23.3+e419edf	dev
master02	Ready	master,worker	42d	v1.23.3+e419edf	prod
master03	Ready	master,worker	42d	v1.23.3+e419edf	

- 4. Switch back to the `developer` user and modify the `hello` application so that it is deployed to the development node. After verifying this change, delete the `schedule-pods` project.

- 4.1. Log in to your OpenShift cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.
```

...output omitted...

```
Using project "schedule-pods".
```

- 4.2. Modify the deployment resource for the `hello` application to select a development node. Make sure to add the node selector in the `spec` group in the `template` section.

```
[student@workstation ~]$ oc edit deployment/hello
```

Add the highlighted lines below to the deployment resource, indenting as shown.

```
...output omitted...
  terminationMessagePath: /dev/termination-log
  terminationMessagePolicy: File
  dnsPolicy: ClusterFirst
  nodeSelector:
    env: dev
  restartPolicy: Always
...output omitted...
```

The following output from `oc edit` is displayed after you save your changes.

```
deployment.apps/hello edited
```

- 4.3. Verify that the application pods are deployed to the node with the `env=dev` label. Although it may take a little time to redeploy, the application pods must be deployed to the `master01` node.

```
[student@workstation ~]$ oc get pods -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE	...
hello-b556ccf8b-8scxd	1/1	Running	0	80s	10.10.0.14	master01	...
hello-b556ccf8b-hb24w	1/1	Running	0	77s	10.10.0.16	master01	...
hello-b556ccf8b-qxlj8	1/1	Running	0	80s	10.10.0.15	master01	...
hello-b556ccf8b-sdpxj	1/1	Running	0	76s	10.10.0.17	master01	...

- 4.4. Remove the `schedule-pods` project.

```
[student@workstation ~]$ oc delete project schedule-pods
project.project.openshift.io "schedule-pods" deleted
```

- ▶ 5. Finish cleaning up this portion of the exercise by removing the `env` label from all nodes.

- 5.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...
```

- 5.2. Remove the `env` label from all nodes that have it.

```
[student@workstation ~]$ oc label node -l env env-
node/master01 labeled
node/master02 labeled
```

- ▶ 6. The `schedule-pods-ts` project contains an application that runs only on nodes that are labeled as `client=ACME`. In the following example, the application pod is pending and you must diagnose the problem using the following steps:

- 6.1. Log in to your OpenShift cluster as the `developer` user and ensure that you are using the `schedule-pods-ts` project.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.

...output omitted...

Using project "schedule-pods-ts".
```

If the output above does not show that you are using the `schedule-pods-ts` project, switch to it.

```
[student@workstation ~]$ oc project schedule-pods-ts
Now using project "schedule-pods-ts" on server ...
```

6.2. Verify that the application pod has a status of `Pending`.

```
[student@workstation ~]$ oc get pods
NAME                                READY   STATUS    RESTARTS   AGE
hello-ts-5dbff9f44-w6csj           0/1     Pending   0           6m19s
```

6.3. Because a pod with a status of `pending` does not have any logs, check the details of the pod using the `oc describe pod` command to see if describing the pod provides any useful information.

```
[student@workstation ~]$ oc describe pod hello-ts-5dbff9f44-8h7c7
...output omitted...
QoS Class:           BestEffort
Node-Selectors:  client=acme
Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s
                   node.kubernetes.io/unreachable:NoExecute op=Exists for 300s
Events:
  Type     Reason             ...           Message
  ----     -
Warning   FailedScheduling   ...   0/3 nodes are available: 3 node(s) didn't match node selector.
```

Based on this information, the pod should be scheduled to a node with the label `client=acme`, but none of the three nodes have this label.

6.4. Log in to your OpenShift cluster as the `admin` user and verify compute node label. To verify it, run the `oc get nodes -L client` to list the details of the available nodes.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.

...output omitted...

[student@workstation ~]$ oc get nodes -L client
```

NAME	STATUS	ROLES	AGE	VERSION	CLIENT
master01	Ready	master,worker	10d	v1.23.3+e419edf	ACME

The information provided indicates that at least one compute node has the label `client=ACME`. You have found the problem. The application must be modified so that it uses the correct node selector.

- 6.5. Log in to your OpenShift cluster as the developer user and edit the deployment resource for the application to use the correct node selector.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.

...output omitted...

[student@workstation ~]$ oc edit deployment/hello-ts
```

Change `acme` to `ACME` as shown below.

```
...output omitted...
  dnsPolicy: ClusterFirst
  nodeSelector:
    client: ACME
  restartPolicy: Always
...output omitted...
```

The following output from `oc edit` is displayed after you save your changes.

```
deployment.apps/hello-ts edited
```

- 6.6. Verify that the a new application pod is deployed and has a status of `Running`.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-ts-69769f64b4-wwhpc	1/1	Running	0	11s

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab schedule-pods finish
```

This concludes the guided exercise.

Limiting Resource Usage by an Application

Objectives

After completing this section, you should be able to limit the resources consumed by containers, pods, and projects.

Defining Resource Requests and Limits for Pods

A pod definition can include both resource requests and resource limits:

Resource requests

Used for scheduling and to indicate that a pod cannot run with less than the specified amount of compute resources. The scheduler tries to find a node with sufficient compute resources to satisfy the pod requests.

Resource limits

Used to prevent a pod from using up all compute resources from a node. The node that runs a pod configures the Linux kernel `cgroups` feature to enforce the pod's resource limits.

Resource request and resource limits should be defined for each container in either a deployment or a deployment configuration resource. If requests and limits have not been defined, then you will find a `resources: {}` line for each container.

Modify the `resources: {}` line to specify the desired requests and or limits. For example:

```
...output omitted...
spec:
  containers:
  - image: quay.io/redhattraining/hello-world-nginx:v1.0
    name: hello-world-nginx
    resources:
      requests:
        cpu: "10m"
        memory: 20Mi
      limits:
        cpu: "80m"
        memory: 100Mi
  status: {}
```

If you use the `oc edit` command to modify a deployment or a deployment configuration, then ensure you use the correct indentation. Indentation mistakes can result in the editor refusing to save changes. To avoid indentation issues, you can use the `oc set resources` command to specify resource requests and limits. The following command sets the same requests and limits as the preceding example:

```
[user@host ~]$ oc set resources deployment hello-world-nginx \
> --requests cpu=10m,memory=20Mi --limits cpu=80m,memory=100Mi
```

If a resource quota applies to a resource request, then the pod should define a resource request. If a resource quota applies to a resource limit, then the pod should also define a resource limit. Red Hat recommends defining resource requests and limits, even if quotas are not used.

Viewing Requests, Limits, and Actual Usage

Using the OpenShift command-line interface, cluster administrators can view compute usage information on individual nodes. The `oc describe node` command displays detailed information about a node, including information about the pods running on the node. For each pod, it shows CPU requests and limits, as well as memory requests and limits. If a request or limit has not been specified, then the pod will show a 0 for that column. A summary of all resource requests and limits is also displayed.

```
[user@host ~]$ oc describe node node1.us-west-1.compute.internal
Name:          node1.us-west-1.compute.internal
Roles:         worker
Labels:        beta.kubernetes.io/arch=amd64
               beta.kubernetes.io/instance-type=m4.xlarge
               beta.kubernetes.io/os=linux
...output omitted...
Non-terminated Pods:          (20 in total)
... Name                      CPU Requests  ... Memory Requests  Memory Limits  AGE
... ----                      -
... tuned-vdwt4              10m (0%)      ... 50Mi (0%)         0 (0%)         8d
... dns-default-2rpwf        110m (3%)     ... 70Mi (0%)         512Mi (3%)     8d
... node-ca-6xwmn            10m (0%)      ... 10Mi (0%)         0 (0%)         8d
...output omitted...
Resource                      Requests      Limits
-----
cpu                            600m (17%)    0 (0%)
memory                        1506Mi (9%)   512Mi (3%)
...output omitted...
```



Note

The summary columns for Requests and Limits display the sum totals of defined requests and limits. In the preceding output, only 1 of the 20 pods running on the node defined a memory limit, and that limit was 512Mi.

The `oc describe node` command displays requests and limits, and the `oc adm top` command shows actual usage. For example, if a pod requests 10m of CPU, then the scheduler will ensure that it places the pod on a node with available capacity. Although the pod requested 10m of CPU, it might use more or less than this value, unless it is also constrained by a CPU limit. Similarly, a pod that does not specify resource requests will still use some amount of resources. The `oc adm top nodes` command shows actual usage for one or more nodes in the cluster, and the `oc adm top pods` command shows actual usage for each pod in a project.

```
[user@host ~]$ oc adm top nodes -l node-role.kubernetes.io/worker
NAME                                CPU(cores)   CPU%   MEMORY(bytes)  MEMORY%
node1.us-west-1.compute.internal    519m         14%    3126Mi         20%
node2.us-west-1.compute.internal    167m         4%     1178Mi         7%
```

Applying Quotas

OpenShift Container Platform can enforce quotas that track and limit the use of two kinds of resources:

Object counts

The number of Kubernetes resources, such as pods, services, and routes.

Compute resources

The number of physical or virtual hardware resources, such as CPU, memory, and storage capacity.

Imposing a quota on the number of Kubernetes resources improves the stability of the OpenShift control plane by avoiding unbounded growth of the Etcd database. Quotas on Kubernetes resources also avoids exhausting other limited software resources, such as IP addresses for services.

In a similar way, imposing a quota on the amount of compute resources avoids exhausting the compute capacity of a single node in an OpenShift cluster. It also avoids having one application starve other applications in a shared cluster by using all the cluster capacity.

OpenShift manages quotas for the number of resources and the use of compute resources in a cluster by using a `ResourceQuota` resource, or a `quota`. A quota specifies hard resource usage limits for a project. All attributes of a quota are optional, meaning that any resource that is not restricted by a quota can be consumed without bounds.



Note

Although a single quota resource can define all of the quotas for a project, a project can also contain multiple quotas. For example, one quota resource might limit compute resources, such as total CPU allowed or total memory allowed. Another quota resource might limit object counts, such as the number of pods allowed or the number of services allowed.

The effect of multiple quotas is cumulative, but it is expected that two different `ResourceQuota` resources for the same project do not limit the use of the same type of Kubernetes or compute resource. For example, two different quotas in a project should not both attempt to limit the maximum number of pods allowed.

The following table describes some resources that a quota can restrict by their count or number:

Resource Name	Quota Description
pods	Total number of pods
replicationcontrollers	Total number of replication controllers
services	Total number of services
secrets	Total number of secrets
persistentvolumeclaims	Total number of persistent volume claims

The following table describes some compute resources that can be restricted by a quota:

Compute Resource Name	Quota Description
requests.cpu	Total CPU use across all containers
requests.memory	Total memory use across all containers
requests.storage	Total storage requests by containers across all persistent volume claims

Quota attributes can track either resource requests or resource limits for all pods in the project. By default, quota attributes track resource requests. Instead, to track resource limits, prefix the compute resource name with `limits`, for example, `limits.cpu`.

The following listing shows a `ResourceQuota` resource defined using YAML syntax. This example specifies quotas for both the number of resources and the use of compute resources:

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: dev-quota
spec:
  hard:
    services: "10"
    cpu: "1300m"
    memory: "1.5Gi"
```

Resource units are the same for pod resource requests and resource limits. For example, `Gi` means GiB, and `m` means millicores. One millicore is the equivalent to 1/1000 of a single CPU core.

Resource quotas can be created in the same way as any other OpenShift Container Platform resource; that is, by passing a YAML or JSON resource definition file to the `oc create` command:

```
[user@host ~]$ oc create --save-config -f dev-quota.yml
```

Another way to create a resource quota is by using the `oc create quota` command, for example:

```
[user@host ~]$ oc create quota dev-quota --hard services=10,cpu=1300,memory=1.5Gi
```

Use the `oc get resourcequota` command to list available quotas, and use the `oc describe resourcequota` command to view usage statistics related to any hard limits defined in the quota, for example:

```
[user@host ~]$ oc get resourcequota
NAME          AGE   REQUEST
compute-quota 51s   cpu: 500m/10, memory: 300Mi/1Gi ...
count-quota   28s   pods: 1/3, replicationcontrollers: 1/5, services: 1/2 ...
```

Without arguments, the `oc describe quota` command displays the cumulative limits set for all `ResourceQuota` resources in the project:

```
[user@host ~]$ oc describe quota
Name:          compute-quota
Namespace:     schedule-demo
Resource       Used    Hard
-----
cpu            500m    10
memory        300Mi   1Gi

Name:          count-quota
Namespace:     schedule-demo
Resource       Used    Hard
-----
pods           1       3
replicationcontrollers 1       5
services       1       2
```

An active quota can be deleted by name using the `oc delete` command:

```
[user@host ~]$ oc delete resourcequota QUOTA
```

When a quota is first created in a project, the project restricts the ability to create any new resources that might violate a quota constraint until it has calculated updated usage statistics. After a quota is created and usage statistics are up-to-date, the project accepts the content creation.

When creating a new resource, the quota usage immediately increments. When deleting a resource, the quota use decrements during the next full recalculation of quota statistics for the project.

Quotas are applied to new resources, but they do not affect existing resources. For example, if you create a quota to limit a project to 15 pods, but there are already 20 pods running, then the quota will not remove the additional 5 pods that exceed the quota.



Important

ResourceQuota constraints are applied for the project as a whole, but many OpenShift processes, such as builds and deployments, create pods inside the project and might fail because starting them would exceed the project quota.

If a modification to a project exceeds the quota for a resource count, then OpenShift denies the action and returns an appropriate error message to the user. However, if the modification exceeds the quota for a compute resource, then the operation does not fail immediately; OpenShift retries the operation several times, giving the administrator an opportunity to increase the quota or to perform another corrective action, such as bringing a new node online.



Important

If a quota that restricts usage of compute resources for a project is set, then OpenShift refuses to create pods that do not specify resource requests or resource limits for that compute resource. To use most templates and builders with a project restricted by quotas, the project must also contain a limit range resource that specifies default values for container resource requests.

Applying Limit Ranges

A `LimitRange` resource, also called a **limit**, defines the default, minimum, and maximum values for compute resource requests, and the limits for a single pod or container defined inside the project. A resource request or limit for a pod is the sum of its containers.

To understand the difference between a limit range and a resource quota, consider that a limit range defines valid ranges and default values for a single pod, and a resource quota defines only top values for the sum of all pods in a project. A cluster administrator concerned about resource usage in an OpenShift cluster usually defines both limits and quotas for a project.

A limit range resource can also define default, minimum, and maximum values for the storage capacity requested by an image, image stream, or persistent volume claim. If a resource that is added to a project does not provide a compute resource request, then it takes the default value provided by the limit ranges for the project.

If a new resource provides compute resource requests or limits that are smaller than the minimum specified by the project limit ranges, then the resource is not created. Similarly, if a new resource provides compute resource requests or limits that are higher than the maximum specified by the project limit ranges, then the resource is not created.

The following listing shows a limit range defined using YAML syntax:

```
apiVersion: "v1"
kind: "LimitRange"
metadata:
  name: "dev-limits"
spec:
  limits:
    - type: "Pod"
      max: ❶
        cpu: "500m"
        memory: "750Mi"
      min: ❷
        cpu: "10m"
        memory: "5Mi"
    - type: "Container"
      max: ❸
        cpu: "500m"
        memory: "750Mi"
      min: ❹
        cpu: "10m"
        memory: "5Mi"
      default: ❺
        cpu: "100m"
        memory: "100Mi"
      defaultRequest: ❻
        cpu: "20m"
        memory: "20Mi"
    - type: openshift.io/Image ❼
      max:
        storage: 1Gi
    - type: openshift.io/ImageStream ❽
      max:
        openshift.io/image-tags: 10
```

```

openshift.io/images: 20
- type: "PersistentVolumeClaim" 9
  min:
    storage: "1Gi"
  max:
    storage: "50Gi"

```

- ❶ The maximum amount of CPU and memory that all containers within a pod can consume. A new pod that exceeds the maximum limits is not created. An existing pod that exceeds the maximum limits is restarted.
- ❷ The minimum amount of CPU and memory consumed across all containers within a pod. A pod that does not satisfy the minimum requirements is not created. Because many pods only have one container, you might set the minimum pod values to the same values as the minimum container values.
- ❸ The maximum amount of CPU and memory that an individual container within a pod can consume. A new container that exceeds the maximum limits does not create the associated pod. An existing container that exceeds the maximum limits restarts the entire pod.
- ❹ The minimum amount of CPU and memory that an individual container within a pod can consume. A container that does not satisfy the minimum requirements prevents the associated pod from being created.
- ❺ The default maximum amount of CPU and memory that an individual container can consume. This is used when a CPU resource limit or a memory limit is not defined for the container.
- ❻ The default CPU and memory an individual container requests. This default is used when a CPU resource request or a memory request is not defined for the container. If CPU and memory quotas are enabled for a namespace, then configuring the `defaultRequest` section allows pods to start, even if the containers do not specify resource requests.
- ❼ The maximum image size that can be pushed to the internal registry.
- ❽ The maximum number of image tags and versions that an image stream resource can reference.
- ❾ The minimum and maximum sizes allowed for a persistent volume claim.

Users can create a limit range resource in the same way as any other OpenShift resource; that is, by passing a YAML or JSON resource definition file to the `oc create` command:

```
[user@host ~]$ oc create --save-config -f dev-limits.yml
```

Red Hat OpenShift Container Platform does not provide an `oc create` command specifically for limit ranges like it does for resource quotas. The only alternative is to use YAML or JSON files.

Use the `oc describe limitrange` command to view the limit constraints enforced in a project:

```

[user@host ~]$ oc describe limitrange dev-limits
Name:          dev-limits
Namespace:     schedule-demo
Type           Resource           Min  Max    Default Request ...
----           -
Pod            cpu                10m  500m  -              ...

```

Pod	memory	5Mi	750Mi	-	...
Container	memory	5Mi	750Mi	20Mi	...
Container	cpu	10m	500m	20m	...
openshift.io/Image	storage	-	1Gi	-	...
openshift.io/ImageStream	openshift.io/image-tags	-	10	-	...
openshift.io/ImageStream	openshift.io/images	-	20	-	...
PersistentVolumeClaim	storage	1Gi	50Gi	-	...

An active limit range can be deleted by name with the `oc delete` command:

```
[user@host ~]$ oc delete limitrange dev-limits
```

After a limit range is created in a project, all requests to create new resources are evaluated against each limit range resource in the project. If the new resource violates the minimum or maximum constraint enumerated by any limit range, then the resource is rejected. If the new resource does not set an explicit value, and the constraint supports a default value, then the default value is applied to the new resource as its usage value.

All resource update requests are also evaluated against each limit range resource in the project. If the updated resource violates any constraint, the update is rejected.



Important

Avoid setting `LimitRange` constraints that are too high, or `ResourceQuota` constraints that are too low. A violation of `LimitRange` constraints prevents pod creation, resulting in error messages. A violation of `ResourceQuota` constraints prevents a pod from being scheduled to any node. The pod might be created but remain in the pending state.

Applying Quotas to Multiple Projects

The `ClusterResourceQuota` resource is created at cluster level, similar to a persistent volume, and specifies resource constraints that apply to multiple projects.

Cluster administrators can specify which projects are subject to cluster resource quotas in two ways:

- Using the `openshift.io/requester` annotation to specify the project owner. All projects with the specified owner are subject to the quota.
- Using a selector. All projects whose labels match the selector are subject to the quota.

The following is an example of creating a cluster resource quota for all projects owned by the `qa` user:

```
[user@host ~]$ oc create clusterquota user-qa \
> --project-annotation-selector openshift.io/requester=qa \
> --hard pods=12,secrets=20
```

The following is an example of creating a cluster resource quota for all projects that have been assigned the `environment=qa` label:

```
[user@host ~]$ oc create clusterquota env-qa \
> --project-label-selector environment=qa \
> --hard pods=10,services=5
```

Project users can use the `oc describe QUOTA` command to view cluster resource quotas that apply to the current project, if any.

Use the `oc delete` command to delete a cluster resource quota:

```
[user@host ~]$ oc delete clusterquota QUOTA
```



Note

To avoid large locking overheads, it is not recommended to have a single cluster resource quota that matches over a hundred projects. When updating or creating project resources, the project is locked while searching for all applicable resource quotas.

Customizing the Default Project Template

Cluster administrators can customize the default project template. Additional resources, such as quotas, limit ranges, and network policies, are created when a user creates a new project.

As a cluster administrator, create a new project template using the `oc adm create-bootstrap-project-template` command, and redirect the output to a file:

```
[user@host ~]$ oc adm create-bootstrap-project-template \
> -o yaml > /tmp/project-template.yaml
```

Customize the template resource file to add additional resources, such as quotas, limit ranges, and network policies. Recall that quotas are configured by cluster administrators and cannot be added, modified, or deleted by project administrators. Project administrators can modify and delete limit ranges and network policies, even if those resources are created by the project template.

New projects create resources specified in the `objects` section. Additional objects should follow the same indentation as the `Project` and `RoleBinding` resources.

