



OpenShift Container Platform 4.19

Building applications

Creating and managing applications on OpenShift Container Platform

OpenShift Container Platform 4.19 Building applications

Creating and managing applications on OpenShift Container Platform

Legal Notice

Copyright © Red Hat.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered 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 and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. 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 an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides instructions for the various ways to create and manage instances of user-provisioned applications running on OpenShift Container Platform. This includes working with projects and provisioning applications using the Open Service Broker API.

Table of Contents

CHAPTER 1. BUILDING APPLICATIONS OVERVIEW	8
1.1. WORKING ON A PROJECT	8
1.2. WORKING ON AN APPLICATION	8
1.2.1. Creating an application	8
1.2.2. Maintaining an application	8
1.2.3. Deploying an application	8
CHAPTER 2. PROJECTS	9
2.1. WORKING WITH PROJECTS	9
2.1.1. Creating a project	9
2.1.1.1. Creating a project by using the web console	9
2.1.1.2. Creating a project by using the CLI	11
2.1.2. Viewing a project	11
2.1.2.1. Viewing a project by using the web console	11
2.1.2.2. Viewing a project using the CLI	12
2.1.3. Providing access permissions to your project using the Developer perspective	12
2.1.4. Customizing the available cluster roles using the web console	13
2.1.5. Adding to a project	14
2.1.6. Checking the project status	15
2.1.6.1. Checking project status by using the web console	15
2.1.6.2. Checking project status by using the CLI	15
2.1.7. Deleting a project	16
2.1.7.1. Deleting a project by using the web console	16
2.1.7.2. Deleting a project by using the CLI	17
2.2. CREATING A PROJECT AS ANOTHER USER	17
2.2.1. API impersonation	17
2.2.2. Impersonating a user when you create a project	17
2.3. CONFIGURING PROJECT CREATION	18
2.3.1. About project creation	18
2.3.2. Modifying the template for new projects	18
2.3.3. Disabling project self-provisioning	19
2.3.4. Customizing the project request message	21
CHAPTER 3. CREATING APPLICATIONS	23
3.1. USING TEMPLATES	23
3.1.1. Understanding templates	23
3.1.2. Uploading a template	23
3.1.3. Creating an application by using the web console	23
3.1.4. Creating objects from templates by using the CLI	24
3.1.4.1. Adding labels	24
3.1.4.2. Listing parameters	24
3.1.4.3. Generating a list of objects	25
3.1.5. Modifying uploaded templates	27
3.1.6. Using instant app and quick start templates	27
3.1.6.1. Quick start templates	27
3.1.6.1.1. Web framework quick start templates	28
3.1.7. Writing templates	28
3.1.7.1. Writing the template description	29
3.1.7.2. Writing template labels	33
3.1.7.3. Writing template parameters	34
3.1.7.4. Writing the template object list	36

3.1.7.5. Marking a template as bindable	37
3.1.7.6. Exposing template object fields	37
3.1.7.7. Waiting for template readiness	39
3.1.7.8. Creating a template from existing objects	41
3.2. CREATING APPLICATIONS BY USING THE DEVELOPER PERSPECTIVE	42
3.2.1. Prerequisites	43
3.2.2. Creating sample applications	43
3.2.3. Creating applications by using Quick Starts	44
3.2.4. Importing a codebase from Git to create an application	44
3.2.5. Creating applications by deploying container image	48
3.2.6. Deploying a Java application by uploading a JAR file	49
3.2.7. Using the Devfile registry to access devfiles	50
3.2.8. Using the Developer Catalog to add services or components to your application	51
3.2.9. Additional resources	51
3.3. CREATING APPLICATIONS FROM INSTALLED OPERATORS	52
3.3.1. Creating an etcd cluster using an Operator	52
3.4. CREATING APPLICATIONS BY USING THE CLI	53
3.4.1. Creating an application from source code	53
3.4.1.1. Local	53
3.4.1.2. Remote	54
3.4.1.3. Build strategy detection	54
3.4.1.4. Language detection	55
3.4.2. Creating an application from an image	56
3.4.2.1. Docker Hub MySQL image	56
3.4.2.2. Image in a private registry	56
3.4.2.3. Existing image stream and optional image stream tag	56
3.4.3. Creating an application from a template	57
3.4.3.1. Template parameters	57
3.4.4. Modifying application creation	57
3.4.4.1. Specifying environment variables	58
3.4.4.2. Specifying build environment variables	59
3.4.4.3. Specifying labels	59
3.4.4.4. Viewing the output without creation	59
3.4.4.5. Creating objects with different names	60
3.4.4.6. Creating objects in a different project	60
3.4.4.7. Creating multiple objects	60
3.4.4.8. Grouping images and source in a single pod	60
3.4.4.9. Searching for images, templates, and other inputs	61
3.4.4.10. Setting the import mode	61
3.5. CREATING APPLICATIONS USING RUBY ON RAILS	61
3.5.1. Prerequisites	61
3.5.2. Setting up the database	62
3.5.3. Writing your application	62
3.5.3.1. Creating a welcome page	63
3.5.3.2. Configuring application for OpenShift Container Platform	64
3.5.3.3. Storing your application in Git	64
3.5.4. Deploying your application to OpenShift Container Platform	65
3.5.4.1. Creating the database service	66
3.5.4.2. Creating the frontend service	66
3.5.4.3. Creating a route for your application	68
CHAPTER 4. VIEWING APPLICATION COMPOSITION BY USING THE TOPOLOGY VIEW	69
4.1. PREREQUISITES	69

4.2. VIEWING THE TOPOLOGY OF YOUR APPLICATION	69
4.3. INTERACTING WITH APPLICATIONS AND COMPONENTS	70
4.4. SCALING APPLICATION PODS AND CHECKING BUILDS AND ROUTES	72
4.5. ADDING COMPONENTS TO AN EXISTING PROJECT	72
4.6. GROUPING MULTIPLE COMPONENTS WITHIN AN APPLICATION	73
4.7. ADDING SERVICES TO YOUR APPLICATION	74
4.8. REMOVING SERVICES FROM YOUR APPLICATION	75
4.9. LABELS AND ANNOTATIONS USED FOR THE TOPOLOGY VIEW	76
4.10. ADDITIONAL RESOURCES	77
CHAPTER 5. EXPORTING APPLICATIONS	78
5.1. PREREQUISITES	78
5.2. PROCEDURE	78
CHAPTER 6. WORKING WITH HELM CHARTS	80
6.1. UNDERSTANDING HELM	80
6.1.1. Key features	80
6.1.2. Red Hat Certification of Helm charts for OpenShift	80
6.1.3. Additional resources	80
6.2. INSTALLING HELM	80
6.2.1. On Linux	81
6.2.2. On Windows 7/8	81
6.2.3. On Windows 10	81
6.2.4. On MacOS	82
6.3. CONFIGURING CUSTOM HELM CHART REPOSITORIES	82
6.3.1. Installing a Helm chart on an OpenShift Container Platform cluster	83
6.3.2. Creating Helm releases using the Developer perspective	84
6.3.3. Using Helm in the web terminal	85
6.3.4. Creating a custom Helm chart on OpenShift Container Platform	85
6.3.5. Adding custom Helm chart repositories	86
6.3.6. Adding namespace-scoped custom Helm chart repositories	87
6.3.7. Creating credentials and CA certificates to add Helm chart repositories	89
6.3.8. Filtering Helm Charts by their certification level	91
6.3.9. Disabling Helm Chart repositories	91
6.4. WORKING WITH HELM RELEASES	92
6.4.1. Prerequisites	92
6.4.2. Upgrading a Helm release	92
6.4.3. Rolling back a Helm release	92
6.4.4. Deleting a Helm release	93
CHAPTER 7. DEPLOYMENTS	94
7.1. UNDERSTANDING DEPLOYMENTS	94
7.1.1. Building blocks of a deployment	94
7.1.1.1. Replica sets	94
7.1.1.2. Replication controllers	95
7.1.2. Deployments	97
7.1.3. DeploymentConfig objects	97
7.1.4. Comparing Deployment and DeploymentConfig objects	99
7.1.4.1. Design	99
7.1.4.2. Deployment-specific features	99
7.1.4.2.1. Rollover	100
7.1.4.2.2. Proportional scaling	100
7.1.4.2.3. Pausing mid-rollout	100
7.1.4.3. DeploymentConfig object-specific features	100

7.1.4.3.1. Automatic rollbacks	100
7.1.4.3.2. Triggers	100
7.1.4.3.3. Lifecycle hooks	100
7.1.4.3.4. Custom strategies	100
7.2. MANAGING DEPLOYMENT PROCESSES	100
7.2.1. Managing DeploymentConfig objects	101
7.2.1.1. Starting a deployment	101
7.2.1.2. Viewing a deployment	101
7.2.1.3. Retrying a deployment	102
7.2.1.4. Rolling back a deployment	102
7.2.1.5. Executing commands inside a container	102
7.2.1.6. Viewing deployment logs	103
7.2.1.7. Deployment triggers	104
7.2.1.7.1. Config change deployment triggers	104
7.2.1.7.2. Image change deployment triggers	104
7.2.1.7.3. Setting deployment triggers	105
7.2.1.8. Setting deployment resources	105
7.2.1.9. Scaling manually	107
7.2.1.10. Accessing private repositories from DeploymentConfig objects	107
7.2.1.11. Assigning pods to specific nodes	107
7.2.1.12. Running a pod with a different service account	108
7.3. USING DEPLOYMENT STRATEGIES	109
7.3.1. Choosing a deployment strategy	109
7.3.2. Rolling strategy	109
7.3.2.1. Canary deployments	111
7.3.2.2. Creating a rolling deployment	111
7.3.2.3. Editing a deployment by using the Developer perspective	112
7.3.2.4. Starting a rolling deployment using the Developer perspective	113
7.3.3. Recreate strategy	114
7.3.3.1. Editing a deployment by using the Developer perspective	115
7.3.3.2. Starting a recreate deployment using the Developer perspective	116
7.3.4. Custom strategy	117
7.3.4.1. Editing a deployment by using the Developer perspective	119
7.3.5. Lifecycle hooks	119
7.3.5.1. Pod-based lifecycle hook	120
7.3.5.2. Setting lifecycle hooks	121
7.4. USING ROUTE-BASED DEPLOYMENT STRATEGIES	121
7.4.1. Proxy shards and traffic splitting	122
7.4.2. N-1 compatibility	122
7.4.3. Graceful termination	122
7.4.4. Blue-green deployments	123
7.4.4.1. Setting up a blue-green deployment	123
7.4.5. A/B deployments	124
7.4.5.1. Load balancing for A/B testing	124
7.4.5.1.1. Managing weights of an existing route using the web console	126
7.4.5.1.2. Managing weights of a new route using the web console	126
7.4.5.1.3. Managing weights using the CLI	126
7.4.5.1.4. One service, multiple Deployment objects	128
7.4.6. Additional resources	129
CHAPTER 8. QUOTAS	130
8.1. RESOURCE QUOTAS PER PROJECT	130
8.1.1. Resources managed by quotas	130

8.1.2. Quota scopes	132
8.1.3. Quota enforcement	132
8.1.4. Requests versus limits	133
8.1.5. Sample resource quota definitions	133
8.1.6. Creating a quota	136
8.1.6.1. Creating object count quotas	137
8.1.6.2. Setting resource quota for extended resources	138
8.1.7. Viewing a quota	140
8.1.8. Configuring explicit resource quotas	141
8.2. RESOURCE QUOTAS ACROSS MULTIPLE PROJECTS	143
8.2.1. Selecting multiple projects during quota creation	144
8.2.2. Viewing applicable cluster resource quotas	145
8.2.3. Selection granularity	146
CHAPTER 9. USING CONFIG MAPS WITH APPLICATIONS	147
9.1. UNDERSTANDING CONFIG MAPS	147
9.1.1. Config map restrictions	148
9.2. USE CASES: CONSUMING CONFIG MAPS IN PODS	148
9.2.1. Populating environment variables in containers by using config maps	148
9.2.2. Setting command-line arguments for container commands with config maps	150
9.2.3. Injecting content into a volume by using config maps	151
CHAPTER 10. MONITORING PROJECT AND APPLICATION METRICS USING THE DEVELOPER PERSPECTIVE	154
10.1. PREREQUISITES	154
10.2. MONITORING YOUR PROJECT METRICS	154
10.3. MONITORING YOUR APPLICATION METRICS	157
10.4. IMAGE VULNERABILITIES BREAKDOWN	158
10.5. MONITORING YOUR APPLICATION AND IMAGE VULNERABILITIES METRICS	158
10.6. ADDITIONAL RESOURCES	159
CHAPTER 11. MONITORING APPLICATION HEALTH BY USING HEALTH CHECKS	160
11.1. UNDERSTANDING HEALTH CHECKS	160
11.1.1. Example probes	161
11.2. CONFIGURING HEALTH CHECKS USING THE CLI	164
11.3. MONITORING APPLICATION HEALTH USING THE DEVELOPER PERSPECTIVE	167
11.4. EDITING HEALTH CHECKS USING THE DEVELOPER PERSPECTIVE	167
11.5. MONITORING HEALTH CHECK FAILURES USING THE DEVELOPER PERSPECTIVE	168
CHAPTER 12. EDITING APPLICATIONS	170
12.1. PREREQUISITES	170
12.2. EDITING THE SOURCE CODE OF AN APPLICATION USING THE DEVELOPER PERSPECTIVE	170
12.3. EDITING THE APPLICATION CONFIGURATION USING THE DEVELOPER PERSPECTIVE	170
CHAPTER 13. PRUNING OBJECTS TO RECLAIM RESOURCES	173
13.1. BASIC PRUNING OPERATIONS	173
13.2. PRUNING GROUPS	173
13.3. PRUNING DEPLOYMENT RESOURCES	174
13.4. PRUNING BUILDS	175
13.5. AUTOMATICALLY PRUNING IMAGES	176
13.6. MANUALLY PRUNING IMAGES	178
13.6.1. Considerations when pruning images	178
13.6.2. Limitations when pruning images	178
13.6.3. Image prune conditions	178

13.6.3.1. Pruning an image by age and tag	179
13.6.3.2. Pruning an image by size limit	179
13.6.4. Running image prune operations	180
13.6.5. Using secure or insecure connections	180
13.6.6. Image pruning CLI options	181
13.6.6.1. Additional information about the --prune-registry flag	182
13.6.7. Image pruning problems	183
13.6.7.1. Images not being pruned	183
13.6.7.2. Using a secure connection against insecure registry	183
13.6.7.3. Using an insecure connection against a secured registry	184
13.6.7.4. Using the wrong certificate authority	184
13.7. HARD PRUNING THE REGISTRY	184
13.8. PRUNING CRON JOBS	187
CHAPTER 14. IDLING APPLICATIONS	188
14.1. IDLING APPLICATIONS	188
14.1.1. Idling a single service	188
14.1.2. Idling multiple services	188
14.2. UNIDLING APPLICATIONS	188
CHAPTER 15. DELETING APPLICATIONS	190
15.1. DELETING APPLICATIONS USING THE DEVELOPER PERSPECTIVE	190
CHAPTER 16. USING THE RED HAT MARKETPLACE	191
16.1. RED HAT MARKETPLACE FEATURES	191
16.1.1. Connect OpenShift Container Platform clusters to the Marketplace	191
16.1.2. Install applications	191
16.1.3. Deploy applications from different perspectives	191

CHAPTER 1. BUILDING APPLICATIONS OVERVIEW

Using OpenShift Container Platform, you can create, edit, delete, and manage applications using the web console or command-line interface (CLI).

1.1. WORKING ON A PROJECT

Using projects, you can organize and manage applications in isolation. You can manage the entire project lifecycle, including [creating, viewing, and deleting a project](#) in OpenShift Container Platform.

After you create the project, you can [grant or revoke access to a project](#) and [manage cluster roles](#) for the users using the Developer perspective. You can also [edit the project configuration resource](#) while creating a project template that is used for automatic provisioning of new projects.

Using the CLI, you can [create a project as a different user](#) by impersonating a request to the OpenShift Container Platform API. When you make a request to create a new project, the OpenShift Container Platform uses an endpoint to provision the project according to a customizable template. As a cluster administrator, you can choose to [prevent an authenticated user group from self-provisioning new projects](#).

1.2. WORKING ON AN APPLICATION

1.2.1. Creating an application

To create applications, you must have created a project or have access to a project with the appropriate roles and permissions. You can create an application by using either [the Developer perspective in the web console](#), [installed Operators](#), or [the OpenShift CLI \(oc\)](#). You can source the applications to be added to the project from Git, JAR files, devfiles, or the developer catalog.

You can also use components that include source or binary code, images, and templates to create an application by using the OpenShift CLI (**oc**). With the OpenShift Container Platform web console, you can create an application from an Operator installed by a cluster administrator.

1.2.2. Maintaining an application

After you create the application, you can use the web console to [monitor your project or application metrics](#). You can also [edit](#) or [delete](#) the application using the web console.

When the application is running, not all applications resources are used. As a cluster administrator, you can choose to [idle these scalable resources](#) to reduce resource consumption.

1.2.3. Deploying an application

You can deploy your application using [Deployment](#) or [DeploymentConfig](#) objects and [manage](#) them from the web console. You can create [deployment strategies](#) that help reduce downtime during a change or an upgrade to the application.

You can also use [Helm](#), a software package manager that simplifies deployment of applications and services to OpenShift Container Platform clusters.

CHAPTER 2. PROJECTS

2.1. WORKING WITH PROJECTS

A *project* allows a community of users to organize and manage their content in isolation from other communities.



NOTE

Projects starting with **openshift-** and **kube-** are [default projects](#). These projects host cluster components that run as pods and other infrastructure components. As such, OpenShift Container Platform does not allow you to create projects starting with **openshift-** or **kube-** using the **oc new-project** command. Cluster administrators can create these projects using the **oc adm new-project** command.



IMPORTANT

Do not run workloads in or share access to default projects. Default projects are reserved for running core cluster components.

The following default projects are considered highly privileged: **default**, **kube-public**, **kube-system**, **openshift**, **openshift-infra**, **openshift-node**, and other system-created projects that have the **openshift.io/run-level** label set to **0** or **1**. Functionality that relies on admission plugins, such as pod security admission, security context constraints, cluster resource quotas, and image reference resolution, does not work in highly privileged projects.

2.1.1. Creating a project

You can use the OpenShift Container Platform web console or the OpenShift CLI (**oc**) to create a project in your cluster.

2.1.1.1. Creating a project by using the web console

You can use the OpenShift Container Platform web console to create a project in your cluster.



NOTE

Projects starting with **openshift-** and **kube-** are considered critical by OpenShift Container Platform. As such, OpenShift Container Platform does not allow you to create projects starting with **openshift-** using the web console.

Prerequisites


- Ensure that you have the appropriate roles and permissions to create projects, applications, and other workloads in OpenShift Container Platform.

Procedure

- If you are using the **Administrator** perspective:
 - a. Navigate to **Home → Projects**.

- b. Click **Create Project**:
 - i. In the **Create Project** dialog box, enter a unique name, such as **myproject**, in the **Name** field.
 - ii. Optional: Add the **Display name** and **Description** details for the project.
 - iii. Click **Create**.
The dashboard for your project is displayed.
 - c. Optional: Select the **Details** tab to view the project details.
 - d. Optional: If you have adequate permissions for a project, you can use the **Project Access** tab to provide or revoke admin, edit, and view privileges for the project.
- If you are using the **Developer** perspective:
 - a. Click the **Project** menu and select **Create Project**:

Figure 2.1. Create project



- i. In the **Create Project** dialog box, enter a unique name, such as **myproject**, in the **Name** field.
 - ii. Optional: Add the **Display name** and **Description** details for the project.
 - iii. Click **Create**.
- b. Optional: Use the left navigation panel to navigate to the **Project** view and see the dashboard for your project.
 - c. Optional: In the project dashboard, select the **Details** tab to view the project details.
 - d. Optional: If you have adequate permissions for a project, you can use the **Project Access** tab of the project dashboard to provide or revoke admin, edit, and view privileges for the project.

Additional resources

- [Customizing the available cluster roles using the web console](#)

2.1.1.2. Creating a project by using the CLI

If allowed by your cluster administrator, you can create a new project.



NOTE

Projects starting with **openshift-** and **kube-** are considered critical by OpenShift Container Platform. As such, OpenShift Container Platform does not allow you to create Projects starting with **openshift-** or **kube-** using the **oc new-project** command. Cluster administrators can create these projects using the **oc adm new-project** command.

Procedure

- Run:

```
$ oc new-project <project_name> \
  --description="<description>" --display-name="<display_name>"
```

For example:

```
$ oc new-project hello-openshift \
  --description="This is an example project" \
  --display-name="Hello OpenShift"
```



NOTE

The number of projects you are allowed to create might be limited by the system administrator. After your limit is reached, you might have to delete an existing project in order to create a new one.

2.1.2. Viewing a project

You can use the OpenShift Container Platform web console or the OpenShift CLI (**oc**) to view a project in your cluster.

2.1.2.1. Viewing a project by using the web console

You can view the projects that you have access to by using the OpenShift Container Platform web console.



IMPORTANT

Starting with OpenShift Container Platform 4.19, the perspectives in the web console have unified. The **Developer** perspective is no longer enabled by default.

All users can interact with all OpenShift Container Platform web console features. However, if you are not the cluster owner, you might need to request permission to access certain features from the cluster owner.

You can still enable the **Developer** perspective. On the **Getting Started** pane in the web console, you can take a tour of the console, find information on setting up your cluster, view a quick start for enabling the **Developer** perspective, and follow links to explore new features and capabilities.

Procedure

- If you are logged in as an administrator:
 - a. Navigate to **Home** → **Projects** in the navigation menu.
 - b. Select a project to view. The **Overview** tab includes a dashboard for your project.
 - c. Select the **Details** tab to view the project details.
 - d. Select the **YAML** tab to view and update the YAML configuration for the project resource.
 - e. Select the **Workloads** tab to see workloads in the project.
 - f. Select the **RoleBindings** tab to view and create role bindings for your project.
- If you are logged in as a developer:
 - a. Navigate to the **Project** page in the navigation menu.
 - b. Select **All Projects** from the **Project** drop-down menu at the top of the screen to list all of the projects in your cluster.
 - c. Select a project to view.
 - d. Select the **Details** tab to view the project details.
 - e. If you have adequate permissions for a project, select the **Project access** tab view and update the privileges for the project.

2.1.2.2. Viewing a project using the CLI

When viewing projects, you are restricted to seeing only the projects you have access to view based on the authorization policy.

Procedure

1. To view a list of projects, run:

```
$ oc get projects
```

2. You can change from the current project to a different project for CLI operations. The specified project is then used in all subsequent operations that manipulate project-scoped content:

```
$ oc project <project_name>
```

2.1.3. Providing access permissions to your project using the Developer perspective

You can use the **Project** view in the **Developer** perspective to grant or revoke access permissions to your project.

Prerequisites

- You have created a project.

Procedure

To add users to your project and provide **Admin**, **Edit**, or **View** access to them:

1. In the **Developer** perspective, navigate to the **Project** page.
2. Select your project from the **Project** menu.
3. Select the **Project Access** tab.
4. Click **Add access** to add a new row of permissions to the default ones.

Figure 2.2. Project permissions

The screenshot shows the 'Project Access' page in the Developer perspective. The left sidebar has a 'Project' menu item highlighted. The main content area shows the 'Project Access' tab for project 'tw'. It lists existing users 'kube:admin' (Admin role) and 'pipeline' (Edit role). A new row is being added with 'Name' and 'Select a role' dropdown. An 'Add Access' button is visible. A message at the bottom says 'You made changes to this page. Click Save to save changes or Reload to cancel changes.' with 'Save' and 'Reload' buttons.

5. Enter the user name, click the **Select a role** drop-down list, and select an appropriate role.
6. Click **Save** to add the new permissions.

You can also use:

- The **Select a role** drop-down list, to modify the access permissions of an existing user.
- The **Remove Access** icon, to completely remove the access permissions of an existing user to the project.



NOTE

Advanced role-based access control is managed in the **Roles** and **Roles Binding** views in the **Administrator** perspective.

2.1.4. Customizing the available cluster roles using the web console

In the **Developer** perspective of the web console, the **Project → Project access** page enables a project administrator to grant roles to users in a project. By default, the available cluster roles that can be granted to users in a project are **admin**, **edit**, and **view**.

As a cluster administrator, you can define which cluster roles are available in the **Project access** page for all projects cluster-wide. You can specify the available roles by customizing the **spec.customization.projectAccess.availableClusterRoles** object in the **Console** configuration resource.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. In the **Administrator** perspective, navigate to **Administration → Cluster settings**.
2. Click the **Configuration** tab.
3. From the **Configuration resource** list, select **Console operator.openshift.io**.
4. Navigate to the **YAML** tab to view and edit the YAML code.
5. In the YAML code under **spec**, customize the list of available cluster roles for project access. The following example specifies the default **admin**, **edit**, and **view** roles:

```
apiVersion: operator.openshift.io/v1
kind: Console
metadata:
  name: cluster
# ...
spec:
  customization:
    projectAccess:
      availableClusterRoles:
        - admin
        - edit
        - view
```

6. Click **Save** to save the changes to the **Console** configuration resource.

Verification

1. In the **Developer** perspective, navigate to the **Project** page.
2. Select a project from the **Project** menu.
3. Select the **Project access** tab.
4. Click the menu in the **Role** column and verify that the available roles match the configuration that you applied to the **Console** resource configuration.

2.1.5. Adding to a project

You can add items to your project by using the **+Add** page.

Prerequisites

- You have created a project.

Procedure

1. Navigate to the **+Add** page.
2. Select your project from the **Project** menu.
3. Click on an item on the **+Add** page and then follow the workflow.



NOTE

You can also use the search feature in the **Add*** page to find additional items to add to your project. Click ***** under **Add** at the top of the page and type the name of a component in the search field.

2.1.6. Checking the project status

You can use the OpenShift Container Platform web console or the OpenShift CLI (**oc**) to view the status of your project.

2.1.6.1. Checking project status by using the web console

You can review the status of your project by using the web console.

Prerequisites

- You have created a project.

Procedure

1. Navigate to **Home → Projects**.
2. Select a project from the list.
3. Review the project status in the **Overview** page.

2.1.6.2. Checking project status by using the CLI

You can review the status of your project by using the OpenShift CLI (**oc**).

Prerequisites

- You have installed the OpenShift CLI (**oc**).
- You have created a project.

Procedure

1. Switch to your project:

```
$ oc project <project_name> 1
```

- 1 Replace **<project_name>** with the name of your project.

2. Obtain a high-level overview of the project:

```
$ oc status
```

2.1.7. Deleting a project

You can use the OpenShift Container Platform web console or the OpenShift CLI (**oc**) to delete a project.

When you delete a project, the server updates the project status to **Terminating** from **Active**. Then, the server clears all content from a project that is in the **Terminating** state before finally removing the project. While a project is in **Terminating** status, you cannot add new content to the project. Projects can be deleted from the CLI or the web console.

2.1.7.1. Deleting a project by using the web console

You can delete a project by using the web console.

Prerequisites

- You have created a project.
- You have the required permissions to delete the project.

Procedure

- If you are using the **Administrator** perspective:
 - a. Navigate to **Home → Projects**.
 - b. Select a project from the list.
 - c. Click the **Actions** drop-down menu for the project and select **Delete Project**.



NOTE

The **Delete Project** option is not available if you do not have the required permissions to delete the project.

1. In the **Delete Project?** pane, confirm the deletion by entering the name of your project.
 2. Click **Delete**.
- If you are using the **Developer** perspective:
 - a. Navigate to the **Project** page.
 - b. Select the project that you want to delete from the **Project** menu.
 - c. Click the **Actions** drop-down menu for the project and select **Delete Project**.

**NOTE**

If you do not have the required permissions to delete the project, the **Delete Project** option is not available.

1. In the **Delete Project?** pane, confirm the deletion by entering the name of your project.
2. Click **Delete**.

2.1.7.2. Deleting a project by using the CLI

You can delete a project by using the OpenShift CLI (**oc**).

Prerequisites

- You have installed the OpenShift CLI (**oc**).
- You have created a project.
- You have the required permissions to delete the project.

Procedure

1. Delete your project:

```
$ oc delete project <project_name> 1
```

- 1** Replace **<project_name>** with the name of the project that you want to delete.

2.2. CREATING A PROJECT AS ANOTHER USER

Impersonation allows you to create a project as a different user.

2.2.1. API impersonation

You can configure a request to the OpenShift Container Platform API to act as though it originated from another user. For more information, see [User impersonation](#) in the Kubernetes documentation.

2.2.2. Impersonating a user when you create a project

You can impersonate a different user when you create a project request. Because **system:authenticated:oauth** is the only bootstrap group that can create project requests, you must impersonate that group.

Procedure

- To create a project request on behalf of a different user:

```
$ oc new-project <project> --as=<user> \
  --as-group=system:authenticated --as-group=system:authenticated:oauth
```

2.3. CONFIGURING PROJECT CREATION

In OpenShift Container Platform, *projects* are used to group and isolate related objects. When a request is made to create a new project using the web console or **oc new-project** command, an endpoint in OpenShift Container Platform is used to provision the project according to a template, which can be customized.

As a cluster administrator, you can allow and configure how developers and service accounts can create, or *self-provision*, their own projects.

2.3.1. About project creation

The OpenShift Container Platform API server automatically provisions new projects based on the project template that is identified by the **projectRequestTemplate** parameter in the cluster's project configuration resource. If the parameter is not defined, the API server creates a default template that creates a project with the requested name, and assigns the requesting user to the **admin** role for that project.