The following example creates a quota using the name of the project; the quota imposes a limit of 3 CPUs, 10 GB of memory, and 10 pods on the project:

```
apiVersion: template.openshift.io/v1
kind: Template
metadata:
  creationTimestamp: null
  name: project-request
objects:
- apiVersion: project.openshift.io/v1
  kind: Project
  ...output omitted...
- apiVersion: rbac.authorization.k8s.io/v1
  kind: RoleBinding
  ...output omitted...
```

```

- apiVersion: v1
  kind: ResourceQuota
  metadata:
    name: ${PROJECT_NAME}-quota
  spec:
    hard:
      cpu: "3"
      memory: 10Gi
      pods: "10"
  parameters:
    - name: PROJECT_NAME
    - name: PROJECT_DISPLAYNAME
    - name: PROJECT_DESCRIPTION
    - name: PROJECT_ADMIN_USER
    - name: PROJECT_REQUESTING_USER

```

Use the `oc create` command to create a new template resource in the `openshift-config` namespace:

```

[user@host ~]$ oc create -f project-template.yaml -n openshift-config
template.template.openshift.io/project-request created

```

Update the `projects.config.openshift.io/cluster` resource to use the new project template. Modify the `spec` section. By default, the name of the project template is `project-request`.

```

apiVersion: config.openshift.io/v1
kind: Project
metadata:
  ...output omitted...
  name: cluster
  ...output omitted...
spec:
  projectRequestTemplate:
    name: project-request

```

A successful update to the `projects.config.openshift.io/cluster` resource creates new `apiserver` pods in the `openshift-apiserver` namespace. After the new `apiserver` pods are running, new projects create the resources specified in the customized project template.

To revert to the original project template, modify the `projects.config.openshift.io/cluster` resource to clear the `spec` resource so that it matches: `spec: {}`



References

For more information, refer to the *Quotas* chapter in the Red Hat OpenShift Container Platform 4.10 *Building applications* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/building_applications/index#quotas

For more information about limit ranges, refer to the *Restrict resource consumption with limit ranges* section in the *Working with clusters* chapter in the Red Hat OpenShift Container Platform 4.10 *Nodes* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-cluster-limit-ranges

Customizing OpenShift project creation

<https://developers.redhat.com/blog/2020/02/05/customizing-openshift-project-creation/>

► Guided Exercise

Limiting Resource Usage by an Application

In this exercise, you will configure an application so that it does not consume all computing resources from the cluster and its compute nodes.

Outcomes

You should be able to use the OpenShift command-line interface to:

- Configure an application to specify resource requests for CPU and memory usage.
- Modify an application to work within existing cluster restrictions.
- Create a quota to limit the total amount of CPU, memory, and configuration maps available to a project.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates the resource files that you will be using in the activity.

```
[student@workstation ~]$ lab schedule-limit start
```

Instructions

- 1. As the developer user, create a new project named `schedule-limit`.

- 1.1. Log in to your OpenShift cluster as the developer user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create a new project for this guided exercise named `schedule-limit`.

```
[student@workstation ~]$ oc new-project schedule-limit
Now using project "schedule-limit" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

- 2. Deploy a test application for this exercise that explicitly requests container resources for CPU and memory.

- 2.1. Create a deployment resource file and save it to `~/D0280/labs/schedule-limit/hello-limit.yaml`. Name the application `hello-limit` and use the container image located at `quay.io/redhattraining/hello-world-nginx:v1.0`.

```
[student@workstation ~]$ oc create deployment hello-limit \
> --image quay.io/redhattraining/hello-world-nginx:v1.0 \
> --dry-run=client -o yaml > ~/D0280/labs/schedule-limit/hello-limit.yaml
```

- 2.2. Edit `~/D0280/labs/schedule-limit/hello-limit.yaml` to replace the `resources: {}` line with the highlighted lines below.

```
[student@workstation ~]$ vim ~/D0280/labs/schedule-limit/hello-limit.yaml
```

Ensure that you have proper indentation before saving the file.

```
...output omitted...
spec:
  containers:
  - image: quay.io/redhattraining/hello-world-nginx:v1.0
    name: hello-world-nginx
    resources:
      requests:
        cpu: "3"
        memory: 20Mi
status: {}
```

- 2.3. Create the new application using your resource file.

```
[student@workstation ~]$ oc create --save-config \
> -f ~/D0280/labs/schedule-limit/hello-limit.yaml
deployment.apps/hello-limit created
```

- 2.4. Although a new deployment was created for the application, the application pod should have a status of `Pending`.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-limit-d86874d86b-fpmrt	0/1	Pending	0	10s

- 2.5. The pod cannot be scheduled because none of the compute nodes have sufficient CPU resources. This can be verified by viewing warning events.

```
[student@workstation ~]$ oc get events --field-selector type=Warning
```

LAST SEEN	TYPE	REASON	OBJECT	MESSAGE
88s	Warning	FailedScheduling	pod/hello-limit-d86874d86b-fpmrt	0/3 nodes are available: 3 Insufficient cpu.

- ▶ 3. Redeploy your application so that it requests fewer CPU resources.

- 3.1. Edit `~/D0280/labs/schedule-limit/hello-limit.yaml` to request 1.2 CPUs for the container.

```
[student@workstation ~]$ vim ~/D0280/labs/schedule-limit/hello-limit.yaml
```

Change the `cpu: "3"` line to match the highlighted line below.

```
...output omitted...
resources:
  requests:
    cpu: "1200m"
    memory: 20Mi
```

3.2. Apply the changes to your application.

```
[student@workstation ~]$ oc apply -f ~/D0280/labs/schedule-limit/hello-limit.yaml
deployment.apps/hello-limit configured
```

3.3. Verify that your application deploys successfully. You might need to run `oc get pods` multiple times until you see a running pod. The previous pod with a pending status will terminate and eventually disappear.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-limit-d86874d86b-fpmrt	0/1	Terminating	0	2m19s
hello-limit-7c7998ff6b-ctsjp	1/1	Running	0	6s



Note

If your application pod does not get scheduled, modify the `~/D0280/labs/schedule-limit/hello-limit.yaml` file to reduce the CPU request to `1100m`. Apply the changes again and verify the pod status is `Running`.

- ▶ 4. Attempt to scale your application from one pod to four pods. After verifying that this change would exceed the capacity of your cluster, delete the resources associated with the `hello-limit` application.

4.1. Manually scale the `hello-limit` application up to four pods.

```
[student@workstation ~]$ oc scale --replicas 4 deployment/hello-limit
deployment.apps/hello-limit scaled
```

4.2. Check to see if all four pods are running. You might need to run `oc get pods` multiple times until you see that at least one pod is pending. Depending on the current cluster load, multiple pods might be in a pending state.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-limit-d55cd65c-765s9	1/1	Running	0	12s
hello-limit-d55cd65c-hmblw	0/1	Pending	0	12s
hello-limit-d55cd65c-r8lvw	1/1	Running	0	12s
hello-limit-d55cd65c-vkrhk	0/1	Pending	0	12s

**Note**

If your application pod still does not deploy, scale the number of application pods to zero and then modify `~/D0280/labs/schedule-limit/hello-limit.yaml` to reduce the CPU request to 1000m.

Run `oc apply -f ~/D0280/labs/schedule-limit/hello-limit.yaml` to apply the changes then rerun the `oc scale` command to scale out to four pods.

- 4.3. Warning events indicate that one or more pods cannot be scheduled because none of the nodes has sufficient CPU resources. Your warning messages might be slightly different.

```
[student@workstation ~]$ oc get events --field-selector type=Warning
LAST SEEN   TYPE      REASON            OBJECT
MESSAGE
...output omitted...
76s         Warning   FailedScheduling  pod/hello-limit-d55cd65c-vkrhk    0/3
nodes are available: 3 Insufficient cpu.
```

- 4.4. Delete all of the resources associated with the `hello-limit` application.

```
[student@workstation ~]$ oc delete all -l app=hello-limit
pod "hello-limit-d55cd65c-765s9" deleted
pod "hello-limit-d55cd65c-hmblw" deleted
pod "hello-limit-d55cd65c-r8lvw" deleted
pod "hello-limit-d55cd65c-vkrhk" deleted
deployment.apps "hello-limit" deleted
replicaset.apps "hello-limit-5cc86ff6b8" deleted
replicaset.apps "hello-limit-7d6bdcc99b" deleted
```

- 5. Deploy a second application to test memory usage. This second application sets a memory limit of 200MB per container.

- 5.1. Use the resource file located at `~/D0280/labs/schedule-limit/loadtest.yaml` to create the `loadtest` application. In addition to creating a deployment, this resource file also creates a service and a route.

```
[student@workstation ~]$ oc create --save-config \
> -f ~/D0280/labs/schedule-limit/loadtest.yaml
deployment.apps/loadtest created
service/loadtest created
route.route.openshift.io/loadtest created
```

- 5.2. The `loadtest` container image is designed to increase either CPU or memory load on the container by making a request to the application API. Identify the fully-qualified domain name used in the route.

```
[student@workstation ~]$ oc get routes
NAME      HOST/PORT          ...
loadtest  loadtest.apps.ocp4.example.com ...
```

- **6.** Generate additional memory load that can be handled by the container.
- 6.1. Open two additional terminal windows to continuously monitor the load of your application. Access the application API from the first terminal to simulate additional memory pressure on the container.
 - From the second terminal window, run `watch oc get pods` to continuously monitor the status of each pod.
 - From the third terminal, run `watch oc adm top pod` to continuously monitor the CPU and memory usage of each pod.
 - 6.2. In the first terminal window, use the application API to increase the memory load by 150MB for 60 seconds. Use the fully-qualified domain name previously identified in the route. While you wait for the `curl` command to complete, observe the output in the other two terminal windows.

```
[student@workstation ~]$ curl -X GET \
> http://loadtest.apps.ocp4.example.com/api/loadtest/v1/mem/150/60
curl: (52) Empty reply from server
```

- 6.3. In the second terminal window, observe the output of `watch oc get pods`. Because the container can handle the additional load, you should see that the single application pod has a status of **Running** for the entire `curl` request.

```
Every 2.0s: oc get pods          ...

NAME                                READY   STATUS    RESTARTS   AGE
loadtest-f7495948-tlxgm             1/1     Running   0           7m34s
```

- 6.4. In the third terminal window, observe the output of `watch oc adm top pod`. The starting memory usage for the pod is about 20-30Mi.

```
Every 2.0s: oc adm top pod      ...

NAME                                CPU(cores)   MEMORY(bytes)
loadtest-f7495948-tlxgm             0m           20Mi
```

As the API request is made, you should see memory usage for the pod increase to about 170-180Mi.

```
Every 2.0s: oc adm top pod      ...

NAME                                CPU(cores)   MEMORY(bytes)
loadtest-f7495948-tlxgm             0m           172Mi
```

A short while after the `curl` request completes, you should see memory usage drop back down to about 20-30Mi.

```
Every 2.0s: oc adm top pod      ...

NAME                                CPU(cores)   MEMORY(bytes)
loadtest-f7495948-tlxgm             0m           20Mi
```

► 7. Generate additional memory load that cannot be handled by the container.

- 7.1. Use the application API to increase the memory load by 200MB for 60 seconds. Observe the output in the other two terminal windows.

```
[student@workstation ~]$ curl -X GET \
> http://loadtest.apps.ocp4.example.com/api/loadtest/v1/mem/200/60
<html><body><h1>502 Bad Gateway</h1>
The server returned an invalid or incomplete response.
</body></html>
```

- 7.2. In the second terminal window, observe the output of `watch oc get pods`. Almost immediately after running the `curl` command, the status of the pod will transition to **OOMKilled**. You might even see a status of **Error**. The pod is out of memory and needs to be killed and restarted. The status might change to **CrashLoopBackOff** before returning to a **Running** status. The restart count will also increment.

```
Every 2.0s: oc get pods ...
```

NAME	READY	STATUS	RESTARTS	AGE
loadtest-f7495948-tlxgm	0/1	OOMKilled	0	9m13s

In some cases the pod might have restarted and changed to a status of **Running** before you have time to switch to the second terminal window. The restart count will have incremented from 0 to 1.

```
Every 2.0s: oc get pods ...
```

NAME	READY	STATUS	RESTARTS	AGE
loadtest-f7495948-tlxgm	1/1	Running	1	9m33s

- 7.3. In the third terminal window, observe the output of `watch oc adm top pod`. After the pod is killed, metrics will show that the pod is using little to no resources for a brief period of time.

```
Every 2.0s: oc adm top pod ...
```

NAME	CPU(cores)	MEMORY(bytes)
loadtest-f7495948-tlxgm	8m	0Mi

- 7.4. In the first terminal window, delete all of the resources associated with the second application.

```
[student@workstation ~]$ oc delete all -l app=loadtest
pod "loadtest-f7495948-tlxgm" deleted
service "loadtest" deleted
deployment.apps "loadtest" deleted
route.route.openshift.io "loadtest" deleted
```

In the second and third terminal windows, press `Ctrl+C` to end the `watch` command. Optionally, close the second and third terminal windows.

► 8. As a cluster administrator, create quotas for the `schedule-limit` project.

- 8.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.
```

...output omitted...

- 8.2. Create a quota named `project-quota` that limits the `schedule-limit` project to 3 CPUs, 1 GB of memory, and 2 configuration maps.

```
[student@workstation ~]$ oc create quota project-quota \
> --hard cpu="3",memory="1G",configmaps="2" \
> -n schedule-limit
resourcequota/project-quota created
```



Note

This exercise places a quota on configuration maps to demonstrate what happens when a user tries to exceed the quota.

- ▶ 9. As a developer, attempt to exceed the configuration map quota for the project.

- 9.1. Log in to your OpenShift cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.
```

...output omitted...

- 9.2. Attempt to create a configuration map. All namespaces in the cluster have two configuration maps containing certificates required for the operation of the cluster. Creating a configuration map fails because a third configuration map exceeds the quota of two configuration maps.

```
[student@workstation ~]$ oc create configmap my-config
error: failed to create configmap: configmaps "my-config" is forbidden: exceeded
quota: project-quota, requested: configmaps=1, used: configmaps=2, limited:
configmaps=2
```

- ▶ 10. As a cluster administrator, configure all new projects with default quota and limit range resources.

- 10.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.
```

...output omitted...

- 10.2. Redirect the output of the `oc adm create-bootstrap-project-template` command so that you can customize project creation.

```
[student@workstation ~]$ oc adm create-bootstrap-project-template \
> -o yaml > /tmp/project-template.yaml
```

- 10.3. Edit the `/tmp/project-template.yaml` file to add the quota and limit range resources defined in the `~/D0280/labs/schedule-limit/quota-limits.yaml` file. Add the lines before the parameters section.

```
...output omitted...
- apiVersion: v1
  kind: ResourceQuota
  metadata:
    name: ${PROJECT_NAME}-quota
  spec:
    hard:
      cpu: 3
      memory: 10G
- apiVersion: v1
  kind: LimitRange
  metadata:
    name: ${PROJECT_NAME}-limits
  spec:
    limits:
      - type: Container
        defaultRequest:
          cpu: 30m
          memory: 30M
parameters:
- name: PROJECT_NAME
- name: PROJECT_DISPLAYNAME
- name: PROJECT_DESCRIPTION
- name: PROJECT_ADMIN_USER
- name: PROJECT_REQUESTING_USER
```

**Note**

The `~/D0280/solutions/schedule-limit/project-template.yaml` file contains the correct configuration and can be used for comparison.

- 10.4. Use the `/tmp/project-template.yaml` file to create a new template resource in the `openshift-config` namespace.

```
[student@workstation ~]$ oc create -f /tmp/project-template.yaml \
> -n openshift-config
template.template.openshift.io/project-request created
```

- 10.5. Update the cluster to use the custom project template.

```
[student@workstation ~]$ oc edit projects.config.openshift.io/cluster
```

Modify the `spec` section to use the following lines in bold.


```
...output omitted...
spec:
  projectRequestTemplate:
    name: project-request
```

Ensure proper indentation, and then save your changes.

```
project.config.openshift.io/cluster edited
```

10.6. After a successful change, the `apiserver` pods in the `openshift-apiserver` namespace are recreated.

```
[student@workstation ~]$ watch oc get pods -n openshift-apiserver
```

Wait until all three new `apiserver` pods are ready and running.

```
Every 2.0s: oc get pods -n openshift-apiserver ...
```

NAME	READY	STATUS	RESTARTS	AGE
apiserver-868dccf5fb-885dz	2/2	Running	0	63s
apiserver-868dccf5fb-8j4vh	2/2	Running	0	39s
apiserver-868dccf5fb-r4j9b	2/2	Running	0	24s

Press `Ctrl+C` to end the `watch` command.

10.7. Create a test project to confirm that the custom project template works as expected.

```
[student@workstation ~]$ oc new-project template-test
Now using project "template-test" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

10.8. List the resource quota and limit range resources in the test project.

```
[student@workstation ~]$ oc get resourcequotas,limitranges
```

NAME	AGE	REQUEST	LIMIT
resourcequota/template-test-quota	87s	cpu: 0/3, memory: 0/10G	

NAME	CREATED AT
limitrange/template-test-limits	2021-06-02T15:46:37Z

► 11. Clean up the lab environment by deleting the projects created in this exercise.

11.1. Delete the `schedule-limit` project.

```
[student@workstation ~]$ oc delete project schedule-limit
project.project.openshift.io "schedule-limit" deleted
```

11.2. Delete the `template-test` project.

```
[student@workstation ~]$ oc delete project template-test  
project.project.openshift.io "template-test" deleted
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab schedule-limit finish
```

This concludes the guided exercise.

Scaling an Application

Objectives

After completing this section, you should be able to control the number of replicas of a pod, specify the number of replicas in a deployment, manually scale the number of replicas, and create a horizontal pod autoscaler (HPA) resource.

Specifying Pod Replicas in Configuration Workloads

The number of pod replicas for a specific deployment can be increased or decreased to meet your needs. Despite the `ReplicaSet` and `ReplicationController` resources, the number of replicas needed for an application is typically defined in a deployment resource. A replica set or replication controller (managed by a deployment) guarantees that the specified number of replicas of a pod are running at all times. The replica set or replication controller can add or remove pods as necessary to conform to the desired replica count.

Deployment resources contain:

- The desired number of replicas
- A selector for identifying managed pods
- A pod definition, or template, for creating a replicated pod (including labels to apply to the pod)

The following deployment resource (created using the `oc create deployment` command) displays the following items:

```
apiVersion: apps/v1
kind: Deployment
...output omitted...
spec:
  replicas: 1 ❶
  selector:
    matchLabels:
      deployment: scale ❷
  strategy: {}
  template: ❸
    metadata:
      labels:
        deployment: scale ❹
    spec:
      containers:
...output omitted...
```

- ❶ Specifies the desired number of copies (or replicas) of the pod to run.
- ❷ A replica set uses a selector to count the number of running pods, in the same way that a service uses a selector to find the pods to load balance.
- ❸ A template for pods that the replica set or replication controller creates.
- ❹ Labels on pods created by the template must match those used for the selector.

If the deployment resource is under version control, then modify the `replicas` line in the resource file and apply the changes using the `oc apply` command.

In a deployment resource, a selector is a set of labels that all of the pods managed by the replica set must match. The same set of labels must be included in the pod definition that the deployment instantiates. This selector is used to determine how many instances of the pod are already running in order to adjust as needed.



Note

The replication controller does not perform autoscaling, because it does not track load or traffic. The horizontal pod autoscaler resource, presented later in this section, manages autoscaling.

Manually Scaling the Number of Pod Replicas

Developers and administrators can choose to manually scale the number of pod replicas in a project. More pods may be needed for an anticipated surge in traffic, or the pod count may be reduced to reclaim resources that can be used elsewhere by the cluster. Whether increasing or decreasing the pod replica count, the first step is to identify the appropriate deployment to scale using the `oc get` command:

```
[user@host ~]$ oc get deployment
NAME      READY   UP-TO-DATE   AVAILABLE   AGE
scale     1/1     1             1           8h
```

The number of replicas in a deployment resource can be changed manually using the `oc scale` command:

```
[user@host ~]$ oc scale --replicas 5 deployment/scale
```

The deployment resource propagates the change to the replica set; it reacts to the change by creating new pods (replicas) or deleting existing ones, depending on whether the new desired replica count is less than or greater than the existing count.

Although it is possible to manipulate a replica set resource directly, the recommended practice is to manipulate the deployment resource instead. A new deployment creates either a new replica set or a new replication controller and changes made directly to a previous replica set or replication controller are ignored.

Autoscaling Pods

OpenShift can autoscale a deployment based on current load on the application pods, by means of a `HorizontalPodAutoscaler` resource type.

A horizontal pod autoscaler resource uses performance metrics collected by the OpenShift Metrics subsystem. The Metrics subsystem comes pre-installed in OpenShift 4, rather than requiring a separate install, as in OpenShift 3. To autoscale a deployment, you must specify resource requests for pods so that the horizontal pod autoscaler can calculate the percentage of usage.

The recommended way to create a horizontal pod autoscaler resource is using the `oc autoscale` command, for example:

```
[user@host ~]$ oc autoscale deployment/hello \
> --min 1 --max 10 --cpu-percent 80
```

The previous command creates a horizontal pod autoscaler resource that changes the number of replicas on the `hello` deployment to keep its pods under 80% of their total requested CPU usage.

The `oc autoscale` command creates a horizontal pod autoscaler resource using the name of the deployment as an argument (`hello` in the previous example).

The maximum and minimum values for the horizontal pod autoscaler resource serve to accommodate bursts of load and avoid overloading the OpenShift cluster. If the load on the application changes too quickly, then it might be advisable to keep a number of spare pods to cope with sudden bursts of user requests. Conversely, too many pods can use up all cluster capacity and impact other applications sharing the same OpenShift cluster.

To get information about horizontal pod autoscaler resources in the current project, use the `oc get` command. For example:

```
[user@host ~]$ oc get hpa
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	...
hello	Deployment/hello	<unknown>/80%	1	10	1	...
scale	Deployment/scale	60%/80%	2	10	2	...



Important

The horizontal pod autoscaler initially has a value of `<unknown>` in the `TARGETS` column. It might take up to five minutes before `<unknown>` changes to display a percentage for current usage.

A persistent value of `<unknown>` in the `TARGETS` column might indicate that the deployment does not define resource requests for the metric. The horizontal pod autoscaler will not scale these pods.

Most of the pods created using the `oc new-app` command do not define resource requests. Using the OpenShift autoscaler may therefore require editing the deployment resources, creating custom YAML or JSON resource files for your application, or adding limit range resources to your project that define default resource requests.



References

For more information, refer to the *Automatically scaling pods with the Horizontal Pod Autoscaler* section in the *Working with pods* chapter in the Red Hat OpenShift Container Platform 4.10 Nodes documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/nodes/index#nodes-pods-autoscaling

► Guided Exercise

Scaling an Application

In this exercise, you will scale an application manually and automatically.

Outcomes

You should be able to use the OpenShift command-line interface to:

- Manually scale an application.
- Configure an application to automatically scale based on usage.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates the resource files that you will be using in the activity.

```
[student@workstation ~]$ lab schedule-scale start
```

Instructions

- 1. As the `developer` user, create a new project named `schedule-scale`.

- 1.1. Log in to your OpenShift cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p developer \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Create a new project for this guided exercise named `schedule-scale`.

```
[student@workstation ~]$ oc new-project schedule-scale
Now using project "schedule-scale" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

- 2. Deploy a test application for this exercise which explicitly requests container resources for CPU and memory.

- 2.1. Modify the resource file located at `~/D0280/labs/schedule-scale/loadtest.yaml` to set both requests and limits for CPU and memory usage.

```
[student@workstation ~]$ vim ~/D0280/labs/schedule-scale/loadtest.yaml
```

- 2.2. Replace the `resources: {}` line with the highlighted lines listed below. Ensure that you have proper indentation before saving the file.

```
...output omitted...
spec:
  containers:
  - image: quay.io/redhattraining/loadtest:v1.0
    name: loadtest
    resources:
      requests:
        cpu: "25m"
        memory: 25Mi
      limits:
        cpu: "100m"
        memory: 100Mi
status: {}
```

- 2.3. Create the new application using your resource file.

```
[student@workstation ~]$ oc create --save-config \
> -f ~/DO280/labs/schedule-scale/loadtest.yaml
deployment.apps/loadtest created
service/loadtest created
route.route.openshift.io/loadtest created
```

- 2.4. Verify that your application pod has a status of `Running`. You might need to run the `oc get pods` command multiple times.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
loadtest-5f9565dbfb-jm9md	1/1	Running	0	23s

- 2.5. Verify that your application pod specifies resource limits and requests.

```
[student@workstation ~]$ oc describe pod/loadtest-5f9565dbfb-jm9md | \
> grep -A2 -E "Limits|Requests"
```

```
Limits:
  cpu:    100m
  memory: 100Mi
Requests:
  cpu:    25m
  memory: 25Mi
```

- ▶ 3. Manually scale the `loadtest` deployment by first increasing and then decreasing the number of running pods.

- 3.1. Scale the `loadtest` deployment up to five pods.

```
[student@workstation ~]$ oc scale --replicas 5 deployment/loadtest
deployment.apps/loadtest scaled
```

- 3.2. Verify that all five application pods are running. You might need to run the `oc get pods` command multiple times.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
loadtest-5f9565dbfb-22f9s	1/1	Running	0	54s
loadtest-5f9565dbfb-8l2rn	1/1	Running	0	54s
loadtest-5f9565dbfb-jm9md	1/1	Running	0	3m17s
loadtest-5f9565dbfb-lfhns	1/1	Running	0	54s
loadtest-5f9565dbfb-prjkl	1/1	Running	0	54s

- 3.3. Scale the `loadtest` deployment back down to one pod.

```
[student@workstation ~]$ oc scale --replicas 1 deployment/loadtest
deployment.apps/loadtest scaled
```

- 3.4. Verify that only one application pod is running. You might need to run the `oc get pods` command multiple times while waiting for the other pods to terminate.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
loadtest-5f9565dbfb-prjkl	1/1	Running	0	72s

- 4. Configure the `loadtest` application to automatically scale based on load, and then test the application by applying load.

- 4.1. Create a horizontal pod autoscaler that ensures the `loadtest` application always has 2 application pods running; that number can be increased to a maximum of 10 pods when CPU load exceeds 50%.

```
[student@workstation ~]$ oc autoscale deployment/loadtest \
> --min 2 --max 10 --cpu-percent 50
horizontalpodautoscaler.autoscaling/loadtest autoscaled
```

- 4.2. Wait until the `loadtest` horizontal pod autoscaler reports usage in the `TARGETS` column.

```
[student@workstation ~]$ watch oc get hpa/loadtest
```

Press `Ctrl+C` to exit the `watch` command after `<unknown>` changes to a percentage.

```
Every 2.0s: oc get hpa/loadtest ...
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	...
loadtest	Deployment/loadtest	0%/50%	2	10	2	...

**Note**

It can take up to five minutes before the <unknown> entry in the TARGETS column changes to 0%. If the <unknown> entry does not change, then use the `oc describe` command to verify that the containers for the `loadtest` application request CPU resources.

- 4.3. The `loadtest` container image is designed to either increase CPU or memory load on the container by making a request to the application API. Identify the fully-qualified domain name used in the route.

```
[student@workstation ~]$ oc get route/loadtest
NAME          HOST/PORT          ...
loadtest      loadtest-schedule-scale.apps.ocp4.example.com  ...
```

- 4.4. Access the application API to simulate additional CPU pressure on the container. Move on to the next step while you wait for the `curl` command to complete.

```
[student@workstation ~]$ curl -X GET \
> http://loadtest-schedule-scale.apps.ocp4.example.com/api/loadtest/v1/cpu/1
curl: (52) Empty reply from server
```

- 4.5. Open a second terminal window and continuously monitor the status of the horizontal pod autoscaler.

**Note**

The increased activity of the application does not immediately trigger the autoscaler. Wait a few moments if you do not see any changes to the number of replicas.

```
[student@workstation ~]$ watch oc get hpa/loadtest
```

As the load increases (visible in the TARGETS column), you should see the count under REPLICAS increase. Observe the output for a minute or two before moving on to the next step. Your output will likely be different from what is displayed below.

```
Every 2.0s: oc get hpa/loadtest          ...

NAME          REFERENCE          TARGETS  MINPODS  MAXPODS  REPLICAS  ...
loadtest      Deployment/loadtest 172%/50%    2         10         9         ...
```

**Note**

Although the horizontal pod autoscaler resource can be quick to scale out, it is slower to scale in. Rather than waiting for the `loadtest` application to scale back down to two pods, continue with the rest of the exercise.

- 5. Back in the first terminal window, create a second application named `scaling`. Scale the application, and then verify the responses coming from the application pods.

- 5.1. Create a new application with the `oc new-app` command using the container image located at `quay.io/redhattraining/scaling:v1.0`.

```
[student@workstation ~]$ oc new-app --name scaling \
> --image quay.io/redhattraining/scaling:v1.0
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "scaling" created
    deployment.apps "scaling" created
    service "scaling" created
--> Success
    Application is not exposed.
    You can expose services to the outside world by executing one or more of the
    commands below:
    'oc expose svc/scaling'
    Run 'oc status' to view your app.
```

- 5.2. Create a route to the application by exposing the service for the application.

```
[student@workstation ~]$ oc expose svc/scaling
route.route.openshift.io/scaling exposed
```

- 5.3. Scale the application up to three pods using the deployment resource for the application.

```
[student@workstation ~]$ oc scale --replicas 3 deployment/scaling
deployment.apps/scaling scaled
```

- 5.4. Verify that all three pods for the `scaling` application are running, and identify their associated IP addresses.

```
[student@workstation ~]$ oc get pods -o wide -l deployment=scaling
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE	...
scaling-1-bm4m2	1/1	Running	0	45s	10.10.0.29	master01	...
scaling-1-w7whl	1/1	Running	0	45s	10.8.0.45	master03	...
scaling-1-xqvs2	1/1	Running	0	6m1s	10.9.0.58	master02	...

- 5.5. Display the host name used to route requests to the `scaling` application.

```
[student@workstation ~]$ oc get route/scaling
```

NAME	HOST/PORT	...
scaling	scaling-schedule-scale.apps.ocp4.example.com	...

- 5.6. When you access the host name for your application, the PHP page will output the IP address of the pod that replied to the request. Send several requests to your application, and then sort the responses to count the number of requests sent to each pod. Run the script located at `~/D0280/labs/schedule-scale/curl-route.sh`.

```
[student@workstation ~]$ ~/D0280/labs/schedule-scale/curl-route.sh
34 Server IP: 10.10.0.29
34 Server IP: 10.8.0.45
32 Server IP: 10.9.0.58
```

- **6. Optional:** Check the status of the horizontal pod autoscaler running for the `loadtest` application.

If the `watch oc get hpa/loadtest` command is still running in the second terminal window, switch to it and observe the output. Provided enough time has passed, the replica count should be back down to two.

```
Every 2.0s: oc get hpa/loadtest          ...
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	...
loadtest	Deployment/loadtest	0%/50%	2	10	2	...

When finished, press `Ctrl+C` to exit the `watch` command, and then close the second terminal window.

- **7.** Clean up the lab environment by deleting the `schedule-scale` project.

```
[student@workstation ~]$ oc delete project schedule-scale
project.project.openshift.io "schedule-scale" deleted
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab schedule-scale finish
```

This concludes the guided exercise.

► Lab

Controlling Pod Scheduling

Performance Checklist

In this lab, you will configure an application to run on a subset of the cluster nodes and to scale with load.

Outcomes

You should be able to use the OpenShift command-line interface to:

- Add a new label to nodes.
- Deploy pods to nodes that match a specified label.
- Request CPU and memory resources for pods.
- Configure an application to scale automatically.
- Create a quota to limit the amount of resources a project can consume.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates a directory for exercise files.

```
[student@workstation ~]$ lab schedule-review start
```

Instructions

1. As the `admin` user, label two nodes with the `tier` label. Give the `master01` node the label of `tier=gold` and the `master02` node the label of `tier=silver`.
2. Switch to the `developer` user and create a new project named `schedule-review`.
3. Create a new application named `loadtest` using the container image located at `quay.io/redhattraining/loadtest:v1.0`.
 - The `loadtest` application should be deployed to nodes labeled with `tier=silver`.
 - Ensure that each container requests `100m` of CPU and `20Mi` of memory.
4. Create a route to your application named `loadtest` using the default (automatically generated) host name. Depending on how you created your application, you might need to create a service before creating the route. Your application works as expected if the following URL returns `{"health":"ok"}`.
 - `http://loadtest-schedule-review.apps.ocp4.example.com/api/loadtest/v1/healthz`

5. Create a horizontal pod autoscaler named `loadtest` for the `loadtest` application that will scale from 2 pods to a maximum of 40 pods if CPU load exceeds 70%. You can test the horizontal pod autoscaler by accessing the following URL to trigger the load test actions.
 - `http://loadtest-schedule-review.apps.ocp4.example.com/api/loadtest/v1/cpu/3`

**Note**

Although the horizontal pod autoscaler will scale the `loadtest` application, your OpenShift cluster will run out of resources before it reaches a maximum of 40 pods.

The load test trigger URL does not return any output.

6. As the `admin` user, implement a quota named `review-quota` on the `schedule-review` project. Limit the `schedule-review` project to a maximum of 1 full CPU, 2G of memory, and 20 pods.

Evaluation

Run the following `lab` command to verify your work. If the `lab` command reports any errors, review your changes, make corrections, and run the `lab` command again until successful.

```
[student@workstation ~]$ lab schedule-review grade
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab schedule-review finish
```

This concludes the lab.

► Solution

Controlling Pod Scheduling

Performance Checklist

In this lab, you will configure an application to run on a subset of the cluster nodes and to scale with load.

Outcomes

You should be able to use the OpenShift command-line interface to:

- Add a new label to nodes.
- Deploy pods to nodes that match a specified label.
- Request CPU and memory resources for pods.
- Configure an application to scale automatically.
- Create a quota to limit the amount of resources a project can consume.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command ensures that the cluster API is reachable and creates a directory for exercise files.

```
[student@workstation ~]$ lab schedule-review start
```

Instructions

1. As the `admin` user, label two nodes with the `tier` label. Give the `master01` node the label of `tier=gold` and the `master02` node the label of `tier=silver`.

- 1.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.
```

...output omitted...

- 1.2. Identify if any nodes already use the `tier` label.

```
[student@workstation ~]$ oc get nodes -L tier
```

NAME	STATUS	ROLES	AGE	VERSION	TIER
master01	Ready	master,worker	12d	v1.23.3+e419edf	
master02	Ready	master,worker	12d	v1.23.3+e419edf	
master03	Ready	master,worker	12d	v1.23.3+e419edf	

- 1.3. Label the master01 node with the label tier=gold.

```
[student@workstation ~]$ oc label node master01 tier=gold
node/master01 labeled
```

- 1.4. Label the master02 node with the label tier=silver.

```
[student@workstation ~]$ oc label node master02 tier=silver
node/master02 labeled
```

- 1.5. Confirm that the labels have been added correctly.

```
[student@workstation ~]$ oc get nodes -L tier
```

NAME	STATUS	ROLES	AGE	VERSION	TIER
master01	Ready	master,worker	12d	v1.23.3+e419edf	gold
master02	Ready	master,worker	12d	v1.23.3+e419edf	silver
master03	Ready	master,worker	12d	v1.23.3+e419edf	

2. Switch to the developer user and create a new project named schedule-review.

- 2.1. Log in to your OpenShift cluster as the developer user.

```
[student@workstation ~]$ oc login -u developer -p developer
Login successful.
...output omitted...
```

- 2.2. Create the schedule-review project.

```
[student@workstation ~]$ oc new-project schedule-review
Now using project "schedule-review" on server "https://api.ocp4.example.com:6443".
...output omitted...
```

3. Create a new application named loadtest using the container image located at quay.io/redhattraining/loadtest:v1.0.

- The loadtest application should be deployed to nodes labeled with tier=silver.
- Ensure that each container requests 100m of CPU and 20Mi of memory.

- 3.1. In order to make the upcoming adjustments easier, create a deployment resource file without actually creating the application.

```
[student@workstation ~]$ oc create deployment loadtest --dry-run=client \
> --image quay.io/redhattraining/loadtest:v1.0 \
> -o yaml > ~/D0280/labs/schedule-review/loadtest.yaml
```

- 3.2. Edit `~/D0280/labs/schedule-review/loadtest.yaml` to specify a node selector.

```
[student@workstation ~]$ vim ~/D0280/labs/schedule-review/loadtest.yaml
```

Add the highlighted lines listed below and ensure that you have proper indentation.

```
...output omitted...
spec:
  nodeSelector:
    tier: silver
  containers:
  - image: quay.io/redhattraining/loadtest:v1.0
    name: loadtest
    resources: {}
status: {}
```

- 3.3. Continue editing `~/D0280/labs/schedule-review/loadtest.yaml`. Replace the `resources: {}` line with the highlighted lines listed below. Ensure that you have proper indentation before saving the file.

```
...output omitted...
spec:
  nodeSelector:
    tier: silver
  containers:
  - image: quay.io/redhattraining/loadtest:v1.0
    name: loadtest
    resources:
      requests:
        cpu: "100m"
        memory: 20Mi
status: {}
```

- 3.4. Create the `loadtest` application.

```
[student@workstation ~]$ oc create --save-config \
> -f ~/D0280/labs/schedule-review/loadtest.yaml
deployment.apps/loadtest created
```

- 3.5. Verify that your application pod is running. You might need to run the `oc get pods` command multiple times.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
loadtest-85f7669897-z4mq7	1/1	Running	0	53s

- 3.6. Verify that your application pod specifies resource requests.


```
[student@workstation ~]$ oc describe pod/loadtest-85f7669897-z4mq7 | \
> grep -A2 Requests
Requests:
  cpu:          100m
  memory:       20Mi
```

4. Create a route to your application named `loadtest` using the default (automatically generated) host name. Depending on how you created your application, you might need to create a service before creating the route. Your application works as expected if the following URL returns `{"health": "ok"}`.

- `http://loadtest-schedule-review.apps.ocp4.example.com/api/loadtest/v1/healthz`

- 4.1. Create a service by exposing the deployment for the `loadtest` application.

```
[student@workstation ~]$ oc expose deployment/loadtest \
> --port 80 --target-port 8080
service/loadtest exposed
```

- 4.2. Create a route named `loadtest` by exposing the `loadtest` service.

```
[student@workstation ~]$ oc expose service/loadtest --name loadtest
route.route.openshift.io/loadtest exposed
```

- 4.3. Identify the host name created by the `loadtest` route.

```
[student@workstation ~]$ oc get route/loadtest
NAME          HOST/PORT                                     ...
loadtest      loadtest-schedule-review.apps.ocp4.example.com ...
```

- 4.4. Verify access to the `loadtest` application using the host name identified in the previous step.

```
[student@workstation ~]$ curl \
> http://loadtest-schedule-review.apps.ocp4.example.com/api/loadtest/v1/healthz
{"health": "ok"}
```

5. Create a horizontal pod autoscaler named `loadtest` for the `loadtest` application that will scale from 2 pods to a maximum of 40 pods if CPU load exceeds 70%. You can test the horizontal pod autoscaler by accessing the following URL to trigger the load test actions.

- `http://loadtest-schedule-review.apps.ocp4.example.com/api/loadtest/v1/cpu/3`

**Note**

Although the horizontal pod autoscaler will scale the `loadtest` application, your OpenShift cluster will run out of resources before it reaches a maximum of 40 pods.

The load test trigger URL does not return any output.

- 5.1. Create the horizontal pod autoscaler for the `loadtest` application.

```
[student@workstation ~]$ oc autoscale deployment/loadtest --name loadtest \
> --min 2 --max 40 --cpu-percent 70
horizontalpodautoscaler.autoscaling/loadtest autoscaled
```

- 5.2. Wait until the `loadtest` horizontal pod autoscaler reports default usage in the `TARGETS` column.

```
[student@workstation ~]$ watch oc get hpa/loadtest
```

Press `Ctrl+C` to exit the `watch` command after `<unknown>` changes to `0%`.

```
Every 2.0s: oc get hpa/loadtest      ...
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	...
loadtest	Deployment/loadtest	0%/70%	2	40	2	...



Note

It can take up to five minutes before the `<unknown>` entry in the `TARGETS` column changes to `0%`.

If the `<unknown>` entry does not change, then use the `oc describe` command to verify that the containers for the `loadtest` application request CPU resources.

- 5.3. Test the horizontal pod autoscaler by applying CPU load. Use the previously identified host name for the `loadtest` route. Wait for the `curl` command to complete.

```
[student@workstation ~]$ curl \
> http://loadtest-schedule-review.apps.ocp4.example.com/api/loadtest/v1/cpu/3
```



Note

The command does not generate any output.

- 5.4. Verify that additional pods have been added. Run the `oc get hpa/loadtest` command multiple times until you see the changes reflected. Your output will likely differ, but check that the replica count is greater than 2.

```
[student@workstation ~]$ oc get hpa/loadtest
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	...
loadtest	Deployment/loadtest	1043%/70%	2	40	21	...

6. As the `admin` user, implement a quota named `review-quota` on the `schedule-review` project. Limit the `schedule-review` project to a maximum of 1 full CPU, 2G of memory, and 20 pods.

- 6.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat
Login successful.
...output omitted...
```

6.2. Create the resource quota.

```
[student@workstation ~]$ oc create quota review-quota \
> --hard cpu="1",memory="2G",pods="20"
resourcequota/review-quota created
```



Note

The quota will not impact existing pods, but the scheduler will evaluate the quota if new resources, such as pods, are requested.

Evaluation

Run the following `lab` command to verify your work. If the `lab` command reports any errors, review your changes, make corrections, and run the `lab` command again until successful.

```
[student@workstation ~]$ lab schedule-review grade
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab schedule-review finish
```

This concludes the lab.

Summary

In this chapter, you learned:

- The default pod scheduler uses regions and zones to achieve both performance and redundancy.
- Labeling nodes and using node selectors influences pod placement.
- Resource requests define the minimum amount of resources a pod needs in order to be scheduled.
- Quotas restrict the amount of resources a project is allowed to consume.
- A customized project template can automatically create quotas and limit ranges for new projects.
- The `oc scale` command manually scales the number of replicas of a pod.
- Horizontal pod autoscalers dynamically scale pod replicas based on load.

Chapter 7

Describing Cluster Updates

Goal	Describe how to perform a cluster update.
Objectives	Describe the cluster update process.
Sections	Describing the Cluster Update Process (and Quiz)

Describing the Cluster Update Process

Objectives

After completing this section, you should be able to describe the cluster update process.

Introducing Cluster Updates

Red Hat OpenShift Container Platform 4 adds many new features by using Red Hat Enterprise Linux CoreOS. Red Hat released a new software distribution system that provides the best upgrade path to update your cluster and the underlying operating system. This new distribution system is one of the significant benefits of OpenShift 4 architectural changes, enabling clusters to upgrade Over-the-Air (OTA).