When a project request is submitted, the API substitutes the following parameters into the template:

Table 2.1. Default project template parameters

Parameter	Description
PROJECT_NAME	The name of the project. Required.
PROJECT_DISPLAYNAME	The display name of the project. May be empty.
PROJECT_DESCRIPTION	The description of the project. May be empty.
PROJECT_ADMIN_USER	The user name of the administrating user.
PROJECT_REQUESTING_USER	The user name of the requesting user.

Access to the API is granted to developers with the **self-provisioner** role and the **self-provisioners** cluster role binding. This role is available to all authenticated developers by default.

2.3.2. Modifying the template for new projects

As a cluster administrator, you can modify the default project template so that new projects are created using your custom requirements.

To create your own custom project template:

Prerequisites

- You have access to an OpenShift Container Platform cluster using an account with **cluster-admin** permissions.

Procedure

1. Log in as a user with **cluster-admin** privileges.

2. Generate the default project template:

```
$ oc adm create-bootstrap-project-template -o yaml > template.yaml
```

3. Use a text editor to modify the generated **template.yaml** file by adding objects or modifying existing objects.

4. The project template must be created in the **openshift-config** namespace. Load your modified template:

```
$ oc create -f template.yaml -n openshift-config
```

5. Edit the project configuration resource using the web console or CLI.

- Using the web console:
 - i. Navigate to the **Administration → Cluster Settings** page.
 - ii. Click **Configuration** to view all configuration resources.
 - iii. Find the entry for **Project** and click **Edit YAML**.
- Using the CLI:
 - i. Edit the **project.config.openshift.io/cluster** resource:

```
$ oc edit project.config.openshift.io/cluster
```

6. Update the **spec** section to include the **projectRequestTemplate** and **name** parameters, and set the name of your uploaded project template. The default name is **project-request**.

Project configuration resource with custom project template

```
apiVersion: config.openshift.io/v1
kind: Project
metadata:
  # ...
spec:
  projectRequestTemplate:
    name: <template_name>
  # ...
```

7. After you save your changes, create a new project to verify that your changes were successfully applied.

2.3.3. Disabling project self-provisioning

You can prevent an authenticated user group from self-provisioning new projects.

Procedure

1. Log in as a user with **cluster-admin** privileges.

- View the **self-provisioners** cluster role binding usage by running the following command:

```
$ oc describe clusterrolebinding.rbac self-provisioners
```

Example output

```
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
```

Review the subjects in the **self-provisioners** section.

- Remove the **self-provisioner** cluster role from the group **system:authenticated:oauth**.

- If the **self-provisioners** cluster role binding binds only the **self-provisioner** role to the **system:authenticated:oauth** group, run the following command:

```
$ oc patch clusterrolebinding.rbac self-provisioners -p '{"subjects": null}'
```

- If the **self-provisioners** cluster role binding binds the **self-provisioner** role to more users, groups, or service accounts than the **system:authenticated:oauth** group, run the following command:

```
$ oc adm policy \
  remove-cluster-role-from-group self-provisioner \
  system:authenticated:oauth
```

- Edit the **self-provisioners** cluster role binding to prevent automatic updates to the role. Automatic updates reset the cluster roles to the default state.

- To update the role binding using the CLI:

- Run the following command:

```
$ oc edit clusterrolebinding.rbac self-provisioners
```

- In the displayed role binding, set the **rbac.authorization.kubernetes.io/autoupdate** parameter value to **false**, as shown in the following example:

```
apiVersion: authorization.openshift.io/v1
kind: ClusterRoleBinding
metadata:
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "false"
# ...
```


- To update the role binding by using a single command:

```
$ oc patch clusterrolebinding.rbac self-provisioners -p '{ "metadata": { "annotations": {  
  "rbac.authorization.kubernetes.io/autoupdate": "false" } } }'
```

5. Log in as an authenticated user and verify that it can no longer self-provision a project:

```
$ oc new-project test
```

Example output

```
Error from server (Forbidden): You may not request a new project via this API.
```

Consider customizing this project request message to provide more helpful instructions specific to your organization.

2.3.4. Customizing the project request message

When a developer or a service account that is unable to self-provision projects makes a project creation request using the web console or CLI, the following error message is returned by default:

```
You may not request a new project via this API.
```

Cluster administrators can customize this message. Consider updating it to provide further instructions on how to request a new project specific to your organization. For example:

- To request a project, contact your system administrator at **projectname@example.com**.
- To request a new project, fill out the project request form located at **<https://internal.example.com/openshift-project-request>**.

To customize the project request message:

Procedure

1. Edit the project configuration resource using the web console or CLI.
 - Using the web console:
 - i. Navigate to the **Administration → Cluster Settings** page.
 - ii. Click **Configuration** to view all configuration resources.
 - iii. Find the entry for **Project** and click **Edit YAML**.
 - Using the CLI:
 - i. Log in as a user with **cluster-admin** privileges.
 - ii. Edit the **project.config.openshift.io/cluster** resource:

```
$ oc edit project.config.openshift.io/cluster
```

2. Update the **spec** section to include the **projectRequestMessage** parameter and set the value to your custom message:

Project configuration resource with custom project request message

```
apiVersion: config.openshift.io/v1
kind: Project
metadata:
  # ...
spec:
  projectRequestMessage: <message_string>
  # ...
```

For example:

```
apiVersion: config.openshift.io/v1
kind: Project
metadata:
  # ...
spec:
  projectRequestMessage: To request a project, contact your system administrator at
projectname@example.com.
  # ...
```

3. After you save your changes, attempt to create a new project as a developer or service account that is unable to self-provision projects to verify that your changes were successfully applied.

CHAPTER 3. CREATING APPLICATIONS

3.1. USING TEMPLATES

The following sections provide an overview of templates, as well as how to use and create them.

3.1.1. Understanding templates

A template describes a set of objects that can be parameterized and processed to produce a list of objects for creation by OpenShift Container Platform. A template can be processed to create anything you have permission to create within a project, for example services, build configurations, and deployment configurations. A template can also define a set of labels to apply to every object defined in the template.

You can create a list of objects from a template using the CLI or, if a template has been uploaded to your project or the global template library, using the web console.

3.1.2. Uploading a template

If you have a JSON or YAML file that defines a template, you can upload the template to projects using the CLI. This saves the template to the project for repeated use by any user with appropriate access to that project. Instructions about writing your own templates are provided later in this topic.

Procedure

- Upload a template using one of the following methods:
 - Upload a template to your current project's template library, pass the JSON or YAML file with the following command:

```
$ oc create -f <filename>
```

- Upload a template to a different project using the **-n** option with the name of the project:

```
$ oc create -f <filename> -n <project>
```

The template is now available for selection using the web console or the CLI.

3.1.3. Creating an application by using the web console

You can use the web console to create an application from a template.

Procedure

1. Navigate to your project and click **+Add**
2. Click **All services** in the **Developer Catalog** tile.
3. Click **Builder Images** under **Type** to see the available builder images.

**NOTE**

Only image stream tags that have the **builder** tag listed in their annotations appear in this list, as demonstrated here:

```
kind: "ImageStream"
apiVersion: "image.openshift.io/v1"
metadata:
  name: "ruby"
  creationTimestamp: null
spec:
  # ...
  tags:
    - name: "2.6"
      annotations:
        description: "Build and run Ruby 2.6 applications"
        iconClass: "icon-ruby"
        tags: "builder,ruby" 1
        supports: "ruby:2.6,ruby"
        version: "2.6"
  # ...
```

- 1** Including **builder** here ensures this image stream tag appears in the web console as a builder.

4. Modify the settings in the new application screen to configure the objects to support your application.

3.1.4. Creating objects from templates by using the CLI

You can use the CLI to process templates and use the configuration that is generated to create objects.

3.1.4.1. Adding labels

Labels are used to manage and organize generated objects, such as pods. The labels specified in the template are applied to every object that is generated from the template.

Procedure

- Add labels in the template from the command line:

```
$ oc process -f <filename> -l name=otherLabel
```

3.1.4.2. Listing parameters

The list of parameters that you can override are listed in the **parameters** section of the template.

Procedure

1. You can list parameters with the CLI by using the following command and specifying the file to be used:

```
$ oc process --parameters -f <filename>
```

Alternatively, if the template is already uploaded:

```
$ oc process --parameters -n <project> <template_name>
```

For example, the following shows the output when listing the parameters for one of the quick start templates in the default **openshift** project:

```
$ oc process --parameters -n openshift rails-postgresql-example
```

Example output

NAME	DESCRIPTION
GENERATOR	VALUE
SOURCE_REPOSITORY_URL	The URL of the repository with your application source code https://github.com/sclorg/rails-ex.git
SOURCE_REPOSITORY_REF	Set this to a branch name, tag or other ref of your repository if you are not using the default branch
CONTEXT_DIR	Set this to the relative path to your project if it is not in the root of your repository
APPLICATION_DOMAIN	The exposed hostname that will route to the Rails service rails-postgresql-example.openshiftapps.com
GITHUB_WEBHOOK_SECRET	A secret string used to configure the GitHub webhook expression [a-zA-Z0-9]{40}
SECRET_KEY_BASE	Your secret key for verifying the integrity of signed cookies expression [a-z0-9]{127}
APPLICATION_USER	The application user that is used within the sample application to authorize access on pages openshift
APPLICATION_PASSWORD	The application password that is used within the sample application to authorize access on pages secret
DATABASE_SERVICE_NAME	Database service name postgresql
POSTGRESQL_USER	database username expression user[A-Z0-9]{3}
POSTGRESQL_PASSWORD	database password expression [a-zA-Z0-9]{8}
POSTGRESQL_DATABASE	database name root
POSTGRESQL_MAX_CONNECTIONS	database max connections 10
POSTGRESQL_SHARED_BUFFERS	database shared buffers 12MB

The output identifies several parameters that are generated with a regular expression-like generator when the template is processed.

3.1.4.3. Generating a list of objects

Using the CLI, you can process a file defining a template to return the list of objects to standard output.

Procedure

1. Process a file defining a template to return the list of objects to standard output:

```
$ oc process -f <filename>
```

Alternatively, if the template has already been uploaded to the current project:

```
$ oc process <template_name>
```

2. Create objects from a template by processing the template and piping the output to **oc create**:

```
$ oc process -f <filename> | oc create -f -
```

Alternatively, if the template has already been uploaded to the current project:

```
$ oc process <template> | oc create -f -
```

3. You can override any parameter values defined in the file by adding the **-p** option for each **<name>=<value>** pair you want to override. A parameter reference appears in any text field inside the template items.

For example, in the following the **POSTGRESQL_USER** and **POSTGRESQL_DATABASE** parameters of a template are overridden to output a configuration with customized environment variables:

- a. Creating a List of objects from a template

```
$ oc process -f my-rails-postgresql \
  -p POSTGRESQL_USER=bob \
  -p POSTGRESQL_DATABASE=mydatabase
```

- b. The JSON file can either be redirected to a file or applied directly without uploading the template by piping the processed output to the **oc create** command:

```
$ oc process -f my-rails-postgresql \
  -p POSTGRESQL_USER=bob \
  -p POSTGRESQL_DATABASE=mydatabase \
  | oc create -f -
```

- c. If you have large number of parameters, you can store them in a file and then pass this file to **oc process**:

```
$ cat postgres.env
POSTGRESQL_USER=bob
POSTGRESQL_DATABASE=mydatabase
```

```
$ oc process -f my-rails-postgresql --param-file=postgres.env
```

- d. You can also read the environment from standard input by using **"-"** as the argument to **--param-file**:

```
$ sed s/bob/alice/ postgres.env | oc process -f my-rails-postgresql --param-file=-
```

3.1.5. Modifying uploaded templates

You can edit a template that has already been uploaded to your project.

Procedure

- Modify a template that has already been uploaded:

```
$ oc edit template <template>
```

3.1.6. Using instant app and quick start templates

OpenShift Container Platform provides a number of default instant app and quick start templates to make it easy to quickly get started creating a new application for different languages. Templates are provided for Rails (Ruby), Django (Python), Node.js, CakePHP (PHP), and Dancer (Perl). Your cluster administrator must create these templates in the default, global **openshift** project so you have access to them.

By default, the templates build using a public source repository on GitHub that contains the necessary application code.

Procedure

1. You can list the available default instant app and quick start templates with:

```
$ oc get templates -n openshift
```

2. To modify the source and build your own version of the application:
 - a. Fork the repository referenced by the template's default **SOURCE_REPOSITORY_URL** parameter.
 - b. Override the value of the **SOURCE_REPOSITORY_URL** parameter when creating from the template, specifying your fork instead of the default value.
By doing this, the build configuration created by the template now points to your fork of the application code, and you can modify the code and rebuild the application at will.



NOTE

Some of the instant app and quick start templates define a database deployment configuration. The configuration they define uses ephemeral storage for the database content. These templates should be used for demonstration purposes only as all database data is lost if the database pod restarts for any reason.

3.1.6.1. Quick start templates

A quick start template is a basic example of an application running on OpenShift Container Platform. Quick starts come in a variety of languages and frameworks, and are defined in a template, which is constructed from a set of services, build configurations, and deployment configurations. This template references the necessary images and source repositories to build and deploy the application.

To explore a quick start, create an application from a template. Your administrator must have already installed these templates in your OpenShift Container Platform cluster, in which case you can simply select it from the web console.

Quick starts refer to a source repository that contains the application source code. To customize the quick start, fork the repository and, when creating an application from the template, substitute the default source repository name with your forked repository. This results in builds that are performed using your source code instead of the provided example source. You can then update the code in your source repository and launch a new build to see the changes reflected in the deployed application.

3.1.6.1.1. Web framework quick start templates

These quick start templates provide a basic application of the indicated framework and language:

- CakePHP: a PHP web framework that includes a MySQL database
- Dancer: a Perl web framework that includes a MySQL database
- Django: a Python web framework that includes a PostgreSQL database
- NodeJS: a NodeJS web application that includes a MongoDB database
- Rails: a Ruby web framework that includes a PostgreSQL database

3.1.7. Writing templates

You can define new templates to make it easy to recreate all the objects of your application. The template defines the objects it creates along with some metadata to guide the creation of those objects.

The following is an example of a simple template object definition (YAML):

```
apiVersion: template.openshift.io/v1
kind: Template
metadata:
  name: redis-template
  annotations:
    description: "Description"
    iconClass: "icon-redis"
    tags: "database,nosql"
objects:
- apiVersion: v1
  kind: Pod
  metadata:
    name: redis-master
  spec:
    containers:
    - env:
      - name: REDIS_PASSWORD
        value: ${REDIS_PASSWORD}
      image: dockerfile/redis
      name: master
      ports:
      - containerPort: 6379
        protocol: TCP
  parameters:
  - description: Password used for Redis authentication
    from: '[A-Z0-9]{8}'
    generate: expression
```



```
name: REDIS_PASSWORD
labels:
  redis: master
```

3.1.7.1. Writing the template description

The template description informs you what the template does and helps you find it when searching in the web console. Additional metadata beyond the template name is optional, but useful to have. In addition to general descriptive information, the metadata also includes a set of tags. Useful tags include the name of the language the template is related to for example, Java, PHP, Ruby, and so on.

The following is an example of template description metadata:

```
kind: Template
apiVersion: template.openshift.io/v1
metadata:
  name: cakephp-mysql-example ❶
  annotations:
    openshift.io/display-name: "CakePHP MySQL Example (Ephemeral)" ❷
  description: >-
    An example CakePHP application with a MySQL database. For more information
    about using this template, including OpenShift considerations, see
    https://github.com/sclorg/cakephp-ex/blob/master/README.md.

    WARNING: Any data stored will be lost upon pod destruction. Only use this
    template for testing." ❸
  openshift.io/long-description: >-
    This template defines resources needed to develop a CakePHP application,
    including a build configuration, application DeploymentConfig, and
    database DeploymentConfig. The database is stored in
    non-persistent storage, so this configuration should be used for
    experimental purposes only. ❹
  tags: "quickstart,php,cakephp" ❺
  iconClass: icon-php ❻
  openshift.io/provider-display-name: "Red Hat, Inc." ❼
  openshift.io/documentation-url: "https://github.com/sclorg/cakephp-ex" ❽
  openshift.io/support-url: "https://access.redhat.com" ❾
  message: "Your admin credentials are ${ADMIN_USERNAME}:${ADMIN_PASSWORD}" ❿
```

- ❶ The unique name of the template.
- ❷ A brief, user-friendly name, which can be employed by user interfaces.
- ❸ A description of the template. Include enough detail that users understand what is being deployed and any caveats they must know before deploying. It should also provide links to additional information, such as a README file. Newlines can be included to create paragraphs.
- ❹ Additional template description. This may be displayed by the service catalog, for example.
- ❺ Tags to be associated with the template for searching and grouping. Add tags that include it into one of the provided catalog categories. Refer to the **id** and **categoryAliases** in **CATALOG_CATEGORIES** in the console constants file. The categories can also be customized for the whole cluster.

- 6** An icon to be displayed with your template in the web console.

Example 3.1. Available icons

- **icon-3scale**
- **icon-aerogear**
- **icon-amq**
- **icon-angularjs**
- **icon-ansible**
- **icon-apache**
- **icon-beaker**
- **icon-camel**
- **icon-capedwarf**
- **icon-cassandra**
- **icon-catalog-icon**
- **icon-clojure**
- **icon-codeigniter**
- **icon-cordova**
- **icon-datagrid**
- **icon-datavirt**
- **icon-debian**
- **icon-decisionserver**
- **icon-django**
- **icon-dotnet**
- **icon-drupal**
- **icon-eap**
- **icon-elastic**
- **icon-erlang**
- **icon-fedora**
- **icon-freebsd**

- **icon-git**
- **icon-github**
- **icon-gitlab**
- **icon-glassfish**
- **icon-go-gopher**
- **icon-golang**
- **icon-grails**
- **icon-hadoop**
- **icon-haproxy**
- **icon-helm**
- **icon-infinispan**
- **icon-jboss**
- **icon-jenkins**
- **icon-jetty**
- **icon-joomla**
- **icon-jruby**
- **icon-js**
- **icon-knative**
- **icon-kubevirt**
- **icon-laravel**
- **icon-load-balancer**
- **icon-mariadb**
- **icon-mediawiki**
- **icon-memcached**
- **icon-mongodb**
- **icon-mssql**
- **icon-mysql-database**
- **icon-nginx**

- **icon-nodejs**
- **icon-openjdk**
- **icon-openliberty**
- **icon-openshift**
- **icon-openstack**
- **icon-other-linux**
- **icon-other-unknown**
- **icon-perl**
- **icon-phalcon**
- **icon-php**
- **icon-play**
- **iconpostgresql**
- **icon-processserver**
- **icon-python**
- **icon-quarkus**
- **icon-rabbitmq**
- **icon-rails**
- **icon-redhat**
- **icon-redis**
- **icon-rh-integration**
- **icon-rh-spring-boot**
- **icon-rh-tomcat**
- **icon-ruby**
- **icon-scala**
- **icon-serverlessfx**
- **icon-shadowman**
- **icon-spring-boot**
- **icon-spring**

- **icon-sso**
- **icon-stackoverflow**
- **icon-suse**
- **icon-symfony**
- **icon-tomcat**
- **icon-ubuntu**
- **icon-vertx**
- **icon-wildfly**
- **icon-windows**
- **icon-wordpress**
- **icon-xamarin**
- **icon-zend**

- 7 The name of the person or organization providing the template.
- 8 A URL referencing further documentation for the template.
- 9 A URL where support can be obtained for the template.
- 10 An instructional message that is displayed when this template is instantiated. This field should inform the user how to use the newly created resources. Parameter substitution is performed on the message before being displayed so that generated credentials and other parameters can be included in the output. Include links to any next-steps documentation that users should follow.

3.1.7.2. Writing template labels

Templates can include a set of labels. These labels are added to each object created when the template is instantiated. Defining a label in this way makes it easy for users to find and manage all the objects created from a particular template.

The following is an example of template object labels:

```
kind: "Template"
apiVersion: "v1"
...
labels:
  template: "cakephp-mysql-example" 1
  app: "${NAME}" 2
```

- 1 A label that is applied to all objects created from this template.
- 2 A parameterized label that is also applied to all objects created from this template. Parameter expansion is carried out on both label keys and values.

expansion is carried out on both keys and values.

3.1.7.3. Writing template parameters

Parameters allow a value to be supplied by you or generated when the template is instantiated. Then, that value is substituted wherever the parameter is referenced. References can be defined in any field in the objects list field. This is useful for generating random passwords or allowing you to supply a hostname or other user-specific value that is required to customize the template. Parameters can be referenced in two ways:

- As a string value by placing values in the form **`${PARAMETER_NAME}`** in any string field in the template.
- As a JSON or YAML value by placing values in the form **`${{PARAMETER_NAME}}`** in place of any field in the template.

When using the **`${PARAMETER_NAME}`** syntax, multiple parameter references can be combined in a single field and the reference can be embedded within fixed data, such as **`"http://${PARAMETER_1}${PARAMETER_2}"`**. Both parameter values are substituted and the resulting value is a quoted string.

When using the **`${{PARAMETER_NAME}}`** syntax only a single parameter reference is allowed and leading and trailing characters are not permitted. The resulting value is unquoted unless, after substitution is performed, the result is not a valid JSON object. If the result is not a valid JSON value, the resulting value is quoted and treated as a standard string.

A single parameter can be referenced multiple times within a template and it can be referenced using both substitution syntaxes within a single template.

A default value can be provided, which is used if you do not supply a different value:

The following is an example of setting an explicit value as the default value:

```
parameters:
- name: USERNAME
  description: "The user name for Joe"
  value: joe
```

Parameter values can also be generated based on rules specified in the parameter definition, for example generating a parameter value:

```
parameters:
- name: PASSWORD
  description: "The random user password"
  generate: expression
  from: "[a-zA-Z0-9]{12}"
```

In the previous example, processing generates a random password 12 characters long consisting of all upper and lowercase alphabet letters and numbers.

The syntax available is not a full regular expression syntax. However, you can use **`\w`**, **`\d`**, **`\a`**, and **`\A`** modifiers:

- **`[w]{10}`** produces 10 alphabet characters, numbers, and underscores. This follows the PCRE standard and is equal to **`[a-zA-Z0-9_]{10}`**.

- `[d]{10}` produces 10 numbers. This is equal to `[0-9]{10}`.
- `[a]{10}` produces 10 alphabetical characters. This is equal to `[a-zA-Z]{10}`.
- `[A]{10}` produces 10 punctuation or symbol characters. This is equal to `[~!@#$$%^&*()\- _+={}\[\] \<, > . ? / ' " ; : ' ] { 10 }`.

NOTE

Depending on if the template is written in YAML or JSON, and the type of string that the modifier is embedded within, you might need to escape the backslash with a second backslash. The following examples are equivalent:

Example YAML template with a modifier

```
parameters:
- name: singlequoted_example
  generate: expression
  from: '[A]{10}'
- name: doublequoted_example
  generate: expression
  from: "[\A]{10}"
```

Example JSON template with a modifier

```
{
  "parameters": [
    {
      "name": "json_example",
      "generate": "expression",
      "from": "[\A]{10}"
    }
  ]
}
```

Here is an example of a full template with parameter definitions and references:

```
kind: Template
apiVersion: template.openshift.io/v1
metadata:
  name: my-template
objects:
- kind: BuildConfig
  apiVersion: build.openshift.io/v1
  metadata:
    name: cakephp-mysql-example
  annotations:
    description: Defines how to build the application
  spec:
    source:
      type: Git
      git:
        uri: "${SOURCE_REPOSITORY_URL}" 1
        ref: "${SOURCE_REPOSITORY_REF}"
```

```

    contextDir: "${CONTEXT_DIR}"
  - kind: DeploymentConfig
    apiVersion: apps.openshift.io/v1
    metadata:
      name: frontend
    spec:
      replicas: "${REPLICA_COUNT}" ❷
  parameters:
  - name: SOURCE_REPOSITORY_URL ❸
    displayName: Source Repository URL ❹
    description: The URL of the repository with your application source code ❺
    value: https://github.com/sclorg/cakephp-ex.git ❻
    required: true ❼
  - name: GITHUB_WEBHOOK_SECRET
    description: A secret string used to configure the GitHub webhook
    generate: expression ❽
    from: "[a-zA-Z0-9]{40}" ❾
  - name: REPLICA_COUNT
    description: Number of replicas to run
    value: "2"
    required: true
  message: "... The GitHub webhook secret is ${GITHUB_WEBHOOK_SECRET} ..." ❿

```

- ❶ This value is replaced with the value of the **SOURCE_REPOSITORY_URL** parameter when the template is instantiated.
- ❷ This value is replaced with the unquoted value of the **REPLICA_COUNT** parameter when the template is instantiated.
- ❸ The name of the parameter. This value is used to reference the parameter within the template.
- ❹ The user-friendly name for the parameter. This is displayed to users.
- ❺ A description of the parameter. Provide more detailed information for the purpose of the parameter, including any constraints on the expected value. Descriptions should use complete sentences to follow the console's text standards. Do not make this a duplicate of the display name.
- ❻ A default value for the parameter which is used if you do not override the value when instantiating the template. Avoid using default values for things like passwords, instead use generated parameters in combination with secrets.
- ❼ Indicates this parameter is required, meaning you cannot override it with an empty value. If the parameter does not provide a default or generated value, you must supply a value.
- ❽ A parameter which has its value generated.
- ❾ The input to the generator. In this case, the generator produces a 40 character alphanumeric value including upper and lowercase characters.
- ❿ Parameters can be included in the template message. This informs you about generated values.

3.1.7.4. Writing the template object list

The main portion of the template is the list of objects which is created when the template is instantiated.

This can be any valid API object, such as a build configuration, deployment configuration, or service. The object is created exactly as defined here, with any parameter values substituted in prior to creation. The definition of these objects can reference parameters defined earlier.

The following is an example of an object list:

```
kind: "Template"
apiVersion: "v1"
metadata:
  name: my-template
objects:
- kind: "Service" ❶
  apiVersion: "v1"
  metadata:
    name: "cakephp-mysql-example"
    annotations:
      description: "Exposes and load balances the application pods"
  spec:
    ports:
      - name: "web"
        port: 8080
        targetPort: 8080
    selector:
      name: "cakephp-mysql-example"
```

❶ The definition of a service, which is created by this template.



NOTE

If an object definition metadata includes a fixed **namespace** field value, the field is stripped out of the definition during template instantiation. If the **namespace** field contains a parameter reference, normal parameter substitution is performed and the object is created in whatever namespace the parameter substitution resolved the value to, assuming the user has permission to create objects in that namespace.

3.1.7.5. Marking a template as bindable

The Template Service Broker advertises one service in its catalog for each template object of which it is aware. By default, each of these services is advertised as being bindable, meaning an end user is permitted to bind against the provisioned service.

Procedure

Template authors can prevent end users from binding against services provisioned from a given template.

- Prevent end user from binding against services provisioned from a given template by adding the annotation **template.openshift.io/bindable: "false"** to the template.

3.1.7.6. Exposing template object fields

Template authors can indicate that fields of particular objects in a template should be exposed. The Template Service Broker recognizes exposed fields on **ConfigMap**, **Secret**, **Service**, and **Route** objects, and returns the values of the exposed fields when a user binds a service backed by the broker.

To expose one or more fields of an object, add annotations prefixed by **template.openshift.io/expose-** or **template.openshift.io/base64-expose-** to the object in the template.

Each annotation key, with its prefix removed, is passed through to become a key in a **bind** response.

Each annotation value is a Kubernetes JSONPath expression, which is resolved at bind time to indicate the object field whose value should be returned in the **bind** response.



NOTE

Bind response key-value pairs can be used in other parts of the system as environment variables. Therefore, it is recommended that every annotation key with its prefix removed should be a valid environment variable name – beginning with a character **A-Z**, **a-z**, or **_** and being followed by zero or more characters **A-Z**, **a-z**, **0-9**, or **_**.



NOTE

Unless escaped with a backslash, Kubernetes' JSONPath implementation interprets characters such as **.**, **@**, and others as metacharacters, regardless of their position in the expression. Therefore, for example, to refer to a **ConfigMap** datum named **my.key**, the required JSONPath expression would be **{.data['my\\.key']}**. Depending on how the JSONPath expression is then written in YAML, an additional backslash might be required, for example **"{.data['my\\.key']}"**.

The following is an example of different objects' fields being exposed:

```
kind: Template
apiVersion: template.openshift.io/v1
metadata:
  name: my-template
objects:
- kind: ConfigMap
  apiVersion: v1
  metadata:
    name: my-template-config
    annotations:
      template.openshift.io/expose-username: "{.data['my\\.username']}"
  data:
    my.username: foo
- kind: Secret
  apiVersion: v1
  metadata:
    name: my-template-config-secret
    annotations:
      template.openshift.io/base64-expose-password: "{.data['password']}"
  stringData:
    password: <password>
- kind: Service
  apiVersion: v1
  metadata:
    name: my-template-service
    annotations:
      template.openshift.io/expose-service_ip_port: "{.spec.clusterIP}:{.spec.ports[?
        (.name==\"web\")].port}"
```

```
spec:
  ports:
    - name: "web"
      port: 8080
- kind: Route
  apiVersion: route.openshift.io/v1
  metadata:
    name: my-template-route
    annotations:
      template.openshift.io/expose-uri: "http://{.spec.host}{.spec.path}"
  spec:
    path: mypath
```

An example response to a **bind** operation given the above partial template follows:

```
{
  "credentials": {
    "username": "foo",
    "password": "YmFy",
    "service_ip_port": "172.30.12.34:8080",
    "uri": "http://route-test.router.default.svc.cluster.local/mypath"
  }
}
```

Procedure

- Use the **template.openshift.io/expose-** annotation to return the field value as a string. This is convenient, although it does not handle arbitrary binary data.
- If you want to return binary data, use the **template.openshift.io/base64-expose-** annotation instead to base64 encode the data before it is returned.