This software distribution system for OTA manages the controller manifests, cluster roles, and any other resources necessary to update a cluster to a particular version.

This feature ensures that a cluster runs the latest available version seamlessly. OTA also enables a cluster to use new features as they become available, including the latest bug fixes and security patches. OTA substantially decreases downtime due to upgrades.

Red Hat hosts and manages this service at <https://console.redhat.com/openshift>, and hosts cluster images at <https://quay.io/>.



Important

As of OpenShift 4.10, the OTA system requires a persistent connection to the Internet. It is not possible to deploy this feature on disconnected environments.

For more information on how to update disconnected clusters, consult the *Update* guide and the *Installation configuration* chapter listed in the references section.

OTA enables faster updates by allowing the skipping of intermediary versions. For example, you can update from 4.10.1 to 4.10.3, thus bypassing 4.10.2.

You use a single interface (<https://console.redhat.com/openshift>) to manage the lifecycle of all your OpenShift clusters.




Figure 7.1: Managing clusters at cloud.redhat.com

The service defines *upgrade paths* that correspond to cluster eligibility for certain updates.

Upgrade paths belong to update channels. A channel can be visualized as a representation of the upgrade path. The channel controls the frequency and stability of updates. The OTA policy engine represents channels as a series of pointers to particular versions within the upgrade path.

A channel name consists of three parts: the tier (release candidate, fast, and stable), the major version (4), and the minor version (.2). Example channel names include: **stable-4.10**, **fast-4.10**, **eus-4.10**, and **candidate-4.10**. Each channel delivers patches for a given cluster version.

Describing the Candidate Channel

The *candidate* channel delivers updates for testing feature acceptance in the next version of OpenShift Container Platform. The release candidate versions are subject to further checks and are promoted to the fast or stable channels when they meet the quality standards.



Note

The updates listed in the *candidate* channel are not supported by Red Hat.

Describing the Fast Channel

The *fast* channel delivers updates as soon as they are available. This channel is best suited for development and QA environments. You can use the **fast-4.10** channel to upgrade from a previous minor version of OpenShift Container Platform.

**Note**

Customers can help to improve OpenShift by joining the Red Hat connected customers program. If you join this program, then your cluster is registered to the fast channel.

Describing the Stable Channel

The stable channel contains delayed updates, which means that it delivers only minor updates for a given cluster version and is better suited for production environments.

Red Hat support and site reliability engineering (SRE) teams monitor operational clusters with new fast updates. If operational clusters pass additional testing and validation, updates in the fast channel are enabled in the stable channel.

If Red Hat observes operational issues from a fast channel update, then that update is skipped in the stable channel. The stable channel delay provides time to observe unforeseen problems in actual OpenShift clusters that testing did not reveal.

Describing the Extended Update Support Channel

The extended update support (EUS) channel is offered in certain minor versions of OpenShift Container Platform to customers with Premium Subscriptions. The EUS versions extend the maintenance phase to 18 months. There is no difference between `stable-4.10` and `eus-4.10` channels and you can switch to the EUS channel as soon as it becomes available.

Starting with OpenShift Container Platform 4.8, Red Hat will denote all even numbered minor releases (eg. 4.8, 4.10, 4.12) as Extended Update Support (EUS) releases.

**Note**

After you upgrade to a version that is exclusive to the EUS channel, that cluster will no longer be eligible for minor version upgrades until upgrades to the next EUS version become available.

Describing Upgrade Paths

The following describes how these upgrade paths would apply to Red Hat OpenShift Container Platform version 4.10:

- When using the `stable-4.10` channel, you can upgrade your cluster from 4.10.0 to 4.10.1 or 4.10.2. If an issue is discovered in the 4.10.3 release, then you cannot upgrade to that version. When a patch becomes available in the 4.10.4 release, you can update your cluster to that version.

This channel is suited for production environments, as the releases in that channel are tested by Red Hat SREs and support services.

- The `fast-4.10` channel can deliver 4.10.1 and 4.10.2 updates but not 4.11.1. This channel is also supported by Red Hat and can be applied to production environments.

Administrators must specifically choose a different minor version channel, such as `fast-4.11`, in order to upgrade to a new release in a new minor version.

- The `candidate-4.10` channel allows you to install the latest features of OpenShift. With this channel, you can upgrade to all *z-stream* releases, such as 4.10.1, 4.10.2, 4.10.3, and so on.

You use this channel to have access to the latest features of the product as they get released. This channel is suited for development and pre-production environments.

- When switching to the `eus-4.10` channel, the `stable-4.10` channel will not receive *z-stream* updates until the next EUS version becomes available.

**Note**

Starting with OpenShift Container Platform 4.8, Red Hat will denote all even numbered minor releases (eg 4.8, 4.10, 4.12) as Extended Update Support (EUS) releases.

The following graphic describes the update graphs for stable and candidate channels:




Figure 7.2: Update graphs for stable and candidate channels

The *stable* and *fast* channels are classified as General Availability (GA), whereas the *candidate* channel (release candidate channel) is not supported by Red Hat.

To ensure the stability of the cluster and the proper level of support, you should only switch from a *stable* channel to a *fast* channel, and vice versa. Although it is possible to switch from a *stable* channel or *fast* channel to a *candidate* channel, it is not recommended. The *candidate* channel is best suited for testing feature acceptance and assisting in qualifying the next version of OpenShift Container Platform.

**Note**

The release of updates for patch and CVE fixes ranges from several hours to a day. This delay provides time to assess any operational impacts to OpenShift clusters.

Changing the Update Channel

You can change the update channel to `stable-4.10`, `fast-4.10`, or `candidate-4.10` using the web console or the OpenShift CLI client:

- In the web console, navigate to the **Administration > Cluster Settings** page on the details tab, and then click the **pencil** icon.




Figure 7.3: Current update channel in the web console

A window displays options to select an update channel.




Figure 7.4: Changing the update channel in the web console

- Execute the following command to switch to another update channel using the `oc` client. You can also switch to another update channel, such as `stable-4.10`, to update to the next minor version of OpenShift Container Platform.

```
[user@host ~]$ oc patch clusterversion version --type="merge" \
> --patch '{"spec":{"channel":"fast-4.10"}}'
clusterversion.config.openshift.io/version patched
```

Describing OTA

OTA follows a client-server approach. Red Hat hosts the cluster images and the update infrastructure. One feature of OTA is the generation of all possible update paths for your cluster. OTA gathers information about the cluster and your entitlement to determine the available upgrade paths. The web console sends a notification when a new update is available.

The following diagram describes the updates architecture: Red Hat hosts both the cluster images and a "watcher", which automatically detects new images that are pushed to Quay. The *Cluster Version Operator* (CVO) receives its update status from that watcher. The CVO starts by updating the cluster components via their operators, and then updates any extra components that the *Operator Lifecycle Manager* (OLM) manages.




Figure 7.5: OpenShift Container Platform Updates Architecture

Telemetry allows Red Hat to determine the update path. The cluster uses Prometheus-based telemetry to report on the state of each cluster operator. The data is anonymized and sent back to Red Hat servers that advise cluster administrators about potential new releases.



Note

Red Hat values customer privacy. For a complete list of the data that Telemeter gathers, consult the *Sample Metrics* document listed in the references section.

In the future, Red Hat intends to extend the list of updated operators that are included in the upgrade path to include independent software vendors (ISV) operators.




Figure 7.6: Managing cluster updates using telemetry

Discussing the Update Process

There are two components involved in the cluster update process:

Machine Config Operator

The Machine Config Operator applies the desired machine state to each of the nodes. This component also handles the rolling upgrade of nodes in the cluster, and uses CoreOS Ignition as the configuration format.

Operator Lifecycle Manager (OLM)

The Operator Lifecycle Manager (OLM) orchestrates updates to any operators running in the cluster.

Updating the Cluster

You can update the cluster via the web console, or from the command-line. Updating via the web console is easier than using the command-line. The **Administration > Cluster Settings** page displays an **Update Status** of *Available updates* when a new update is available. From this page, click **Update now** to begin the process:




Figure 7.7: Updating the cluster using the web console

**Important**

Red Hat does not support reverting your cluster to a previous version, or rollback. Red Hat only supports upgrading to a newer version.

The update process also updates the underlying operating system when there are updates available. It uses the `rpm-ostree` technology for managing transactional upgrades. Updates are delivered via container images and are part of the OpenShift update process. When the update deploys, the nodes pull the new image, extract it, write the packages to the disk, and then modify the bootloader to boot into the new version. The machine reboots and implements a rolling update to ensure that the cluster capacity is minimally impacted.

Update the Cluster Using the Command Line

The following steps describe the procedure for updating a cluster as a cluster administrator using the command-line interface:

1. Be sure to update all operators installed through the Operator Lifecycle Manager (OLM) to the latest version before updating the OpenShift cluster.
2. Retrieve the cluster version and review the current update channel information and confirm the channel. If you are running the cluster in production, then ensure that the channel reads `stable`.

```
[user@host ~]$ oc get clusterversion
NAME      VERSION  AVAILABLE  PROGRESSING  SINCE   STATUS
version   4.10.3   True       False        43d     Cluster version is 4.10.3

[user@host ~]$ oc get clusterversion -o jsonpath='{.items[0].spec.channel}'
stable-4.10
```

3. View the available updates and note the version number of the update that you want to apply.

```
[user@host ~]$ oc adm upgrade
Cluster version is 4.10.3

Updates:

VERSION IMAGE
4.10.4    quay.io/openshift-release-dev/ocp-release@sha256:...
...output omitted...
```

4. Apply the latest update to your cluster or update to a specific version:

- Run the following command to install the latest available update for your cluster.

```
[user@host ~]$ oc adm upgrade --to-latest=true
```

- Run the following command to install a specific version. *VERSION* corresponds to one of the available versions that the `oc adm upgrade` command returns.

```
[user@host ~]$ oc adm upgrade --to=VERSION
```



Important

OpenShift Container Platform 4.9 uses Kubernetes 1.22, which removed a significant number of deprecated `v1beta1` APIs.

OpenShift Container Platform 4.8.14 introduced a requirement that an administrator must provide a manual acknowledgment before the cluster can be upgraded from version 4.8 to 4.9. This is to help prevent issues after upgrading to OpenShift Container Platform 4.9, where APIs that have been removed are still in use by workloads, tools, or other components running on or interacting with the cluster.

Administrators must evaluate their cluster for any APIs in use that will be removed and migrate the affected components to use the appropriate new API version. After this is done, the administrator can provide the administrator acknowledgment.

All clusters require this administrator acknowledgment before they can be upgraded to OpenShift Container Platform 4.9.

```
[user@host ~]$ oc patch configmap admin-acks -n openshift-config \
> --type=merge \
> --patch '{"data":{"ack-4.8-kube-1.22-api-removals-in-4.9":"true"}}'
configmap/admin-acks patched
```

5. The previous command initializes the update process. Run the following command to review the status of the Cluster Version Operator (CVO) and the installed cluster operators.

```
[user@host ~]$ oc get clusterversion
NAME      VERSION  AVAILABLE  PROGRESSING  SINCE   STATUS
version  4.10.3   True       True         30m    Working towards 4.10.4 ...

[user@host ~]$ oc get clusteroperators
```

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
authentication	4.10.3	True	False	False
cloud-credential	4.10.4	False	True	False
openshift-apiserver	4.10.4	True	False	True
...output omitted...				

6. The following command allows you to review the cluster version status history to monitor the status of the update. It might take some time for all the objects to finish updating.

The history contains a list of the most recent versions applied to the cluster. This value is updated when the CVO applies an update. The list is ordered by date, where the newest update is first in the list.

If the rollout completed successfully, updates in the history have a state of **Completed**. Otherwise, the update has a state of **Partial** if the update failed or did not complete.

```
[user@host ~]$ oc describe clusterversion
...output omitted...
History:
  Completion Time: 2022-04-28T04:38:12Z
  Image:          quay.io/openshift-release-dev/ocp-release@sha256:...
  Started Time:   2022-04-28T03:35:05Z
  State:          Partial
  Verified:       true
  Version:        4.10.4
  Completion Time: 2022-03-15T12:39:02Z
  Image:          quay.io/openshift-release-dev/ocp-release@sha256:...
  Started Time:   2022-03-15T12:23:14Z
  State:          Completed
  Verified:       false
  Version:        4.10.3
```



Important

If an upgrade fails, the operator stops and reports the status of the failing component. Rolling your cluster back to a previous version is not supported.

If your upgrade fails, contact Red Hat support.

7. After the process completes, you can confirm that the cluster has been updated to the new version.

```
[user@host ~]$ oc get clusterversion
NAME      VERSION  AVAILABLE  PROGRESSING  SINCE   STATUS
version  4.10.4  True       True          30m    Cluster version is 4.10.4
```




References

For more information about installing Red Hat OpenShift Container Platform in a disconnected environment, refer to the *Installation configuration* chapter in the Red Hat OpenShift Container Platform 4.10 *Installing* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/installing/index#installation-configuration

For more information about update channels, update prerequisites, and updating clusters in disconnected environments, refer to the *Updating a restricted network cluster*, and *Updating a cluster between minor versions* chapters in the the Red Hat OpenShift Container Platform 4.10 *Updating clusters* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/updating_clusters/index#updating-restricted-network-cluster

For more information about the update of operators installed through the Operator Lifecycle Manager (OLM), refer to the *Upgrading installed Operators* section in the *Administrator tasks* chapter in the Red Hat OpenShift Container Platform 4.10 *Working with Operators* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/operators/index#olm-upgrading-operators

For more information about the OpenShift Container Platform update graph, visit the following page in the customer portal: https://access.redhat.com/labs/ocpupgradegraph/update_path

For more information about the deprecated APIs in OpenShift Container Platform 4.9, visit the following page in the customer portal: <https://access.redhat.com/articles/6329921>

For more information about the OpenShift Container Platform lifecycle policy, visit the following page in the customer portal: <https://access.redhat.com/support/policy/updates/openshift>

For more information about OpenShift Extended Update Support (EUS), visit the following page in the customer portal: <https://access.redhat.com/support/policy/updates/openshift-eus>

For more information about performing an EUS-to-EUS update, visit the following page: <https://docs.openshift.com/container-platform/4.10/updating/preparing-eus-eus-upgrade.html>

► Quiz

Describing the Cluster Update Process

Choose the correct answers to the following questions:

- 1. Which two of the following updates would be available in the fast - 4.10 update channel? (Choose two.)
 - a. 4.9.2
 - b. 4.10.1
 - c. 4.11.1
 - d. 4.10.5
- 2. Which of the following components retrieves the updated cluster images from Quay.io?
 - a. Cluster Monitoring (Prometheus)
 - b. Operator Lifecycle Manager (OLM)
 - c. Cluster Version Operator (CVO)
 - d. Telemetry client (Telemeter)
- 3. Which of the following components manages the updates of operators that are not cluster operators?
 - a. Operator Lifecycle Manager (OLM)
 - b. Telemetry client (Telemeter)
 - c. Cluster Version Operator (CVO)
- 4. Which two of the following commands allow you to retrieve the version of the cluster that is currently running? (Choose two.)
 - a. `oc adm upgrade`
 - b. `oc get clusterchannel`
 - c. `oc get clusterversion`
- 5. Which of the following statements is true regarding the OTA feature?
 - a. The *stable* channel is classified as General Availability (GA), whereas the *fast* channel is classified as Release Candidate (RC).
 - b. When using the stable channel, you cannot skip intermediary versions. For example, when updating from 4.9.27 to 4.9.29, OpenShift must install the 4.9.28 version first.
 - c. It is not recommended to switch from a stable channel or a fast channel to a candidate channel. However, you can switch from a fast channel to a stable channel and vice versa.
 - d. Rolling back a failed update is only supported by Red Hat when you are attempting to update from a z-stream version to one another (for example, from 4.9.2 to 4.9.3, but not from 4.9.3 to 4.10).

- 6. Which two of the following channels are classified as general availability? (Choose two.)
- a. candidate-4.10.stable
 - b. stable-4.10
 - c. candidate-stable-4.10
 - d. fast-4.10
 - e. fast-4.10.1

► Solution

Describing the Cluster Update Process

Choose the correct answers to the following questions:

- 1. Which two of the following updates would be available in the fast - 4.10 update channel? (Choose two.)
 - a. 4.9.2
 - b. 4.10.1
 - c. 4.11.1
 - d. 4.10.5
- 2. Which of the following components retrieves the updated cluster images from Quay.io?
 - a. Cluster Monitoring (Prometheus)
 - b. Operator Lifecycle Manager (OLM)
 - c. Cluster Version Operator (CVO)
 - d. Telemetry client (Telemeter)
- 3. Which of the following components manages the updates of operators that are not cluster operators?
 - a. Operator Lifecycle Manager (OLM)
 - b. Telemetry client (Telemeter)
 - c. Cluster Version Operator (CVO)
- 4. Which two of the following commands allow you to retrieve the version of the cluster that is currently running? (Choose two.)
 - a. `oc adm upgrade`
 - b. `oc get clusterchannel`
 - c. `oc get clusterversion`
- 5. Which of the following statements is true regarding the OTA feature?
 - a. The *stable* channel is classified as General Availability (GA), whereas the *fast* channel is classified as Release Candidate (RC).
 - b. When using the stable channel, you cannot skip intermediary versions. For example, when updating from 4.9.27 to 4.9.29, OpenShift must install the 4.9.28 version first.
 - c. It is not recommended to switch from a stable channel or a fast channel to a candidate channel. However, you can switch from a fast channel to a stable channel and vice versa.
 - d. Rolling back a failed update is only supported by Red Hat when you are attempting to update from a z-stream version to one another (for example, from 4.9.2 to 4.9.3, but not from 4.9.3 to 4.10).

- 6. Which two of the following channels are classified as general availability? (Choose two.)
- a. candidate-4.10.stable
 - b. stable-4.10
 - c. candidate-stable-4.10
 - d. fast-4.10
 - e. fast-4.10.1

Summary

In this chapter, you learned:

- One of the major benefits of OpenShift 4 architectural changes is that you can update your clusters *Over-the-Air* (OTA).
- Red Hat provides a new software distribution system that ensures the best path for updating your cluster and the underlying operating system.
- There are several distribution channels:
 - The *stable* channel delivers delayed updates.
 - The *fast* channel delivers updates as soon as they are available.
 - The *candidate* channel delivers updates for testing feature acceptance in the next version of OpenShift Container Platform.
 - The *eus* channel (only available when running 4.10) extends the maintenance phase for customers with Premium Subscriptions.
- Red Hat does not support reverting your cluster to a previous version. Red Hat only supports upgrading to a newer version.

Chapter 8

Managing a Cluster with the Web Console

Goal

Manage a Red Hat OpenShift cluster using the web console.

Objectives

- Perform cluster administration with the web console.
- Manage applications and Kubernetes Operators with the web console.
- Examine performance and health metrics for cluster nodes and applications.

Sections

- Performing Cluster Administration (and Guided Exercise)
- Managing Workloads and Operators (and Guided Exercise)
- Examining Cluster Metrics (and Guided Exercise)

Lab

Managing a Cluster with the Web Console

Performing Cluster Administration

Objectives

After completing this section, you should be able to perform cluster administration with the web console.

Describing the Web Console

The Red Hat OpenShift web console provides a graphical user interface to perform administrative, management, and troubleshooting tasks. It supports both **Administrator** and **Developer** perspectives. This course explores the **Administrator** perspective.

The following list outlines some of the most important parts of the web console, grouped by the main navigation menu items. The ability to view individual items depends upon the roles and role bindings associated with a user. Users with the `cluster-admin` cluster role can view and modify everything. Users with the `view` cluster role can view most items, but cannot make changes. Additional roles can provide access to specific items.

Home

The **Home > Overview** page provides a quick overview of the cluster, including health metrics, resource counts, and a streaming list of events, such as machine updates or pod failures.

You can navigate to **Home > Search** page to find or create resources of any type. This is also a useful starting point to navigate to resources that do not have dedicated navigation in the menu.

The **Home > Events** page displays a filterable stream of events that occur in the cluster and is a good starting point for troubleshooting.

Operators

Explore and install operators curated by Red Hat using OperatorHub, then navigate to the **Installed Operators** page to managed the operators.

Workloads, Networking, and Storage

Manage common resources such as deployments, services, and persistent volumes. The ability to view pod logs and connect to a terminal can be used for troubleshooting.

Builds

Manage build configurations, builds, and image streams.

Monitoring

View alerts and perform ad-hoc Prometheus queries.

Compute

View and manage compute resources such as nodes, machines, and machine autoscalers.

User Management

View and manage users, groups, service accounts, roles, and role bindings.

Administration

View and manage a wide variety of settings that are of particular interest to cluster administrators, such as cluster updates, cluster operators, CRDs, and resource quotas.

Accessing the OpenShift Web Console

The OpenShift web console runs as pods in the `openshift-console` project and is managed by the operator running in the `openshift-console-operator` project. You can discover the console URL by using the `oc whoami` command.


```
[user@host ~]$ oc whoami --show-console
https://console-openshift-console.apps.cluster.example.com
```

In non-production systems, self-signed certificates are commonly used for the HTTPS endpoint. Web browsers will warn you about the certificate, and you will need to add a security exception when navigating to the web console for the first time.


Finding Resources

The web UI provides multiple ways to locate resources. Many common resources, such as deployments and services, are available in the main menu on the left. You can use the **Home > Search** page to find other resource types. This page provides a complete menu of resource types and a label selector field.

Use the name filter to quickly locate resources on pages with long lists such as the **Projects** page:



It may be useful to filter pods by state to identify potential issues or problematic deployments:



The details page of a resource displays common useful information. The contents of this page vary for different types. For example, the **Pod Details** page displays metrics and status information

and the **Secret Details** page allows you to reveal or copy data stored in the secret. Detail pages provide a YAML editor to view and modify the resource specification from the web console. Some resource types, such as **secrets** and **role bindings**, provide more advanced UIs tailored to the resource type.

Creating Users and Groups

The **User Management > Users** page displays users who have previously logged in to OpenShift. As discussed in *Chapter 3, Configuring Authentication and Authorization*, OpenShift supports a variety of identity providers (IdPs), including HTPasswd, LDAP, and OpenID Connect.

When using the HTPasswd identity provider, the **secrets** editor can simplify adding, updating, or removing entries in the HTPasswd secret. After using a terminal to generate a new or updated HTPasswd entry, switch to the web console to modify the secret.

In the web UI, locate the secret in the **openshift-config** project and then click **Actions > Edit Secret**. The **Edit Key/Value Secret** tool handles the base64 encoding for you. Add a line to allow a new user to log in to OpenShift. Update a line to change the password for a user. Delete a line so that a user cannot log in to OpenShift.

The **User Management > Groups** page displays existing groups, and provides the ability to create new groups.

Creating a Project

The web UI features a variety of pages and forms for configuring projects. To create a project:

1. Navigate to the **Home > Projects** page to display the full list of projects. Then click **Create Project** and complete the form to create a new project.
2. After you have created your new project, you can navigate to the **Role Bindings** tab on the project details page.
3. Red Hat recommends that administrators responsible for multitenant clusters configure **resource quotas** and **limit ranges**, which enforce total project limits and container limits, respectively. Navigate to either **Administration > Resource Quotas** or **Administration > Limit Ranges** to access the appropriate YAML editor, where you can configure these limits.

Discussing Limitations

The OpenShift web console is a powerful tool for graphically administering OpenShift clusters, however some administrative tasks are not currently available in the web console. For example, viewing node logs and executing node debug sessions requires the **oc** command-line tool.



References

For more information, refer to the Red Hat OpenShift Container Platform 4.10 *Web console* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/web_console/index

► Guided Exercise

Performing Cluster Administration

In this exercise, you will perform cluster administration with the web console.

Outcomes

You should be able to use the OpenShift web console to:

- Find resources associated with an operator.
- Review the status of a pod, YAML definition, and logs.
- View and edit cluster configuration resources.
- Create a new project and configure its resource quotas, limit ranges, and role-based access control (RBAC).

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates the resources required for this exercise.

```
[student@workstation ~]$ lab console-admin start
```

Instructions

- 1. As the `admin` user, locate and navigate to the OpenShift web console.

- 1.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.


...output omitted...
```

- 1.2. Identify the URL for the web console.


```
[student@workstation ~]$ oc whoami --show-console
https://console-openshift-console.apps.ocp4.example.com
```

- 1.3. Open a web browser and navigate to the web console URL:
 - `https://console-openshift-console.apps.ocp4.example.com`
- 1.4. Click `localusers` and log in as the `admin` user with the `redhat` password.


2. Review the `openshift-console-operator` and `openshift-console` pod logs.
 - 2.1. In the Red Hat OpenShift Container Platform web UI, click **Home** > **Projects** to display the **Projects** page.
 - 2.2. Type `console` in the **Search by name** field, and then click the `openshift-console-operator` link.



- 2.3. Click **Workloads** and then click **1 of 1 Pod** to navigate to the `console-operator` replica set.




- 2.4. Click the pod name marked by the **P** icon to navigate to the `console-operator` pod.



- 2.5. Review the **Pod Details** page and notice the pod metrics, running status, and volumes.

- 2.6. Click **YAML** to navigate to the pod resource editor.
- 2.7. Click **Logs** to view the console operator logs.
- 2.8. Open the **Project** list and type **openshift-console** to switch to the **openshift-console** project. Click the **Show default projects** option if it is disabled.




- 2.9. Click the first pod in the table and then click **Logs** to view the console pod logs.

► 3. Review the Console, Image, and OAuth cluster settings.

- 3.1. Click **Administration** > **Cluster Settings** on the left menu to view the **Cluster Settings** page. Notice that information about the cluster's update channel and current version are listed at the top and a section for the cluster's update history is listed further below.
- 3.2. Click **Configuration** to navigate to the list of cluster configuration resources.
- 3.3. Click **Console** (config.openshift.io) and then click **YAML** to review the **Console** resource.
- 3.4. Return to the **Cluster Settings** Configuration page. Click **Image** and then click **YAML**. Notice the **internalRegistryHostname** is configured to use the internal image registry.
- 3.5. Return to the **Cluster Settings** Configuration page and click **OAuth**. The **OAuth Details** page has a special section for listing and adding identity providers. Navigate to the **YAML** page to view additional configuration details.

► 4. Review the admin, edit, and view cluster roles.

- 4.1. Click **User Management** > **Roles** on the left menu to view the **Roles** page.
- 4.2. Click **admin** next to the **CR** icon. Review the **Rules** table which describes the allowed actions for various resources.



- 4.3. Return to the **Cluster Roles** page and click the cluster role named **edit** to view the **edit** cluster role details.
- 4.4. Return to the **Cluster Roles** page and type **view** in the **Search by name** field. Click the cluster role named **view** to navigate to the **view** cluster role details. Notice that this role only allows **get**, **list**, and **watch** actions on the listed resources.
- ▶ 5. Add a **tester** user entry to the **localusers** secret.
 - 5.1. Click **Workloads > Secrets** on the left menu, and then select **openshift-config** from the **Project** filter list to display the secrets for the **openshift-config** project. Click the **Show default projects** option if it is disabled.
 - 5.2. Use the filter or scroll to the bottom of the page to locate and then click the **localusers** link.
 - 5.3. Click **Actions > Edit Secret** to navigate to the **Edit Key/Value Secret** tool.
 - 5.4. Use the workstation terminal to generate an **htpasswd** entry using **redhat** as the password.

```
[student@workstation ~]$ htpasswd -n -b tester redhat
tester:$apr1$0Q3BtW0p.HtW97.$wVbJBofBNSNd4sd
```

- 5.5. Append the terminal output from the **htpasswd** command to the **htpasswd** value in the OpenShift web console's secrets editor, and then click **Save**.

```
admin:$apr1$Au9.fFr$0k5wvUBd3eeBt0baa77.dae
leader:$apr1$/abo4Hybn7a.tG5Zo0Bn.QwefXckiy1
developer:$apr1$RjqTY4cv$xl3.BQfg42moSxwnTNkh.
tester:$apr1$0Q3BtW0p.HtW97.$wVbJBofBNSNd4sd
```

- ▶ 6. Create and configure a new project named **console-apps**.
 - 6.1. Click **Home > Projects** on the left menu to view the **Projects** page, and then click **Create Project**.
 - 6.2. Use the following information for the new project, and then click **Create**.

Create Project Form

Field	Value
Name	console-apps
Display Name	Console chapter applications
Description	Example project

- 6.3. Click **Administration > Resource Quotas** on the left menu, and then click **Create Resource Quota**. Modify the YAML document as follows and click **Create**.

```
apiVersion: v1
kind: ResourceQuota
metadata:
```

```

name: quota
namespace: console-apps
spec:
  hard:
    pods: '10'
    requests.cpu: '2'
    requests.memory: 8Gi
    limits.cpu: '4'
    limits.memory: 12Gi

```

- 6.4. Click **Administration > Limit Ranges** on the left menu, and then click **Create Limit Range**. Modify the YAML document to specify a name for the limit range. Specify default memory and CPU container limits and requests, and click **Create**.

```

apiVersion: v1
kind: LimitRange
metadata:
  name: limit-range
  namespace: console-apps
spec:
  limits:
    - default:
        cpu: 500m
        memory: 5Gi
      defaultRequest:
        cpu: 10m
        memory: 100Mi
      type: Container

```

- 6.5. Click **User Management > Groups** on the left menu, and then click **Create Group**. Use the editor to define a Group resource as follows, and click **Create**.

```

apiVersion: user.openshift.io/v1
kind: Group
metadata:
  name: project-team
users:
  - developer
  - tester

```

- 6.6. Click **User Management > Role Bindings** on the left menu, and then click **Create Binding**. Fill out the form as follows to create a role binding for the **project-team** group.

Team Role Binding Form

Field	Value
Binding Type	Namespace Role Binding (RoleBinding)
Name	team
Namespace	console-apps
Role Name	edit
Subject	Group
Subject Name	project-team

Click **Create** to create the RoleBinding in the console-apps namespace.

- 6.7. Return to the **Role Bindings** page and click **Create Binding** to create a role binding for the leader user. Fill out the form as follows:

Leader Role Binding Form

Field	Value
Binding Type	Namespace Role Binding (RoleBinding)
Name	leader
Namespace	console-apps
Role Name	admin
Subject	User
Subject Name	leader

Click **Create** to create the RoleBinding in the console-apps namespace.

- 6.8. Click **admin > Log out**, and then log back in as the developer user with the developer password.
Ensure that the developer account can only see the console-apps project.

**Note**

Previous projects from guided exercises that were not deleted upon completion may also display in the list.

- 6.9. You will continue to use the new console-apps project in the next section, so you do not need to delete it.

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab console-admin finish
```



Important

Do not delete the `console-apps` project. It will be used in the upcoming sections.

This concludes the guided exercise.

Managing Workloads and Operators


Objectives

After completing this section, you should be able to manage applications and Kubernetes Operators with the web console.

Exploring Workload Resources

Workload resources such as Pods, Deployments, Stateful Sets, and Config Maps are listed under the **Workloads** menu. Click a resource type to see a list of resources, and then click the name of the resource to navigate to the details page for that resource.

For example, to navigate to the OpenShift DNS operator pod, click **Workloads** > **Pods**, click **Show default projects**, select `openshift-dns-operator` from the Project list at the top of the page, and click the name of the pod listed in the table:



There are often multiple ways to navigate to common resources. Throughout the web UI, associated resources will often link to each other. The deployment details page displays a list of pods. Click any pod name in that list to display the pod details page for that pod.

Managing Workloads

The web console provides specialized editor pages for many workload resources. Use the **Actions** menu on the resource's details page to navigate to the specialized editor pages:




Figure 8.9: Using the Actions menu to modify a deployment

Some useful action pages are described below:

- All resources feature **Edit Labels** and **Edit Annotations** editors.
- Click **Actions > Add Storage** to add a Persistent Volume Claim (PVC) to a deployment.
- To edit the replica count, navigate to the **Deployment Details** page and click **Actions > Edit Pod Count**.
- To modify the update strategy for a deployment, such as changing rolling update parameters, navigate to the **Deployment Details** page and click **Actions > Edit Update Strategy**. To modify the update strategy for a deployment configuration, navigate to the **Deployment Config Details** page and click **Actions > Edit Deployment Config**.
- Navigate to the **Secret Details** page and click **Actions > Edit Secret** to display the **Edit Key/Value Secret** tool, which automatically uses Base64 to encode and decode values.

You can also use the embedded YAML editor to create or modify workload resources. Drag and drop a JSON or YAML file into the browser-based editor to update the resource from a file without using the `oc` command:




Figure 8.10: Editing a resource using the embedded YAML editor

Along with the ability to edit resources in a dedicated page or the embedded YAML editor, you can perform many other common operations directly from the OpenShift web console. For example, to delete a resource, navigate to the resource's details page and click **Actions** > **Delete Resource Type**.

There is often more than one way to perform a particular task. For example, to manually scale a deployment you can navigate to the **Deployment Details** page and then click **Actions** > **Edit Pod Count**, or you can click the arrows next to the pod count without leaving the page.

Deploying Applications

You can create deployment resources from the **Workloads** > **Deployments** page. This section provides a YAML editor with a prepopulated specification to define your YAML resource.

The **Builds** section contains tools for:

- Creating build configurations for Source-to-Image (S2I), Dockerfile, or custom builds.
- Listing and inspecting builds.
- Managing image streams.

After you initiate a deployment or build, use the resource's details and events pages to verify success, or start investigating the cause of a deployment failure.

Installing and Using Operators


Explore community and partner operators on the OpenShift web console's **Operators** > **OperatorHub** page. Many operators are available for installation from the web console. This includes community operators, which Red Hat does not support.

Operators add features and services to your cluster along with automation traditionally performed by human operators, such as deployment coordination or automatic backups. Operators cover a broad range of categories including:

- Traditional databases such as PostgreSQL and MySQL.
- Popular big data frameworks such as Apache Spark.
- Kafka-based streaming platforms such as Red Hat AMQ streams.
- The Knative serverless framework OpenShift Serverless Operator.

Click the operator listing to view details about the operator, such as its version and where to find documentation. When you are ready to install an operator, click **Install** to start the installation. Complete the **Operator Installation** form to select the target namespace and operator approval strategy. You can install operators to target all namespaces or only specific namespaces. Be aware, however, that not all operators support all installation target options.

After you have installed an operator, it appears on the **Operators > Installed Operators** page. If it is installed for a specific namespace, make sure you select the correct project using the project filter at the top of the page:



The operator details page lists the APIs provided by the operator and allows you to create instances of those resources. For example, from the local storage operator page you can create instances of local volume, local volume set, and local volume discovery.



References

For more information, refer to the Red Hat OpenShift Container Platform 4.10 *Web console* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/web_console/index

For more information, refer to the Red Hat OpenShift Container Platform 4.10 *Operators* documentation at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/operators/index

► Guided Exercise

Managing Workloads and Operators

In this exercise, you will manage cluster workloads with the web console.

Outcomes

You should be able to use the OpenShift web console to:

- Install an operator from OperatorHub.
- Use a custom functionalities provided by an operator.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates the resources required for this activity.

```
[student@workstation ~]$ lab console-workloads start
```

Instructions

- 1. As the `admin` user, locate and navigate to the OpenShift web console.

- 1.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Identify the URL for the web console.

```
[student@workstation ~]$ oc whoami --show-console
https://console-openshift-console.apps.ocp4.example.com
```

- 1.3. Open a web browser and navigate to the web console URL:


- `https://console-openshift-console.apps.ocp4.example.com`

- 1.4. Click `localusers` and log in as the `admin` user with the password of `redhat`.

- 2. Inspect the `openshift-console-operator` and `openshift-console` deployments, replica sets, and pods.

- 2.1. On the left pane click **Workloads > Deployments**, and select all projects from the project list at the top. Type `console` in the **Search by name** field. In case is not activated, click **Show default projects** to show all the projects in the cluster.

Notice that OpenShift has a deployment named `console-operator` with a single pod in the `openshift-console-operator` namespace, which operates a deployment named `console` in the `openshift-console` namespace.




- 2.2. On the left pane click **Workloads > Replica Sets**, and type `console` in the **Search by name** field.

Deployments declare a `ReplicaSet` to ensure that a specified number of pods are always running.

- 2.3. In the status column, click **2 of 2 pods** to display the `console` `ReplicaSet` pod list.

► **3. Install the container security operator (CSO) from the OperatorHub.**

- 3.1. On the left pane click **Operators > OperatorHub**, and then click **Security** to display the list of security operators available from OperatorHub.
- 3.2. Scroll down and click **Quay Container Security** to view the operator page, and then click **Install**.



- 3.3. Select the **All namespaces on the cluster** option, and then click **Install** to install the operator. Leave the other form fields unchanged.

- ▶ **4.** Review the resources created by the operator.
 - 4.1. On the left pane click **Workloads > Deployments**, and inspect the list of deployments. You will notice a **container-security-operator** deployment.
 - 4.2. Click the **container-security-operator** deployment, and then click the **Pods** tab to see the pod created by this deployment. Click the pod name to display the **Pod Details** page.
- ▶ **5.** Create a deployment for a simple web application.
 - 5.1. On the **Project** dropdown select the **console-apps** project.
 - 5.2. On the left pane click **Workloads > Deployments**, and then click **Create Deployment** to display the web console YAML editor. Update the YAML as follows and then click **Create**.

**Note**

You can copy the YAML from the `~/D0280/labs/console-workloads/deployment.yaml` file on the workstation machine.

```
kind: Deployment
apiVersion: apps/v1
metadata:
  name: hello-world
  namespace: console-apps
spec:
  selector:
    matchLabels:
      app: hello-world
  replicas: 1
  template:
    metadata:
      labels:
        app: hello-world
    spec:
      containers:
        - name: hello-world
          image: quay.io/redhattraining/hello-world-nginx:v1.0
          ports:
            - containerPort: 8080
              protocol: TCP
```

- ▶ **6.** Verify the deployment's pod using the security operator.
 - 6.1. Switch to the **console-apps** project.

```
[student@workstation ~]$ oc project console-apps
Now using project "console-apps" on server "https://api.ocp4.example.com:6443".
```

- 6.2. List the pods in the **console-apps** project.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-world-685d5cf88-f4h9l	1/1	Running	0	16m

- 6.3. Check the pod `hello-world-685d5cf88-f4h9l` security by getting the `imagemanifestvulns.secsan.quay.redhat.com` resource with the `console-apps/hello-world-685d5cf88-f4h9l` selector

```
[student@workstation ~]$ oc get imagemanifestvulns.secsan.quay.redhat.com \
> --selector=console-apps/hello-world-685d5cf88-f4h9l \
> -o jsonpath='{.items[*].spec.features[*].vulnerabilities[*].name}'
RHSA-2019:3352: gdb security, bug fix, and enhancement update (Low)
RHSA-2020:1635: gdb security and bug fix update (Moderate)
RHSA-2019:3352: gdb security, bug fix, and enhancement update (Low)
RHSA-2020:1635: gdb security and bug fix update (Moderate)
RHSA-2020:1864: gcc security and bug fix update (Moderate)
RHSA-2021:4386: gcc security and bug fix update (Low)
RHSA-2021:4587: gcc security update (Moderate)
...output omitted...
```

Finish

On the workstation machine, use the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab console-workloads finish
```



Important

Do not delete the `console-apps` project or any of the work you performed in this section. It will be used in the next section.

This concludes the guided exercise.

Examining Cluster Metrics

Objectives


After completing this section, you should be able to examine performance and health metrics for cluster nodes and applications.

Viewing Cluster Metrics

The OpenShift web console incorporates useful graphs to visualize cluster and resource analytics. Cluster administrators and users with either the `view` or the `cluster-monitoring-view` cluster role can access the **Home > Overview** page. The Overview page displays a collection of cluster-wide metrics, provides a high-level view of the overall health of the cluster.

The overview includes:


- Current cluster capacity based on CPU, memory, storage, and network usage.
- A time-series graph of total CPU, memory, and disk utilization.
- The ability to display the top consumers of CPU, memory, and storage.



For any of the resources listed in the **Cluster Utilization** section, administrators can click the link for current resource usage. The link displays a window with a breakdown of top consumers for that resource. Top consumers can be sorted by project, by pod, or by node. The list of top consumers can be useful for identifying problematic pods or nodes. For example, a pod with an unexpected memory leak may appear on the top of the list.

Viewing Project Metrics

The **Project Details** page displays metrics that provide an overview of the resources used within the scope of a specific project. The **Utilization** section displays usage information about resources such as CPU and memory along with the ability to display the top consumers for each resource.



All metrics are pulled from Prometheus. Click any graph to navigate to the **Metrics** page. View the executed query, and inspect the data further.

If a resource quota is created for the project, the current project request and limits appear on the **Project Details** page.

Viewing Resource Metrics

When troubleshooting, it is often useful to view metrics at a smaller granularity than the entire cluster or whole project. The **Pod Details** page displays time-series graphs of the CPU, memory, and file system usage for a specific pod. A sudden change in these critical metrics, such as a CPU spike caused by high load, will be visible on this page.




Figure 8.16: Time-series graphs showing various metrics for a pod

Performing Prometheus Queries in the Web Console

The Prometheus UI is a feature-rich tool for visualizing metrics and configuring alerts. The OpenShift web console provides an interface for executing Prometheus queries directly from the web console.

To perform a query, navigate to **Observe > Metrics**, enter a Prometheus Query Language expression in the text field, and click **Run Queries**. The results of the query display as a time-series graph:




Figure 8.17: Using a Prometheus query to display a time-series graph



Note

The Prometheus Query Language is not discussed in detail in this course. See the references below for a link to the official documentation.



References

For more information, refer to the Red Hat OpenShift Container Platform 4.10 *Monitoring* documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/html-single/monitoring/index

Querying Prometheus

<https://prometheus.io/docs/prometheus/latest/querying/basics/>

► Guided Exercise

Examining Cluster Metrics

In this exercise, you will examine the metrics page and dashboard within the web console.

Outcomes

You should be able to use the Red Hat OpenShift web console to:

- View cluster, project, pod, and node metrics.
- Identify a pod consuming large amounts of memory or CPU.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates the resources required for this exercise.

```
[student@workstation ~]$ lab console-metrics start
```

Instructions

- 1. As the `admin` user, locate and navigate to the OpenShift web console.

- 1.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Identify the URL for the web console.