3.1.7.7. Waiting for template readiness

Template authors can indicate that certain objects within a template should be waited for before a template instantiation by the service catalog, Template Service Broker, or **TemplateInstance** API is considered complete.

Before starting the procedure, read the following considerations:

- Set memory, CPU, and storage default sizes to make sure your application is given enough resources to run smoothly.
- Avoid referencing the **latest** tag from images if that tag is used across major versions. This can cause running applications to break when new images are pushed to that tag.
- A good template builds and deploys cleanly without requiring modifications after the template is deployed.

Procedure

- To use the template feature, mark one or more objects of kind **Build**, **BuildConfig**, **Deployment**, **DeploymentConfig**, **Job**, or **StatefulSet** in a template with the following annotation:

■

"template.alpha.openshift.io/wait-for-ready": "true"

Template instantiation is not complete until all objects marked with the annotation report ready. Similarly, if any of the annotated objects report failed, or if the template fails to become ready within a fixed timeout of one hour, the template instantiation fails.

For the purposes of instantiation, readiness and failure of each object kind are defined as follows:

Kind	Readiness	Failure
Build	Object reports phase complete.	Object reports phase canceled, error, or failed.
BuildConfig	Latest associated build object reports phase complete.	Latest associated build object reports phase canceled, error, or failed.
Deployment	Object reports new replica set and deployment available. This honors readiness probes defined on the object.	Object reports progressing condition as false.
DeploymentConfig	Object reports new replication controller and deployment available. This honors readiness probes defined on the object.	Object reports progressing condition as false.
Job	Object reports completion.	Object reports that one or more failures have occurred.
StatefulSet	Object reports all replicas ready. This honors readiness probes defined on the object.	Not applicable.

The following is an example template extract, which uses the **wait-for-ready** annotation. Further examples can be found in the OpenShift Container Platform quick start templates.

```
kind: Template
apiVersion: template.openshift.io/v1
metadata:
  name: my-template
objects:
- kind: BuildConfig
  apiVersion: build.openshift.io/v1
  metadata:
    name: ...
  annotations:
    # wait-for-ready used on BuildConfig ensures that template instantiation
    # will fail immediately if build fails
    template.alpha.openshift.io/wait-for-ready: "true"
  spec:
    ...
```

```

- kind: DeploymentConfig
  apiVersion: apps.openshift.io/v1
  metadata:
    name: ...
  annotations:
    template.alpha.openshift.io/wait-for-ready: "true"
  spec:
    ...
- kind: Service
  apiVersion: v1
  metadata:
    name: ...
  spec:
    ...

```

3.1.7.8. Creating a template from existing objects

Rather than writing an entire template from scratch, you can export existing objects from your project in YAML form, and then modify the YAML from there by adding parameters and other customizations as template form.

Procedure

- Export objects in a project in YAML form:

```
$ oc get -o yaml all > <yaml_filename>
```

You can also substitute a particular resource type or multiple resources instead of **all**. Run **oc get -h** for more examples.

The object types included in **oc get -o yaml all** are:

- **BuildConfig**
- **Build**
- **DeploymentConfig**
- **ImageStream**
- **Pod**
- **ReplicationController**
- **Route**
- **Service**



NOTE

Using the **all** alias is not recommended because the contents might vary across different clusters and versions. Instead, specify all required resources.

3.2. CREATING APPLICATIONS BY USING THE DEVELOPER PERSPECTIVE

The **Developer** perspective in the web console provides you the following options from the **+Add** view to create applications and associated services and deploy them on OpenShift Container Platform:

- **Getting started resources:** Use these resources to help you get started with Developer

Console. You can choose to hide the header using the Options menu



- **Creating applications using samples:** Use existing code samples to get started with creating applications on the OpenShift Container Platform.
- **Build with guided documentation:** Follow the guided documentation to build applications and familiarize yourself with key concepts and terminologies.
- **Explore new developer features:** Explore the new features and resources within the **Developer** perspective.
- **Developer catalog:** Explore the Developer Catalog to select the required applications, services, or source to image builders, and then add it to your project.
 - **All Services:** Browse the catalog to discover services across OpenShift Container Platform.
 - **Database:** Select the required database service and add it to your application.
 - **Operator Backed:** Select and deploy the required Operator-managed service.
 - **Helm chart:** Select the required Helm chart to simplify deployment of applications and services.
 - **Devfile:** Select a devfile from the **Devfile registry** to declaratively define a development environment.
 - **Event Source:** Select an event source to register interest in a class of events from a particular system.



NOTE

The Managed services option is also available if the RHOAS Operator is installed.

- **Git repository:** Import an existing codebase, Devfile, or Dockerfile from your Git repository using the **From Git**, **From Devfile**, or **From Dockerfile** options respectively, to build and deploy an application on OpenShift Container Platform.
- **Container images:** Use existing images from an image stream or registry to deploy it on to the OpenShift Container Platform.
- **Pipelines:** Use Tekton pipeline to create CI/CD pipelines for your software delivery process on the OpenShift Container Platform.
- **Serverless:** Explore the **Serverless** options to create, build, and deploy stateless and serverless applications on the OpenShift Container Platform.

- **Channel:** Create a Knative channel to create an event forwarding and persistence layer with in-memory and reliable implementations.
- **Samples:** Explore the available sample applications to create, build, and deploy an application quickly.
- **Quick Starts:** Explore the quick start options to create, import, and run applications with step-by-step instructions and tasks.
- **From Local Machine:** Explore the **From Local Machine** tile to import or upload files on your local machine for building and deploying applications easily.
 - **Import YAML:** Upload a YAML file to create and define resources for building and deploying applications.
 - **Upload JAR file:** Upload a JAR file to build and deploy Java applications.
- **Share my Project:** Use this option to add or remove users to a project and provide accessibility options to them.
- **Helm Chart repositories:** Use this option to add Helm Chart repositories in a namespace.
- **Re-ordering of resources:** Use these resources to re-order pinned resources added to your navigation pane. The drag-and-drop icon is displayed on the left side of the pinned resource when you hover over it in the navigation pane. The dragged resource can be dropped only in the section where it resides.

Note that certain options, such as **Pipelines**, **Event Source**, and **Import Virtual Machines**, are displayed only when the [OpenShift Pipelines Operator](#), [OpenShift Serverless Operator](#), and [OpenShift Virtualization Operator](#) are installed, respectively.

3.2.1. Prerequisites

To create applications using the **Developer** perspective ensure that:

- You have [logged in to the web console](#).
- You have created a project or have access to a project with the appropriate [roles and permissions](#) to create applications and other workloads in OpenShift Container Platform.

To create serverless applications, in addition to the preceding prerequisites, ensure that:

- You have [installed the OpenShift Serverless Operator](#).
- You have [created a KnativeServing resource in the knative-serving namespace](#).

3.2.2. Creating sample applications

You can use the sample applications in the **+Add** flow of the **Developer** perspective to create, build, and deploy applications quickly.

Prerequisites

- You have logged in to the OpenShift Container Platform web console and are in the **Developer** perspective.

Procedure

1. In the **+Add** view, click the **Samples** tile to see the **Samples** page.
2. On the **Samples** page, select one of the available sample applications to see the **Create Sample Application** form.
3. In the **Create Sample Application Form**
 - In the **Name** field, the deployment name is displayed by default. You can modify this name as required.
 - In the **Builder Image Version**, a builder image is selected by default. You can modify this image version by using the **Builder Image Version** drop-down list.
 - A sample Git repository URL is added by default.
4. Click **Create** to create the sample application. The build status of the sample application is displayed on the **Topology** view. After the sample application is created, you can see the deployment added to the application.

3.2.3. Creating applications by using Quick Starts

The **Quick Starts** page shows you how to create, import, and run applications on OpenShift Container Platform, with step-by-step instructions and tasks.

Prerequisites

- You have logged in to the OpenShift Container Platform web console and are in the **Developer** perspective.

Procedure

1. In the **+Add** view, click the **Getting Started resources** → **Build with guided documentation** → **View all quick starts** link to view the **Quick Starts** page.
2. In the **Quick Starts** page, click the tile for the quick start that you want to use.
3. Click **Start** to begin the quick start.
4. Perform the steps that are displayed.

3.2.4. Importing a codebase from Git to create an application

You can use the **Developer** perspective to create, build, and deploy an application on OpenShift Container Platform using an existing codebase in GitHub.

The following procedure walks you through the **From Git** option in the **Developer** perspective to create an application.

Procedure

1. In the **+Add** view, click **From Git** in the **Git Repository** tile to see the **Import from git** form.

2. In the **Git** section, enter the Git repository URL for the codebase you want to use to create an application. For example, enter the URL of this sample Node.js application <https://github.com/sclorg/nodejs-ex>. The URL is then validated.
3. Optional: You can click **Show Advanced Git Options** to add details such as:
 - **Git Reference** to point to code in a specific branch, tag, or commit to be used to build the application.
 - **Context Dir** to specify the subdirectory for the application source code you want to use to build the application.
 - **Source Secret** to create a **Secret Name** with credentials for pulling your source code from a private repository.
4. Optional: You can import a **Devfile**, a **Dockerfile**, a **Builder Image**, or a **Serverless Function** through your Git repository to further customize your deployment.
 - If your Git repository contains a **Devfile**, a **Dockerfile**, a **Builder Image**, or a **func.yaml**, it is automatically detected and populated on the respective path fields.
 - If a **Devfile**, a **Dockerfile**, or a **Builder Image** are detected in the same repository, the **Devfile** is selected by default.
 - If **func.yaml** is detected in the Git repository, the **Import Strategy** changes to **Serverless Function**.
 - Alternatively, you can create a serverless function by clicking **Create Serverless function** in the **+Add** view using the Git repository URL.
 - To edit the file import type and select a different strategy, click **Edit import strategy** option.
 - If multiple **Devfiles**, a **Dockerfiles**, or a **Builder Images** are detected, to import a specific instance, specify the respective paths relative to the context directory.
5. After the Git URL is validated, the recommended builder image is selected and marked with a star. If the builder image is not auto-detected, select a builder image. For the <https://github.com/sclorg/nodejs-ex> Git URL, by default the Node.js builder image is selected.
 - a. Optional: Use the **Builder Image Version** drop-down to specify a version.
 - b. Optional: Use the **Edit import strategy** to select a different strategy.
 - c. Optional: For the Node.js builder image, use the **Run command** field to override the command to run the application.
6. In the **General** section:
 - a. In the **Application** field, enter a unique name for the application grouping, for example, **myapp**. Ensure that the application name is unique in a namespace.
 - b. The **Name** field to identify the resources created for this application is automatically populated based on the Git repository URL if there are no existing applications. If there are existing applications, you can choose to deploy the component within an existing application, create a new application, or keep the component unassigned.

**NOTE**

The resource name must be unique in a namespace. Modify the resource name if you get an error.

7. In the **Resources** section, select:

- **Deployment**, to create an application in plain Kubernetes style.
- **Deployment Config**, to create an OpenShift Container Platform style application.
- **Serverless Deployment**, to create a Knative service.

**NOTE**

To set the default resource preference for importing an application, go to **User Preferences → Applications → Resource type** field. The **Serverless Deployment** option is displayed in the **Import from Git** form only if the OpenShift Serverless Operator is installed in your cluster. The **Resources** section is not available while creating a serverless function. For further details, refer to the OpenShift Serverless documentation.

8. In the **Pipelines** section, select **Add Pipeline**, and then click **Show Pipeline Visualization** to see the pipeline for the application. A default pipeline is selected, but you can choose the pipeline you want from the list of available pipelines for the application.

**NOTE**

The **Add pipeline** checkbox is checked and **Configure PAC** is selected by default if the following criterias are fulfilled:

- Pipeline operator is installed
- **pipelines-as-code** is enabled
- **.tekton** directory is detected in the Git repository

9. Add a webhook to your repository. If **Configure PAC** is checked and the GitHub App is set up, you can see the **Use GitHub App** and **Setup a webhook** options. If GitHub App is not set up, you can only see the **Setup a webhook** option:

- Go to **Settings → Webhooks** and click **Add webhook**.
- Set the **Payload URL** to the Pipelines as Code controller public URL.
- Select the content type as **application/json**.
- Add a webhook secret and note it in an alternate location. With **openssl** installed on your local machine, generate a random secret.
- Click **Let me select individual events** and select these events: **Commit comments**, **Issue comments**, **Pull request**, and **Pushes**.
- Click **Add webhook**.

10. Optional: In the **Advanced Options** section, the **Target port** and the **Create a route to the application** is selected by default so that you can access your application using a publicly available URL.
If your application does not expose its data on the default public port, 80, clear the check box, and set the target port number you want to expose.
11. Optional: You can use the following advanced options to further customize your application:

Routing

By clicking the **Routing** link, you can perform the following actions:

- Customize the hostname for the route.
- Specify the path the router watches.
- Select the target port for the traffic from the drop-down list.
- Secure your route by selecting the **Secure Route** check box. Select the required TLS termination type and set a policy for insecure traffic from the respective drop-down lists.



NOTE

For serverless applications, the Knative service manages all the routing options above. However, you can customize the target port for traffic, if required. If the target port is not specified, the default port of **8080** is used.

Domain mapping

If you are creating a **Serverless Deployment**, you can add a custom domain mapping to the Knative service during creation.

- In the **Advanced options** section, click **Show advanced Routing options**.
 - If the domain mapping CR that you want to map to the service already exists, you can select it from the **Domain mapping** drop-down menu.
 - If you want to create a new domain mapping CR, type the domain name into the box, and select the **Create** option. For example, if you type in **example.com**, the **Create** option is **Create "example.com"**.

Health Checks

Click the **Health Checks** link to add Readiness, Liveness, and Startup probes to your application. All the probes have prepopulated default data; you can add the probes with the default data or customize it as required.

To customize the health probes:

- Click **Add Readiness Probe**, if required, modify the parameters to check if the container is ready to handle requests, and select the check mark to add the probe.
- Click **Add Liveness Probe**, if required, modify the parameters to check if a container is still running, and select the check mark to add the probe.

- Click **Add Startup Probe**, if required, modify the parameters to check if the application within the container has started, and select the check mark to add the probe. For each of the probes, you can specify the request type - **HTTP GET**, **Container Command**, or **TCP Socket**, from the drop-down list. The form changes as per the selected request type. You can then modify the default values for the other parameters, such as the success and failure thresholds for the probe, number of seconds before performing the first probe after the container starts, frequency of the probe, and the timeout value.

Build Configuration and Deployment

Click the **Build Configuration** and **Deployment** links to see the respective configuration options. Some options are selected by default; you can customize them further by adding the necessary triggers and environment variables.

For serverless applications, the **Deployment** option is not displayed as the Knative configuration resource maintains the desired state for your deployment instead of a **DeploymentConfig** resource.

Scaling

Click the **Scaling** link to define the number of pods or instances of the application you want to deploy initially.

If you are creating a serverless deployment, you can also configure the following settings:

- **Min Pods** determines the lower limit for the number of pods that must be running at any given time for a Knative service. This is also known as the **minScale** setting.
- **Max Pods** determines the upper limit for the number of pods that can be running at any given time for a Knative service. This is also known as the **maxScale** setting.
- **Concurrency target** determines the number of concurrent requests desired for each instance of the application at a given time.
- **Concurrency limit** determines the limit for the number of concurrent requests allowed for each instance of the application at a given time.
- **Concurrency utilization** determines the percentage of the concurrent requests limit that must be met before Knative scales up additional pods to handle additional traffic.
- **Autoscale window** defines the time window over which metrics are averaged to provide input for scaling decisions when the autoscaler is not in panic mode. A service is scaled-to-zero if no requests are received during this window. The default duration for the autoscale window is **60s**. This is also known as the stable window.

Resource Limit

Click the **Resource Limit** link to set the amount of **CPU** and **Memory** resources a container is guaranteed or allowed to use when running.

Labels

Click the **Labels** link to add custom labels to your application.

12. Click **Create** to create the application and a success notification is displayed. You can see the build status of the application in the **Topology** view.

3.2.5. Creating applications by deploying container image

You can use an external image registry or an image stream tag from an internal registry to deploy an application on your cluster.

Prerequisites

- You have logged in to the OpenShift Container Platform web console and are in the **Developer** perspective.

Procedure

1. In the **+Add** view, click **Container images** to view the **Deploy Images** page.
2. In the **Image** section:
 - a. Select **Image name from external registry** to deploy an image from a public or a private registry, or select **Image stream tag from internal registry** to deploy an image from an internal registry.
 - b. Select an icon for your image in the **Runtime icon** tab.
3. In the **General** section:
 - a. In the **Application name** field, enter a unique name for the application grouping.
 - b. In the **Name** field, enter a unique name to identify the resources created for this component.
4. In the **Resource type** section, select the resource type to generate:
 - a. Select **Deployment** to enable declarative updates for **Pod** and **ReplicaSet** objects.
 - b. Select **DeploymentConfig** to define the template for a **Pod** object, and manage deploying new images and configuration sources.
 - c. Select **Serverless Deployment** to enable scaling to zero when idle.
5. Click **Create**. You can view the build status of the application in the **Topology** view.

3.2.6. Deploying a Java application by uploading a JAR file

You can use the web console **Developer** perspective to upload a JAR file by using the following options:

- Navigate to the **+Add** view of the **Developer** perspective, and click **Upload JAR file** in the **From Local Machine** tile. Browse and select your JAR file, or drag a JAR file to deploy your application.
- Navigate to the **Topology** view and use the **Upload JAR file** option, or drag a JAR file to deploy your application.
- Use the in-context menu in the **Topology** view, and then use the **Upload JAR file** option to upload your JAR file to deploy your application.

Prerequisites

- The Cluster Samples Operator must be installed by a cluster administrator.

- You have access to the OpenShift Container Platform web console and are in the **Developer** perspective.

Procedure

1. In the **Topology** view, right-click anywhere to view the **Add to Project** menu.
2. Hover over the **Add to Project** menu to see the menu options, and then select the **Upload JAR file** option to see the **Upload JAR file** form. Alternatively, you can drag the JAR file into the **Topology** view.
3. In the **JAR file** field, browse for the required JAR file on your local machine and upload it. Alternatively, you can drag the JAR file on to the field. A toast alert is displayed at the top right if an incompatible file type is dragged into the **Topology** view. A field error is displayed if an incompatible file type is dropped on the field in the upload form.
4. The runtime icon and builder image are selected by default. If a builder image is not auto-detected, select a builder image. If required, you can change the version using the **Builder Image Version** drop-down list.
5. Optional: In the **Application Name** field, enter a unique name for your application to use for resource labelling.
6. In the **Name** field, enter a unique component name for the associated resources.
7. Optional: Use the **Resource type** drop-down list to change the resource type.
8. In the **Advanced options** menu, click **Create a Route to the Application** to configure a public URL for your deployed application.
9. Click **Create** to deploy the application. A toast notification is shown to notify you that the JAR file is being uploaded. The toast notification also includes a link to view the build logs.



NOTE

If you attempt to close the browser tab while the build is running, a web alert is displayed.

After the JAR file is uploaded and the application is deployed, you can view the application in the **Topology** view.

3.2.7. Using the Devfile registry to access devfiles

You can use the devfiles in the **+Add** flow of the **Developer** perspective to create an application. The **+Add** flow provides a complete integration with the [devfile community registry](#). A devfile is a portable YAML file that describes your development environment without needing to configure it from scratch. Using the **Devfile registry**, you can use a preconfigured devfile to create an application.

Procedure

1. Navigate to **Developer Perspective** → **+Add** → **Developer Catalog** → **All Services**. A list of all the available services in the **Developer Catalog** is displayed.
2. Under **Type**, click **Devfiles** to browse for devfiles that support a particular language or framework. Alternatively, you can use the keyword filter to search for a particular devfile using their name, tag, or description.

3. Click the devfile you want to use to create an application. The devfile tile displays the details of the devfile, including the name, description, provider, and the documentation of the devfile.
4. Click **Create** to create an application and view the application in the **Topology** view.


3.2.8. Using the Developer Catalog to add services or components to your application

You use the Developer Catalog to deploy applications and services based on Operator backed services such as Databases, Builder Images, and Helm Charts. The Developer Catalog contains a collection of application components, services, event sources, or source-to-image builders that you can add to your project. Cluster administrators can customize the content made available in the catalog.

Procedure

1. In the **Developer** perspective, navigate to the **+Add** view and from the **Developer Catalog** tile, click **All Services** to view all the available services in the **Developer Catalog**.
2. Under **All Services**, select the kind of service or the component you need to add to your project. For this example, select **Databases** to list all the database services and then click **MariaDB** to see the details for the service.
3. Click **Instantiate Template** to see an automatically populated template with details for the **MariaDB** service, and then click **Create** to create and view the MariaDB service in the **Topology** view.

Figure 3.1. MariaDB in Topology



3.2.9. Additional resources

- For more information about Knative routing settings for OpenShift Serverless, see [Routing](#).

- For more information about domain mapping settings for OpenShift Serverless, see [Configuring a custom domain for a Knative service](#).
- For more information about Knative autoscaling settings for OpenShift Serverless, see [Autoscaling](#).
- For more information about adding a new user to a project, see [Working with projects](#).
- For more information about creating a Helm Chart repository, see [Creating Helm Chart repositories](#).

3.3. CREATING APPLICATIONS FROM INSTALLED OPERATORS

Operators are a method of packaging, deploying, and managing a Kubernetes application. You can create applications on OpenShift Container Platform using Operators that have been installed by a cluster administrator.

This guide walks developers through an example of creating applications from an installed Operator using the OpenShift Container Platform web console.

Additional resources

- See the [Operators](#) guide for more on how Operators work and how the Operator Lifecycle Manager is integrated in OpenShift Container Platform.

3.3.1. Creating an etcd cluster using an Operator

This procedure walks through creating a new etcd cluster using the etcd Operator, managed by Operator Lifecycle Manager (OLM).

Prerequisites

- Access to an OpenShift Container Platform 4.19 cluster.
- The etcd Operator already installed cluster-wide by an administrator.

Procedure

1. Create a new project in the OpenShift Container Platform web console for this procedure. This example uses a project called **my-etcd**.
2. Navigate to the **Operators → Installed Operators** page. The Operators that have been installed to the cluster by the cluster administrator and are available for use are shown here as a list of cluster service versions (CSVs). CSVs are used to launch and manage the software provided by the Operator.

TIP

You can get this list from the CLI using:

```
$ oc get csv
```

3. On the **Installed Operators** page, click the etcd Operator to view more details and available actions.

As shown under **Provided APIs**, this Operator makes available three new resource types, including one for an **etcd Cluster** (the **EtcdCluster** resource). These objects work similar to the built-in native Kubernetes ones, such as **Deployment** or **ReplicaSet**, but contain logic specific to managing etcd.

4. Create a new etcd cluster:
 - a. In the **etcd Cluster** API box, click **Create instance**.
 - b. The next page allows you to make any modifications to the minimal starting template of an **EtcdCluster** object, such as the size of the cluster. For now, click **Create** to finalize. This triggers the Operator to start up the pods, services, and other components of the new etcd cluster.
5. Click the **example** etcd cluster, then click the **Resources** tab to see that your project now contains a number of resources created and configured automatically by the Operator. Verify that a Kubernetes service has been created that allows you to access the database from other pods in your project.
6. All users with the **edit** role in a given project can create, manage, and delete application instances (an etcd cluster, in this example) managed by Operators that have already been created in the project, in a self-service manner, just like a cloud service. If you want to enable additional users with this ability, project administrators can add the role using the following command:

```
$ oc policy add-role-to-user edit <user> -n <target_project>
```

You now have an etcd cluster that will react to failures and rebalance data as pods become unhealthy or are migrated between nodes in the cluster. Most importantly, cluster administrators or developers with proper access can now easily use the database with their applications.

3.4. CREATING APPLICATIONS BY USING THE CLI

You can create an OpenShift Container Platform application from components that include source or binary code, images, and templates by using the OpenShift Container Platform CLI.

The set of objects created by **new-app** depends on the artifacts passed as input: source repositories, images, or templates.

3.4.1. Creating an application from source code

With the **new-app** command you can create applications from source code in a local or remote Git repository.

The **new-app** command creates a build configuration, which itself creates a new application image from your source code. The **new-app** command typically also creates a **Deployment** object to deploy the new image, and a service to provide load-balanced access to the deployment running your image.

OpenShift Container Platform automatically detects whether the pipeline, source, or docker build strategy should be used, and in the case of source build, detects an appropriate language builder image.

3.4.1.1. Local

To create an application from a Git repository in a local directory:

■

```
$ oc new-app /<path to source code>
```



NOTE

If you use a local Git repository, the repository must have a remote named **origin** that points to a URL that is accessible by the OpenShift Container Platform cluster. If there is no recognized remote, running the **new-app** command will create a binary build.

3.4.1.2. Remote

To create an application from a remote Git repository:

```
$ oc new-app https://github.com/sclorg/cakephp-ex
```

To create an application from a private remote Git repository:

```
$ oc new-app https://github.com/youruser/yourprivaterepo --source-secret=yoursecret
```



NOTE

If you use a private remote Git repository, you can use the **--source-secret** flag to specify an existing source clone secret that will get injected into your build config to access the repository.

You can use a subdirectory of your source code repository by specifying a **--context-dir** flag. To create an application from a remote Git repository and a context subdirectory:

```
$ oc new-app https://github.com/sclorg/s2i-ruby-container.git \
  --context-dir=2.0/test/puma-test-app
```

Also, when specifying a remote URL, you can specify a Git branch to use by appending **#<branch_name>** to the end of the URL:

```
$ oc new-app https://github.com/openshift/ruby-hello-world.git#beta4
```

3.4.1.3. Build strategy detection

OpenShift Container Platform automatically determines which build strategy to use by detecting certain files:

- If a Jenkins file exists in the root or specified context directory of the source repository when creating a new application, OpenShift Container Platform generates a pipeline build strategy.



NOTE


The **pipeline** build strategy is deprecated; consider using Red Hat OpenShift Pipelines instead.

- If a Dockerfile exists in the root or specified context directory of the source repository when creating a new application, OpenShift Container Platform generates a docker build strategy.

- If neither a Jenkins file nor a Dockerfile is detected, OpenShift Container Platform generates a source build strategy.

Override the automatically detected build strategy by setting the **--strategy** flag to **docker**, **pipeline**, or **source**.

```
$ oc new-app /home/user/code/myapp --strategy=docker
```



NOTE

The **oc** command requires that files containing build sources are available in a remote Git repository. For all source builds, you must use **git remote -v**.

3.4.1.4. Language detection

If you use the source build strategy, **new-app** attempts to determine the language builder to use by the presence of certain files in the root or specified context directory of the repository:

Table 3.1. Languages detected by **new-app**

Language	Files
dotnet	project.json, *.csproj
jee	pom.xml
nodejs	app.json, package.json
perl	cpanfile, index.pl
php	composer.json, index.php
python	requirements.txt, setup.py
ruby	Gemfile, Rakefile, config.ru
scala	build.sbt
golang	Godeps, main.go

After a language is detected, **new-app** searches the OpenShift Container Platform server for image stream tags that have a **supports** annotation matching the detected language, or an image stream that matches the name of the detected language. If a match is not found, **new-app** searches the [Docker Hub registry](#) for an image that matches the detected language based on name.

You can override the image the builder uses for a particular source repository by specifying the image, either an image stream or container specification, and the repository with a ~ as a separator. Note that if this is done, build strategy detection and language detection are not carried out.

For example, to use the **myproject/my-ruby** imagestream with the source in a remote repository:

```
$ oc new-app myproject/my-ruby~https://github.com/openshift/ruby-hello-world.git
```

To use the **openshift/ruby-20-centos7:latest** container image stream with the source in a local repository:

```
$ oc new-app openshift/ruby-20-centos7:latest~/home/user/code/my-ruby-app
```



NOTE

Language detection requires the Git client to be locally installed so that your repository can be cloned and inspected. If Git is not available, you can avoid the language detection step by specifying the builder image to use with your repository with the **<image>~<repository>** syntax.


The **-i <image> <repository>** invocation requires that **new-app** attempt to clone **repository** to determine what type of artifact it is, so this will fail if Git is not available.

The **-i <image> --code <repository>** invocation requires **new-app** clone **repository** to determine whether **image** should be used as a builder for the source code, or deployed separately, as in the case of a database image.

3.4.2. Creating an application from an image

You can deploy an application from an existing image. Images can come from image streams in the OpenShift Container Platform server, images in a specific registry, or images in the local Docker server.

The **new-app** command attempts to determine the type of image specified in the arguments passed to it. However, you can explicitly tell **new-app** whether the image is a container image using the **--docker-image** argument or an image stream using the **-i|--image-stream** argument.



NOTE

If you specify an image from your local Docker repository, you must ensure that the same image is available to the OpenShift Container Platform cluster nodes.

3.4.2.1. Docker Hub MySQL image

Create an application from the Docker Hub MySQL image, for example:

```
$ oc new-app mysql
```

3.4.2.2. Image in a private registry

Create an application using an image in a private registry, specify the full container image specification:

```
$ oc new-app myregistry:5000/example/myimage
```

3.4.2.3. Existing image stream and optional image stream tag

Create an application from an existing image stream and optional image stream tag:

```
$ oc new-app my-stream:v1
```

3.4.3. Creating an application from a template

You can create an application from a previously stored template or from a template file, by specifying the name of the template as an argument. For example, you can store a sample application template and use it to create an application.