```
[student@workstation ~]$ oc whoami --show-console
https://console-openshift-console.apps.ocp4.example.com
```

- 1.3. Open a web browser and navigate to the web console URL:

- `https://console-openshift-console.apps.ocp4.example.com`

- 1.4. Click `localusers` and log in as the `admin` user with the `redhat` password.

- 2. In this guided exercise, you will see how changes in load are displayed in the web console. Start by observing baseline healthy metrics on the **Overview**, **Pod Details**, and **Project Details** pages.

- 2.1. Click **Home** > **Overview** to display the **Overview** page. Scroll down to the **Cluster Utilization** section, which displays a time-series historical graph of the cluster's CPU, memory, and disk usage.
- 2.2. For each resource in the table, such as **CPU**, **Memory**, or **Filesystem**, click the usage link on the right to view the **Top Consumers** of that resource. By default, the window filters top consumers by project, but you can filter by pod or by node instead.
- 2.3. Click the usage link for **Memory**, filter the top consumers by pod, and then click the name of the pod that consumes the most memory resources.




Figure 8.18: Memory breakdown: Top consumers by pod

- 2.4. Click **Metrics** to display memory usage, CPU usage, and file system usage time-series historical graphs.
 - 2.5. Click **Home** > **Projects**, and then click **console-apps** to display the **Project Details** page.
Notice the **Utilization** section, which displays the metrics for the workloads running in the **console-apps** project. The links in the **Usage** column open windows displaying the pods that consume the most resources. The workloads are running safely within their limits.
 - 2.6. Scroll down to the **Resource Quotas** section, which displays the current CPU and memory usage compared to the allotted quota.
- 3. Find and review the baseline health metrics of a compute node.
- 3.1. Click **Compute** > **Nodes**, then click any of the nodes in the list.
 - 3.2. On the **Utilization** section, notice the time-series graphs that display the metrics for the individual node that you selected.




Figure 8.19: Time-series graphs showing various metrics for a node

- 4. On the workstation machine, execute the `load.sh` script to generate load on the example `books` deployment. The application intentionally contains a memory leak that consumes multiple megabytes of RAM with every request to its `/leak` path.

4.1. In a terminal on the `workstation` machine, run the following command.

```
[student@workstation ~]$ ~/D0280/labs/console-metrics/load.sh
```

- 5. In the OpenShift web console, observe the change in metrics and identify the problematic pod. The data displayed in the web console automatically refreshes, so there is no need to reload the page.
- 5.1. Click **Home** > **Projects**, and then click `console-apps` to display the **Project Details** page. Watch the **Memory Usage** time-series graph to monitor for changes.




Figure 8.20: Utilization graphs indicating a possible memory leak

The memory leak may take a minute or two before it is significant enough to be visible. Although both CPU and memory increase, the total CPU usage remains low.

- 5.2. Click **Home** > **Overview** to display the **Overview** page. The memory consumed by the load test may be too small to notice across a large cluster, but the **Memory breakdown** window (sorted by pod) provides a convenient list of pods using the most memory. View the **Memory breakdown** window by clicking the usage link for **Memory**. Sort the top consumers by pod.





Figure 8.21: The books pod is a top memory consumer

The **books** pod appears at or near the top of the list. If it is not on the list, you may need to wait a minute longer for the load script to complete.


- 5.3. Click the **books** pod link in the **Memory breakdown** window to navigate to the **Pod Details** page. Click the **Metrics** link to see the pod metrics, notice the climbing memory leak visible in the **Memory Usage** time-series graph.



- 5.4. Click **Observe** > **Metrics** to display the web console **Metrics** page. Type the following Prometheus query in the expression input field:

```
avg(container_memory_working_set_bytes{namespace='console-apps'}) BY (pod)
```

Click **Run Queries** to view the results in the OpenShift web console.



► **6.** Delete the `console-apps` project and stop the load test.

6.1. Click **Home > Projects**, and then click **Delete Project** in the menu at the end of the `console-apps` row.

<code>console-apps</code>	✓ Active	admin	No labels	⋮
<code>default</code>	✓ Active	No requester	No labels	Edit Project
<code>kube-node-lease</code>	✓ Active	No requester	No labels	Delete Project
<code>kube-public</code>	✓ Active	No requester	No labels	⋮

6.2. In the **Delete Project** dialog box, type `console-apps` and then click **Delete**.

6.3. If the `load.sh` script is still running on the workstation terminal, press **Ctrl+C** in the terminal to stop the load test.

Finish

On the workstation machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab console-metrics finish
```

This concludes the guided exercise.

► Lab

Managing a Cluster with the Web Console

In this lab, you will manage the OpenShift cluster using the web console.

Outcomes

You should be able to use the OpenShift web console to:

- Modify a secret to add `htpasswd` entries for new users.
- Configure a new project with role-based access controls and resource quotas.
- Use an OperatorHub operator to deploy a database.
- Create a deployment, service, and route for a web application.
- Troubleshoot an application using events and logs.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates a directory for the exercise files.

```
[student@workstation ~]$ lab console-review start
```

Instructions

1. Log in to the OpenShift web console as the `admin` user.
2. Add `htpasswd` entries to the `localusers` secret for users named `dba` and `tester` using `redhat` as the password.
3. Create a new `app-team` group that contains the `developer` and `dba` users.
4. Create a new `console-review` project with a `view` role binding for the `tester` user and an `edit` role binding for the `app-team` group. Set a resource quota that limits the project to two pods.
5. As the `dba` user, deploy a MySQL Database instance into the `console-review` project. Set `database` as the app name and `famous` as the username, password and database name. Use the `registry.redhat.io/rhel8/mysql-80:1` image.
6. As the `developer` user, create a deployment, service, and route in the `console-review` project with issues that you will troubleshoot in the next step. Use the `quay.io/redhattraining/famous-quotes:2.1` image, one replica, and name all of the new resources `famous-quotes`. When correctly configured, the `famous-quotes` application connects to the MySQL database and displays a list of famous quotes.

**Note**

You can copy the deployment and service YAML resources from the `~/DO280/labs/console-review` directory on the `workstation` machine.

Specify the following environment variables in the deployment:

Deployment Environment Variables

Name	Value
QUOTES_HOSTNAME	database
QUOTES_USER	famous
QUOTES_DATABASE	famous
QUOTES_PASSWORD	famous

**Important**

You will troubleshoot issues with the deployment in the next step.

7. Troubleshoot and fix the deployment issues.
8. Navigate to the famous-quotes website in a browser and observe the working application.

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab console-review grade
```

Finish

As the student user on the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab console-review finish
```

This concludes the lab.

► Solution

Managing a Cluster with the Web Console

In this lab, you will manage the OpenShift cluster using the web console.

Outcomes

You should be able to use the OpenShift web console to:

- Modify a secret to add `htpasswd` entries for new users.
- Configure a new project with role-based access controls and resource quotas.
- Use an OperatorHub operator to deploy a database.
- Create a deployment, service, and route for a web application.
- Troubleshoot an application using events and logs.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that the cluster API is reachable and creates a directory for the exercise files.

```
[student@workstation ~]$ lab console-review start
```

Instructions

1. Log in to the OpenShift web console as the `admin` user.
 - 1.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Identify the URL for the web console.

```
[student@workstation ~]$ oc whoami --show-console
https://console-openshift-console.apps.ocp4.example.com
```

- 1.3. Open a web browser and navigate to the web console URL:
 - `https://console-openshift-console.apps.ocp4.example.com`
- 1.4. Click `localusers` and log in as the `admin` user with the `redhat` password.

2. Add `htpasswd` entries to the `localusers` secret for users named `dba` and `tester` using `redhat` as the password.
 - 2.1. In the Red Hat OpenShift Container Platform web UI, click **Workloads > Secrets**, and then select `openshift-config` from the **Project** search list to display the secrets for the `openshift-config` project. If default projects are hidden, then click **Show default projects** to show all the projects in the cluster.
 - 2.2. Scroll to the bottom of the page and click the `localusers` link to display the `localusers` **Secret Details**.
 - 2.3. Click **Actions > Edit Secret** at the top of the page to navigate to the **Edit Key/Value Secret** tool.
 - 2.4. Use a terminal on the `workstation` machine to generate an encrypted `htpasswd` entry for both users.

```
[student@workstation ~]$ htpasswd -n -b dba redhat
dba:$apr1$YF4aCK.9$gho0THlWTC.cLByNEHdaV
```

```
[student@workstation ~]$ htpasswd -n -b tester redhat
tester:$apr1$XdTSqET7$ihkC5bIs7PhYUm2KhiI.0
```

- 2.5. Append the terminal output from the `htpasswd` commands to the `htpasswd` value in the OpenShift web console's secrets editor and then click **Save**.

```
admin:$apr1$Au9.fFr$0k5wvUBd3eeBt0baa77.dae
leader:$apr1$/abo4Hybn7a.tG5Zo0Bn.QWefXckiy1
developer:$apr1$RjqTY4cv$xl3.BQfg42moSxwnTNkh.
dba:$apr1$YF4aCK.9$gho0THlWTC.cLByNEHdaV
tester:$apr1$XdTSqET7$ihkC5bIs7PhYUm2KhiI.0
```

3. Create a new `app-team` group that contains the `developer` and `dba` users.
 - 3.1. Click **User Management > Groups**, and then click **Create Group**. Use the YAML editor to define a Group resource as follows, and then click **Create** to add the new `app-team` group. Update the YAML as follows and then click **Create**.

```
apiVersion: user.openshift.io/v1
kind: Group
metadata:
  name: app-team
users:
  - developer
  - dba
```

4. Create a new `console-review` project with a `view` role binding for the `tester` user and an `edit` role binding for the `app-team` group. Set a resource quota that limits the project to two pods.
 - 4.1. Click **Home > Projects** to view the Projects page, and then click **Create Project**. Type `console-review` in the **Name** field, and then provide an optional **Display Name** and **Description**. Click **Create**.

- 4.2. Click **User Management > RoleBindings** and then click **Create binding**. Complete the form as follows to create a namespaced Role Binding for the **app-team** group. Click **Create** to create the namespaced RoleBinding.

App Team Role Binding Form

Field	Value
Binding Type	Namespace Role Binding (RoleBinding)
Name	app-team
Namespace	console-review
Role Name	edit
Subject	Group
Subject Name	app-team

- 4.3. Click the **RoleBindings** link to return to the **RoleBindings** page, and then click **Create Binding**. Complete the form as follows to create a namespaced Role Binding for the **tester** user. Click **Create** to create the namespaced RoleBinding.

Tester Role Binding Form

Field	Value
Binding Type	Namespace Role Binding (RoleBinding)
Name	tester
Namespace	console-review
Role Name	view
Subject	User
Subject Name	tester

- 4.4. Click **Administration > Resource Quotas**, and then click **Create Resource Quota**. Modify the YAML document to specify a limit of four pods as follows. Remove the CPU and memory requests and limits, and then click **Create**.

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: quota
  namespace: console-review
spec:
  hard:
    pods: '2'
```


5. As the `dba` user, deploy a MySQL Database instance into the `console-review` project. Set `database` as the app name and `famous` as the username, password and database name. Use the `registry.redhat.io/rhel8/mysql-80:1` image.

5.1. Log in as the `dba` user with the `redhat` password.

```
[student@workstation ~]$ oc login -u dba -p redhat \
> https://api.ocp4.example.com:6443
...output omitted...
```

5.2. Switch to the `console-review` project.

```
[student@workstation ~]$ oc project console-review
Now using project "console-review" on server "https://api.ocp4.example.com:6443".
```

5.3. Create a new database deployment using the container image located at `registry.redhat.io/rhel8/mysql-80:1`.

```
[student@workstation ~]$ oc new-app --name database \
> --image registry.redhat.io/rhel8/mysql-80:1 \
> -e MYSQL_USER=famous \
> -e MYSQL_PASSWORD=famous \
> -e MYSQL_DATABASE=famous
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "database" created
    deployment.apps "database" created
    service "database" created
--> Success
...output omitted...
```

6. As the `developer` user, create a deployment, service, and route in the `console-review` project with issues that you will troubleshoot in the next step. Use the `quay.io/redhattraining/famous-quotes:2.1` image, one replica, and name all of the new resources `famous-quotes`. When correctly configured, the `famous-quotes` application connects to the MySQL database and displays a list of famous quotes.



Note

You can copy the deployment and service YAML resources from the `~/D0280/labs/console-review` directory on the `workstation` machine.

Specify the following environment variables in the deployment:

Deployment Environment Variables

Name	Value
QUOTES_HOSTNAME	database
QUOTES_USER	famous
QUOTES_DATABASE	famous
QUOTES_PASSWORD	famous

**Important**

You will troubleshoot issues with the deployment in the next step.

- 6.1. Click **admin > Log out**, and then log in as the **developer** user with the **developer** password.
- 6.2. Switch to the **Administrator** perspective. Then click the **console-review** project to switch to the **console-review** project.
- 6.3. Click **Workloads > Deployments**, and then click **Create Deployment** to display the web console YAML editor. Update the YAML as follows and then click **Create**.

```
kind: Deployment
apiVersion: apps/v1
metadata:
  name: famous-quotes
  namespace: console-review
spec:
  selector:
    matchLabels:
      app: famous-quotes
  replicas: 2
  template:
    metadata:
      labels:
        app: famous-quotes
    spec:
      containers:
        - name: famous-quotes
          image: 'quay.io/redhattraining/famous-quotes:2.1'
          ports:
            - containerPort: 8000
              protocol: TCP
          env:
            - name: QUOTES_HOSTNAME
              value: database
            - name: QUOTES_USER
              value: famous
            - name: QUOTES_DATABASE
```

```

value: famous
- name: QUOTES_PASSWORD
  value: famous

```

**Note**

You can use the YAML manifest from the `~/D0280/labs/console-review/deployment.yaml` to create the resource.

- 6.4. Click **Networking** > **Services**, and then click **Create Service** to display the web console YAML editor. Update the YAML as follows and then click **Create**.

```

kind: Service
apiVersion: v1
metadata:
  name: famous-quotes
  namespace: console-review
spec:
  selector:
    app: famous-quotes
  ports:
    - protocol: TCP
      port: 8000
      targetPort: 8000

```

**Note**

You can use the YAML manifest from the `~/D0280/labs/console-review/service.yaml` to create the resource.

- 6.5. Click **Networking** > **Routes**, and then click **Create Route**. Complete the form as follows, leaving the other fields unchanged, and then click **Create**:

Create Route Form

Field	Value
Name	famous
Service	famous-quotes
Target Port	8000 → 8000 (TCP)

7. Troubleshoot and fix the deployment issues.
- 7.1. Click **developer** > **Log out**, and then log in as the **admin** user with the password **redhat**.
 - 7.2. Click **Home** > **Events**, and then select **console-review** from the project list filter at the top. Notice the famous-quotes quota error.

```
(combined from similar events): Error creating: pods "famous-quotes-8449f7dc6d-
k94xq" is forbidden: exceeded quota: quota, requested: pods=1, used: pods=2,
limited: pods=2
```

- 7.3. Click **Administration** > **Resource Quotas**, and then select **console-review** from the **Project** filter list.
- 7.4. Click the **quota** link in the list of resource quotas, and then click the **YAML** tab. Modify the **spec** to specify a limit of four pods as follows, and then click **Save**.

```
kind: ResourceQuota
apiVersion: v1
metadata:
  name: quota
  namespace: console-review
...output omitted...
spec:
  hard:
    pods: '4'
...output omitted...
```

**Note**

The project requires a pod for the famous-quotes's specified replica and an additional pod in order to roll out a change.

- 7.5. Click **Workloads** > **Pods**, and review the list of pods. The **famous-quotes** pod takes a minute or two to appear on the list.
8. Navigate to the famous-quotes website in a browser and observe the working application.
- 8.1. Click **Networking** > **Routes**, click the **famous** route name, and then click the link in the **Location** column. Firefox will open a new tab rendering a list of famous quotes.

Evaluation

As the student user on the **workstation** machine, use the **lab** command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab console-review grade
```

Finish

As the student user on the **workstation** machine, use the **lab** command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab console-review finish
```

This concludes the lab.

Summary

In this chapter, you learned that:

- The OpenShift web console provides a GUI for visualizing and managing OpenShift resources.
- Some resources feature a specialized page that makes creating and editing resources more convenient than writing YAML by hand, such as the **Edit Key/Value Secret** editor, which automatically handles Base64 encoding and decoding.
- You can install partner and community operators from the embedded **OperatorHub** page.
- Cluster-wide metrics such as CPU, memory, and storage usage are displayed on the **Dashboards** page.
- **Project Details** pages display metrics specific to the project, such as the top ten memory consumers by pod and the current resource quota usage.

Chapter 9

Comprehensive Review

Goal

Review tasks from *Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster*

Objectives

- Review tasks from *Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster*

Sections

- Comprehensive Review

Labs

- Troubleshoot an OpenShift Cluster and Applications
- Configure a Project Template with Resource and Network Restrictions

Comprehensive Review

Objectives

After completing this section, you should be able to demonstrate knowledge and skills learned in *Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster*.

Reviewing Red Hat OpenShift Administration II: Operating a Production Kubernetes Cluster

Before beginning the comprehensive review for this course, you should be comfortable with the topics covered in each chapter.

You can refer to earlier sections in the textbook for extra study.

Chapter 1, Describing the Red Hat OpenShift Container Platform

Describe the architecture of OpenShift Container Platform.

- Describe the typical usage of the product and its features.
- Describe the architecture of Red Hat OpenShift Container Platform.
- Describe what a cluster operator is, how it works, and name the major cluster operators.

Chapter 2, Verifying the Health of a Cluster

Describe OpenShift installation methods and verify the health of a newly installed cluster.

- Describe the OpenShift installation process, full-stack automation, and pre-existing infrastructure installation methods.
- Execute commands that assist in troubleshooting, verify that the OpenShift nodes are healthy, and troubleshoot common issues with OpenShift and Kubernetes deployments.
- Identify the components and resources of persistent storage and deploy an application that uses a persistent volume claim.

Chapter 3, Configuring Authentication and Authorization

Configure authentication with the HTPasswd identity provider and assign roles to users and groups.

- Configure the HTPasswd identity provider for OpenShift authentication.
- Define role-based access controls and apply permissions to users.

Chapter 4, Configuring Application Security

Restrict permissions of applications using security context constraints and protect access credentials using secrets.

- Create and apply secrets to manage sensitive information and share secrets between applications.
- Create service accounts and apply permissions, and manage security context constraints.

Chapter 5, Configuring OpenShift Networking for Applications

Troubleshoot OpenShift software-defined networking (SDN) and configure network policies.

- Troubleshoot OpenShift software-defined networking using the command-line interface.
- Allow and protect network connections to applications inside an OpenShift cluster.
- Restrict network traffic between projects and pods.

Chapter 6, Controlling Pod Scheduling

Control the nodes on which a pod runs.

- Describe pod scheduling algorithms, the methods used to influence scheduling, and apply these methods.
- Limit the resources consumed by containers, pods, and projects.
- Control the number of replicas of a pod, specify the number of replicas in a deployment, manually scale the number of replicas, and create a horizontal pod autoscaler (HPA) resource.

Chapter 7, Describing Cluster Updates

Describe how to perform a cluster update.

Describe the cluster update process.

Chapter 8, Managing a Cluster with the Web Console

Manage a Red Hat OpenShift cluster using the web console.

- Perform cluster administration with the web console.
- Manage applications and Kubernetes Operators with the web console.
- Examine performance and health metrics for cluster nodes and applications.

► Lab

Troubleshoot an OpenShift Cluster and Applications

In this review, you will enable developers to access a cluster and troubleshoot application deployments.

Outcomes

You should be able to:

- Create a new project.
- Perform a smoke test of the OpenShift cluster by creating an application using the source-to-image process.
- Create applications using a deployment resource.
- Use the HTPasswd identity provider for managing users.
- Create and manage groups.
- Manage RBAC and SCC for users and groups.
- Manage secrets for databases and applications.
- Troubleshoot common problems.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command:

- Ensures that the cluster API is reachable.
- Removes existing users and groups.
- Removes existing identity providers.
- Removes the `cluster-admin` cluster role binding from the `admin` user.
- Ensures that authenticated users can create new projects.