Upload an application template to your current project's template library. The following example uploads an application template from a file called **examples/sample-app/application-template-stibuild.json**:

```
$ oc create -f examples/sample-app/application-template-stibuild.json
```

Then create a new application by referencing the application template. In this example, the template name is **ruby-helloworld-sample**:

```
$ oc new-app ruby-helloworld-sample
```

To create a new application by referencing a template file in your local file system, without first storing it in OpenShift Container Platform, use the **-f|--file** argument. For example:

```
$ oc new-app -f examples/sample-app/application-template-stibuild.json
```

3.4.3.1. Template parameters

When creating an application based on a template, use the **-p|--param** argument to set parameter values that are defined by the template:

```
$ oc new-app ruby-helloworld-sample \
  -p ADMIN_USERNAME=admin -p ADMIN_PASSWORD=mypassword
```

You can store your parameters in a file, then use that file with **--param-file** when instantiating a template. If you want to read the parameters from standard input, use **--param-file=-**. The following is an example file called **helloworld.params**:

```
ADMIN_USERNAME=admin
ADMIN_PASSWORD=mypassword
```

Reference the parameters in the file when instantiating a template:

```
$ oc new-app ruby-helloworld-sample --param-file=helloworld.params
```

3.4.4. Modifying application creation

The **new-app** command generates OpenShift Container Platform objects that build, deploy, and run the application that is created. Normally, these objects are created in the current project and assigned names that are derived from the input source repositories or the input images. However, with **new-app** you can modify this behavior.

Table 3.2. **new-app** output objects

Object	Description
BuildConfig	A BuildConfig object is created for each source repository that is specified in the command line. The BuildConfig object specifies the strategy to use, the source location, and the build output location.
ImageStreams	For the BuildConfig object, two image streams are usually created. One represents the input image. With source builds, this is the builder image. With Docker builds, this is the FROM image. The second one represents the output image. If a container image was specified as input to new-app , then an image stream is created for that image as well.
DeploymentConfig	A DeploymentConfig object is created either to deploy the output of a build, or a specified image. The new-app command creates emptyDir volumes for all Docker volumes that are specified in containers included in the resulting DeploymentConfig object.
Service	The new-app command attempts to detect exposed ports in input images. It uses the lowest numeric exposed port to generate a service that exposes that port. To expose a different port, after new-app has completed, simply use the oc expose command to generate additional services.
Other	Other objects can be generated when instantiating templates, according to the template.

3.4.4.1. Specifying environment variables

When generating applications from a template, source, or an image, you can use the **-e|--env** argument to pass environment variables to the application container at run time:

```
$ oc new-app openshift/postgresql-92-centos7 \
  -e POSTGRESQL_USER=user \
  -e POSTGRESQL_DATABASE=db \
  -e POSTGRESQL_PASSWORD=password
```

The variables can also be read from file using the **--env-file** argument. The following is an example file called **postgresql.env**:

```
POSTGRESQL_USER=user
POSTGRESQL_DATABASE=db
POSTGRESQL_PASSWORD=password
```

Read the variables from the file:

```
$ oc new-app openshift/postgresql-92-centos7 --env-file=postgresql.env
```

Additionally, environment variables can be given on standard input by using **--env-file=-**:

```
$ cat postgresql.env | oc new-app openshift/postgresql-92-centos7 --env-file=-
```

**NOTE**

Any **BuildConfig** objects created as part of **new-app** processing are not updated with environment variables passed with the **-e|--env** or **--env-file** argument.

3.4.4.2. Specifying build environment variables

When generating applications from a template, source, or an image, you can use the **--build-env** argument to pass environment variables to the build container at run time:

```
$ oc new-app openshift/ruby-23-centos7 \
  --build-env HTTP_PROXY=http://myproxy.net:1337/ \
  --build-env GEM_HOME=~/.gem
```

The variables can also be read from a file using the **--build-env-file** argument. The following is an example file called **ruby.env**:

```
HTTP_PROXY=http://myproxy.net:1337/
GEM_HOME=~/.gem
```

Read the variables from the file:

```
$ oc new-app openshift/ruby-23-centos7 --build-env-file=ruby.env
```

Additionally, environment variables can be given on standard input by using **--build-env-file=-**:

```
$ cat ruby.env | oc new-app openshift/ruby-23-centos7 --build-env-file=-
```

3.4.4.3. Specifying labels

When generating applications from source, images, or templates, you can use the **-l|--label** argument to add labels to the created objects. Labels make it easy to collectively select, configure, and delete objects associated with the application.

```
$ oc new-app https://github.com/openshift/ruby-hello-world -l name=hello-world
```

3.4.4.4. Viewing the output without creation

To see a dry-run of running the **new-app** command, you can use the **-o|--output** argument with a **yaml** or **json** value. You can then use the output to preview the objects that are created or redirect it to a file that you can edit. After you are satisfied, you can use **oc create** to create the OpenShift Container Platform objects.

To output **new-app** artifacts to a file, run the following:

```
$ oc new-app https://github.com/openshift/ruby-hello-world \
  -o yaml > myapp.yaml
```

Edit the file:

```
$ vi myapp.yaml
```

Create a new application by referencing the file:

```
$ oc create -f myapp.yaml
```

3.4.4.5. Creating objects with different names

Objects created by **new-app** are normally named after the source repository, or the image used to generate them. You can set the name of the objects produced by adding a **--name** flag to the command:

```
$ oc new-app https://github.com/openshift/ruby-hello-world --name=myapp
```

3.4.4.6. Creating objects in a different project

Normally, **new-app** creates objects in the current project. However, you can create objects in a different project by using the **-n|--namespace** argument:

```
$ oc new-app https://github.com/openshift/ruby-hello-world -n myproject
```

3.4.4.7. Creating multiple objects

The **new-app** command allows creating multiple applications specifying multiple parameters to **new-app**. Labels specified in the command line apply to all objects created by the single command. Environment variables apply to all components created from source or images.

To create an application from a source repository and a Docker Hub image:

```
$ oc new-app https://github.com/openshift/ruby-hello-world mysql
```



NOTE

If a source code repository and a builder image are specified as separate arguments, **new-app** uses the builder image as the builder for the source code repository. If this is not the intent, specify the required builder image for the source using the **~** separator.

3.4.4.8. Grouping images and source in a single pod

The **new-app** command allows deploying multiple images together in a single pod. To specify which images to group together, use the **+** separator. The **--group** command-line argument can also be used to specify the images that should be grouped together. To group the image built from a source repository with other images, specify its builder image in the group:

```
$ oc new-app ruby+mysql
```

To deploy an image built from source and an external image together:

```
$ oc new-app \
  ruby~https://github.com/openshift/ruby-hello-world \
```



```
mysql \
--group=ruby+mysql
```

3.4.4.9. Searching for images, templates, and other inputs

To search for images, templates, and other inputs for the **oc new-app** command, add the **--search** and **--list** flags. For example, to find all of the images or templates that include PHP:

```
$ oc new-app --search php
```

3.4.4.10. Setting the import mode

To set the import mode when using **oc new-app**, add the **--import-mode** flag. This flag can be appended with **Legacy** or **PreserveOriginal**, which provides users the option to create image streams using a single sub-manifest, or all manifests, respectively.

```
$ oc new-app --image=registry.redhat.io/ubi8/httpd-24:latest --import-mode=Legacy --name=test
```

```
$ oc new-app --image=registry.redhat.io/ubi8/httpd-24:latest --import-mode=PreserveOriginal --
name=test
```

3.5. CREATING APPLICATIONS USING RUBY ON RAILS

Ruby on Rails is a web framework written in Ruby. This guide covers using Rails 4 on OpenShift Container Platform.



WARNING

Go through the whole tutorial to have an overview of all the steps necessary to run your application on the OpenShift Container Platform. If you experience a problem try reading through the entire tutorial and then going back to your issue. It can also be useful to review your previous steps to ensure that all the steps were run correctly.

3.5.1. Prerequisites

- Basic Ruby and Rails knowledge.
- Locally installed version of Ruby 2.0.0+, Rubygems, Bundler.
- Basic Git knowledge.
- Running instance of OpenShift Container Platform 4.
- Make sure that an instance of OpenShift Container Platform is running and is available. Also make sure that your **oc** CLI client is installed and the command is accessible from your command shell, so you can use it to log in using your email address and password.

3.5.2. Setting up the database

Rails applications are almost always used with a database. For local development use the PostgreSQL database.

Procedure

1. Install the database:

```
$ sudo yum install -y postgresql postgresql-server postgresql-devel
```

2. Initialize the database:

```
$ sudo postgresql-setup initdb
```

This command creates the **/var/lib/pgsql/data** directory, in which the data is stored.

3. Start the database:

```
$ sudo systemctl start postgresql.service
```

4. When the database is running, create your **rails** user:

```
$ sudo -u postgres createuser -s rails
```

Note that the user created has no password.

3.5.3. Writing your application

If you are starting your Rails application from scratch, you must install the Rails gem first. Then you can proceed with writing your application.

Procedure

1. Install the Rails gem:

```
$ gem install rails
```

Example output

```
Successfully installed rails-4.3.0
1 gem installed
```

2. After you install the Rails gem, create a new application with PostgreSQL as your database:

```
$ rails new rails-app --database=postgresql
```

3. Change into your new application directory:

```
$ cd rails-app
```

4. If you already have an application, make sure the **pg** (postgresql) gem is present in your **Gemfile**. If not, edit your **Gemfile** by adding the gem:

```
gem 'pg'
```

5. Generate a new **Gemfile.lock** with all your dependencies:

```
$ bundle install
```

6. In addition to using the **postgresql** database with the **pg** gem, you also must ensure that the **config/database.yml** is using the **postgresql** adapter.

Make sure you updated **default** section in the **config/database.yml** file, so it looks like this:

```
default: &default
  adapter: postgresql
  encoding: unicode
  pool: 5
  host: localhost
  username: rails
  password: <password>
```

7. Create your application's development and test databases:

```
$ rake db:create
```

This creates **development** and **test** database in your PostgreSQL server.

3.5.3.1. Creating a welcome page

Since Rails 4 no longer serves a static **public/index.html** page in production, you must create a new root page.

To have a custom welcome page must do following steps:

- Create a controller with an index action.
- Create a view page for the welcome controller index action.
- Create a route that serves applications root page with the created controller and view.

Rails offers a generator that completes all necessary steps for you.

Procedure

1. Run Rails generator:

```
$ rails generate controller welcome index
```

All the necessary files are created.

2. edit line 2 in **config/routes.rb** file as follows:

```
root 'welcome#index'
```

3. Run the rails server to verify the page is available:

```
$ rails server
```

You should see your page by visiting <http://localhost:3000> in your browser. If you do not see the page, check the logs that are output to your server to debug.

3.5.3.2. Configuring application for OpenShift Container Platform

To have your application communicate with the PostgreSQL database service running in OpenShift Container Platform you must edit the **default** section in your **config/database.yml** to use environment variables, which you must define later, upon the database service creation.

Procedure

- Edit the **default** section in your **config/database.yml** with pre-defined variables as follows:

Sample config/database YAML file

```
<% user = ENV.key?("POSTGRESQL_ADMIN_PASSWORD") ? "root" :
ENV["POSTGRESQL_USER"] %>
<% password = ENV.key?("POSTGRESQL_ADMIN_PASSWORD") ?
ENV["POSTGRESQL_ADMIN_PASSWORD"] : ENV["POSTGRESQL_PASSWORD"] %>
<% db_service = ENV.fetch("DATABASE_SERVICE_NAME", "").upcase %>

default: &default
  adapter: postgresql
  encoding: unicode
  # For details on connection pooling, see rails configuration guide
  # http://guides.rubyonrails.org/configuring.html#database-pooling
  pool: <%= ENV["POSTGRESQL_MAX_CONNECTIONS"] || 5 %>
  username: <%= user %>
  password: <%= password %>
  host: <%= ENV["#{db_service}_SERVICE_HOST"] %>
  port: <%= ENV["#{db_service}_SERVICE_PORT"] %>
  database: <%= ENV["POSTGRESQL_DATABASE"] %>
```

3.5.3.3. Storing your application in Git

Building an application in OpenShift Container Platform usually requires that the source code be stored in a git repository, so you must install **git** if you do not already have it.

Prerequisites

- Install git.

Procedure

1. Make sure you are in your Rails application directory by running the **ls -1** command. The output of the command should look like:

```
$ ls -1
```

Example output

```

app
bin
config
config.ru
db
Gemfile
Gemfile.lock
lib
log
public
Rakefile
README.rdoc
test
tmp
vendor

```

2. Run the following commands in your Rails app directory to initialize and commit your code to git:

```
$ git init
```

```
$ git add .
```

```
$ git commit -m "initial commit"
```

After your application is committed you must push it to a remote repository. GitHub account, in which you create a new repository.

3. Set the remote that points to your **git** repository:

```
$ git remote add origin git@github.com:<namespace/repository-name>.git
```

4. Push your application to your remote git repository.

```
$ git push
```

3.5.4. Deploying your application to OpenShift Container Platform

You can deploy you application to OpenShift Container Platform.

After creating the **rails-app** project, you are automatically switched to the new project namespace.

Deploying your application in OpenShift Container Platform involves three steps:

- Creating a database service from OpenShift Container Platform's PostgreSQL image.
- Creating a frontend service from OpenShift Container Platform's Ruby 2.0 builder image and your Ruby on Rails source code, which are wired with the database service.
- Creating a route for your application.

Procedure

- To deploy your Ruby on Rails application, create a new project for the application:

```
$ oc new-project rails-app --description="My Rails application" --display-name="Rails Application"
```

3.5.4.1. Creating the database service

Your Rails application expects a running database service. For this service use PostgreSQL database image.

To create the database service, use the **oc new-app** command. To this command you must pass some necessary environment variables which are used inside the database container. These environment variables are required to set the username, password, and name of the database. You can change the values of these environment variables to anything you would like. The variables are as follows:

- POSTGRESQL_DATABASE
- POSTGRESQL_USER
- POSTGRESQL_PASSWORD

Setting these variables ensures:

- A database exists with the specified name.
- A user exists with the specified name.
- The user can access the specified database with the specified password.

Procedure

1. Create the database service:

```
$ oc new-app postgresql -e POSTGRESQL_DATABASE=db_name -e  
POSTGRESQL_USER=username -e POSTGRESQL_PASSWORD=password
```

To also set the password for the database administrator, append to the previous command with:

```
-e POSTGRESQL_ADMIN_PASSWORD=admin_pw
```

2. Watch the progress:

```
$ oc get pods --watch
```

3.5.4.2. Creating the frontend service

To bring your application to OpenShift Container Platform, you must specify a repository in which your application lives.

Procedure

1. Create the frontend service and specify database related environment variables that were setup when creating the database service:

```
$ oc new-app path/to/source/code --name=rails-app -e POSTGRESQL_USER=username -e
POSTGRESQL_PASSWORD=password -e POSTGRESQL_DATABASE=db_name -e
DATABASE_SERVICE_NAME=postgresql
```

With this command, OpenShift Container Platform fetches the source code, sets up the builder, builds your application image, and deploys the newly created image together with the specified environment variables. The application is named **rails-app**.

2. Verify the environment variables have been added by viewing the JSON document of the **rails-app** deployment config:

```
$ oc get dc rails-app -o json
```

You should see the following section:

Example output

```
env": [
  {
    "name": "POSTGRESQL_USER",
    "value": "username"
  },
  {
    "name": "POSTGRESQL_PASSWORD",
    "value": "password"
  },
  {
    "name": "POSTGRESQL_DATABASE",
    "value": "db_name"
  },
  {
    "name": "DATABASE_SERVICE_NAME",
    "value": "postgresql"
  }
],
```

3. Check the build process:

```
$ oc logs -f build/rails-app-1
```

4. After the build is complete, look at the running pods in OpenShift Container Platform:

```
$ oc get pods
```

You should see a line starting with **myapp-<number>-<hash>**, and that is your application running in OpenShift Container Platform.

5. Before your application is functional, you must initialize the database by running the database migration script. There are two ways you can do this:

- Manually from the running frontend container:
 - Exec into frontend container with **rsh** command:

–

```
$ oc rsh <frontend_pod_id>
```

- Run the migration from inside the container:

```
$ RAILS_ENV=production bundle exec rake db:migrate
```

If you are running your Rails application in a **development** or **test** environment you do not have to specify the **RAILS_ENV** environment variable.

- By adding pre-deployment lifecycle hooks in your template.

3.5.4.3. Creating a route for your application

You can expose a service to create a route for your application.

Procedure

- To expose a service by giving it an externally-reachable hostname like **www.example.com** use OpenShift Container Platform route. In your case you need to expose the frontend service by typing:

```
$ oc expose service rails-app --hostname=www.example.com
```



WARNING

Ensure the hostname you specify resolves into the IP address of the router.

CHAPTER 4. VIEWING APPLICATION COMPOSITION BY USING THE TOPOLOGY VIEW

The **Topology** view in the **Developer** perspective of the web console provides a visual representation of all the applications within a project, their build status, and the components and services associated with them.

4.1. PREREQUISITES

To view your applications in the **Topology** view and interact with them, ensure that:

- You have [logged in to the web console](#).
- You have the appropriate [roles and permissions](#) in a project to create applications and other workloads in OpenShift Container Platform.
- You are in [the Developer perspective](#).

4.2. VIEWING THE TOPOLOGY OF YOUR APPLICATION

You can navigate to the **Topology** view using the left navigation panel in the **Developer** perspective. After you deploy an application, you are directed automatically to the **Graph view** where you can see the status of the application pods, quickly access the application on a public URL, access the source code to modify it, and see the status of your last build. You can zoom in and out to see more details for a particular application.

The **Topology** view provides you the option to monitor your applications using the **List view**. Use the

List view icon () to see a list of all your applications and use the **Graph view** icon () to switch back to the graph view.

You can customize the views as required using the following:

- Use the **Find by name** field to find the required components. Search results may appear outside of the visible area; click **Fit to Screen** from the lower-left toolbar to resize the **Topology** view to show all components.
- Use the **Display Options** drop-down list to configure the **Topology** view of the various application groupings. The options are available depending on the types of components deployed in the project:
 - **Expand group**
 - Virtual Machines: Toggle to show or hide the virtual machines.
 - Application Groupings: Clear to condense the application groups into cards with an overview of an application group and alerts associated with it.
 - Helm Releases: Clear to condense the components deployed as Helm Release into cards with an overview of a given release.
 - Knative Services: Clear to condense the Knative Service components into cards with an overview of a given component.

- Operator Groupings: Click to condense the components deployed with an Operator into cards with an overview of the given group.
- **Show** elements based on **Pod Count** or **Labels**
 - Pod Count: Select to show the number of pods of a component in the component icon.
 - Labels: Toggle to show or hide the component labels.

The **Topology** view also provides you the **Export application** option to download your application in the ZIP file format. You can then import the downloaded application to another project or cluster. For more details, see *Exporting an application to another project or cluster* in the *Additional resources* section.

4.3. INTERACTING WITH APPLICATIONS AND COMPONENTS

In the **Topology** view in the **Developer** perspective of the web console, the **Graph view** provides the following options to interact with applications and components:

- Click **Open URL** () to see your application exposed by the route on a public URL.
- Click **Edit Source code** to access your source code and modify it.




NOTE

This feature is available only when you create applications using the **From Git**, **From Catalog**, and the **From Dockerfile** options.

- Hover your cursor over the lower left icon on the pod to see the name of the latest build and its status. The status of the application build is indicated as **New** (), **Pending** (), **Running** (), **Completed** (), **Failed** (), and **Canceled** ().
- The status or phase of the pod is indicated by different colors and tooltips as:
 - **Running** (): The pod is bound to a node and all of the containers are created. At least one container is still running or is in the process of starting or restarting.
 - **Not Ready** (): The pods which are running multiple containers, not all containers are ready.
 - **Warning** (): Containers in pods are being terminated, however termination did not succeed. Some containers may be other states.
 - **Failed** (): All containers in the pod terminated but least one container has terminated in failure. That is, the container either exited with non-zero status or was terminated by the system.
 - **Pending** (): The pod is accepted by the Kubernetes cluster, but one or more of the containers has not been set up and made ready to run. This includes time a pod spends waiting to be scheduled as well as the time spent downloading container images over the network.
 - **Succeeded** (): All containers in the pod terminated successfully and will not be restarted.

- **Terminating**(): When a pod is being deleted, it is shown as **Terminating** by some kubectl commands. **Terminating** status is not one of the pod phases. A pod is granted a graceful termination period, which defaults to 30 seconds.
- **Unknown**(): The state of the pod could not be obtained. This phase typically occurs due to an error in communicating with the node where the pod should be running.
- After you create an application and an image is deployed, the status is shown as **Pending**. After the application is built, it is displayed as **Running**.

Figure 4.1. Application topology



The application resource name is appended with indicators for the different types of resource objects as follows:

- **CJ: CronJob**
- **D: Deployment**
- **DC: DeploymentConfig**
- **DS: DaemonSet**
- **J: Job**
- **P: Pod**
- **SS: StatefulSet**
-  (Knative): A serverless application



NOTE

Serverless applications take some time to load and display on the **Graph view**. When you deploy a serverless application, it first creates a service resource and then a revision. After that, it is deployed and displayed on the **Graph view**. If it is the only workload, you might be redirected to the **Add** page. After the revision is deployed, the serverless application is displayed on the **Graph view**.

4.4. SCALING APPLICATION PODS AND CHECKING BUILDS AND ROUTES

The **Topology** view provides the details of the deployed components in the **Overview** panel. You can use the **Overview** and **Details** tabs to scale the application pods, check build status, services, and routes as follows:

- Click on the component node to see the **Overview** panel to the right. Use the **Details** tab to:
 - Scale your pods using the up and down arrows to increase or decrease the number of instances of the application manually. For serverless applications, the pods are automatically scaled down to zero when idle and scaled up depending on the channel traffic.
 - Check the **Labels**, **Annotations**, and **Status** of the application.
- Click the **Resources** tab to:
 - See the list of all the pods, view their status, access logs, and click on the pod to see the pod details.
 - See the builds, their status, access logs, and start a new build if needed.
 - See the services and routes used by the component.

For serverless applications, the **Resources** tab provides information on the revision, routes, and the configurations used for that component.

4.5. ADDING COMPONENTS TO AN EXISTING PROJECT

You can add components to a project.


Procedure

1. Navigate to the **+Add** view.




2. Click **Add to Project** () next to left navigation pane or press **Ctrl+Space**
3. Search for the component and click the **Start/Create/Install** button or click **Enter** to add the component to the project and see it in the topology **Graph view**.

Figure 4.2. Adding component via quick search



Alternatively, you can also use the available options in the context menu, such as **Import from Git**, **Container Image**, **Database**, **From Catalog**, **Operator Backed**, **Helm Charts**, **Samples**, or **Upload JAR file**, by right-clicking in the topology **Graph view** to add a component to your project.

Figure 4.3. Context menu to add services



4.6. GROUPING MULTIPLE COMPONENTS WITHIN AN APPLICATION

You can use the **+Add** view to add multiple components or services to your project and use the topology **Graph view** to group applications and resources within an application group.


Prerequisites

- You have created and deployed minimum two or more components on OpenShift Container Platform using the **Developer** perspective.

Procedure

- To add a service to the existing application group, press **Shift+** drag it to the existing application group. Dragging a component and adding it to an application group adds the required labels to the component.

Figure 4.4. Application grouping



Alternatively, you can also add the component to an application as follows:

- Click the service pod to see the **Overview** panel to the right.
- Click the **Actions** drop-down menu and select **Edit Application Grouping**.
- In the **Edit Application Grouping** dialog box, click the **Application** drop-down list, and select an appropriate application group.
- Click **Save** to add the service to the application group.

You can remove a component from an application group by selecting the component and using **Shift+** drag to drag it out of the application group.

4.7. ADDING SERVICES TO YOUR APPLICATION

To add a service to your application use the **+Add** actions using the context menu in the topology **Graph view**.




NOTE

In addition to the context menu, you can add services by using the sidebar or hovering and dragging the dangling arrow from the application group.

Procedure

1. Right-click an application group in the topology **Graph view** to display the context menu.

Figure 4.5. Add resource context menu



2. Use **Add to Application** to select a method for adding a service to the application group, such as **From Git**, **Container Image**, **From Dockerfile**, **From Devfile**, **Upload JAR file**, **Event Source**, **Channel**, or **Broker**.
3. Complete the form for the method you choose and click **Create**. For example, to add a service based on the source code in your Git repository, choose the **From Git** method, fill in the **Import from Git** form, and click **Create**.


4.8. REMOVING SERVICES FROM YOUR APPLICATION

In the topology **Graph view** remove a service from your application using the context menu.

Procedure

1. Right-click on a service in an application group in the topology **Graph view** to display the context menu.
2. Select **Delete Deployment** to delete the service.

Figure 4.6. Deleting deployment option



4.9. LABELS AND ANNOTATIONS USED FOR THE TOPOLOGY VIEW

The **Topology** view uses the following labels and annotations:

Icon displayed in the node

Icons in the node are defined by looking for matching icons using the **app.openshift.io/runtime** label, followed by the **app.kubernetes.io/name** label. This matching is done using a predefined set of icons.

Link to the source code editor or the source

The **app.openshift.io/vcs-uri** annotation is used to create links to the source code editor.

Node Connector

The **app.openshift.io/connects-to** annotation is used to connect the nodes.

App grouping

The **app.kubernetes.io/part-of=<appname>** label is used to group the applications, services, and components.

For detailed information on the labels and annotations OpenShift Container Platform applications must use, see [Guidelines for labels and annotations for OpenShift applications](#).

4.10. ADDITIONAL RESOURCES


- See [Importing a codebase from Git to create an application](#) for more information on creating an application from Git.
- See [Exporting applications](#).

CHAPTER 5. EXPORTING APPLICATIONS

As a developer, you can export your application in the ZIP file format. Based on your needs, import the exported application to another project in the same cluster or a different cluster by using the **Import YAML** option in the **+Add** view. Exporting your application helps you to reuse your application resources and saves your time.

5.1. PREREQUISITES

- You have installed the gitops-primer Operator from the OperatorHub.




NOTE

The **Export application** option is disabled in the **Topology** view even after installing the gitops-primer Operator.

- You have created an application in the **Topology** view to enable **Export application**.

5.2. PROCEDURE

- In the developer perspective, perform one of the following steps:
 - Navigate to the **+Add** view and click **Export application** in the **Application portability** tile.
 - Navigate to the **Topology** view and click **Export application**.
- Click **OK** in the **Export Application** dialog box. A notification opens to confirm that the export of resources from your project has started.
- Optional steps that you might need to perform in the following scenarios:
 - If you have started exporting an incorrect application, click **Export application → Cancel Export**.
 - If your export is already in progress and you want to start a fresh export, click **Export application → Restart Export**.
 - If you want to view logs associated with exporting an application, click **Export application** and the **View Logs** link.



4. After a successful export, click **Download** in the dialog box to download application resources in ZIP format onto your machine.

CHAPTER 6. WORKING WITH HELM CHARTS

6.1. UNDERSTANDING HELM

Helm is a software package manager that simplifies deployment of applications and services to OpenShift Container Platform clusters.

Helm uses a packaging format called *charts*. A Helm chart is a collection of files that describes the OpenShift Container Platform resources.

Creating a chart in a cluster creates a running instance of the chart known as a *release*.

Each time a chart is created, or a release is upgraded or rolled back, an incremental revision is created.

6.1.1. Key features

Helm provides the ability to:

- Search through a large collection of charts stored in the chart repository.
- Modify existing charts.
- Create your own charts with OpenShift Container Platform or Kubernetes resources.
- Package and share your applications as charts.

6.1.2. Red Hat Certification of Helm charts for OpenShift

You can choose to verify and certify your Helm charts by Red Hat for all the components you will be deploying on the Red Hat OpenShift Container Platform. Charts go through an automated Red Hat OpenShift certification workflow that guarantees security compliance as well as best integration and experience with the platform. Certification assures the integrity of the chart and ensures that the Helm chart works seamlessly on Red Hat OpenShift clusters.

6.1.3. Additional resources

- For more information on how to certify your Helm charts as a Red Hat partner, see [Red Hat Certification of Helm charts for OpenShift](#).
- For more information on OpenShift and Container certification guides for Red Hat partners, see [Partner Guide for OpenShift and Container Certification](#).
- For a list of the charts, see [the Red Hat Helm index file](#).

6.2. INSTALLING HELM

The following section describes how to install Helm on different platforms using the CLI.

You can also find the URL to the latest binaries from the OpenShift Container Platform web console by clicking the ? icon in the upper-right corner and selecting **Command Line Tools**.

Prerequisites

- You have installed Go, version 1.13 or higher.

6.2.1. On Linux

1. Download the Helm binary and add it to your path:

- Linux (x86_64, amd64)

```
# curl -L https://mirror.openshift.com/pub/openshift-v4/clients/helm/latest/helm-linux-amd64 -o /usr/local/bin/helm
```

- Linux on IBM Z® and IBM® LinuxONE (s390x)

```
# curl -L https://mirror.openshift.com/pub/openshift-v4/clients/helm/latest/helm-linux-s390x -o /usr/local/bin/helm
```

- Linux on IBM Power® (ppc64le)

```
# curl -L https://mirror.openshift.com/pub/openshift-v4/clients/helm/latest/helm-linux-ppc64le -o /usr/local/bin/helm
```

2. Make the binary file executable:

```
# chmod +x /usr/local/bin/helm
```

3. Check the installed version:

```
$ helm version
```

Example output

```
version.BuildInfo{Version:"v3.0",
GitCommit:"b31719aab7963acf4887a1c1e6d5e53378e34d93", GitTreeState:"clean",
GoVersion:"go1.13.4"}
```

6.2.2. On Windows 7/8

1. Download the latest [.exe file](#) and put in a directory of your preference.
2. Right click **Start** and click **Control Panel**.
3. Select **System and Security** and then click **System**.
4. From the menu on the left, select **Advanced systems settings** and click **Environment Variables** at the bottom.
5. Select **Path** from the **Variable** section and click **Edit**.
6. Click **New** and type the path to the folder with the **.exe** file into the field or click **Browse** and select the directory, and click **OK**.

6.2.3. On Windows 10

1. Download the latest [.exe file](#) and put in a directory of your preference.

2. Click **Search** and type **env** or **environment**.
3. Select **Edit environment variables for your account**
4. Select **Path** from the **Variable** section and click **Edit**.
5. Click **New** and type the path to the directory with the exe file into the field or click **Browse** and select the directory, and click **OK**.

6.2.4. On MacOS

1. Download the Helm binary and add it to your path:

```
# curl -L https://mirror.openshift.com/pub/openshift-v4/clients/helm/latest/helm-darwin-amd64  
-o /usr/local/bin/helm
```

2. Make the binary file executable:

```
# chmod +x /usr/local/bin/helm
```

3. Check the installed version:

```
$ helm version
```

Example output

```
version.BuildInfo{Version:"v3.0",  
GitCommit:"b31719aab7963acf4887a1c1e6d5e53378e34d93", GitTreeState:"clean",  
GoVersion:"go1.13.4"}
```

6.3. CONFIGURING CUSTOM HELM CHART REPOSITORIES

You can create Helm releases on an OpenShift Container Platform cluster using the following methods:

- The CLI.
- The **Developer** perspective of the web console.

The **Developer Catalog**, in the **Developer** perspective of the web console, displays the Helm charts available in the cluster. By default, it lists the Helm charts from the Red Hat OpenShift Helm chart repository. For a list of the charts, see [the Red Hat Helm index file](#).

As a cluster administrator, you can add multiple cluster-scoped and namespace-scoped Helm chart repositories, separate from the default cluster-scoped Helm repository, and display the Helm charts from these repositories in the **Developer Catalog**.

As a regular user or project member with the appropriate role-based access control (RBAC) permissions, you can add multiple namespace-scoped Helm chart repositories, apart from the default cluster-scoped Helm repository, and display the Helm charts from these repositories in the **Developer Catalog**.

In the **Developer** perspective of the web console, you can use the **Helm** page to:

- Create Helm Releases and Repositories using the **Create** button.

- Create, update, or delete a cluster-scoped or namespace-scoped Helm chart repository.
- View the list of the existing Helm chart repositories in the Repositories tab, which can also be easily distinguished as either cluster scoped or namespace scoped.

6.3.1. Installing a Helm chart on an OpenShift Container Platform cluster

Prerequisites

- You have a running OpenShift Container Platform cluster and you have logged into it.
- You have installed Helm.

Procedure

1. Create a new project:

```
$ oc new-project vault
```

2. Add a repository of Helm charts to your local Helm client:

```
$ helm repo add openshift-helm-charts https://charts.openshift.io/
```

Example output

```
"openshift-helm-charts" has been added to your repositories
```

3. Update the repository:

```
$ helm repo update
```

4. Install an example HashiCorp Vault:

```
$ helm install example-vault openshift-helm-charts/hashicorp-vault
```

Example output

```
NAME: example-vault
LAST DEPLOYED: Fri Mar 11 12:02:12 2022
NAMESPACE: vault
STATUS: deployed
REVISION: 1
NOTES:
Thank you for installing HashiCorp Vault!
```

5. Verify that the chart has installed successfully:

```
$ helm list
```

Example output

NAME	NAMESPACE	REVISION	UPDATED	STATUS	CHART
------	-----------	----------	---------	--------	-------

APP VERSION

example-vault vault 1 2022-03-11 12:02:12.296226673 +0530 IST deployed vault-0.19.0 1.9.2

6.3.2. Creating Helm releases using the Developer perspective

You can use either the **Developer** perspective in the web console or the CLI to select and create a release from the Helm charts listed in the **Developer Catalog**. You can create Helm releases by installing Helm charts and see them in the **Developer** perspective of the web console.

Prerequisites

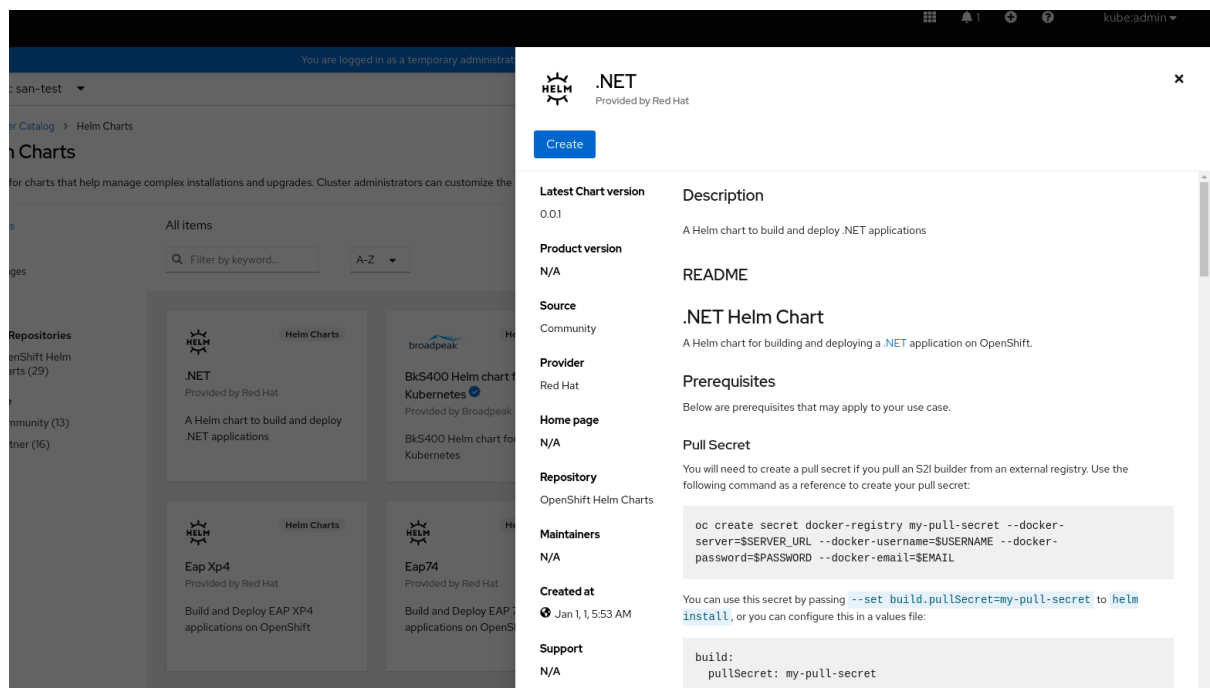
- You have logged in to the web console and have switched to [the Developer perspective](#).

Procedure

To create Helm releases from the Helm charts provided in the **Developer Catalog**:

- In the **Developer** perspective, navigate to the **+Add** view and select a project. Then click **Helm Chart** option to see all the Helm Charts in the **Developer Catalog**.
- Select a chart and read the description, README, and other details about the chart.
- Click **Create**.

Figure 6.1. Helm charts in developer catalog



- In the **Create Helm Release** page:
 - Enter a unique name for the release in the **Release Name** field.
 - Select the required chart version from the **Chart Version** drop-down list.
 - Configure your Helm chart by using the **Form View** or the **YAML View**.

**NOTE**

Where available, you can switch between the **YAML View** and **Form View**. The data is persisted when switching between the views.

- d. Click **Create** to create a Helm release. The web console displays the new release in the **Topology** view.

If a Helm chart has release notes, the web console displays them.

If a Helm chart creates workloads, the web console displays them on the **Topology** or **Helm release details** page. The workloads are **DaemonSet**, **CronJob**, **Pod**, **Deployment**, and **DeploymentConfig**.

- e. View the newly created Helm release in the **Helm Releases** page.

You can upgrade, rollback, or delete a Helm release by using the **Actions** button on the side panel or by right-clicking a Helm release.

6.3.3. Using Helm in the web terminal

You can use Helm by [Accessing the web terminal](#) in the **Developer** perspective of the web console.

6.3.4. Creating a custom Helm chart on OpenShift Container Platform

Procedure

1. Create a new project:

```
$ oc new-project nodejs-ex-k
```

2. Download an example Node.js chart that contains OpenShift Container Platform objects:

```
$ git clone https://github.com/redhat-developer/redhat-helm-charts
```

3. Go to the directory with the sample chart:

```
$ cd redhat-helm-charts/alpha/nodejs-ex-k/
```

4. Edit the **Chart.yaml** file and add a description of your chart:

```
apiVersion: v2 1
name: nodejs-ex-k 2
description: A Helm chart for OpenShift 3
icon: https://static.redhat.com/libs/redhat/brand-assets/latest/corp/logo.svg 4
version: 0.2.1 5
```

1 The chart API version. It should be **v2** for Helm charts that require at least Helm 3.

2 The name of your chart.

3 The description of your chart.

- 4 The URL to an image to be used as an icon.
- 5 The Version of your chart as per the Semantic Versioning (SemVer) 2.0.0 Specification.

5. Verify that the chart is formatted properly:

```
$ helm lint
```

Example output

```
[INFO] Chart.yaml: icon is recommended

1 chart(s) linted, 0 chart(s) failed
```

6. Navigate to the previous directory level:

```
$ cd ..
```

7. Install the chart:

```
$ helm install nodejs-chart nodejs-ex-k
```

8. Verify that the chart has installed successfully:

```
$ helm list
```

Example output

```
NAME      NAMESPACE REVISION UPDATED      STATUS  CHART          APP VERSION
nodejs-chart nodejs-ex-k 1      2019-12-05 15:06:51.379134163 -0500 EST deployed nodejs-0.1.0 1.16.0
```

6.3.5. Adding custom Helm chart repositories

As a cluster administrator, you can add custom Helm chart repositories to your cluster and enable access to the Helm charts from these repositories in the **Developer Catalog**.

Procedure

1. To add a new Helm Chart Repository, you must add the Helm Chart Repository custom resource (CR) to your cluster.

Sample Helm Chart Repository CR

```
apiVersion: helm.openshift.io/v1beta1
kind: HelmChartRepository
metadata:
  name: <name>
spec:
  # optional name that might be used by console
  # name: <chart-display-name>
  connectionConfig:
```

```
url: <helm-chart-repository-url>
```


For example, to add an Azure sample chart repository, run:

```
$ cat <<EOF | oc apply -f -
apiVersion: helm.openshift.io/v1beta1
kind: HelmChartRepository
metadata:
  name: azure-sample-repo
spec:
  name: azure-sample-repo
  connectionConfig:
    url: https://raw.githubusercontent.com/Azure-Samples/helm-charts/master/docs
EOF
```

2. Navigate to the **Developer Catalog** in the web console to verify that the Helm charts from the chart repository are displayed.

For example, use the **Chart repositories** filter to search for a Helm chart from the repository.

Figure 6.2. Chart repositories filter



NOTE

If a cluster administrator removes all of the chart repositories, then you cannot view the Helm option in the **+Add** view, **Developer Catalog**, and left navigation panel.

6.3.6. Adding namespace-scoped custom Helm chart repositories

The cluster-scoped **HelmChartRepository** custom resource definition (CRD) for Helm repository provides the ability for administrators to add Helm repositories as custom resources. The namespace-scoped **ProjectHelmChartRepository** CRD allows project members with the appropriate role-based access control (RBAC) permissions to create Helm repository resources of their choice but scoped to their namespace. Such project members can see charts from both cluster-scoped and namespace-scoped Helm repository resources.



NOTE

- Administrators can limit users from creating namespace-scoped Helm repository resources. By limiting users, administrators have the flexibility to control the RBAC through a namespace role instead of a cluster role. This avoids unnecessary permission elevation for the user and prevents access to unauthorized services or applications.
- The addition of the namespace-scoped Helm repository does not impact the behavior of the existing cluster-scoped Helm repository.

As a regular user or project member with the appropriate RBAC permissions, you can add custom namespace-scoped Helm chart repositories to your cluster and enable access to the Helm charts from these repositories in the **Developer Catalog**.

Procedure

1. To add a new namespace-scoped Helm Chart Repository, you must add the Helm Chart Repository custom resource (CR) to your namespace.

Sample Namespace-scoped Helm Chart Repository CR

```
apiVersion: helm.openshift.io/v1beta1
kind: ProjectHelmChartRepository
metadata:
  name: <name>
spec:
  url: https://my.chart-repo.org/stable

# optional name that might be used by console
name: <chart-repo-display-name>

# optional and only needed for UI purposes
description: <My private chart repo>

# required: chart repository URL
connectionConfig:
  url: <helm-chart-repository-url>
```

For example, to add an Azure sample chart repository scoped to your **my-namespace** namespace, run:

```
$ cat <<EOF | oc apply --namespace my-namespace -f -
apiVersion: helm.openshift.io/v1beta1
kind: ProjectHelmChartRepository
metadata:
  name: azure-sample-repo
spec:
  name: azure-sample-repo
  connectionConfig:
    url: https://raw.githubusercontent.com/Azure-Samples/helm-charts/master/docs
EOF
```


The output verifies that the namespace-scoped Helm Chart Repository CR is created:

Example output

```
projecthelmchartrepository.helm.openshift.io/azure-sample-repo created
```

2. Navigate to the **Developer Catalog** in the web console to verify that the Helm charts from the chart repository are displayed in your **my-namespace** namespace.
For example, use the **Chart repositories** filter to search for a Helm chart from the repository.

Figure 6.3. Chart repositories filter in your namespace



Alternatively, run:

```
$ oc get projecthelmchartrepositories --namespace my-namespace
```

Example output

```
NAME          AGE
azure-sample-repo 1m
```



NOTE

If a cluster administrator or a regular user with appropriate RBAC permissions removes all of the chart repositories in a specific namespace, then you cannot view the Helm option in the **+Add** view, **Developer Catalog**, and left navigation panel for that specific namespace.

6.3.7. Creating credentials and CA certificates to add Helm chart repositories

Some Helm chart repositories need credentials and custom certificate authority (CA) certificates to connect to it. You can use the web console as well as the CLI to add credentials and certificates.

Procedure

To configure the credentials and certificates, and then add a Helm chart repository using the CLI:

1. In the **openshift-config** namespace, create a **ConfigMap** object with a custom CA certificate in PEM encoded format, and store it under the **ca-bundle.crt** key within the config map:

```
■
```

```
$ oc create configmap helm-ca-cert \
--from-file=ca-bundle.crt=/path/to/certs/ca.crt \
-n openshift-config
```

2. In the **openshift-config** namespace, create a **Secret** object to add the client TLS configurations:

```
$ oc create secret tls helm-tls-configs \
--cert=/path/to/certs/client.crt \
--key=/path/to/certs/client.key \
-n openshift-config
```

Note that the client certificate and key must be in PEM encoded format and stored under the keys **tls.crt** and **tls.key**, respectively.

3. Add the Helm repository as follows:

```
$ cat <<EOF | oc apply -f -
apiVersion: helm.openshift.io/v1beta1
kind: HelmChartRepository
metadata:
  name: <helm-repository>
spec:
  name: <helm-repository>
  connectionConfig:
    url: <URL for the Helm repository>
    tlsConfig:
      name: helm-tls-configs
    ca:
      name: helm-ca-cert
EOF
```

The **ConfigMap** and **Secret** are consumed in the HelmChartRepository CR using the **tlsConfig** and **ca** fields. These certificates are used to connect to the Helm repository URL.

4. By default, all authenticated users have access to all configured charts. However, for chart repositories where certificates are needed, you must provide users with read access to the **helm-ca-cert** config map and **helm-tls-configs** secret in the **openshift-config** namespace, as follows:

```
$ cat <<EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  namespace: openshift-config
  name: helm-chartrepos-tls-conf-viewer
rules:
- apiGroups: [""]
  resources: ["configmaps"]
  resourceNames: ["helm-ca-cert"]
  verbs: ["get"]
- apiGroups: [""]
  resources: ["secrets"]
  resourceNames: ["helm-tls-configs"]
  verbs: ["get"]
```

```

---
kind: RoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  namespace: openshift-config
  name: helm-chartrepos-tls-conf-viewer
subjects:
  - kind: Group
    apiGroup: rbac.authorization.k8s.io
    name: 'system:authenticated'
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: helm-chartrepos-tls-conf-viewer
EOF

```

6.3.8. Filtering Helm Charts by their certification level

You can filter Helm charts based on their certification level in the **Developer Catalog**.

Procedure

1. In the **Developer** perspective, navigate to the **+Add** view and select a project.
2. From the **Developer Catalog** tile, select the **Helm Chart** option to see all the Helm charts in the **Developer Catalog**.
3. Use the filters to the left of the list of Helm charts to filter the required charts:
 - Use the **Chart Repositories** filter to filter charts provided by **Red Hat Certification Charts** or **OpenShift Helm Charts**.
 - Use the **Source** filter to filter charts sourced from **Partners**, **Community**, or **Red Hat**. Certified charts are indicated with the () icon.



NOTE

The **Source** filter will not be visible when there is only one provider type.

You can now select the required chart and install it.

6.3.9. Disabling Helm Chart repositories

You can disable Helm Charts from a particular Helm Chart Repository in the catalog by setting the **disabled** property in the **HelmChartRepository** custom resource to **true**.

Procedure

- To disable a Helm Chart repository by using CLI, add the **disabled: true** flag to the custom resource. For example, to remove an Azure sample chart repository, run:

```

$ cat <<EOF | oc apply -f -
apiVersion: helm.openshift.io/v1beta1
kind: HelmChartRepository

```

```

metadata:
  name: azure-sample-repo
spec:
  connectionConfig:
    url:https://raw.githubusercontent.com/Azure-Samples/helm-charts/master/docs
  disabled: true
EOF

```

- To disable a recently added Helm Chart repository by using Web Console:
 1. Go to **Custom Resource Definitions** and search for the **HelmChartRepository** custom resource.
 2. Go to **Instances**, find the repository you want to disable, and click its name.
 3. Go to the **YAML** tab, add the **disabled: true** flag in the **spec** section, and click **Save**.

Example

```

spec:
  connectionConfig:
    url: <url-of-the-repositoru-to-be-disabled>
  disabled: true

```

The repository is now disabled and will not appear in the catalog.

6.4. WORKING WITH HELM RELEASES

You can use the **Developer** perspective in the web console to update, rollback, or delete a Helm release.

6.4.1. Prerequisites

- You have logged in to the web console and have switched to [the Developer perspective](#).

6.4.2. Upgrading a Helm release

You can upgrade a Helm release to upgrade to a new chart version or update your release configuration.

Procedure

1. In the **Topology** view, select the Helm release to see the side panel.
2. Click **Actions → Upgrade Helm Release**.
3. In the **Upgrade Helm Release** page, select the **Chart Version** you want to upgrade to, and then click **Upgrade** to create another Helm release. The **Helm Releases** page displays the two revisions.

6.4.3. Rolling back a Helm release

If a release fails, you can rollback the Helm release to a previous version.

Procedure

To rollback a release using the **Helm** view:

1. In the **Developer** perspective, navigate to the **Helm** view to see the **Helm Releases** in the namespace.
2. Click the Options menu  adjoining the listed release, and select **Rollback**.
3. In the **Rollback Helm Release** page, select the **Revision** you want to rollback to and click **Rollback**.
4. In the **Helm Releases** page, click on the chart to see the details and resources for that release.
5. Go to the **Revision History** tab to see all the revisions for the chart.

Figure 6.4. Helm revision history

Helm Releases > Helm Release Details

 elasticsearch Deployed Actions

Revision	Updated	Status	Chart Name	Chart Version	App Version	Description
1	4 minutes ago	Superseded	elasticsearch	7.6.0	7.6.0	Install complete
2	3 minutes ago	Superseded	elasticsearch	7.6.2	7.6.2	Upgrade complete
3	less than a minute ago	Deployed	elasticsearch	7.6.2	7.6.2	Rollback to 2

6. If required, you can further use the Options menu  adjoining a particular revision and select the revision to rollback to.

6.4.4. Deleting a Helm release

Procedure

1. In the **Topology** view, right-click the Helm release and select **Delete Helm Release**.
2. In the confirmation prompt, enter the name of the chart and click **Delete**.

CHAPTER 7. DEPLOYMENTS

7.1. UNDERSTANDING DEPLOYMENTS

The **Deployment** and **DeploymentConfig** API objects in OpenShift Container Platform provide two similar but different methods for fine-grained management over common user applications. They are composed of the following separate API objects:

- A **Deployment** or **DeploymentConfig** object, either of which describes the desired state of a particular component of the application as a pod template.
- **Deployment** objects involve one or more *replica sets*, which contain a point-in-time record of the state of a deployment as a pod template. Similarly, **DeploymentConfig** objects involve one or more *replication controllers*, which preceded replica sets.
- One or more pods, which represent an instance of a particular version of an application.

Use **Deployment** objects unless you need a specific feature or behavior provided by **DeploymentConfig** objects.



IMPORTANT

As of OpenShift Container Platform 4.14, **DeploymentConfig** objects are deprecated. **DeploymentConfig** objects are still supported, but are not recommended for new installations. Only security-related and critical issues will be fixed.

Instead, use **Deployment** objects or another alternative to provide declarative updates for pods.

7.1.1. Building blocks of a deployment

Deployments and deployment configs are enabled by the use of native Kubernetes API objects **ReplicaSet** and **ReplicationController**, respectively, as their building blocks.

Users do not have to manipulate replica sets, replication controllers, or pods owned by **Deployment** or **DeploymentConfig** objects. The deployment systems ensure changes are propagated appropriately.

TIP

If the existing deployment strategies are not suited for your use case and you must run manual steps during the lifecycle of your deployment, then you should consider creating a custom deployment strategy.

The following sections provide further details on these objects.

7.1.1.1. Replica sets

A **ReplicaSet** is a native Kubernetes API object that ensures a specified number of pod replicas are running at any given time.



NOTE

Only use replica sets if you require custom update orchestration or do not require updates at all. Otherwise, use deployments. Replica sets can be used independently, but are used by deployments to orchestrate pod creation, deletion, and updates. Deployments manage their replica sets automatically, provide declarative updates to pods, and do not have to manually manage the replica sets that they create.

The following is an example **ReplicaSet** definition:

```
apiVersion: apps/v1
kind: ReplicaSet
metadata:
  name: frontend-1
  labels:
    tier: frontend
spec:
  replicas: 3
  selector: 1
    matchLabels: 2
      tier: frontend
    matchExpressions: 3
      - {key: tier, operator: In, values: [frontend]}
  template:
    metadata:
      labels:
        tier: frontend
    spec:
      containers:
        - image: openshift/hello-openshift
          name: helloworld
          ports:
            - containerPort: 8080
              protocol: TCP
      restartPolicy: Always
```

- 1 A label query over a set of resources. The result of **matchLabels** and **matchExpressions** are logically conjoined.
- 2 Equality-based selector to specify resources with labels that match the selector.
- 3 Set-based selector to filter keys. This selects all resources with key equal to **tier** and value equal to **frontend**.

7.1.1.2. Replication controllers

Similar to a replica set, a replication controller ensures that a specified number of replicas of a pod are running at all times. If pods exit or are deleted, the replication controller instantiates more up to the defined number. Likewise, if there are more running than desired, it deletes as many as necessary to match the defined amount. The difference between a replica set and a replication controller is that a replica set supports set-based selector requirements whereas a replication controller only supports equality-based selector requirements.

A replication controller configuration consists of:

- The number of replicas desired, which can be adjusted at run time.
- A **Pod** definition to use when creating a replicated pod.
- A selector for identifying managed pods.

A selector is a set of labels assigned to the pods that are managed by the replication controller. These labels are included in the **Pod** definition that the replication controller instantiates. The replication controller uses the selector to determine how many instances of the pod are already running in order to adjust as needed.

The replication controller does not perform auto-scaling based on load or traffic, as it does not track either. Rather, this requires its replica count to be adjusted by an external auto-scaler.



NOTE

Use a **DeploymentConfig** to create a replication controller instead of creating replication controllers directly.

If you require custom orchestration or do not require updates, use replica sets instead of replication controllers.

The following is an example definition of a replication controller:

```
apiVersion: v1
kind: ReplicationController
metadata:
  name: frontend-1
spec:
  replicas: 1 ①
  selector: ②
    name: frontend
  template: ③
    metadata:
      labels: ④
        name: frontend ⑤
    spec:
      containers:
        - image: openshift/hello-openshift
          name: helloworld
          ports:
            - containerPort: 8080
              protocol: TCP
          restartPolicy: Always
```

- ① The number of copies of the pod to run.
- ② The label selector of the pod to run.
- ③ A template for the pod the controller creates.
- ④ Labels on the pod should include those from the label selector.
- ⑤ The maximum name length after expanding any parameters is 63 characters.

7.1.2. Deployments


Kubernetes provides a first-class, native API object type in OpenShift Container Platform called **Deployment**. **Deployment** objects describe the desired state of a particular component of an application as a pod template. Deployments create replica sets, which orchestrate pod lifecycles.

For example, the following deployment definition creates a replica set to bring up one **hello-openshift** pod:

Deployment definition

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: hello-openshift
spec:
  replicas: 1
  selector:
    matchLabels:
      app: hello-openshift
  template:
    metadata:
      labels:
        app: hello-openshift
    spec:
      containers:
        - name: hello-openshift
          image: openshift/hello-openshift:latest
          ports:
            - containerPort: 80
```

7.1.3. DeploymentConfig objects



IMPORTANT

As of OpenShift Container Platform 4.14, **DeploymentConfig** objects are deprecated. **DeploymentConfig** objects are still supported, but are not recommended for new installations. Only security-related and critical issues will be fixed.

Instead, use **Deployment** objects or another alternative to provide declarative updates for pods.

Building on replication controllers, OpenShift Container Platform adds expanded support for the software development and deployment lifecycle with the concept of **DeploymentConfig** objects. In the simplest case, a **DeploymentConfig** object creates a new replication controller and lets it start up pods.

However, OpenShift Container Platform deployments from **DeploymentConfig** objects also provide the ability to transition from an existing deployment of an image to a new one and also define hooks to be run before or after creating the replication controller.

The **DeploymentConfig** deployment system provides the following capabilities:

- A **DeploymentConfig** object, which is a template for running applications.

- Triggers that drive automated deployments in response to events.
- User-customizable deployment strategies to transition from the previous version to the new version. A strategy runs inside a pod commonly referred as the deployment process.
- A set of hooks (lifecycle hooks) for executing custom behavior in different points during the lifecycle of a deployment.
- Versioning of your application to support rollbacks either manually or automatically in case of deployment failure.
- Manual replication scaling and autoscaling.

When you create a **DeploymentConfig** object, a replication controller is created representing the **DeploymentConfig** object's pod template. If the deployment changes, a new replication controller is created with the latest pod template, and a deployment process runs to scale down the old replication controller and scale up the new one.

Instances of your application are automatically added and removed from both service load balancers and routers as they are created. As long as your application supports graceful shutdown when it receives the **TERM** signal, you can ensure that running user connections are given a chance to complete normally.

The OpenShift Container Platform **DeploymentConfig** object defines the following details:

1. The elements of a **ReplicationController** definition.
2. Triggers for creating a new deployment automatically.
3. The strategy for transitioning between deployments.
4. Lifecycle hooks.

Each time a deployment is triggered, whether manually or automatically, a deployer pod manages the deployment (including scaling down the old replication controller, scaling up the new one, and running hooks). The deployment pod remains for an indefinite amount of time after it completes the deployment to retain its logs of the deployment. When a deployment is superseded by another, the previous replication controller is retained to enable easy rollback if needed.

Example DeploymentConfig definition

```
apiVersion: apps.openshift.io/v1
kind: DeploymentConfig
metadata:
  name: frontend
spec:
  replicas: 5
  selector:
    name: frontend
  template: { ... }
  triggers:
    - type: ConfigChange 1
    - imageChangeParams:
        automatic: true
        containerNames:
          - helloworld
```

```

from:
  kind: ImageStreamTag
  name: hello-openshift:latest
type: ImageChange ❷
strategy:
  type: Rolling ❸

```

- ❶ A configuration change trigger results in a new replication controller whenever changes are detected in the pod template of the deployment configuration.
- ❷ An image change trigger causes a new deployment to be created each time a new version of the backing image is available in the named image stream.
- ❸ The default **Rolling** strategy makes a downtime-free transition between deployments.

7.1.4. Comparing Deployment and DeploymentConfig objects

Both Kubernetes **Deployment** objects and OpenShift Container Platform-provided **DeploymentConfig** objects are supported in OpenShift Container Platform; however, it is recommended to use **Deployment** objects unless you need a specific feature or behavior provided by **DeploymentConfig** objects.

The following sections go into more detail on the differences between the two object types to further help you decide which type to use.



IMPORTANT

As of OpenShift Container Platform 4.14, **DeploymentConfig** objects are deprecated. **DeploymentConfig** objects are still supported, but are not recommended for new installations. Only security-related and critical issues will be fixed.

Instead, use **Deployment** objects or another alternative to provide declarative updates for pods.

7.1.4.1. Design

One important difference between **Deployment** and **DeploymentConfig** objects is the properties of the [CAP theorem](#) that each design has chosen for the rollout process. **DeploymentConfig** objects prefer consistency, whereas **Deployments** objects take availability over consistency.

For **DeploymentConfig** objects, if a node running a deployer pod goes down, it will not get replaced. The process waits until the node comes back online or is manually deleted. Manually deleting the node also deletes the corresponding pod. This means that you can not delete the pod to unstick the rollout, as the kubelet is responsible for deleting the associated pod.

However, deployment rollouts are driven from a controller manager. The controller manager runs in high availability mode on masters and uses leader election algorithms to value availability over consistency. During a failure it is possible for other masters to act on the same deployment at the same time, but this issue will be reconciled shortly after the failure occurs.

7.1.4.2. Deployment-specific features

7.1.4.2.1. Rollover

The deployment process for **Deployment** objects is driven by a controller loop, in contrast to **DeploymentConfig** objects that use deployer pods for every new rollout. This means that the **Deployment** object can have as many active replica sets as possible, and eventually the deployment controller will scale down all old replica sets and scale up the newest one.

DeploymentConfig objects can have at most one deployer pod running, otherwise multiple deployers might conflict when trying to scale up what they think should be the newest replication controller. Because of this, only two replication controllers can be active at any point in time. Ultimately, this results in faster rapid rollouts for **Deployment** objects.

7.1.4.2.2. Proportional scaling

Because the deployment controller is the sole source of truth for the sizes of new and old replica sets owned by a **Deployment** object, it can scale ongoing rollouts. Additional replicas are distributed proportionally based on the size of each replica set.

DeploymentConfig objects cannot be scaled when a rollout is ongoing because the controller will have issues with the deployer process about the size of the new replication controller.

7.1.4.2.3. Pausing mid-rollout

Deployments can be paused at any point in time, meaning you can also pause ongoing rollouts. However, you currently cannot pause deployer pods; if you try to pause a deployment in the middle of a rollout, the deployer process is not affected and continues until it finishes.

7.1.4.3. DeploymentConfig object-specific features

7.1.4.3.1. Automatic rollbacks

Currently, deployments do not support automatically rolling back to the last successfully deployed replica set in case of a failure.

7.1.4.3.2. Triggers

Deployments have an implicit config change trigger in that every change in the pod template of a deployment automatically triggers a new rollout. If you do not want new rollouts on pod template changes, pause the deployment:

```
$ oc rollout pause deployments/<name>
```

7.1.4.3.3. Lifecycle hooks


Deployments do not yet support any lifecycle hooks.

7.1.4.3.4. Custom strategies

Deployments do not support user-specified custom deployment strategies.

7.2. MANAGING DEPLOYMENT PROCESSES

7.2.1. Managing DeploymentConfig objects



IMPORTANT

As of OpenShift Container Platform 4.14, **DeploymentConfig** objects are deprecated. **DeploymentConfig** objects are still supported, but are not recommended for new installations. Only security-related and critical issues will be fixed.

Instead, use **Deployment** objects or another alternative to provide declarative updates for pods.

DeploymentConfig objects can be managed from the OpenShift Container Platform web console's **Workloads** page or using the **oc** CLI. The following procedures show CLI usage unless otherwise stated.

7.2.1.1. Starting a deployment

You can start a rollout to begin the deployment process of your application.

Procedure

1. To start a new deployment process from an existing **DeploymentConfig** object, run the following command:

```
$ oc rollout latest dc/<name>
```



NOTE

If a deployment process is already in progress, the command displays a message and a new replication controller will not be deployed.

7.2.1.2. Viewing a deployment

You can view a deployment to get basic information about all the available revisions of your application.

Procedure

1. To show details about all recently created replication controllers for the provided **DeploymentConfig** object, including any currently running deployment process, run the following command:

```
$ oc rollout history dc/<name>
```

2. To view details specific to a revision, add the **--revision** flag:

```
$ oc rollout history dc/<name> --revision=1
```

3. For more detailed information about a **DeploymentConfig** object and its latest revision, use the **oc describe** command:

```
$ oc describe dc <name>
```

7.2.1.3. Retrying a deployment

If the current revision of your **DeploymentConfig** object failed to deploy, you can restart the deployment process.

Procedure

1. To restart a failed deployment process:

```
$ oc rollout retry dc/<name>
```

If the latest revision of it was deployed successfully, the command displays a message and the deployment process is not retried.



NOTE

Retrying a deployment restarts the deployment process and does not create a new deployment revision. The restarted replication controller has the same configuration it had when it failed.

7.2.1.4. Rolling back a deployment

Rollbacks revert an application back to a previous revision and can be performed using the REST API, the CLI, or the web console.

Procedure

1. To rollback to the last successful deployed revision of your configuration:

```
$ oc rollout undo dc/<name>
```

The **DeploymentConfig** object's template is reverted to match the deployment revision specified in the undo command, and a new replication controller is started. If no revision is specified with **--to-revision**, then the last successfully deployed revision is used.

2. Image change triggers on the **DeploymentConfig** object are disabled as part of the rollback to prevent accidentally starting a new deployment process soon after the rollback is complete. To re-enable the image change triggers:

```
$ oc set triggers dc/<name> --auto
```



NOTE

Deployment configs also support automatically rolling back to the last successful revision of the configuration in case the latest deployment process fails. In that case, the latest template that failed to deploy stays intact by the system and it is up to users to fix their configurations.

7.2.1.5. Executing commands inside a container

You can add a command to a container, which modifies the container's startup behavior by overruling the image's **ENTRYPOINT**. This is different from a lifecycle hook, which instead can be run once per deployment at a specified time.

Procedure

1. Add the **command** parameters to the **spec** field of the **DeploymentConfig** object. You can also add an **args** field, which modifies the **command** (or the **ENTRYPOINT** if **command** does not exist).

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
  template:
# ...
    spec:
      containers:
      - name: <container_name>
        image: 'image'
        command:
        - '<command>'
        args:
        - '<argument_1>'
        - '<argument_2>'
        - '<argument_3>'
```

For example, to execute the **java** command with the **-jar** and **/opt/app-root/springboots2idemo.jar** arguments:

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
  template:
# ...
    spec:
      containers:
      - name: example-spring-boot
        image: 'image'
        command:
        - java
        args:
        - '-jar'
        - /opt/app-root/springboots2idemo.jar
# ...
```

7.2.1.6. Viewing deployment logs

Procedure

1. To stream the logs of the latest revision for a given **DeploymentConfig** object:

```
$ oc logs -f dc/<name>
```

If the latest revision is running or failed, the command returns the logs of the process that is responsible for deploying your pods. If it is successful, it returns the logs from a pod of your application.

2. You can also view logs from older failed deployment processes, if and only if these processes (old replication controllers and their deployer pods) exist and have not been pruned or deleted manually:

```
$ oc logs --version=1 dc/<name>
```

7.2.1.7. Deployment triggers

A **DeploymentConfig** object can contain triggers, which drive the creation of new deployment processes in response to events inside the cluster.



WARNING

If no triggers are defined on a **DeploymentConfig** object, a config change trigger is added by default. If triggers are defined as an empty field, deployments must be started manually.

7.2.1.7.1. Config change deployment triggers

The config change trigger results in a new replication controller whenever configuration changes are detected in the pod template of the **DeploymentConfig** object.



NOTE

If a config change trigger is defined on a **DeploymentConfig** object, the first replication controller is automatically created soon after the **DeploymentConfig** object itself is created and it is not paused.

Config change deployment trigger

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
# ...
triggers:
  - type: "ConfigChange"
```

7.2.1.7.2. Image change deployment triggers

The image change trigger results in a new replication controller whenever the content of an image stream tag changes (when a new version of the image is pushed).

Image change deployment trigger

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
# ...
triggers:
- type: "ImageChange"
  imageChangeParams:
    automatic: true ❶
    from:
      kind: "ImageStreamTag"
      name: "origin-ruby-sample:latest"
      namespace: "myproject"
    containerNames:
      - "helloworld"
```

❶ If the **imageChangeParams.automatic** field is set to **false**, the trigger is disabled.

With the above example, when the **latest** tag value of the **origin-ruby-sample** image stream changes and the new image value differs from the current image specified in the **DeploymentConfig** object's **helloworld** container, a new replication controller is created using the new image for the **helloworld** container.



NOTE

If an image change trigger is defined on a **DeploymentConfig** object (with a config change trigger and **automatic=false**, or with **automatic=true**) and the image stream tag pointed by the image change trigger does not exist yet, the initial deployment process will automatically start as soon as an image is imported or pushed by a build to the image stream tag.

7.2.1.7.3. Setting deployment triggers

Procedure

1. You can set deployment triggers for a **DeploymentConfig** object using the **oc set triggers** command. For example, to set a image change trigger, use the following command:

```
$ oc set triggers dc/<dc_name> \
  --from-image=<project>/<image>:<tag> -c <container_name>
```

7.2.1.8. Setting deployment resources

A deployment is completed by a pod that consumes resources (memory, CPU, and ephemeral storage) on a node. By default, pods consume unbounded node resources. However, if a project specifies default container limits, then pods consume resources up to those limits.



NOTE

The minimum memory limit for a deployment is 12 MB. If a container fails to start due to a **Cannot allocate memory** pod event, the memory limit is too low. Either increase or remove the memory limit. Removing the limit allows pods to consume unbounded node resources.

You can also limit resource use by specifying resource limits as part of the deployment strategy. Deployment resources can be used with the recreate, rolling, or custom deployment strategies.

Procedure

1. In the following example, each of **resources**, **cpu**, **memory**, and **ephemeral-storage** is optional:

```
kind: Deployment
apiVersion: apps/v1
metadata:
  name: hello-openshift
# ...
spec:
# ...
type: "Recreate"
resources:
  limits:
    cpu: "100m" 1
    memory: "256Mi" 2
    ephemeral-storage: "1Gi" 3
```

- 1 **cpu** is in CPU units: **100m** represents 0.1 CPU units ($100 * 1e-3$).
- 2 **memory** is in bytes: **256Mi** represents 268435456 bytes ($256 * 2^{20}$).
- 3 **ephemeral-storage** is in bytes: **1Gi** represents 1073741824 bytes (2^{30}).

However, if a quota has been defined for your project, one of the following two items is required:

- A **resources** section set with an explicit **requests**:

```
kind: Deployment
apiVersion: apps/v1
metadata:
  name: hello-openshift
# ...
spec:
# ...
type: "Recreate"
resources:
  requests: 1
    cpu: "100m"
    memory: "256Mi"
    ephemeral-storage: "1Gi"
```

1 The **requests** object contains the list of resources that correspond to the list of resources in the quota.

- A limit range defined in your project, where the defaults from the **LimitRange** object apply to pods created during the deployment process.

To set deployment resources, choose one of the above options. Otherwise, deploy pod creation fails, citing a failure to satisfy quota.

Additional resources

- For more information about resource limits and requests, see [Understanding managing application memory](#).

7.2.1.9. Scaling manually

In addition to rollbacks, you can exercise fine-grained control over the number of replicas by manually scaling them.



NOTE

Pods can also be auto-scaled using the **oc autoscale** command.

Procedure

1. To manually scale a **DeploymentConfig** object, use the **oc scale** command. For example, the following command sets the replicas in the **frontend DeploymentConfig** object to **3**.

```
$ oc scale dc frontend --replicas=3
```

The number of replicas eventually propagates to the desired and current state of the deployment configured by the **DeploymentConfig** object **frontend**.

7.2.1.10. Accessing private repositories from DeploymentConfig objects

You can add a secret to your **DeploymentConfig** object so that it can access images from a private repository. This procedure shows the OpenShift Container Platform web console method.

Procedure

1. Create a new project.
2. Navigate to **Workloads → Secrets**.
3. Create a secret that contains credentials for accessing a private image repository.
4. Navigate to **Workloads → DeploymentConfigs**.
5. Create a **DeploymentConfig** object.
6. On the **DeploymentConfig** object editor page, set the **Pull Secret** and save your changes.

7.2.1.11. Assigning pods to specific nodes

You can use node selectors in conjunction with labeled nodes to control pod placement.

Cluster administrators can set the default node selector for a project in order to restrict pod placement to specific nodes. As a developer, you can set a node selector on a **Pod** configuration to restrict nodes even further.

Procedure

1. To add a node selector when creating a pod, edit the **Pod** configuration, and add the **nodeSelector** value. This can be added to a single **Pod** configuration, or in a **Pod** template:

```
apiVersion: v1
kind: Pod
metadata:
  name: my-pod
# ...
spec:
  nodeSelector:
    disktype: ssd
# ...
```

Pods created when the node selector is in place are assigned to nodes with the specified labels. The labels specified here are used in conjunction with the labels added by a cluster administrator.

For example, if a project has the **type=user-node** and **region=east** labels added to a project by the cluster administrator, and you add the above **disktype: ssd** label to a pod, the pod is only ever scheduled on nodes that have all three labels.



NOTE

Labels can only be set to one value, so setting a node selector of **region=west** in a **Pod** configuration that has **region=east** as the administrator-set default, results in a pod that will never be scheduled.

7.2.1.12. Running a pod with a different service account

You can run a pod with a service account other than the default.

Procedure

1. Edit the **DeploymentConfig** object:

```
$ oc edit dc/<deployment_config>
```

2. Add the **serviceAccount** and **serviceAccountName** parameters to the **spec** field, and specify the service account you want to use:

```
apiVersion: apps.openshift.io/v1
kind: DeploymentConfig
metadata:
  name: example-dc
# ...
spec:
```



```
# ...
securityContext: {}
serviceAccount: <service_account>
serviceAccountName: <service_account>
```

7.3. USING DEPLOYMENT STRATEGIES

Deployment strategies are used to change or upgrade applications without downtime so that users barely notice a change.

Because users generally access applications through a route handled by a router, deployment strategies can focus on **DeploymentConfig** object features or routing features. Strategies that focus on **DeploymentConfig** object features impact all routes that use the application. Strategies that use router features target individual routes.

Most deployment strategies are supported through the **DeploymentConfig** object, and some additional strategies are supported through router features.

7.3.1. Choosing a deployment strategy

Consider the following when choosing a deployment strategy:

- Long-running connections must be handled gracefully.
- Database conversions can be complex and must be done and rolled back along with the application.
- If the application is a hybrid of microservices and traditional components, downtime might be required to complete the transition.
- You must have the infrastructure to do this.
- If you have a non-isolated test environment, you can break both new and old versions.

A deployment strategy uses readiness checks to determine if a new pod is ready for use. If a readiness check fails, the **DeploymentConfig** object retries to run the pod until it times out. The default timeout is **10m**, a value set in **TimeoutSeconds** in **dc.spec.strategy.*params**.

7.3.2. Rolling strategy

A rolling deployment slowly replaces instances of the previous version of an application with instances of the new version of the application. The rolling strategy is the default deployment strategy used if no strategy is specified on a **DeploymentConfig** object.

A rolling deployment typically waits for new pods to become **ready** via a readiness check before scaling down the old components. If a significant issue occurs, the rolling deployment can be aborted.

When to use a rolling deployment:

- When you want to take no downtime during an application update.
- When your application supports having old code and new code running at the same time.

A rolling deployment means you have both old and new versions of your code running at the same time. This typically requires that your application handle N-1 compatibility.

Example rolling strategy definition

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
# ...
strategy:
  type: Rolling
  rollingParams:
    updatePeriodSeconds: 1 1
    intervalSeconds: 1 2
    timeoutSeconds: 120 3
    maxSurge: "20%" 4
    maxUnavailable: "10%" 5
    pre: {} 6
    post: {}
```

- 1 The time to wait between individual pod updates. If unspecified, this value defaults to **1**.
- 2 The time to wait between polling the deployment status after update. If unspecified, this value defaults to **1**.
- 3 The time to wait for a scaling event before giving up. Optional; the default is **600**. Here, *giving up* means automatically rolling back to the previous complete deployment.
- 4 **maxSurge** is optional and defaults to **25%** if not specified. See the information below the following procedure.
- 5 **maxUnavailable** is optional and defaults to **25%** if not specified. See the information below the following procedure.
- 6 **pre** and **post** are both lifecycle hooks.

The rolling strategy:

1. Executes any **pre** lifecycle hook.
2. Scales up the new replication controller based on the surge count.
3. Scales down the old replication controller based on the max unavailable count.
4. Repeats this scaling until the new replication controller has reached the desired replica count and the old replication controller has been scaled to zero.
5. Executes any **post** lifecycle hook.



IMPORTANT

When scaling down, the rolling strategy waits for pods to become ready so it can decide whether further scaling would affect availability. If scaled up pods never become ready, the deployment process will eventually time out and result in a deployment failure.

The **maxUnavailable** parameter is the maximum number of pods that can be unavailable during the update. The **maxSurge** parameter is the maximum number of pods that can be scheduled above the original number of pods. Both parameters can be set to either a percentage (e.g., **10%**) or an absolute value (e.g., **2**). The default value for both is **25%**.

These parameters allow the deployment to be tuned for availability and speed. For example:

- **maxUnavailable*=0** and **maxSurge*=20%** ensures full capacity is maintained during the update and rapid scale up.
- **maxUnavailable*=10%** and **maxSurge*=0** performs an update using no extra capacity (an in-place update).
- **maxUnavailable*=10%** and **maxSurge*=10%** scales up and down quickly with some potential for capacity loss.

Generally, if you want fast rollouts, use **maxSurge**. If you have to take into account resource quota and can accept partial unavailability, use **maxUnavailable**.



WARNING

The default setting for **maxUnavailable** is **1** for all the machine config pools in OpenShift Container Platform. It is recommended to not change this value and update one control plane node at a time. Do not change this value to **3** for the control plane pool.

7.3.2.1. Canary deployments

All rolling deployments in OpenShift Container Platform are *canary deployments*; a new version (the canary) is tested before all of the old instances are replaced. If the readiness check never succeeds, the canary instance is removed and the **DeploymentConfig** object will be automatically rolled back.

The readiness check is part of the application code and can be as sophisticated as necessary to ensure the new instance is ready to be used. If you must implement more complex checks of the application (such as sending real user workloads to the new instance), consider implementing a custom deployment or using a blue-green deployment strategy.

7.3.2.2. Creating a rolling deployment

Rolling deployments are the default type in OpenShift Container Platform. You can create a rolling deployment using the CLI.

Procedure

1. Create an application based on the example deployment images found in [Quay.io](https://quay.io):

```
$ oc new-app quay.io/openshifttest/deployment-example:latest
```

**NOTE**

This image does not expose any ports. If you want to expose your applications over an external LoadBalancer service or enable access to the application over the public internet, create a service by using the **oc expose dc/deployment-example --port=<port>** command after completing this procedure.

2. If you have the router installed, make the application available via a route or use the service IP directly.

```
$ oc expose svc/deployment-example
```

3. Browse to the application at **deployment-example.<project>.<router_domain>** to verify you see the **v1** image.
4. Scale the **DeploymentConfig** object up to three replicas:

```
$ oc scale dc/deployment-example --replicas=3
```

5. Trigger a new deployment automatically by tagging a new version of the example as the **latest** tag:

```
$ oc tag deployment-example:v2 deployment-example:latest
```

6. In your browser, refresh the page until you see the **v2** image.
7. When using the CLI, the following command shows how many pods are on version 1 and how many are on version 2. In the web console, the pods are progressively added to v2 and removed from v1:

```
$ oc describe dc deployment-example
```

During the deployment process, the new replication controller is incrementally scaled up. After the new pods are marked as **ready** (by passing their readiness check), the deployment process continues.

If the pods do not become ready, the process aborts, and the deployment rolls back to its previous version.

7.3.2.3. Editing a deployment by using the Developer perspective

You can edit the deployment strategy, image settings, environment variables, and advanced options for your deployment by using the **Developer** perspective.

Prerequisites

- You are in the **Developer** perspective of the web console.
- You have created an application.

Procedure

1. Navigate to the **Topology** view.

2. Click your application to see the **Details** panel.
3. In the **Actions** drop-down menu, select **Edit Deployment** to view the **Edit Deployment** page.
4. You can edit the following **Advanced options** for your deployment:
 - a. Optional: You can pause rollouts by clicking **Pause rollouts**, and then selecting the **Pause rollouts for this deployment** checkbox.
By pausing rollouts, you can make changes to your application without triggering a rollout. You can resume rollouts at any time.
 - b. Optional: Click **Scaling** to change the number of instances of your image by modifying the number of **Replicas**.
5. Click **Save**.

7.3.2.4. Starting a rolling deployment using the Developer perspective

You can upgrade an application by starting a rolling deployment.


Prerequisites

- You are in the **Developer** perspective of the web console.
- You have created an application.

Procedure

1. In the **Topology** view, click the application node to see the **Overview** tab in the side panel. Note that the **Update Strategy** is set to the default **Rolling** strategy.
2. In the **Actions** drop-down menu, select **Start Rollout** to start a rolling update. The rolling deployment spins up the new version of the application and then terminates the old one.

Figure 7.1. Rolling update



Additional resources

- [Creating and deploying applications on OpenShift Container Platform using the Developer perspective](#)
- [Viewing the applications in your project](#)

7.3.3. Recreate strategy

The recreate strategy has basic rollout behavior and supports lifecycle hooks for injecting code into the deployment process.

Example recreate strategy definition


```
kind: Deployment
apiVersion: apps/v1
metadata:
  name: hello-openshift
# ...
spec:
# ...
strategy:
  type: Recreate
  recreateParams: 1
```

```
pre: {} 2
mid: {}
post: {}
```

- 1 **recreateParams** are optional.
- 2 **pre**, **mid**, and **post** are lifecycle hooks.

The recreate strategy:

1. Executes any **pre** lifecycle hook.
2. Scales down the previous deployment to zero.
3. Executes any **mid** lifecycle hook.
4. Scales up the new deployment.
5. Executes any **post** lifecycle hook.



IMPORTANT

During scale up, if the replica count of the deployment is greater than one, the first replica of the deployment will be validated for readiness before fully scaling up the deployment. If the validation of the first replica fails, the deployment will be considered a failure.

When to use a recreate deployment:

- When you must run migrations or other data transformations before your new code starts.
- When you do not support having new and old versions of your application code running at the same time.
- When you want to use a RWO volume, which is not supported being shared between multiple replicas.

A recreate deployment incurs downtime because, for a brief period, no instances of your application are running. However, your old code and new code do not run at the same time.

7.3.3.1. Editing a deployment by using the Developer perspective

You can edit the deployment strategy, image settings, environment variables, and advanced options for your deployment by using the **Developer** perspective.

Prerequisites

- You are in the **Developer** perspective of the web console.
- You have created an application.

Procedure

1. Navigate to the **Topology** view.

2. Click your application to see the **Details** panel.
3. In the **Actions** drop-down menu, select **Edit Deployment** to view the **Edit Deployment** page.
4. You can edit the following **Advanced options** for your deployment:
 - a. Optional: You can pause rollouts by clicking **Pause rollouts**, and then selecting the **Pause rollouts for this deployment** checkbox.
By pausing rollouts, you can make changes to your application without triggering a rollout.
You can resume rollouts at any time.
 - b. Optional: Click **Scaling** to change the number of instances of your image by modifying the number of **Replicas**.
5. Click **Save**.

7.3.3.2. Starting a recreate deployment using the Developer perspective

You can switch the deployment strategy from the default rolling update to a recreate update using the **Developer** perspective in the web console.

Prerequisites

- Ensure that you are in the **Developer** perspective of the web console.
- Ensure that you have created an application using the **Add** view and see it deployed in the **Topology** view.

Procedure

To switch to a recreate update strategy and to upgrade an application:

1. Click your application to see the **Details** panel.
2. In the **Actions** drop-down menu, select **Edit Deployment Config** to see the deployment configuration details of the application.
3. In the YAML editor, change the **spec.strategy.type** to **Recreate** and click **Save**.
4. In the **Topology** view, select the node to see the **Overview** tab in the side panel. The **Update Strategy** is now set to **Recreate**.
5. Use the **Actions** drop-down menu to select **Start Rollout** to start an update using the recreate strategy. The recreate strategy first terminates pods for the older version of the application and then spins up pods for the new version.

Figure 7.2. Recreate update

nodejs-ex1 Actions

Overview Resources

0 pods → **0 pods**

Name nodejs-ex1	Latest Version 3
Namespace test-project	Message manual change
Labels app=nodejs-ex1 app.kubernetes.io/com... =nodejs... app.kubernetes.io/inst... =nodejs-... app.kubernetes.io/name=nodejs app.openshift.io/runtime=nodejs app.openshift.io/runtime-... =10-S...	Update Strategy Recreate
Pod Selector app=nodejs-ex1, deploymentconfig=nodejs-ex1	Min Ready Seconds Not Configured
	Triggers ImageChange, ConfigChange

Additional resources

- [Creating and deploying applications on OpenShift Container Platform using the Developer perspective](#)
- [Viewing the applications in your project](#)

7.3.4. Custom strategy

The custom strategy allows you to provide your own deployment behavior.

Example custom strategy definition

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
```

```
# ...
strategy:
  type: Custom
  customParams:
    image: organization/strategy
    command: [ "command", "arg1" ]
  environment:
    - name: ENV_1
      value: VALUE_1
```

In the above example, the **organization/strategy** container image provides the deployment behavior. The optional **command** array overrides any **CMD** directive specified in the image's **Dockerfile**. The optional environment variables provided are added to the execution environment of the strategy process.

Additionally, OpenShift Container Platform provides the following environment variables to the deployment process:

Environment variable	Description
OPENSHIFT_DEPLOYMENT_NAME	The name of the new deployment, a replication controller.
OPENSHIFT_DEPLOYMENT_NAMESPACE	The name space of the new deployment.

The replica count of the new deployment will initially be zero. The responsibility of the strategy is to make the new deployment active using the logic that best serves the needs of the user.

Alternatively, use the **customParams** object to inject the custom deployment logic into the existing deployment strategies. Provide a custom shell script logic and call the **openshift-deploy** binary. Users do not have to supply their custom deployer container image; in this case, the default OpenShift Container Platform deployer image is used instead:

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: example-dc
# ...
spec:
# ...
strategy:
  type: Rolling
  customParams:
    command:
      - /bin/sh
      - -c
      - |
        set -e
        openshift-deploy --until=50%
        echo Halfway there
        openshift-deploy
        echo Complete
```

This results in following deployment:

```
Started deployment #2
--> Scaling up custom-deployment-2 from 0 to 2, scaling down custom-deployment-1 from 2 to 0
(keep 2 pods available, don't exceed 3 pods)
  Scaling custom-deployment-2 up to 1
--> Reached 50% (currently 50%)
Halfway there
--> Scaling up custom-deployment-2 from 1 to 2, scaling down custom-deployment-1 from 2 to 0
(keep 2 pods available, don't exceed 3 pods)
  Scaling custom-deployment-1 down to 1
  Scaling custom-deployment-2 up to 2
  Scaling custom-deployment-1 down to 0
--> Success
Complete
```

If the custom deployment strategy process requires access to the OpenShift Container Platform API or the Kubernetes API the container that executes the strategy can use the service account token available inside the container for authentication.

7.3.4.1. Editing a deployment by using the Developer perspective

You can edit the deployment strategy, image settings, environment variables, and advanced options for your deployment by using the **Developer** perspective.

Prerequisites

- You are in the **Developer** perspective of the web console.
- You have created an application.

Procedure

1. Navigate to the **Topology** view.
2. Click your application to see the **Details** panel.
3. In the **Actions** drop-down menu, select **Edit Deployment** to view the **Edit Deployment** page.
4. You can edit the following **Advanced options** for your deployment:
 - a. Optional: You can pause rollouts by clicking **Pause rollouts**, and then selecting the **Pause rollouts for this deployment** checkbox.
By pausing rollouts, you can make changes to your application without triggering a rollout. You can resume rollouts at any time.
 - b. Optional: Click **Scaling** to change the number of instances of your image by modifying the number of **Replicas**.
5. Click **Save**.

7.3.5. Lifecycle hooks

The rolling and recreate strategies support *lifecycle hooks*, or deployment hooks, which allow behavior to be injected into the deployment process at predefined points within the strategy:

Example pre lifecycle hook

```
pre:
  failurePolicy: Abort
  execNewPod: {} 1
```

1 **execNewPod** is a pod-based lifecycle hook.

Every hook has a *failure policy*, which defines the action the strategy should take when a hook failure is encountered:

Abort	The deployment process will be considered a failure if the hook fails.
Retry	The hook execution should be retried until it succeeds.
Ignore	Any hook failure should be ignored and the deployment should proceed.

Hooks have a type-specific field that describes how to execute the hook. Currently, pod-based hooks are the only supported hook type, specified by the **execNewPod** field.

7.3.5.1. Pod-based lifecycle hook

Pod-based lifecycle hooks execute hook code in a new pod derived from the template in a **DeploymentConfig** object.

The following simplified example deployment uses the rolling strategy. Triggers and some other minor details are omitted for brevity:

```
kind: DeploymentConfig
apiVersion: apps.openshift.io/v1
metadata:
  name: frontend
spec:
  template:
    metadata:
      labels:
        name: frontend
    spec:
      containers:
        - name: helloworld
          image: openshift/origin-ruby-sample
  replicas: 5
  selector:
    name: frontend
  strategy:
    type: Rolling
    rollingParams:
      pre:
        failurePolicy: Abort
        execNewPod:
          containerName: helloworld 1
```

```

command: [ "/usr/bin/command", "arg1", "arg2" ] ❷
env: ❸
  - name: CUSTOM_VAR1
    value: custom_value1
volumes:
  - data ❹

```

- ❶ The **helloworld** name refers to `spec.template.spec.containers[0].name`.
- ❷ This **command** overrides any **ENTRYPOINT** defined by the **openshift/origin-ruby-sample** image.
- ❸ **env** is an optional set of environment variables for the hook container.
- ❹ **volumes** is an optional set of volume references for the hook container.

In this example, the **pre** hook will be executed in a new pod using the **openshift/origin-ruby-sample** image from the **helloworld** container. The hook pod has the following properties:

- The hook command is **/usr/bin/command arg1 arg2**.
- The hook container has the **CUSTOM_VAR1=custom_value1** environment variable.
- The hook failure policy is **Abort**, meaning the deployment process fails if the hook fails.
- The hook pod inherits the **data** volume from the **DeploymentConfig** object pod.

7.3.5.2. Setting lifecycle hooks

You can set lifecycle hooks, or deployment hooks, for a deployment using the CLI.

Procedure

1. Use the **oc set deployment-hook** command to set the type of hook you want: **--pre**, **--mid**, or **--post**. For example, to set a pre-deployment hook:

```

$ oc set deployment-hook dc/frontend \
  --pre -c helloworld -e CUSTOM_VAR1=custom_value1 \
  --volumes data --failure-policy=abort -- /usr/bin/command arg1 arg2

```

7.4. USING ROUTE-BASED DEPLOYMENT STRATEGIES

Deployment strategies provide a way for the application to evolve. Some strategies use **Deployment** objects to make changes that are seen by users of all routes that resolve to the application. Other advanced strategies, such as the ones described in this section, use router features in conjunction with **Deployment** objects to impact specific routes.