```
[student@workstation ~]$ lab review-troubleshoot start
```

Instructions

Complete the following tasks:

1. As the `kubeadmin` user, create the `review-troubleshoot` project. The password for the `kubeadmin` user is located in the `/usr/local/etc/ocp4.config` file on the `RHT_OCP4_KUBEADM_PASSWD` line. Perform all subsequent tasks in the `review-troubleshoot` project.
2. Perform a smoke test of the cluster to verify basic cluster functionality. Use a `deployment` to create an application named `hello-world-nginx`. The application source code is located in the `hello-world-nginx` subdirectory of the repository.
 - <https://github.com/RedHatTraining/DO280-apps>

Create a route for the application, use `hello-world.apps.ocp4.example.com` as the hostname. Verify that the application responds to external requests.

3. Configure the cluster to use an HTPasswd identity provider. The name of the identity provider is `cluster-users`. The identity provider reads `htpasswd` credentials stored in the `comprevue-users` secret.

Ensure that four user accounts exist: `admin`, `leader`, `developer`, and `qa-engineer`. All user accounts must use `review` as the password.

Add the `cluster-admin` role to the `admin` user.

4. As the `admin` user, create three user groups: `leaders`, `developers`, and `qa`.

Assign the `leader` user to the `leaders` group, the `developer` user to the `developers` group, and the `qa-engineer` user to the `qa` group.

Assign roles to each group:

- Assign the `self-provisioner` role to the `leaders` group, which allows members to create projects. For this role to be effective, you must also remove the ability of any authenticated user to create new projects.
- Assign the `edit` role to the `developers` group for the `review-troubleshoot` project only, which allows members to create and delete project resources.
- Assign the `view` role to the `qa` group for the `review-troubleshoot` project only, which provides members with read access to project resources.

5. As the `developer` user, use a deployment to create an application named `mysql` in the `review-troubleshoot` project. Use the image available at `registry.redhat.io/rhel8/mysql-80:1-139`. This application provides a shared database service for other project applications.

Create a generic secret named `mysql` using `password` as the key and `r3dh4t123` as the value.

Set the `MYSQL_ROOT_` environment variables from the values in the `mysql` secret.

Configure the `mysql` database application to mount a persistent volume claim (PVC) to the `/var/lib/mysql/data` directory within the pod. The PVC must be 2 GB in size and must only request the `ReadWriteOnce` access mode.

6. As the `developer` user, use a `deployment` to create an application named `wordpress`. Create the application in the `review-troubleshoot` project. Use the image available at `quay.io/redhattraining/wordpress:5.7-php7.4-apache`.

The Wordpress application requires that you set several environment variables when executing the `oc new-app` command:

- `WORDPRESS_DB_HOST` with a value of `mysql`.
- `WORDPRESS_DB_NAME` to have a value of `wordpress`.
- `WORDPRESS_DB_USER` to have a value of `root`.
- `WORDPRESS_USER` with a value of `wpuser`.
- `WORDPRESS_PASSWORD` with a value of `wppass`.
- `WORDPRESS_TITLE` to have a value of `review-troubleshoot`.
- `WORDPRESS_URL` with a value of `wordpress.apps.ocp4.example.com`.
- `WORDPRESS_EMAIL` with a value of `student@redhat.com`.

Set the `WORDPRESS_DB_*` environment variables to retrieve the database password from the `mysql` secret.

The `wordpress` application requires the `anyuid` security context constraint. Create a service account named `wordpress-sa`, and then assign the `anyuid` security context constraint to it. Configure the `wordpress` deployment to use the `wordpress-sa` service account.

The `wordpress` application also requires that the database `WORDPRESS_DB_NAME` exists on the database server. Create an empty database named `wordpress` by executing the SQL command `CREATE DATABASE wordpress` inside the MySQL pod.

Create a route for the application, use `wordpress.apps.ocp4.example.com` as the hostname. If you correctly deploy the application, then an installation wizard displays when you access the application from a browser.

7. As the `developer` user, deploy the `famous-quotes` application in the `review-troubleshoot` project using the `~/D0280/labs/review-troubleshoot/deploy_famous-quotes.sh` script. This script creates the `defaultdb` database and the resources defined in the `~/D0280/labs/review-troubleshoot/famous-quotes.yaml` file.

Use the `mysql` secret to initialize environment variables for the `famous-quotes` deployment with the prefix `QUOTES_`.

The application pods do not initially deploy after you execute the script. The `famous-quotes` deployment specifies a node selector, and there are no cluster nodes with a matching node label.

Remove the node selector from the deployment, which enables OpenShift to schedule application pods on any available node.

Create a route for the application, use `quotes.apps.ocp4.example.com` as the hostname. Verify that the application responds to external requests.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab review-troubleshoot grade
```

Finish

As the student user on the workstation machine, use the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab review-troubleshoot finish
```

This concludes the lab.

► Solution

Troubleshoot an OpenShift Cluster and Applications

In this review, you will enable developers to access a cluster and troubleshoot application deployments.

Outcomes

You should be able to:

- Create a new project.
- Perform a smoke test of the OpenShift cluster by creating an application using the source-to-image process.
- Create applications using a deployment resource.
- Use the HTPasswd identity provider for managing users.
- Create and manage groups.
- Manage RBAC and SCC for users and groups.
- Manage secrets for databases and applications.
- Troubleshoot common problems.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command:

- Ensures that the cluster API is reachable.
- Removes existing users and groups.
- Removes existing identity providers.
- Removes the `cluster-admin` cluster role binding from the `admin` user.
- Ensures that authenticated users can create new projects.

```
[student@workstation ~]$ lab review-troubleshoot start
```

Instructions

Complete the following tasks:

1. As the `kubeadmin` user, create the `review-troubleshoot` project. The password for the `kubeadmin` user is located in the `/usr/local/etc/ocp4.config` file on the `RHT_OCP4_KUBEADM_PASSWORD` line. Perform all subsequent tasks in the `review-troubleshoot` project.
 - 1.1. Source the classroom configuration file that is accessible at `/usr/local/etc/ocp4.config`, and log in as the `kubeadmin` user.

```
[student@workstation ~]$ source /usr/local/etc/ocp4.config
[student@workstation ~]$ oc login -u kubeadmin -p ${RHT_OCP4_KUBEADM_PASSWD} \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

1.2. Create the review-troubleshoot project.

```
[student@workstation ~]$ oc new-project review-troubleshoot
Now using project "review-troubleshoot" on server
"https://api.ocp4.example.com:6443".

...output omitted...
```

2. Perform a smoke test of the cluster to verify basic cluster functionality. Use a deployment to create an application named `hello-world-nginx`. The application source code is located in the `hello-world-nginx` subdirectory of the repository.

- <https://github.com/RedHatTraining/DO280-apps>

Create a route for the application, use `hello-world.apps.ocp4.example.com` as the hostname. Verify that the application responds to external requests.

2.1. Use the `oc new-app` command to create the `hello-world-nginx` deployment.

```
[student@workstation ~]$ oc new-app --name hello-world-nginx \
> https://github.com/RedHatTraining/DO280-apps \
> --context-dir hello-world-nginx
...output omitted...
--> Creating resources ...
  imagestream.image.openshift.io "ubi8" created
  imagestream.image.openshift.io "hello-world-nginx" created
  buildconfig.build.openshift.io "hello-world-nginx" created
  deployment.apps "hello-world-nginx" created
  service "hello-world-nginx" created
--> Success
...output omitted...
Run 'oc status' to view your app.
```

- 2.2. Create a route to the application by exposing the `hello-world-nginx` service. Set the hostname to `hello-world.apps.ocp4.example.com`.

```
[student@workstation ~]$ oc expose service hello-world-nginx \
> --hostname hello-world.apps.ocp4.example.com
route.route.openshift.io/hello-world-nginx exposed
```

- 2.3. Wait until the application pod is running.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-world-nginx-1-build	0/1	Completed	0	2m
hello-world-nginx-695754d9f7-8rv4x	1/1	Running	0	1m

2.4. Verify access to the application.

```
[student@workstation ~]$ curl -s "http://hello-world.apps.ocp4.example.com" | \
>   grep -i Hello
<h1>Hello, world from nginx!</h1>
```

3. Configure the cluster to use an HTTPBasic identity provider. The name of the identity provider is `cluster-users`. The identity provider reads `htpasswd` credentials stored in the `compreview-users` secret.

Ensure that four user accounts exist: `admin`, `leader`, `developer`, and `qa-engineer`. All user accounts must use `review` as the password.

Add the `cluster-admin` role to the `admin` user.

- 3.1. Create a temporary `htpasswd` authentication file at `/tmp/cluster-users`.

```
[student@workstation ~]$ touch /tmp/cluster-users
```

- 3.2. Populate the `/tmp/cluster-users` file with the required user and password values.

```
[student@workstation ~]$ for USER in admin leader developer qa-engineer
>   do
>   htpasswd -B -b /tmp/cluster-users "${USER}" review
>   done
Adding password for user admin
Adding password for user leader
Adding password for user developer
Adding password for user qa-engineer
```

- 3.3. Create a `compreview-users` secret from the `/tmp/cluster-users` file.

```
[student@workstation ~]$ oc create secret generic compreview-users \
>   --from-file htpasswd=/tmp/cluster-users -n openshift-config
secret/compreview-users created
```

- 3.4. Export the existing OAuth resource to a YAML file.

```
[student@workstation ~]$ oc get oauth cluster -o yaml > /tmp/oauth.yaml
```

- 3.5. Edit the `/tmp/oauth.yaml` file to add the HTTPBasic identity provider definition to the `identityProviders` list. Set the identity provider name to `cluster-users`, and set the `fileName` name to `compreview-users`.

After making these modifications, the file reads as follows:

```

apiVersion: config.openshift.io/v1
kind: OAuth
metadata:
  ...output omitted...
  name: cluster
  ...output omitted...
spec:
  identityProviders:
    - name: cluster-users
      mappingMethod: claim
      type: HTPasswd
      httppasswd:
        fileData:
          name: compreview-users

```

**Note**

The `name`, `mappingMethod`, `type`, and `httppasswd` keys all use the same indentation.

3.6. Replace the existing OAuth resource with the resource definition in the modified file:

```

[student@workstation ~]$ oc replace -f /tmp/oauth.yaml
oauth.config.openshift.io/cluster replaced

```

3.7. Assign the `admin` user the `cluster-admin` role.

```

[student@workstation ~]$ oc adm policy add-cluster-role-to-user \
> cluster-admin admin
Warning: User 'admin' not found
clusterrole.rbac.authorization.k8s.io/cluster-admin added: "admin"

```

**Note**

You can safely ignore the warning about the `admin` user not being found. The `user/admin` resource does not exist in your OpenShift cluster until after the `admin` user logs in for the first time.

3.8. Wait until the pods in the `openshift-authentication` namespace are restarted.

```

[student@workstation ~]$ oc get pods -n openshift-authentication
NAME                                READY   STATUS    RESTARTS   AGE
oauth-openshift-768c6476d4-tg8kc    1/1     Running   0           40s
oauth-openshift-768c6476d4-9dc48    1/1     Running   0           60s
oauth-openshift-768c6476d4-qngxw    1/1     Running   0           80s

```

**Note**

You might need to run the command multiple times until the desired condition is reached.

4. As the `admin` user, create three user groups: `leaders`, `developers`, and `qa`.

Assign the `leader` user to the `leaders` group, the `developer` user to the `developers` group, and the `qa-engineer` user to the `qa` group.

Assign roles to each group:

- Assign the `self-provisioner` role to the `leaders` group, which allows members to create projects. For this role to be effective, you must also remove the ability of any authenticated user to create new projects.
- Assign the `edit` role to the `developers` group for the `review-troubleshoot` project only, which allows members to create and delete project resources.
- Assign the `view` role to the `qa` group for the `review-troubleshoot` project only, which provides members with read access to project resources.

4.1. Log in as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p review
Login successful.
```

...output omitted...



Note

If the login command fails, then check that the pods in the `openshift-authentication` namespace are restarted and running.

4.2. Create the three user groups.

```
[student@workstation ~]$ for GROUP in leaders developers qa
> do
> oc adm groups new "${GROUP}"
> done
group.user.openshift.io/leaders created
group.user.openshift.io/developers created
group.user.openshift.io/qa created
```

4.3. Add each user to the appropriate group.

```
[student@workstation ~]$ oc adm groups add-users leaders leader
group.user.openshift.io/leaders added: "leader"

[student@workstation ~]$ oc adm groups add-users developers developer
group.user.openshift.io/developers added: "developer"

[student@workstation ~]$ oc adm groups add-users qa qa-engineer
group.user.openshift.io/qa added: "qa-engineer"
```

4.4. Allow members of the `leaders` group to create new projects:

```
[student@workstation ~]$ oc adm policy add-cluster-role-to-group \
> self-provisioner leaders
clusterrole.rbac.authorization.k8s.io/self-provisioner added: "leaders"
```

- 4.5. Remove the `self-provisioner` cluster role from the `system:authenticated:oauth` group.

```
[student@workstation ~]$ oc adm policy remove-cluster-role-from-group \
> self-provisioner system:authenticated:oauth
Warning: Your changes may get lost whenever a master is restarted, unless you
prevent reconciliation of this rolebinding using the following command:
oc annotate clusterrolebinding.rbac self-provisioners
'rbac.authorization.kubernetes.io/autoupdate=false' --overwrite
clusterrole.rbac.authorization.k8s.io/self-provisioner removed:
"system:authenticated:oauth"
```

**Note**

You can safely ignore the warning about your changes being lost.

- 4.6. Allow members of the `developers` group to create and delete resources in the `review-troubleshoot` project:

```
[student@workstation ~]$ oc policy add-role-to-group edit developers
clusterrole.rbac.authorization.k8s.io/edit added: "developers"
```

- 4.7. Allow members of the `qa` group to view project resources:

```
[student@workstation ~]$ oc policy add-role-to-group view qa
clusterrole.rbac.authorization.k8s.io/view added: "qa"
```

5. As the `developer` user, use a deployment to create an application named `mysql` in the `review-troubleshoot` project. Use the image available at `registry.redhat.io/rhel8/mysql-80:1-139`. This application provides a shared database service for other project applications.
- Create a generic secret named `mysql` using `password` as the key and `r3dh4t123` as the value.
- Set the `MYSQL_ROOT_` environment variables from the values in the `mysql` secret.
- Configure the `mysql` database application to mount a persistent volume claim (PVC) to the `/var/lib/mysql/data` directory within the pod. The PVC must be 2 GB in size and must only request the `ReadWriteOnce` access mode.

- 5.1. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p review
Login successful.

...output omitted...
```

- 5.2. Create a new application to deploy a `mysql` database server. Use the `oc new-app` command to create a deployment.

```
[student@workstation ~]$ oc new-app --name mysql \
> --image registry.redhat.io/rhel8/mysql-80:1-139
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "mysql" created
    deployment.apps "mysql" created
    service "mysql" created
--> Success
...output omitted...
    Run 'oc status' to view your app.
```

- 5.3. Create a generic secret for the MySQL database named `mysql` using a key of `password` and a value of `r3dh4t123`.

```
[student@workstation ~]$ oc create secret generic mysql \
> --from-literal password=r3dh4t123
secret/mysql created
```

**Note**

You will use this secret to provide the database password for the application.

- 5.4. Use the `mysql` secret to initialize environment variables for the `mysql` deployment.

```
[student@workstation ~]$ oc set env deployment mysql \
> --prefix MYSQL_ROOT_ --from secret/mysql
deployment.apps/mysql updated
```

- 5.5. Use the `oc set volumes` command to configure persistent storage for the `mysql` deployment. The command automatically creates a persistent volume claim with the specified size and access mode. Mounting the volume to the `/var/lib/mysql/data` directory provides access to stored data, even if one database pod is deleted and another database pod is created.

```
[student@workstation ~]$ oc set volumes deployment/mysql --name mysql-storage \
> --add --type pvc --claim-size 2Gi --claim-mode rwo \
> --mount-path /var/lib/mysql/data
deployment.apps/mysql volume updated
```

- 5.6. Verify that the `mysql` pod successfully redeploys after configuring the deployment to use a secret and a volume.

```
[student@workstation ~]$ oc get pods -l deployment=mysql
NAME                READY   STATUS    RESTARTS   AGE
mysql-bbb6b5fbb-dmq9x 1/1     Running   0           60s
```

**Note**

You might need to run the command multiple times until the `mysql` pod displays both 1/1 and Running.

5.7. Verify that a persistent volume claim exists with the correct size and access mode.

```
[student@workstation ~]$ oc get pvc
```

NAME	STATUS	...	CAPACITY	ACCESS MODES	STORAGECLASS	AGE
pvc-ks52v	Bound	...	2Gi	RWO	nfs-storage	2m33s

5.8. Verify that the running `mysql` pod mounts a volume to the `/var/lib/mysql/data` directory.

```
[student@workstation ~]$ oc exec -it mysql-bbb6b5fbb-dmq9x -- \
> df -h /var/lib/mysql/data
```

Filesystem	Size	Used	Avail	Use%	Mounted on
...:/exports/review-troubleshoot-pvc-...	40G	859M	40G	3%	/var/lib/mysql/data

6. As the `developer` user, use a deployment to create an application named `wordpress`. Create the application in the `review-troubleshoot` project. Use the image available at quay.io/redhattraining/wordpress:5.7-php7.4-apache.

The Wordpress application requires that you set several environment variables when executing the `oc new-app` command:

- `WORDPRESS_DB_HOST` with a value of `mysql`.
- `WORDPRESS_DB_NAME` to have a value of `wordpress`.
- `WORDPRESS_DB_USER` to have a value of `root`.
- `WORDPRESS_USER` with a value of `wpuser`.
- `WORDPRESS_PASSWORD` with a value of `wppass`.
- `WORDPRESS_TITLE` to have a value of `review-troubleshoot`.
- `WORDPRESS_URL` with a value of `wordpress.apps.ocp4.example.com`.
- `WORDPRESS_EMAIL` with a value of `student@redhat.com`.

Set the `WORDPRESS_DB_*` environment variables to retrieve the database password from the `mysql` secret.

The `wordpress` application requires the `anyuid` security context constraint. Create a service account named `wordpress-sa`, and then assign the `anyuid` security context constraint to it. Configure the `wordpress` deployment to use the `wordpress-sa` service account.

The `wordpress` application also requires that the database `WORDPRESS_DB_NAME` exists on the database server. Create an empty database named `wordpress` by executing the SQL command `CREATE DATABASE wordpress` inside the MySQL pod.

Create a route for the application, use `wordpress.apps.ocp4.example.com` as the hostname. If you correctly deploy the application, then an installation wizard displays when you access the application from a browser.

6.1. Deploy a `wordpress` application as a deployment.

```
[student@workstation ~]$ oc new-app --name wordpress \
> --image quay.io/redhattraining/wordpress:5.7-php7.4-apache \
> -e WORDPRESS_DB_HOST=mysql \
> -e WORDPRESS_DB_NAME=wordpress \
> -e WORDPRESS_DB_USER=root \
> -e WORDPRESS_USER=wpuser \
> -e WORDPRESS_PASSWORD=wppass \
> -e WORDPRESS_TITLE=review-troubleshoot \
> -e WORDPRESS_URL=wordpress.apps.ocp4.example.com \
> -e WORDPRESS_EMAIL=student@redhat.com
...output omitted...
--> Creating resources ...
    imagestream.image.openshift.io "wordpress" created
    deployment.apps "wordpress" created
    service "wordpress" created
--> Success
...output omitted...
    Run 'oc status' to view your app.
```

- 6.2. Add the `WORDPRESS_DB_*` environment variables to the `wordpress` deployment. The secret contains the database password.

```
[student@workstation ~]$ oc set env deployment/wordpress \
> --prefix WORDPRESS_DB_ --from secret/mysql
deployment.apps/wordpress updated
```

- 6.3. Create the `wordpress-sa` service account.

```
[student@workstation ~]$ oc create serviceaccount wordpress-sa
serviceaccount/wordpress-sa created
```

- 6.4. Log in to the cluster as the `admin` user and grant `anyuid` privileges to the `wordpress-sa` service account.

```
[student@workstation ~]$ oc login -u admin -p review
Login successful.

...output omitted...
[student@workstation ~]$ oc adm policy add-scc-to-user anyuid -z wordpress-sa
clusterrole.rbac.authorization.k8s.io/system:openshift:scc:anyuid added:
"wordpress-sa"
```

- 6.5. Switch back to the `developer` user to perform the remaining steps. Log in to the cluster as the `developer` user.

```
[student@workstation ~]$ oc login -u developer -p review
Login successful.

...output omitted...
```

- 6.6. Configure the `wordpress` deployment to use the `wordpress-sa` service account.

```
[student@workstation ~]$ oc set serviceaccount deployment/wordpress \
>   wordpress-sa
deployment.apps/wordpress serviceaccount updated
```

6.7. Create the `wordpress` database. Use the `mysql` client tool on the running pod.

```
[student@workstation ~]$ oc get pods -l deployment=mysql
NAME                READY   STATUS    RESTARTS   AGE
mysql-fbf67ff96-c4sq7 1/1     Running   0           10m


[student@workstation ~]$ oc exec mysql-fbf67ff96-c4sq7 -- \
> /usr/bin/mysql -uroot -e "CREATE DATABASE wordpress"
```

6.8. Create a route for the `wordpress` application. Set the hostname to `wordpress.apps.ocp4.example.com`.

```
[student@workstation ~]$ oc expose service wordpress \
>   --hostname wordpress.apps.ocp4.example.com
route.route.openshift.io/wordpress exposed
```

6.9. Use a web browser to verify access to the application URL. When you correctly deploy the application, a setup wizard displays in the browser.

- `http://wordpress.apps.ocp4.example.com`



7. As the `developer` user, deploy the `famous-quotes` application in the `review-troubleshoot` project using the `~/D0280/labs/review-troubleshoot/deploy_famous-quotes.sh` script. This script creates the `defaultdb` database and the resources defined in the `~/D0280/labs/review-troubleshoot/famous-quotes.yaml` file.

Use the `mysql` secret to initialize environment variables for the `famous-quotes` deployment with the prefix `QUOTES_`.

The application pods do not initially deploy after you execute the script. The `famous-quotes` deployment specifies a node selector, and there are no cluster nodes with a matching node label.

Remove the node selector from the deployment, which enables OpenShift to schedule application pods on any available node.

Create a route for the application, use `quotes.apps.ocp4.example.com` as the hostname. Verify that the application responds to external requests.