The most common route-based strategy is to use a *blue-green deployment*. The new version (the green version) is brought up for testing and evaluation, while the users still use the stable version (the blue version). When ready, the users are switched to the green version. If a problem arises, you can switch back to the blue version.

Alternatively, you can use an *A/B versions* strategy in which both versions are active at the same time. With this strategy, some users can use *version A*, and other users can use *version B*. You can use this

strategy to experiment with user interface changes or other features in order to get user feedback. You can also use it to verify proper operation in a production context where problems impact a limited number of users.

A canary deployment tests the new version but when a problem is detected it quickly falls back to the previous version. This can be done with both of the above strategies.

The route-based deployment strategies do not scale the number of pods in the services. To maintain desired performance characteristics the deployment configurations might have to be scaled.

7.4.1. Proxy shards and traffic splitting

In production environments, you can precisely control the distribution of traffic that lands on a particular shard. When dealing with large numbers of instances, you can use the relative scale of individual shards to implement percentage based traffic. That combines well with a *proxy shard*, which forwards or splits the traffic it receives to a separate service or application running elsewhere.

In the simplest configuration, the proxy forwards requests unchanged. In more complex setups, you can duplicate the incoming requests and send to both a separate cluster as well as to a local instance of the application, and compare the result. Other patterns include keeping the caches of a DR installation warm, or sampling incoming traffic for analysis purposes.

Any TCP (or UDP) proxy could be run under the desired shard. Use the **oc scale** command to alter the relative number of instances serving requests under the proxy shard. For more complex traffic management, consider customizing the OpenShift Container Platform router with proportional balancing capabilities.

7.4.2. N-1 compatibility

Applications that have new code and old code running at the same time must be careful to ensure that data written by the new code can be read and handled (or gracefully ignored) by the old version of the code. This is sometimes called *schema evolution* and is a complex problem.

This can take many forms: data stored on disk, in a database, in a temporary cache, or that is part of a user's browser session. While most web applications can support rolling deployments, it is important to test and design your application to handle it.

For some applications, the period of time that old code and new code is running side by side is short, so bugs or some failed user transactions are acceptable. For others, the failure pattern may result in the entire application becoming non-functional.

One way to validate N-1 compatibility is to use an A/B deployment: run the old code and new code at the same time in a controlled way in a test environment, and verify that traffic that flows to the new deployment does not cause failures in the old deployment.

7.4.3. Graceful termination

OpenShift Container Platform and Kubernetes give application instances time to shut down before removing them from load balancing rotations. However, applications must ensure they cleanly terminate user connections as well before they exit.

On shutdown, OpenShift Container Platform sends a **TERM** signal to the processes in the container. Application code, on receiving **SIGTERM**, stop accepting new connections. This ensures that load balancers route traffic to other active instances. The application code then waits until all open

connections are closed, or gracefully terminate individual connections at the next opportunity, before exiting.

After the graceful termination period expires, a process that has not exited is sent the **KILL** signal, which immediately ends the process. The **terminationGracePeriodSeconds** attribute of a pod or pod template controls the graceful termination period (default 30 seconds) and can be customized per application as necessary.

7.4.4. Blue-green deployments

Blue-green deployments involve running two versions of an application at the same time and moving traffic from the in-production version (the blue version) to the newer version (the green version). You can use a rolling strategy or switch services in a route.

Because many applications depend on persistent data, you must have an application that supports *N-1 compatibility*, which means it shares data and implements live migration between the database, store, or disk by creating two copies of the data layer.

Consider the data used in testing the new version. If it is the production data, a bug in the new version can break the production version.

7.4.4.1. Setting up a blue-green deployment

Blue-green deployments use two **Deployment** objects. Both are running, and the one in production depends on the service the route specifies, with each **Deployment** object exposed to a different service.



NOTE

Routes are intended for web (HTTP and HTTPS) traffic, so this technique is best suited for web applications.

You can create a new route to the new version and test it. When ready, change the service in the production route to point to the new service and the new (green) version is live.

If necessary, you can roll back to the older (blue) version by switching the service back to the previous version.

Procedure

1. Create two independent application components.
 - a. Create a copy of the example application running the **v1** image under the **example-blue** service:

```
$ oc new-app openshift/deployment-example:v1 --name=example-blue
```

- b. Create a second copy that uses the **v2** image under the **example-green** service:

```
$ oc new-app openshift/deployment-example:v2 --name=example-green
```

2. Create a route that points to the old service:

```
$ oc expose svc/example-blue --name=bluegreen-example
```

3. Browse to the application at **bluegreen-example-`<project>`.`<router_domain>`** to verify you see the **v1** image.
4. Edit the route and change the service name to **example-green**:

```
$ oc patch route/bluegreen-example -p '{"spec":{"to":{"name":"example-green"}}}'
```

5. To verify that the route has changed, refresh the browser until you see the **v2** image.

7.4.5. A/B deployments

The A/B deployment strategy lets you try a new version of the application in a limited way in the production environment. You can specify that the production version gets most of the user requests while a limited fraction of requests go to the new version.

Because you control the portion of requests to each version, as testing progresses you can increase the fraction of requests to the new version and ultimately stop using the previous version. As you adjust the request load on each version, the number of pods in each service might have to be scaled as well to provide the expected performance.

In addition to upgrading software, you can use this feature to experiment with versions of the user interface. Since some users get the old version and some the new, you can evaluate the user's reaction to the different versions to inform design decisions.

For this to be effective, both the old and new versions must be similar enough that both can run at the same time. This is common with bug fix releases and when new features do not interfere with the old. The versions require N-1 compatibility to properly work together.

OpenShift Container Platform supports N-1 compatibility through the web console as well as the CLI.

7.4.5.1. Load balancing for A/B testing

The user sets up a route with multiple services. Each service handles a version of the application.

Each service is assigned a **weight** and the portion of requests to each service is the **service_weight** divided by the **sum_of_weights**. The **weight** for each service is distributed to the service's endpoints so that the sum of the endpoint **weights** is the service **weight**.

The route can have up to four services. The **weight** for the service can be between **0** and **256**. When the **weight** is **0**, the service does not participate in load balancing but continues to serve existing persistent connections. When the service **weight** is not **0**, each endpoint has a minimum **weight** of **1**. Because of this, a service with a lot of endpoints can end up with higher **weight** than intended. In this case, reduce the number of pods to get the expected load balance **weight**.

Procedure

To set up the A/B environment:

1. Create the two applications and give them different names. Each creates a **Deployment** object. The applications are versions of the same program; one is usually the current production version and the other the proposed new version.
 - a. Create the first application. The following example creates an application called **ab-example-a**:


```
$ oc new-app openshift/deployment-example --name=ab-example-a
```

b. Create the second application:

```
$ oc new-app openshift/deployment-example:v2 --name=ab-example-b
```

Both applications are deployed and services are created.

2. Make the application available externally via a route. At this point, you can expose either. It can be convenient to expose the current production version first and later modify the route to add the new version.

```
$ oc expose svc/ab-example-a
```

Browse to the application at **ab-example-a.<project>.<router_domain>** to verify that you see the expected version.

3. When you deploy the route, the router balances the traffic according to the **weights** specified for the services. At this point, there is a single service with default **weight=1** so all requests go to it. Adding the other service as an **alternateBackends** and adjusting the **weights** brings the A/B setup to life. This can be done by the **oc set route-backends** command or by editing the route.



NOTE

When using **alternateBackends**, also use the **roundrobin** load balancing strategy to ensure requests are distributed as expected to the services based on weight. **roundrobin** can be set for a route by using a route annotation. See the *Additional resources* section for more information about route annotations.

Setting the **oc set route-backend** to **0** means the service does not participate in load balancing, but continues to serve existing persistent connections.



NOTE

Changes to the route just change the portion of traffic to the various services. You might have to scale the deployment to adjust the number of pods to handle the anticipated loads.

To edit the route, run:

```
$ oc edit route <route_name>
```

Example output

```
apiVersion: route.openshift.io/v1
kind: Route
metadata:
  name: route-alternate-service
annotations:
  haproxy.router.openshift.io/balance: roundrobin
# ...
spec:
```

```

host: ab-example.my-project.my-domain
to:
  kind: Service
  name: ab-example-a
  weight: 10
alternateBackends:
- kind: Service
  name: ab-example-b
  weight: 15
# ...

```

7.4.5.1.1. Managing weights of an existing route using the web console

Procedure

1. Navigate to the **Networking → Routes** page.



2. Click the Options menu next to the route you want to edit and select **Edit Route**.
3. Edit the YAML file. Update the **weight** to be an integer between **0** and **256** that specifies the relative weight of the target against other target reference objects. The value **0** suppresses requests to this back end. The default is **100**. Run **oc explain routes.spec.alternateBackends** for more information about the options.
4. Click **Save**.

7.4.5.1.2. Managing weights of a new route using the web console

1. Navigate to the **Networking → Routes** page.
2. Click **Create Route**.
3. Enter the route **Name**.
4. Select the **Service**.
5. Click **Add Alternate Service**.
6. Enter a value for **Weight** and **Alternate Service Weight**. Enter a number between **0** and **255** that depicts relative weight compared with other targets. The default is **100**.
7. Select the **Target Port**.
8. Click **Create**.

7.4.5.1.3. Managing weights using the CLI

Procedure

1. To manage the services and corresponding weights load balanced by the route, use the **oc set route-backends** command:

```
$ oc set route-backends ROUTENAME \
  [--zero|--equal] [--adjust] SERVICE=WEIGHT[%] [...] [options]
```

For example, the following sets **ab-example-a** as the primary service with **weight=198** and **ab-example-b** as the first alternate service with a **weight=2**:

```
$ oc set route-backends ab-example ab-example-a=198 ab-example-b=2
```

This means 99% of traffic is sent to service **ab-example-a** and 1% to service **ab-example-b**.

This command does not scale the deployment. You might be required to do so to have enough pods to handle the request load.

2. Run the command with no flags to verify the current configuration:

```
$ oc set route-backends ab-example
```

Example output

NAME	KIND	TO	WEIGHT
routes/ab-example	Service	ab-example-a	198 (99%)
routes/ab-example	Service	ab-example-b	2 (1%)

3. To override the default values for the load balancing algorithm, adjust the annotation on the route by setting the algorithm to **roundrobin**. For a route on OpenShift Container Platform, the default load balancing algorithm is set to **random** or **source** values.
To set the algorithm to **roundrobin**, run the command:

```
$ oc annotate routes/<route-name> haproxy.router.openshift.io/balance=roundrobin
```

For Transport Layer Security (TLS) passthrough routes, the default value is **source**. For all other routes, the default is **random**.

4. To alter the weight of an individual service relative to itself or to the primary service, use the **--adjust** flag. Specifying a percentage adjusts the service relative to either the primary or the first alternate (if you specify the primary). If there are other backends, their weights are kept proportional to the changed.
The following example alters the weight of **ab-example-a** and **ab-example-b** services:

```
$ oc set route-backends ab-example --adjust ab-example-a=200 ab-example-b=10
```

Alternatively, alter the weight of a service by specifying a percentage:

```
$ oc set route-backends ab-example --adjust ab-example-b=5%
```

By specifying **+** before the percentage declaration, you can adjust a weighting relative to the current setting. For example:


```
$ oc set route-backends ab-example --adjust ab-example-b=+15%
```

The **--equal** flag sets the **weight** of all services to **100**:

■

```
$ oc set route-backends ab-example --equal
```

The **--zero** flag sets the **weight** of all services to **0**. All requests then return with a 503 error.



NOTE

Not all routers may support multiple or weighted backends.

7.4.5.1.4. One service, multiple Deployment objects

Procedure

1. Create a new application, adding a label **ab-example=true** that will be common to all shards:

```
$ oc new-app openshift/deployment-example --name=ab-example-a --as-deployment-
config=true --labels=ab-example=true --env=SUBTITLE\=shardA
```

```
$ oc delete svc/ab-example-a
```

The application is deployed and a service is created. This is the first shard.

2. Make the application available via a route, or use the service IP directly:

```
$ oc expose deployment ab-example-a --name=ab-example --selector=ab-example\=true
```

```
$ oc expose service ab-example
```

3. Browse to the application at **ab-example-`<project_name>`.`<router_domain>`** to verify you see the **v1** image.
4. Create a second shard based on the same source image and label as the first shard, but with a different tagged version and unique environment variables:

```
$ oc new-app openshift/deployment-example:v2 \
  --name=ab-example-b --labels=ab-example=true \
  SUBTITLE="shard B" COLOR="red" --as-deployment-config=true
```

```
$ oc delete svc/ab-example-b
```

5. At this point, both sets of pods are being served under the route. However, because both browsers (by leaving a connection open) and the router (by default, through a cookie) attempt to preserve your connection to a back-end server, you might not see both shards being returned to you.

To force your browser to one or the other shard:

- a. Use the **oc scale** command to reduce replicas of **ab-example-a** to **0**.

```
$ oc scale dc/ab-example-a --replicas=0
```

Refresh your browser to show **v2** and **shard B** (in red).

- b. Scale **ab-example-a** to **1** replica and **ab-example-b** to **0**:

```
$ oc scale dc/ab-example-a --replicas=1; oc scale dc/ab-example-b --replicas=0
```

Refresh your browser to show **v1** and **shard A** (in blue).

6. If you trigger a deployment on either shard, only the pods in that shard are affected. You can trigger a deployment by changing the **SUBTITLE** environment variable in either **Deployment** object:

```
$ oc edit dc/ab-example-a
```

or

```
$ oc edit dc/ab-example-b
```

7.4.6. Additional resources

- [Route-specific annotations](#)

CHAPTER 8. QUOTAS

8.1. RESOURCE QUOTAS PER PROJECT

A *resource quota*, defined by a **ResourceQuota** object, provides constraints that limit aggregate resource consumption per project. It can limit the quantity of objects that can be created in a project by type, as well as the total amount of compute resources and storage that might be consumed by resources in that project.

This guide describes how resource quotas work, how cluster administrators can set and manage resource quotas on a per project basis, and how developers and cluster administrators can view them.

8.1.1. Resources managed by quotas

The following describes the set of compute resources and object types that can be managed by a quota.



NOTE

A pod is in a terminal state if **status.phase** in (**Failed**, **Succeeded**) is true.

Table 8.1. Compute resources managed by quota

Resource Name	Description
cpu	The sum of CPU requests across all pods in a non-terminal state cannot exceed this value. cpu and requests.cpu are the same value and can be used interchangeably.
memory	The sum of memory requests across all pods in a non-terminal state cannot exceed this value. memory and requests.memory are the same value and can be used interchangeably.
requests.cpu	The sum of CPU requests across all pods in a non-terminal state cannot exceed this value. cpu and requests.cpu are the same value and can be used interchangeably.
requests.memory	The sum of memory requests across all pods in a non-terminal state cannot exceed this value. memory and requests.memory are the same value and can be used interchangeably.
limits.cpu	The sum of CPU limits across all pods in a non-terminal state cannot exceed this value.
limits.memory	The sum of memory limits across all pods in a non-terminal state cannot exceed this value.

Table 8.2. Storage resources managed by quota

Resource Name	Description
requests.storage	The sum of storage requests across all persistent volume claims in any state cannot exceed this value.
persistentvolumeclaims	The total number of persistent volume claims that can exist in the project.
<storage-class-name>.storageclass.storage.k8s.io/requests.storage	The sum of storage requests across all persistent volume claims in any state that have a matching storage class, cannot exceed this value.
<storage-class-name>.storageclass.storage.k8s.io/persistentvolumeclaims	The total number of persistent volume claims with a matching storage class that can exist in the project.
ephemeral-storage	The sum of local ephemeral storage requests across all pods in a non-terminal state cannot exceed this value. ephemeral-storage and requests.ephemeral-storage are the same value and can be used interchangeably.
requests.ephemeral-storage	The sum of ephemeral storage requests across all pods in a non-terminal state cannot exceed this value. ephemeral-storage and requests.ephemeral-storage are the same value and can be used interchangeably.
limits.ephemeral-storage	The sum of ephemeral storage limits across all pods in a non-terminal state cannot exceed this value.

Table 8.3. Object counts managed by quota

Resource Name	Description
pods	The total number of pods in a non-terminal state that can exist in the project.
replicationcontrollers	The total number of ReplicationControllers that can exist in the project.
resourcequotas	The total number of resource quotas that can exist in the project.
services	The total number of services that can exist in the project.
services.loadbalancers	The total number of services of type LoadBalancer that can exist in the project.
services.nodeports	The total number of services of type NodePort that can exist in the project.

Resource Name	Description
secrets	The total number of secrets that can exist in the project.
configmaps	The total number of ConfigMap objects that can exist in the project.
persistentvolumeclaims	The total number of persistent volume claims that can exist in the project.
openshift.io/imagestreams	The total number of imagestreams that can exist in the project.

8.1.2. Quota scopes

Each quota can have an associated set of *scopes*. A quota only measures usage for a resource if it matches the intersection of enumerated scopes.

Adding a scope to a quota restricts the set of resources to which that quota can apply. Specifying a resource outside of the allowed set results in a validation error.

Scope	Description
BestEffort	Match pods that have best effort quality of service for either cpu or memory .
NotBestEffort	Match pods that do not have best effort quality of service for cpu and memory .

A **BestEffort** scope restricts a quota to limiting the following resources:

- **pods**

A **NotBestEffort** scope restricts a quota to tracking the following resources:

- **pods**
- **memory**
- **requests.memory**
- **limits.memory**
- **cpu**
- **requests.cpu**
- **limits.cpu**

8.1.3. Quota enforcement

After a resource quota for a project is first created, the project restricts the ability to create any new resources that may violate a quota constraint until it has calculated updated usage statistics.

After a quota is created and usage statistics are updated, the project accepts the creation of new content. When you create or modify resources, your quota usage is incremented immediately upon the request to create or modify the resource.

When you delete a resource, your quota use is decremented during the next full recalculation of quota statistics for the project. A configurable amount of time determines how long it takes to reduce quota usage statistics to their current observed system value.

If project modifications exceed a quota usage limit, the server denies the action, and an appropriate error message is returned to the user explaining the quota constraint violated, and what their currently observed usage statistics are in the system.

8.1.4. Requests versus limits

When allocating compute resources, each container might specify a request and a limit value each for CPU, memory, and ephemeral storage. Quotas can restrict any of these values.

If the quota has a value specified for **requests.cpu** or **requests.memory**, then it requires that every incoming container make an explicit request for those resources. If the quota has a value specified for **limits.cpu** or **limits.memory**, then it requires that every incoming container specify an explicit limit for those resources.

8.1.5. Sample resource quota definitions

core-object-counts.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: core-object-counts
spec:
  hard:
    configmaps: "10" ❶
    persistentvolumeclaims: "4" ❷
    replicationcontrollers: "20" ❸
    secrets: "10" ❹
    services: "10" ❺
    services.loadbalancers: "2" ❻
```

- ❶ The total number of **ConfigMap** objects that can exist in the project.
- ❷ The total number of persistent volume claims (PVCs) that can exist in the project.
- ❸ The total number of replication controllers that can exist in the project.
- ❹ The total number of secrets that can exist in the project.
- ❺ The total number of services that can exist in the project.
- ❻ The total number of services of type **LoadBalancer** that can exist in the project.

openshift-object-counts.yaml

```

apiVersion: v1
kind: ResourceQuota
metadata:
  name: openshift-object-counts
spec:
  hard:
    openshift.io/imagestreams: "10" ❶

```

- ❶ The total number of image streams that can exist in the project.

compute-resources.yaml

```

apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources
spec:
  hard:
    pods: "4" ❶
    requests.cpu: "1" ❷
    requests.memory: 1Gi ❸
    limits.cpu: "2" ❹
    limits.memory: 2Gi ❺

```

- ❶ The total number of pods in a non-terminal state that can exist in the project.
- ❷ Across all pods in a non-terminal state, the sum of CPU requests cannot exceed 1 core.
- ❸ Across all pods in a non-terminal state, the sum of memory requests cannot exceed 1Gi.
- ❹ Across all pods in a non-terminal state, the sum of CPU limits cannot exceed 2 cores.
- ❺ Across all pods in a non-terminal state, the sum of memory limits cannot exceed 2Gi.

besteffort.yaml

```

apiVersion: v1
kind: ResourceQuota
metadata:
  name: besteffort
spec:
  hard:
    pods: "1" ❶
  scopes:
    - BestEffort ❷

```

- ❶ The total number of pods in a non-terminal state with **BestEffort** quality of service that can exist in the project.

- 2 Restricts the quota to only matching pods that have **BestEffort** quality of service for either memory or CPU.

compute-resources-long-running.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources-long-running
spec:
  hard:
    pods: "4" 1
    limits.cpu: "4" 2
    limits.memory: "2Gi" 3
  scopes:
    - NotTerminating 4
```

- 1 The total number of pods in a non-terminal state.
- 2 Across all pods in a non-terminal state, the sum of CPU limits cannot exceed this value.
- 3 Across all pods in a non-terminal state, the sum of memory limits cannot exceed this value.
- 4 Restricts the quota to only matching pods where **spec.activeDeadlineSeconds** is set to **nil**. Build pods fall under **NotTerminating** unless the **RestartNever** policy is applied.

compute-resources-time-bound.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources-time-bound
spec:
  hard:
    pods: "2" 1
    limits.cpu: "1" 2
    limits.memory: "1Gi" 3
  scopes:
    - Terminating 4
```

- 1 The total number of pods in a terminating state.
- 2 Across all pods in a terminating state, the sum of CPU limits cannot exceed this value.
- 3 Across all pods in a terminating state, the sum of memory limits cannot exceed this value.
- 4 Restricts the quota to only matching pods where **spec.activeDeadlineSeconds** ≥ 0 . For example, this quota charges for build or deployer pods, but not long running pods like a web server or database.

storage-consumption.yaml

```

apiVersion: v1
kind: ResourceQuota
metadata:
  name: storage-consumption
spec:
  hard:
    persistentvolumeclaims: "10" ❶
    requests.storage: "50Gi" ❷
    gold.storageclass.storage.k8s.io/requests.storage: "10Gi" ❸
    silver.storageclass.storage.k8s.io/requests.storage: "20Gi" ❹
    silver.storageclass.storage.k8s.io/persistentvolumeclaims: "5" ❺
    bronze.storageclass.storage.k8s.io/requests.storage: "0" ❻
    bronze.storageclass.storage.k8s.io/persistentvolumeclaims: "0" ❼
    requests.ephemeral-storage: 2Gi ❽
    limits.ephemeral-storage: 4Gi ❾

```

- ❶ The total number of persistent volume claims in a project
- ❷ Across all persistent volume claims in a project, the sum of storage requested cannot exceed this value.
- ❸ Across all persistent volume claims in a project, the sum of storage requested in the gold storage class cannot exceed this value.
- ❹ Across all persistent volume claims in a project, the sum of storage requested in the silver storage class cannot exceed this value.
- ❺ Across all persistent volume claims in a project, the total number of claims in the silver storage class cannot exceed this value.
- ❻ Across all persistent volume claims in a project, the sum of storage requested in the bronze storage class cannot exceed this value. When this is set to **0**, it means bronze storage class cannot request storage.
- ❼ Across all persistent volume claims in a project, the sum of storage requested in the bronze storage class cannot exceed this value. When this is set to **0**, it means bronze storage class cannot create claims.
- ❽ Across all pods in a non-terminal state, the sum of ephemeral storage requests cannot exceed 2Gi.
- ❾ Across all pods in a non-terminal state, the sum of ephemeral storage limits cannot exceed 4Gi.

8.1.6. Creating a quota

You can create a quota to constrain resource usage in a given project.

Procedure

1. Define the quota in a file.
2. Use the file to create the quota and apply it to a project:

```
$ oc create -f <file> [-n <project_name>]
```

For example:

```
$ oc create -f core-object-counts.yaml -n demoproject
```

8.1.6.1. Creating object count quotas

You can create an object count quota for all standard namespaced resource types on OpenShift Container Platform, such as **BuildConfig** and **DeploymentConfig** objects. An object quota count places a defined quota on all standard namespaced resource types.

When using a resource quota, an object is charged against the quota upon creation. These types of quotas are useful to protect against exhaustion of resources. The quota can only be created if there are enough spare resources within the project.

Procedure

To configure an object count quota for a resource:

1. Run the following command:

```
$ oc create quota <name> \
  --hard=count/<resource>.<group>=<quota>,count/<resource>.<group>=<quota> 1
```

- 1 The **<resource>** variable is the name of the resource, and **<group>** is the API group, if applicable. Use the **oc api-resources** command for a list of resources and their associated API groups.

For example:

```
$ oc create quota test \
  --
  hard=count/deployments.apps=2,count/replicasets.apps=4,count/pods=3,count/secrets=4
```

Example output

```
resourcequota "test" created
```

This example limits the listed resources to the hard limit in each project in the cluster.

2. Verify that the quota was created:

```
$ oc describe quota test
```

Example output

```
Name:          test
Namespace:     quota
Resource       Used  Hard
-----
count/deployments.apps  0  2
count/pods              0  3
count/replicasets.apps  0  4
count/secrets           0  4
```



8.1.6.2. Setting resource quota for extended resources

Overcommitment of resources is not allowed for extended resources, so you must specify **requests** and **limits** for the same extended resource in a quota. Currently, only quota items with the prefix **requests.** is allowed for extended resources. The following is an example scenario of how to set resource quota for the GPU resource **nvidia.com/gpu**.

Procedure

1. Determine how many GPUs are available on a node in your cluster. For example:

```
# oc describe node ip-172-31-27-209.us-west-2.compute.internal | egrep
'Capacity|Allocatable|gpu'
```

Example output

```
openshift.com/gpu-accelerator=true
Capacity:
  nvidia.com/gpu: 2
Allocatable:
  nvidia.com/gpu: 2
  nvidia.com/gpu 0      0
```

In this example, 2 GPUs are available.

2. Create a **ResourceQuota** object to set a quota in the namespace **nvidia**. In this example, the quota is **1**:

Example output

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: gpu-quota
  namespace: nvidia
spec:
  hard:
    requests.nvidia.com/gpu: 1
```

3. Create the quota:

```
# oc create -f gpu-quota.yaml
```

Example output

```
resourcequota/gpu-quota created
```

4. Verify that the namespace has the correct quota set:

```
# oc describe quota gpu-quota -n nvidia
```

Example output

```

Name:          gpu-quota
Namespace:     nvidia
Resource       Used Hard
-----
requests.nvidia.com/gpu 0   1

```

5. Define a pod that asks for a single GPU. The following example definition file is called **gpu-pod.yaml**:

```

apiVersion: v1
kind: Pod
metadata:
  generateName: gpu-pod-
  namespace: nvidia
spec:
  restartPolicy: OnFailure
  containers:
  - name: rhel7-gpu-pod
    image: rhel7
    env:
      - name: NVIDIA_VISIBLE_DEVICES
        value: all
      - name: NVIDIA_DRIVER_CAPABILITIES
        value: "compute,utility"
      - name: NVIDIA_REQUIRE_CUDA
        value: "cuda>=5.0"
    command: ["sleep"]
    args: ["infinity"]
  resources:
    limits:
      nvidia.com/gpu: 1

```

6. Create the pod:

```
# oc create -f gpu-pod.yaml
```

7. Verify that the pod is running:

```
# oc get pods
```

Example output

```

NAME          READY  STATUS   RESTARTS  AGE
gpu-pod-s46h7  1/1    Running  0          1m

```

8. Verify that the quota **Used** counter is correct:

```
# oc describe quota gpu-quota -n nvidia
```

Example output

```
■
```

```

Name:          gpu-quota
Namespace:     nvidia
Resource       Used Hard
-----
requests.nvidia.com/gpu 1   1

```

- Attempt to create a second GPU pod in the **nvidia** namespace. This is technically available on the node because it has 2 GPUs:

```
# oc create -f gpu-pod.yaml
```

Example output

```

Error from server (Forbidden): error when creating "gpu-pod.yaml": pods "gpu-pod-f7z2w" is
forbidden: exceeded quota: gpu-quota, requested: requests.nvidia.com/gpu=1, used:
requests.nvidia.com/gpu=1, limited: requests.nvidia.com/gpu=1

```

This **Forbidden** error message is expected because you have a quota of 1 GPU and this pod tried to allocate a second GPU, which exceeds its quota.

8.1.7. Viewing a quota

You can view usage statistics related to any hard limits defined in a quota for a project by navigating in the web console to the project's **Quota** page.

You can also use the CLI to view quota details.

Procedure

- Get the list of quotas defined in the project. For example, for a project called **demoproject**:

```
$ oc get quota -n demoproject
```

Example output

```

NAME          AGE  REQUEST
LIMIT
besteffort    4s   pods: 1/2
compute-resources-time-bound 10m   pods: 0/2
limits.cpu: 0/1, limits.memory: 0/1Gi
core-object-counts      109s  configmaps: 2/10, persistentvolumeclaims: 1/4,
replicationcontrollers: 1/20, secrets: 9/10, services: 2/10

```

- Describe the quota you are interested in, for example the **core-object-counts** quota:

```
$ oc describe quota core-object-counts -n demoproject
```

Example output

```

Name: core-object-counts
Namespace: demoproject
Resource Used Hard

```



```

-----
configmaps 3 10
persistentvolumeclaims 0 4
replicationcontrollers 3 20
secrets 9 10
services 2 10

```

8.1.8. Configuring explicit resource quotas

Configure explicit resource quotas in a project request template to apply specific resource quotas in new projects.

Prerequisites

- Access to the cluster as a user with the cluster-admin role.
- Install the OpenShift CLI (**oc**).