- 7.1. Run the `~/D0280/labs/review-troubleshoot/deploy_famous-quotes.sh` script.

```
[student@workstation ~]$ ~/D0280/labs/review-troubleshoot/deploy_famous-quotes.sh
Creating famous-quotes database
Deploying famous-quotes application
deployment.apps/famous-quotes created
service/famous-quotes created
```

- 7.2. Use the `mysql` secret to initialize environment variables with the prefix `QUOTES_`.

```
[student@workstation ~]$ oc set env deployment famous-quotes \
> --prefix QUOTES_ --from secret/mysql
```

- 7.3. Verify that the `famous-quotes` application pod is not scheduled for deployment.

```
[student@workstation ~]$ oc get pods -l app=famous-quotes
```

NAME	READY	STATUS	RESTARTS	AGE
famous-quotes-85ff8679d7-vhvbk	0/1	Pending	0	41s

- 7.4. See if any project events provide information about the problem.

```
[student@workstation ~]$ oc get events --sort-by='{.lastTimestamp}' | \
> grep -i Scheduling
...output omitted...
30s Warning FailedScheduling pod/famous-quotes-1-deploy 0/3 nodes are
available: 3 node(s) didn't match node selector.
...output omitted...
```

- 7.5. Save the `famous-quotes` deployment resource to a file.

```
[student@workstation ~]$ oc get deployment/famous-quotes \
> -o yaml > /tmp/famous-deploy.yaml
```

- 7.6. Use an editor to remove the node selector from the `/tmp/famous-deploy.yaml` file. Search for the `nodeSelector` line in the file. Then, delete the following two lines from the `/tmp/famous-deploy.yaml` file.

```
nodeSelector:
  env: quotes
```

- 7.7. Replace the `famous-quotes` deployment with the modified file.

```
[student@workstation ~]$ oc replace -f /tmp/famous-deploy.yaml
deployment.apps/famous-quotes replaced
```

- 7.8. Wait a few moments and then run the `oc get pods` command to ensure that the `famous-quotes` application is now running.

```
[student@workstation ~]$ oc get pods -l app=famous-quotes
```

NAME	READY	STATUS	RESTARTS	AGE
famous-quotes-2-gmz2j	1/1	Running	0	53s

**Note**

You might need to run the command multiple times until the desired condition is reached.

- 7.9. Create a route for the `famous-quotes` application. Set the hostname to `quotes.apps.ocp4.example.com`.

```
[student@workstation ~]$ oc expose service famous-quotes \
> --hostname quotes.apps.ocp4.example.com
route.route.openshift.io/famous-quotes exposed
```

- 7.10. Verify that the `famous-quotes` application responds to requests sent to the URL.

- `http://quotes.apps.ocp4.example.com`

```
[student@workstation ~]$ curl -s "http://quotes.apps.ocp4.example.com" | \
> egrep '</?title>'
<title>Quotes</title>
```

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab review-troubleshoot grade
```

Finish

As the student user on the `workstation` machine, use the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab review-troubleshoot finish
```

This concludes the lab.

► Lab

Configure a Project Template with Resource and Network Restrictions

In this review, you will configure the cluster's default project template to ensure that new projects enforce default quotas, resource limits, and network policies.

Outcomes

You should be able to:

- Modify the default project template to automatically create limit ranges, resource quotas, and network policies.
- Create a TLS secret using the provided files.
- Mount a secret as a volume within an application.
- Create a passthrough route to an application.
- Configure an application to automatically scale.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

The command:

- Ensures that the cluster API is reachable.
- Configures the HTPasswd identity provider and provides access to the `admin`, `leader`, and `developer` users.
- Downloads sample YAML files to `~/DO280/labs/review-template/sample-files`.
- Creates the `review-template-test` project and deploys an application to the project that you can use to test your network policies.
- Adds the `network.openshift.io/policy-group=ingress` label to the default namespace.
- Generates certificate files needed for the secure application.

```
[student@workstation ~]$ lab review-template start
```



Note

If you need help getting started, consider using the *Post-installation network configuration* chapter of the Red Hat OpenShift Container Platform Post-installation configuration guide as a reference.

Instructions

Complete the following tasks:

1. As the `admin` user, update the OpenShift cluster to use a new project template. The project template must automatically create the network policy, limit range, and add quota resources for new projects. New projects must automatically have a label matching the name of the project. For example, a project named `test` has the `name=test` label.

The following table guides you to the needed resources.

Resource	Requirements
Project	Includes a label with the name of the project.
NetworkPolicy	Policy 1: - Routes are accessible to external traffic; this means that traffic is allowed from pods in namespaces with the <code>network.openshift.io/policy-group=ingress</code> label. Policy 2: - Pods in the same namespace can communicate with each other. - Pods do not respond to pods that exist in a different namespace, except namespaces with the <code>network.openshift.io/policy-group=ingress</code> label.
LimitRange	- Each container requests 30 millicores of CPU. - Each container requests 30 MiB of memory. - Each container is limited to 100 millicores of CPU. - Each container is limited to 100 MiB of memory.
ResourceQuota	- Projects are limited to 10 pods. - Projects can request a maximum of 1 GiB of memory. - Projects can request a maximum of 2 CPUs. - Projects can use a maximum of 4 GiB of memory. - Projects can use a maximum of 4 CPUs.

2. As the `developer` user, create a project named `review-template`. Ensure that the `review-template` project inherits the settings specified in the new project template. In the `review-template` project, create a deployment named `hello-secure` using the container image located at `quay.io/redhattraining/hello-world-secure:v1.0`.



Note

The `hello-secure` pod does not run successfully until after you provide access to the TLS certificate and key required by the NGINX server.

3. As the `developer` user, create a TLS secret using the `hello-secure-combined.pem` certificate and the `hello-secure-key.pem` key located in the `~/D0280/labs/review-template` directory. Use the logs from the failed `hello-secure` pod to determine the expected mount point for the certificate. Mount the TLS secret as a volume in the pod using the identified directory. Verify that the `hello-secure` pod successfully redeploys.
4. The `hello-secure-combined.pem` certificate is valid for a single host name. Use the `openssl x509` command with the `-noout` and `-ext 'subjectAltName'` options to

read the `hello-secure-combined.pem` certificate and identify the host name. As the `developer` user, create a passthrough route to the `hello-secure` service using the identified host name. Verify that the route responds to external requests.

**Note**

The `x509(1)` man page provides information on the `openssl x509` command.

5. As the `developer` user, configure the `hello-secure` deployment to scale automatically. The deployment must have at least one pod running. If the average CPU utilization exceeds 80%, then the deployment scales to a maximum of five pods.

**Note**

You can use the script located at `~/D0280/solutions/review-template/test-hpa.sh` to test that your deployment scales as expected.

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab review-template grade
```

Finish

As the student user on the `workstation` machine, use the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab review-template finish
```

This concludes the lab.

► Solution

Configure a Project Template with Resource and Network Restrictions

In this review, you will configure the cluster's default project template to ensure that new projects enforce default quotas, resource limits, and network policies.

Outcomes

You should be able to:

- Modify the default project template to automatically create limit ranges, resource quotas, and network policies.
- Create a TLS secret using the provided files.
- Mount a secret as a volume within an application.
- Create a passthrough route to an application.
- Configure an application to automatically scale.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

The command:

- Ensures that the cluster API is reachable.
- Configures the HTPasswd identity provider and provides access to the `admin`, `leader`, and `developer` users.
- Downloads sample YAML files to `~/DO280/labs/review-template/sample-files`.
- Creates the `review-template-test` project and deploys an application to the project that you can use to test your network policies.
- Adds the `network.openshift.io/policy-group=ingress` label to the default namespace.
- Generates certificate files needed for the secure application.

```
[student@workstation ~]$ lab review-template start
```



Note

If you need help getting started, consider using the *Post-installation network configuration* chapter of the Red Hat OpenShift Container Platform Post-installation configuration guide as a reference.

Instructions

Complete the following tasks:

1. As the `admin` user, update the OpenShift cluster to use a new project template. The project template must automatically create the network policy, limit range, and add quota resources for new projects. New projects must automatically have a label matching the name of the project. For example, a project named `test` has the `name=test` label.

The following table guides you to the needed resources.

Resource	Requirements
Project	Includes a label with the name of the project.
NetworkPolicy	Policy 1: - Routes are accessible to external traffic; this means that traffic is allowed from pods in namespaces with the <code>network.openshift.io/policy-group=ingress</code> label. Policy 2: - Pods in the same namespace can communicate with each other. - Pods do not respond to pods that exist in a different namespace, except namespaces with the <code>network.openshift.io/policy-group=ingress</code> label.
LimitRange	- Each container requests 30 millicores of CPU. - Each container requests 30 MiB of memory. - Each container is limited to 100 millicores of CPU. - Each container is limited to 100 MiB of memory.
ResourceQuota	- Projects are limited to 10 pods. - Projects can request a maximum of 1 GiB of memory. - Projects can request a maximum of 2 CPUs. - Projects can use a maximum of 4 GiB of memory. - Projects can use a maximum of 4 CPUs.

- 1.1. Log in to your OpenShift cluster as the `admin` user.

```
[student@workstation ~]$ oc login -u admin -p redhat \
> https://api.ocp4.example.com:6443
Login successful.

...output omitted...
```

- 1.2. Use the `oc adm create-bootstrap-project-template` command to create a new YAML file that you will customize for this exercise. Save the file to `~/D0280/labs/review-template/project-template.yaml`.

```
[student@workstation ~]$ oc adm create-bootstrap-project-template \
> -o yaml > ~/D0280/labs/review-template/project-template.yaml
```

- 1.3. Change to the `~/D0280/labs/review-template` directory.

```
[student@workstation ~]$ cd ~/D0280/labs/review-template
```

- 1.4. Edit the `project-template.yaml` file to add a label for new projects. Add the bold lines, ensuring proper indentation, and then save the file. The `PROJECT_NAME` variable takes the value of the project name.

```
...output omitted...
- apiVersion: project.openshift.io/v1
  kind: Project
  metadata:
    labels:
      name: ${PROJECT_NAME}
    annotations:
...output omitted...
```

- 1.5. List the files in the `~/D0280/labs/review-template/sample-files` directory. The directory provides two sample network policy files, a sample limit range file, and a sample resource quota file.

```
[student@workstation review-template]$ ls sample-files
allow-from-openshift-ingress.yaml    limitrange.yaml
allow-same-namespace.yaml            resourcequota.yaml
```

- 1.6. Add the content of the `sample-files/allow-*.yaml` files to the `project-template.yaml` file. Because the `project-template.yaml` file expects a list of resources, the first line of each resource must begin with a `-`, and you must indent the remainder of the content accordingly.

Edit the `project-template.yaml` file to add the network policy resources. Add the bold lines, ensuring proper indentation, and then save the file.

```
...output omitted...
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: ${PROJECT_ADMIN_USER}
- apiVersion: networking.k8s.io/v1
  kind: NetworkPolicy
  metadata:
    name: allow-from-openshift-ingress
  spec:
    podSelector: {}
    ingress:
      - from:
          - namespaceSelector:
              matchLabels:
                network.openshift.io/policy-group: ingress
- apiVersion: networking.k8s.io/v1
  kind: NetworkPolicy
  metadata:
    name: allow-same-namespace
  spec:
```

```

    podSelector: {}
  ingress:
    - from:
        - podSelector: {}
  parameters:
    - name: PROJECT_NAME
    ...output omitted...

```

- 1.7. Add the content of the `sample-files/limitrange.yaml` and `sample-files/resourcequota.yaml` files to the `project-template.yaml` file. Because the `project-template.yaml` file expects a list of resources, the first line of each resource must begin with a `-`, and you must indent the remainder of the content accordingly.

Edit the `project-template.yaml` file to add the limit range and resource quota resources. Add the bold lines, ensuring proper indentation, and then save the file.

```

    ...output omitted...
    ingress:
      - from:
          - podSelector: {}
- apiVersion: v1
  kind: LimitRange
  metadata:
    name: project-limitrange
  spec:
    limits:
      - default:
          memory: 100Mi
          cpu: 100m
          defaultRequest:
            memory: 30Mi
            cpu: 30m
          type: Container
- apiVersion: v1
  kind: ResourceQuota
  metadata:
    name: project-quota
  spec:
    hard:
      pods: '10'
      requests.cpu: '2'
      requests.memory: 1Gi
      limits.cpu: '4'
      limits.memory: 4Gi
  parameters:
    - name: PROJECT_NAME
    ...output omitted...

```

- 1.8. Use the `oc create` command to create a new template resource in the `openshift-config` namespace using the `project-template.yaml` file.

**Note**

The `~/D0280/solutions/review-template/project-template.yaml` file contains the correct configuration and can be used for comparison.

```
[student@workstation review-template]$ oc create -f project-template.yaml \
> -n openshift-config
template.template.openshift.io/project-request created
```

- 1.9. List the templates in the `openshift-config` namespace. You use the template name in the next step.

```
[student@workstation review-template]$ oc get templates -n openshift-config
```

NAME	DESCRIPTION	PARAMETERS	OBJECTS
project-request		5 (5 blank)	6

- 1.10. Update the `projects.config.openshift.io/cluster` resource to use the new template.

```
[student@workstation review-template]$ oc edit \
> projects.config.openshift.io/cluster
```

Modify the `spec: {}` line to match the following bold lines. Use the template name identified in the `openshift-config` namespace. Ensure proper indentation, and then save your changes.

```
...output omitted...
spec:
  projectRequestTemplate:
    name: project-request
```

- 1.11. A successful change redeployes the `apiserver` pods in the `openshift-apiserver` namespace. Monitor the redeployment.

```
[student@workstation review-template]$ watch oc get pods -n openshift-apiserver
```

Press `Ctrl+C` to end the `watch` command after all three new pods are running.

```
Every 2.0s: oc get pods -n openshift-apiserver
```

NAME	READY	STATUS	RESTARTS	AGE
apiserver-75cfdcf877-257vs	2/2	Running	0	61s
apiserver-75cfdcf877-l2xnv	2/2	Running	0	29s
apiserver-75cfdcf877-rn9fs	2/2	Running	0	47s

2. As the `developer` user, create a project named `review-template`. Ensure that the `review-template` project inherits the settings specified in the new project template. In the `review-template` project, create a deployment named `hello-secure` using the container image located at `quay.io/redhattraining/hello-world-secure:v1.0`.

**Note**

The `hello-secure` pod does not run successfully until after you provide access to the TLS certificate and key required by the NGINX server.

- 2.1. Log in to your OpenShift cluster as the `developer` user.

```
[student@workstation review-template]$ oc login -u developer -p developer
Login successful.

...output omitted...
```

- 2.2. Create the `review-template` project.

```
[student@workstation review-template]$ oc new-project review-template
Now using project "review-template" on server "https://api.ocp4.example.com:6443".

...output omitted...
```

- 2.3. List the network policy, limit range, and quota resources in the `review-template` project.

```
[student@workstation review-template]$ oc get networkpolicies
NAME                                POD-SELECTOR  AGE
allow-from-openshift-ingress       <none>        94s
allow-same-namespace               <none>        94s

[student@workstation review-template]$ oc get limitranges
NAME          CREATED AT
project-limitrange  2022-05-04T21:53:09Z

[student@workstation review-template]$ oc get resourcequotas
NAME          AGE    REQUEST
project-quota  103s  pods: 0/10, requests.cpu: 0/2, requests.memory: 0/1Gi

LIMIT
limits.cpu: 0/4, limits.memory: 0/4Gi
```

- 2.4. Verify that the `review-template` project has the `name=review-template` label.

```
[student@workstation review-template]$ oc get project review-template \
> --show-labels
NAME          DISPLAY NAME  STATUS  LABELS
review-template              Active  ...,name=review-template
```

- 2.5. Use the `oc new-app` command to create the `hello-secure` deployment using the `quay.io/redhattraining/hello-world-secure:v1.0` container image.

```
[student@workstation review-template]$ oc new-app --name hello-secure \
> --image quay.io/redhattraining/hello-world-secure:v1.0
...output omitted...
```

```
--> Creating resources ...
    imagestream.image.openshift.io "hello-secure" created
    deployment.apps "hello-secure" created
    service "hello-secure" created
--> Success
...output omitted...
    Run 'oc status' to view your app.
```

2.6. Verify that the `hello-secure` pod does not start successfully.

```
[student@workstation review-template]$ watch oc get pods
```

Press `Ctrl+C` to end the `watch` command after the pod has a status of either `CrashLoopBackOff` or `Error`.

```
Every 2.0s: oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-secure-6475f657c9-rmgsr	0/1	CrashLoopBackOff	1	14s

- As the `developer` user, create a TLS secret using the `hello-secure-combined.pem` certificate and the `hello-secure-key.pem` key located in the `~/D0280/labs/review-template` directory. Use the logs from the failed `hello-secure` pod to determine the expected mount point for the certificate. Mount the TLS secret as a volume in the pod using the identified directory. Verify that the `hello-secure` pod successfully redeploys.

- Create a TLS secret using the `hello-secure-combined.pem` certificate and the `hello-secure-key.pem` key.

```
[student@workstation review-template]$ oc create secret tls hello-tls \
> --cert hello-secure-combined.pem --key hello-secure-key.pem
secret/hello-tls created
```

- Identify the name of the failed pod.

```
[student@workstation review-template]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-secure-6475f657c9-rmgsr	0/1	CrashLoopBackOff	4	2m45s

- Examine the logs of the failed pod. The logs indicate that the pod attempts to use the `/run/secrets/nginx/tls.crt` file, but that the file does not exist.

```
[student@workstation review-template]$ oc logs hello-secure-6475f657c9-rmgsr
...output omitted...
nginx: [emerg] BIO_new_file("/run/secrets/nginx/tls.crt") failed (SSL:
error:02001002:system library:fopen:No such file or directory:fopen('/run/
secrets/nginx/tls.crt','r') error:2006D080:BIO routines:BIO_new_file:no such file)
```

- Use the `oc set volumes` command to mount the secret into the `/run/secrets/nginx` directory.

```
[student@workstation review-template]$ oc set volumes deployment/hello-secure \
> --add --type secret --secret-name hello-tls --mount-path /run/secrets/nginx
info: Generated volume name: volume-hlrtf
deployment.apps/hello-secure volume updated
```

3.5. Verify that the `hello-secure` pod successfully redeploy.

```
[student@workstation review-template]$ watch oc get pods
```

Press `Ctrl+C` to end the `watch` command after the pod displays `1/1` and `Running`.

```
Every 2.0s: oc get pods ...
```

NAME	READY	STATUS	RESTARTS	AGE
hello-secure-6bd8fcccb4-nhwr2	1/1	Running	0	25s

4. The `hello-secure-combined.pem` certificate is valid for a single host name. Use the `openssl x509` command with the `-noout` and `-ext 'subjectAltName'` options to read the `hello-secure-combined.pem` certificate and identify the host name. As the `developer` user, create a passthrough route to the `hello-secure` service using the identified host name. Verify that the route responds to external requests.



Note

The `x509(1)` man page provides information on the `openssl x509` command.

- 4.1. Examine the `hello-secure-combined.pem` certificate using the `openssl x509` command.

```
[student@workstation review-template]$ openssl x509 \
> -in hello-secure-combined.pem -noout -ext 'subjectAltName'
X509v3 Subject Alternative Name:
    DNS:hello-secure.apps.ocp4.example.com
```

- 4.2. Create a passthrough route to the `hello-secure` service pointing to `hello-secure.apps.ocp4.example.com`.

```
[student@workstation review-template]$ oc create route passthrough \
> --service hello-secure --hostname hello-secure.apps.ocp4.example.com
route.route.openshift.io/hello-secure created
```

- 4.3. Return to the `/home/student` directory.

```
[student@workstation review-template]$ cd
```

Use the `curl` command to verify that the route responds to external requests.

```
[student@workstation ~]$ curl -s https://hello-secure.apps.ocp4.example.com \
> | grep -i Hello
<h1>Hello, world from nginx!</h1>
```

5. As the `developer` user, configure the `hello-secure` deployment to scale automatically. The deployment must have at least one pod running. If the average CPU utilization exceeds 80%, then the deployment scales to a maximum of five pods.

**Note**

You can use the script located at `~/D0280/solutions/review-template/test-hpa.sh` to test that your deployment scales as expected.

- 5.1. Use the `oc autoscale` command to create the horizontal pod autoscaler resource.

```
[student@workstation ~]$ oc autoscale deployment/hello-secure \
> --min 1 --max 5 --cpu-percent 80
horizontalpodautoscaler.autoscaling/hello-secure autoscaled
```

- 5.2. Run the provided `test-hpa.sh` script. The script uses the `ab` command to apply load using the Apache benchmarking tool.

```
[student@workstation ~]$ ~/D0280/solutions/review-template/test-hpa.sh
...output omitted...
Benchmarking hello-secure.apps.ocp4.example.com (be patient)
Completed 10000 requests
Completed 20000 requests
...output omitted...
Finished 100000 requests
...output omitted...
```

- 5.3. Verify that at least two, but not more than five, `hello-secure` pods are running.

```
[student@workstation ~]$ oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
hello-secure-6bd8fcccb4-7qjcz	1/1	Running	0	96s
hello-secure-6bd8fcccb4-d67xd	1/1	Running	0	66s
hello-secure-6bd8fcccb4-m8vxp	1/1	Running	0	96s
hello-secure-6bd8fcccb4-nhwr2	1/1	Running	0	9m36s

Evaluation

As the student user on the workstation machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab review-template grade
```

Finish

As the student user on the workstation machine, use the `lab` command to complete this exercise.

```
[student@workstation ~]$ lab review-template finish
```

This concludes the lab.