Procedure

1. Add a resource quota definition to a project request template:

- If a project request template does not exist in a cluster:
 - a. Create a bootstrap project template and output it to a file called **template.yaml**:

```
$ oc adm create-bootstrap-project-template -o yaml > template.yaml
```

- b. Add a resource quota definition to **template.yaml**. The following example defines a resource quota named 'storage-consumption'. The definition must be added before the **parameters:** section in the template:

```

- apiVersion: v1
  kind: ResourceQuota
  metadata:
    name: storage-consumption
    namespace: ${PROJECT_NAME}
  spec:
    hard:
      persistentvolumeclaims: "10" 1
      requests.storage: "50Gi" 2
      gold.storageclass.storage.k8s.io/requests.storage: "10Gi" 3
      silver.storageclass.storage.k8s.io/requests.storage: "20Gi" 4
      silver.storageclass.storage.k8s.io/persistentvolumeclaims: "5" 5
      bronze.storageclass.storage.k8s.io/requests.storage: "0" 6
      bronze.storageclass.storage.k8s.io/persistentvolumeclaims: "0" 7

```

- 1 The total number of persistent volume claims in a project.
- 2 Across all persistent volume claims in a project, the sum of storage requested cannot exceed this value.
- 3 Across all persistent volume claims in a project, the sum of storage requested in the gold storage class cannot exceed this value.

- 4 Across all persistent volume claims in a project, the sum of storage requested in the silver storage class cannot exceed this value.
- 5 Across all persistent volume claims in a project, the total number of claims in the silver storage class cannot exceed this value.
- 6 Across all persistent volume claims in a project, the sum of storage requested in the bronze storage class cannot exceed this value. When this value is set to **0**, the bronze storage class cannot request storage.
- 7 Across all persistent volume claims in a project, the sum of storage requested in the bronze storage class cannot exceed this value. When this value is set to **0**, the bronze storage class cannot create claims.

- c. Create a project request template from the modified **template.yaml** file in the **openshift-config** namespace:

```
$ oc create -f template.yaml -n openshift-config
```



NOTE

To include the configuration as a **kubectrl.kubernetes.io/last-applied-configuration** annotation, add the **--save-config** option to the **oc create** command.

By default, the template is called **project-request**.

- If a project request template already exists within a cluster:



NOTE

If you declaratively or imperatively manage objects within your cluster by using configuration files, edit the existing project request template through those files instead.

- a. List templates in the **openshift-config** namespace:

```
$ oc get templates -n openshift-config
```

- b. Edit an existing project request template:

```
$ oc edit template <project_request_template> -n openshift-config
```

- c. Add a resource quota definition, such as the preceding **storage-consumption** example, into the existing template. The definition must be added before the **parameters:** section in the template.

- 2. If you created a project request template, reference it in the cluster's project configuration resource:

- a. Access the project configuration resource for editing:

- By using the web console:
 - i. Navigate to the **Administration → Cluster Settings** page.
 - ii. Click **Configuration** to view all configuration resources.
 - iii. Find the entry for **Project** and click **Edit YAML**.
- By using the CLI:
 - i. Edit the **project.config.openshift.io/cluster** resource:

```
$ oc edit project.config.openshift.io/cluster
```

- b. Update the **spec** section of the project configuration resource to include the **projectRequestTemplate** and **name** parameters. The following example references the default project request template name **project-request**:

```
apiVersion: config.openshift.io/v1
kind: Project
metadata:
# ...
spec:
  projectRequestTemplate:
    name: project-request
```

3. Verify that the resource quota is applied when projects are created:

- a. Create a project:

```
$ oc new-project <project_name>
```

- b. List the project's resource quotas:

```
$ oc get resourcequotas
```


- c. Describe the resource quota in detail:

```
$ oc describe resourcequotas <resource_quota_name>
```

8.2. RESOURCE QUOTAS ACROSS MULTIPLE PROJECTS

A multi-project quota, defined by a **ClusterResourceQuota** object, allows quotas to be shared across multiple projects. Resources used in each selected project are aggregated and that aggregate is used to limit resources across all the selected projects.

This guide describes how cluster administrators can set and manage resource quotas across multiple projects.



IMPORTANT

Do not run workloads in or share access to default projects. Default projects are reserved for running core cluster components.

The following default projects are considered highly privileged: **default**, **kube-public**, **kube-system**, **openshift**, **openshift-infra**, **openshift-node**, and other system-created projects that have the **openshift.io/run-level** label set to **0** or **1**. Functionality that relies on admission plugins, such as pod security admission, security context constraints, cluster resource quotas, and image reference resolution, does not work in highly privileged projects.

8.2.1. Selecting multiple projects during quota creation

When creating quotas, you can select multiple projects based on annotation selection, label selection, or both.

Procedure

1. To select projects based on annotations, run the following command:

```
$ oc create clusterquota for-user \
  --project-annotation-selector openshift.io/requester=<user_name> \
  --hard pods=10 \
  --hard secrets=20
```

This creates the following **ClusterResourceQuota** object:

```
apiVersion: quota.openshift.io/v1
kind: ClusterResourceQuota
metadata:
  name: for-user
spec:
  quota: 1
  hard:
    pods: "10"
    secrets: "20"
  selector:
    annotations: 2
    openshift.io/requester: <user_name>
    labels: null 3
status:
  namespaces: 4
  - namespace: ns-one
    status:
      hard:
        pods: "10"
        secrets: "20"
      used:
        pods: "1"
        secrets: "9"
  total: 5
  hard:
    pods: "10"
```

```

    secrets: "20"
  used:
    pods: "1"
    secrets: "9"

```

- 1 The **ResourceQuotaSpec** object that will be enforced over the selected projects.
- 2 A simple key-value selector for annotations.
- 3 A label selector that can be used to select projects.
- 4 A per-namespace map that describes current quota usage in each selected project.
- 5 The aggregate usage across all selected projects.

This multi-project quota document controls all projects requested by **<user_name>** using the default project request endpoint. You are limited to 10 pods and 20 secrets.

2. Similarly, to select projects based on labels, run this command:

```

$ oc create clusterresourcequota for-name \
  --project-label-selector=name=frontend \
  --hard=pods=10 --hard=secrets=20

```

- 1 Both **clusterresourcequota** and **clusterquota** are aliases of the same command. **for-name** is the name of the **ClusterResourceQuota** object.
- 2 To select projects by label, provide a key-value pair by using the format **--project-label-selector=key=value**.

This creates the following **ClusterResourceQuota** object definition:

```

apiVersion: quota.openshift.io/v1
kind: ClusterResourceQuota
metadata:
  creationTimestamp: null
  name: for-name
spec:
  quota:
    hard:
      pods: "10"
      secrets: "20"
  selector:
    annotations: null
    labels:
      matchLabels:
        name: frontend

```

8.2.2. Viewing applicable cluster resource quotas

A project administrator is not allowed to create or modify the multi-project quota that limits his or her project, but the administrator is allowed to view the multi-project quota documents that are applied to his or her project. The project administrator can do this via the **AppliedClusterResourceQuota**

resource.

Procedure

1. To view quotas applied to a project, run:

```
$ oc describe AppliedClusterResourceQuota
```

Example output

```
Name: for-user
Namespace: <none>
Created: 19 hours ago
Labels: <none>
Annotations: <none>
Label Selector: <null>
AnnotationSelector: map[openshift.io/requester:<user-name>]
Resource Used Hard
-----
pods      1    10
secrets   9    20
```

8.2.3. Selection granularity

Because of the locking consideration when claiming quota allocations, the number of active projects selected by a multi-project quota is an important consideration. Selecting more than 100 projects under a single multi-project quota can have detrimental effects on API server responsiveness in those projects.

CHAPTER 9. USING CONFIG MAPS WITH APPLICATIONS

Config maps allow you to decouple configuration artifacts from image content to keep containerized applications portable.

The following sections define config maps and how to create and use them.

9.1. UNDERSTANDING CONFIG MAPS

Many applications require configuration by using some combination of configuration files, command-line arguments, and environment variables. In OpenShift Container Platform, these configuration artifacts are decoupled from image content to keep containerized applications portable.

The **ConfigMap** object provides mechanisms to inject containers with configuration data while keeping containers agnostic of OpenShift Container Platform. A config map can be used to store fine-grained information like individual properties or coarse-grained information like entire configuration files or JSON blobs.

The **ConfigMap** object holds key-value pairs of configuration data that can be consumed in pods or used to store configuration data for system components such as controllers. For example:

ConfigMap Object Definition

```
kind: ConfigMap
apiVersion: v1
metadata:
  creationTimestamp: 2016-02-18T19:14:38Z
  name: example-config
  namespace: my-namespace
data: ❶
  example.property.1: hello
  example.property.2: world
  example.property.file: |-
    property.1=value-1
    property.2=value-2
    property.3=value-3
binaryData:
  bar: L3Jvb3QvMTAw ❷
```

❶ Contains the configuration data.

❷ Points to a file that contains non-UTF8 data, for example, a binary Java keystore file. Enter the file data in Base 64.



NOTE

You can use the **binaryData** field when you create a config map from a binary file, such as an image.

Configuration data can be consumed in pods in a variety of ways. A config map can be used to:

- Populate environment variable values in containers

- Set command-line arguments in a container
- Populate configuration files in a volume

Users and system components can store configuration data in a config map.

A config map is similar to a secret, but designed to more conveniently support working with strings that do not contain sensitive information.

9.1.1. Config map restrictions

A config map must be created before its contents can be consumed in pods.

Controllers can be written to tolerate missing configuration data. Consult individual components configured by using config maps on a case-by-case basis.

ConfigMap objects reside in a project.

They can only be referenced by pods in the same project.

The Kubelet only supports the use of a config map for pods it gets from the API server.

This includes any pods created by using the CLI, or indirectly from a replication controller. It does not include pods created by using the OpenShift Container Platform node's **--manifest-url** flag, its **--config** flag, or its REST API because these are not common ways to create pods.

Additional resources

- [Creating and using config maps](#)

9.2. USE CASES: CONSUMING CONFIG MAPS IN PODS

The following sections describe some uses cases when consuming **ConfigMap** objects in pods.

9.2.1. Populating environment variables in containers by using config maps

You can use config maps to populate individual environment variables in containers or to populate environment variables in containers from all keys that form valid environment variable names.

As an example, consider the following config map:

ConfigMap with two environment variables

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config ❶
  namespace: default ❷
data:
  special.how: very ❸
  special.type: charm ❹
```

- ❶ Name of the config map.

- 2 The project in which the config map resides. Config maps can only be referenced by pods in the same project.
- 3 4 Environment variables to inject.

ConfigMap with one environment variable

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: env-config 1
  namespace: default
data:
  log_level: INFO 2
```

- 1 Name of the config map.
- 2 Environment variable to inject.

Procedure

- You can consume the keys of this **ConfigMap** in a pod using **configMapKeyRef** sections.

Sample Pod specification configured to inject specific environment variables

```
apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  securityContext:
    runAsNonRoot: true
    seccompProfile:
      type: RuntimeDefault
  containers:
    - name: test-container
      image: gcr.io/google_containers/busybox
      command: [ "/bin/sh", "-c", "env" ]
      env: 1
        - name: SPECIAL_LEVEL_KEY 2
          valueFrom:
            configMapKeyRef:
              name: special-config 3
              key: special.how 4
        - name: SPECIAL_TYPE_KEY
          valueFrom:
            configMapKeyRef:
              name: special-config 5
              key: special.type 6
              optional: true 7
      envFrom: 8
        - configMapRef:
```

```

    name: env-config 9
    securityContext:
      allowPrivilegeEscalation: false
      capabilities:
        drop: [ALL]
    restartPolicy: Never

```

- 1 Stanza to pull the specified environment variables from a **ConfigMap**.
- 2 Name of a pod environment variable that you are injecting a key's value into.
- 3 5 Name of the **ConfigMap** to pull specific environment variables from.
- 4 6 Environment variable to pull from the **ConfigMap**.
- 7 Makes the environment variable optional. As optional, the pod will be started even if the specified **ConfigMap** and keys do not exist.
- 8 Stanza to pull all environment variables from a **ConfigMap**.
- 9 Name of the **ConfigMap** to pull all environment variables from.

When this pod is run, the pod logs will include the following output:

```

SPECIAL_LEVEL_KEY=very
log_level=INFO

```



NOTE

SPECIAL_TYPE_KEY=charm is not listed in the example output because **optional: true** is set.

9.2.2. Setting command-line arguments for container commands with config maps

You can use a config map to set the value of the commands or arguments in a container by using the Kubernetes substitution syntax **\$(VAR_NAME)**.

As an example, consider the following config map:

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config
  namespace: default
data:
  special.how: very
  special.type: charm

```

Procedure

- To inject values into a command in a container, you must consume the keys you want to use as environment variables. Then you can refer to them in a container's command using the **\$(VAR_NAME)** syntax.

Sample pod specification configured to inject specific environment variables

```

apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  securityContext:
    runAsNonRoot: true
    seccompProfile:
      type: RuntimeDefault
  containers:
    - name: test-container
      image: gcr.io/google_containers/busybox
      command: [ "/bin/sh", "-c", "echo $(SPECIAL_LEVEL_KEY) $(SPECIAL_TYPE_KEY)" ]
      env:
        - name: SPECIAL_LEVEL_KEY
          valueFrom:
            configMapKeyRef:
              name: special-config
              key: special.how
        - name: SPECIAL_TYPE_KEY
          valueFrom:
            configMapKeyRef:
              name: special-config
              key: special.type
      securityContext:
        allowPrivilegeEscalation: false
        capabilities:
          drop: [ALL]
      restartPolicy: Never

```

- 1 Inject the values into a command in a container using the keys you want to use as environment variables.

When this pod is run, the output from the echo command run in the test-container container is as follows:

```
very charm
```

9.2.3. Injecting content into a volume by using config maps

You can inject content into a volume by using config maps.

Example ConfigMap custom resource (CR)

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config
  namespace: default

```

```
data:
  special.how: very
  special.type: charm
```

Procedure

You have a couple different options for injecting content into a volume by using config maps.

- The most basic way to inject content into a volume by using a config map is to populate the volume with files where the key is the file name and the content of the file is the value of the key:

```
apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  securityContext:
    runAsNonRoot: true
    seccompProfile:
      type: RuntimeDefault
  containers:
  - name: test-container
    image: gcr.io/google_containers/busybox
    command: [ "/bin/sh", "-c", "cat", "/etc/config/special.how" ]
    volumeMounts:
    - name: config-volume
      mountPath: /etc/config
      securityContext:
        allowPrivilegeEscalation: false
        capabilities:
          drop: [ALL]
  volumes:
  - name: config-volume
    configMap:
      name: special-config 1
  restartPolicy: Never
```

- 1** File containing key.

When this pod is run, the output of the cat command will be:

```
very
```

- You can also control the paths within the volume where config map keys are projected:

```
apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  securityContext:
    runAsNonRoot: true
    seccompProfile:
      type: RuntimeDefault
```

```
containers:
  - name: test-container
    image: gcr.io/google_containers/busybox
    command: [ "/bin/sh", "-c", "cat", "/etc/config/path/to/special-key" ]
    volumeMounts:
      - name: config-volume
        mountPath: /etc/config
    securityContext:
      allowPrivilegeEscalation: false
      capabilities:
        drop: [ALL]
  volumes:
    - name: config-volume
      configMap:
        name: special-config
        items:
          - key: special.how
            path: path/to/special-key ❶
      restartPolicy: Never
```

❶ Path to config map key.

When this pod is run, the output of the cat command will be:

```
very
```

CHAPTER 10. MONITORING PROJECT AND APPLICATION METRICS USING THE DEVELOPER PERSPECTIVE

The **Observe** view in the **Developer** perspective provides options to monitor your project or application metrics, such as CPU, memory, and bandwidth usage, and network related information.

10.1. PREREQUISITES

- You have [created and deployed applications on OpenShift Container Platform](#) .
- You have [logged in to the web console](#) and have switched to [the Developer perspective](#).

10.2. MONITORING YOUR PROJECT METRICS

After you create applications in your project and deploy them, you can use the **Developer** perspective in the web console to see the metrics for your project.

Procedure

1. Go to **Observe** to see the **Dashboard**, **Metrics**, **Alerts**, and **Events** for your project.
2. Optional: Use the **Dashboard** tab to see graphs depicting the following application metrics:
 - CPU usage
 - Memory usage
 - Bandwidth consumption
 - Network-related information such as the rate of transmitted and received packets and the rate of dropped packets.

In the **Dashboard** tab, you can access the Kubernetes compute resources dashboards.



NOTE


In the **Dashboard** list, the **Kubernetes / Compute Resources / Namespace (Pods)** dashboard is selected by default.

Use the following options to see further details:

- Select a dashboard from the **Dashboard** list to see the filtered metrics. All dashboards produce additional sub-menus when selected, except **Kubernetes / Compute Resources / Namespace (Pods)**.
- Select an option from the **Time Range** list to determine the time frame for the data being captured.
- Set a custom time range by selecting **Custom time range** from the **Time Range** list. You can input or select the **From** and **To** dates and times. Click **Save** to save the custom time range.
- Select an option from the **Refresh Interval** list to determine the time period after which the data is refreshed.


- Hover your cursor over the graphs to see specific details for your pod.
 - Click **Inspect** located in the upper-right corner of every graph to see any particular graph details. The graph details appear in the **Metrics** tab.
3. Optional: Use the **Metrics** tab to query for the required project metric.

Figure 10.1. Monitoring metrics



- In the **Select Query** list, select an option to filter the required details for your project. The filtered metrics for all the application pods in your project are displayed in the graph. The pods in your project are also listed below.
 - From the list of pods, clear the colored square boxes to remove the metrics for specific pods to further filter your query result.
 - Click **Show PromQL** to see the Prometheus query. You can further modify this query with the help of prompts to customize the query and filter the metrics you want to see for that namespace.
 - Use the drop-down list to set a time range for the data being displayed. You can click **Reset Zoom** to reset it to the default time range.
 - Optional: In the **Select Query** list, select **Custom Query** to create a custom Prometheus query and filter relevant metrics.
4. Optional: Use the **Alerts** tab to do the following tasks:
- See the rules that trigger alerts for the applications in your project.
 - Identify the alerts firing in the project.
 - Silence such alerts if required.

Figure 10.2. Monitoring alerts



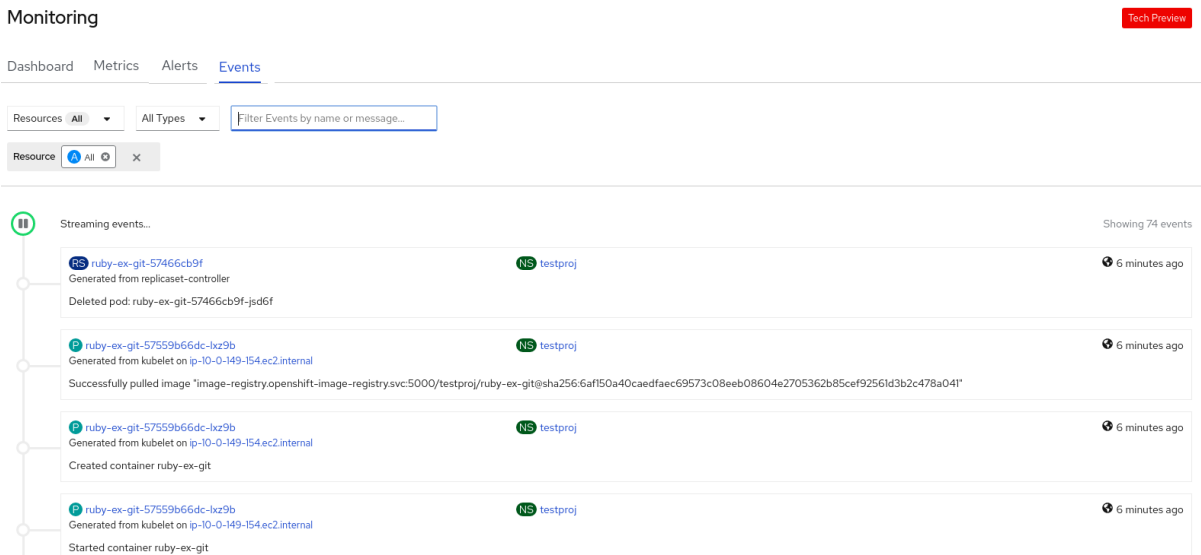
Use the following options to see further details:

- Use the **Filter** list to filter the alerts by their **Alert State** and **Severity**.
- Click on an alert to go to the details page for that alert. In the **Alerts Details** page, you can click **View Metrics** to see the metrics for the alert.
- Use the **Notifications** toggle adjoining an alert rule to silence all the alerts for that rule, and then select the duration for which the alerts will be silenced from the **Silence for** list. You must have the permissions to edit alerts to see the **Notifications** toggle.

- Use the Options menu adjoining an alert rule to see the details of the alerting rule.

5. Optional: Use the **Events** tab to see the events for your project.

Figure 10.3. Monitoring events



You can filter the displayed events using the following options:

- In the **Resources** list, select a resource to see events for that resource.
- In the **All Types** list, select a type of event to see events relevant to that type.
- Search for specific events using the **Filter events by names or messages** field.


10.3. MONITORING YOUR APPLICATION METRICS

After you create applications in your project and deploy them, you can use the **Topology** view in the **Developer** perspective to see the alerts and metrics for your application. Critical and warning alerts for your application are indicated on the workload node in the **Topology** view.

Procedure

To see the alerts for your workload:


1. In the **Topology** view, click the workload to see the workload details in the right panel.
2. Click the **Observe** tab to see the critical and warning alerts for the application; graphs for metrics, such as CPU, memory, and bandwidth usage; and all the events for the application.



NOTE

Only critical and warning alerts in the **Firing** state are displayed in the **Topology** view. Alerts in the **Silenced**, **Pending** and **Not Firing** states are not displayed.

Figure 10.4. Monitoring application metrics



- a. Click the alert listed in the right panel to see the alert details in the **Alert Details** page.
- b. Click any of the charts to go to the **Metrics** tab to see the detailed metrics for the application.
- c. Click **View monitoring dashboard** to see the monitoring dashboard for that application.


10.4. IMAGE VULNERABILITIES BREAKDOWN

In the **Developer** perspective, the project dashboard shows the **Image Vulnerabilities** link in the **Status** section. Using this link, you can view the **Image Vulnerabilities breakdown** window, which includes details regarding vulnerable container images and fixable container images. The icon color indicates severity:

- Red: High priority. Fix immediately.
- Orange: Medium priority. Can be fixed after high-priority vulnerabilities.
- Yellow: Low priority. Can be fixed after high and medium-priority vulnerabilities.

Based on the severity level, you can prioritize vulnerabilities and fix them in an organized manner.

Figure 10.5. Viewing image vulnerabilities



10.5. MONITORING YOUR APPLICATION AND IMAGE VULNERABILITIES METRICS

After you create applications in your project and deploy them, use the **Developer** perspective in the web console to see the metrics for your application dependency vulnerabilities across your cluster. The metrics help you to analyze the following image vulnerabilities in detail:

- Total count of vulnerable images in a selected project
- Severity-based counts of all vulnerable images in a selected project
- Drilldown into severity to obtain the details, such as count of vulnerabilities, count of fixable vulnerabilities, and number of affected pods for each vulnerable image

Prerequisites

- You have installed the Red Hat Quay Container Security operator from the Operator Hub.

**NOTE**

The Red Hat Quay Container Security operator detects vulnerabilities by scanning the images that are in the quay registry.

Procedure

1. For a general overview of the image vulnerabilities, on the navigation panel of the **Developer** perspective, click **Project** to see the project dashboard.
2. Click **Image Vulnerabilities** in the **Status** section. The window that opens displays details such as **Vulnerable Container Images** and **Fixable Container Images**.
3. For a detailed vulnerabilities overview, click the **Vulnerabilities** tab on the project dashboard.
 - a. To get more detail about an image, click its name.
 - b. View the default graph with all types of vulnerabilities in the **Details** tab.
 - c. Optional: Click the toggle button to view a specific type of vulnerability. For example, click **App dependency** to see vulnerabilities specific to application dependency.
 - d. Optional: You can filter the list of vulnerabilities based on their **Severity** and **Type** or sort them by **Severity**, **Package**, **Type**, **Source**, **Current Version**, and **Fixed in Version**.
 - e. Click a **Vulnerability** to get its associated details:
 - **Base image** vulnerabilities display information from a Red Hat Security Advisory (RHSA).
 - **App dependency** vulnerabilities display information from the Snyk security application.

10.6. ADDITIONAL RESOURCES

- [About OpenShift Container Platform monitoring](#)

CHAPTER 11. MONITORING APPLICATION HEALTH BY USING HEALTH CHECKS

In software systems, components can become unhealthy due to transient issues such as temporary connectivity loss, configuration errors, or problems with external dependencies. OpenShift Container Platform applications have a number of options to detect and handle unhealthy containers.

11.1. UNDERSTANDING HEALTH CHECKS

A health check periodically performs diagnostics on a running container using any combination of the readiness, liveness, and startup health checks.

You can include one or more probes in the specification for the pod that contains the container which you want to perform the health checks.



NOTE

If you want to add or edit health checks in an existing pod, you must edit the pod **DeploymentConfig** object or use the web console. You cannot use the CLI to add or edit health checks for an existing pod.

Readiness probe

A *readiness probe* determines if a container is ready to accept service requests. If the readiness probe fails for a container, the kubelet removes the pod from the list of available service endpoints. After a failure, the probe continues to examine the pod. If the pod becomes available, the kubelet adds the pod to the list of available service endpoints.

Liveness health check

A *liveness probe* determines if a container is still running. If the liveness probe fails due to a condition such as a deadlock, the kubelet kills the container. The pod then responds based on its restart policy. For example, a liveness probe on a pod with a **restartPolicy** of **Always** or **OnFailure** kills and restarts the container.

Startup probe

A *startup probe* indicates whether the application within a container is started. All other probes are disabled until the startup succeeds. If the startup probe does not succeed within a specified time period, the kubelet kills the container, and the container is subject to the pod **restartPolicy**. Some applications can require additional startup time on their first initialization. You can use a startup probe with a liveness or readiness probe to delay that probe long enough to handle lengthy start-up time using the **failureThreshold** and **periodSeconds** parameters.

For example, you can add a startup probe, with a **failureThreshold** of 30 failures and a **periodSeconds** of 10 seconds (30 * 10s = 300s) for a maximum of 5 minutes, to a liveness probe. After the startup probe succeeds the first time, the liveness probe takes over.

You can configure liveness, readiness, and startup probes with any of the following types of tests:

- **HTTP GET**: When using an HTTP **GET** test, the test determines the healthiness of the container by using a web hook. The test is successful if the HTTP response code is between **200** and **399**. You can use an HTTP **GET** test with applications that return HTTP status codes when completely initialized.

- **Container Command:** When using a container command test, the probe executes a command inside the container. The probe is successful if the test exits with a **0** status.
- **TCP socket:** When using a TCP socket test, the probe attempts to open a socket to the container. The container is only considered healthy if the probe can establish a connection. You can use a TCP socket test with applications that do not start listening until initialization is complete.

You can configure several fields to control the behavior of a probe:

- **initialDelaySeconds:** The time, in seconds, after the container starts before the probe can be scheduled. The default is 0.
- **periodSeconds:** The delay, in seconds, between performing probes. The default is **10**. This value must be greater than **timeoutSeconds**.
- **timeoutSeconds:** The number of seconds of inactivity after which the probe times out and the container is assumed to have failed. The default is **1**. This value must be lower than **periodSeconds**.
- **successThreshold:** The number of times that the probe must report success after a failure to reset the container status to successful. The value must be **1** for a liveness probe. The default is **1**.
- **failureThreshold:** The number of times that the probe is allowed to fail. The default is 3. After the specified attempts:
 - for a liveness probe, the container is restarted
 - for a readiness probe, the pod is marked **Unready**
 - for a startup probe, the container is killed and is subject to the pod's **restartPolicy**

11.1.1. Example probes

The following are samples of different probes as they would appear in an object specification.

Sample readiness probe with a container command readiness probe in a pod spec

```
apiVersion: v1
kind: Pod
metadata:
  labels:
    test: health-check
  name: my-application
# ...
spec:
  containers:
  - name: goproxy-app ❶
    args:
    image: registry.k8s.io/goproxy:0.1 ❷
    readinessProbe: ❸
      exec: ❹
        command: ❺
```

```

- cat
- /tmp/healthy
# ...

```

- 1 The container name.
- 2 The container image to deploy.
- 3 A readiness probe.
- 4 A container command test.
- 5 The commands to execute on the container.

Sample container command startup probe and liveness probe with container command tests in a pod spec

```

apiVersion: v1
kind: Pod
metadata:
  labels:
    test: health-check
  name: my-application
# ...
spec:
  containers:
  - name: goproxy-app 1
    args:
    image: registry.k8s.io/goproxy:0.1 2
    livenessProbe: 3
      httpGet: 4
        scheme: HTTPS 5
        path: /healthz
        port: 8080 6
        httpHeaders:
        - name: X-Custom-Header
          value: Awesome
    startupProbe: 7
      httpGet: 8
        path: /healthz
        port: 8080 9
      failureThreshold: 30 10
      periodSeconds: 10 11
# ...

```

- 1 The container name.
- 2 Specify the container image to deploy.
- 3 A liveness probe.
- 4 An HTTP **GET** test.

- 5 The internet scheme: **HTTP** or **HTTPS**. The default value is **HTTP**.
- 6 The port on which the container is listening.
- 7 A startup probe.
- 8 An HTTP **GET** test.
- 9 The port on which the container is listening.
- 10 The number of times to try the probe after a failure.
- 11 The number of seconds to perform the probe.

Sample liveness probe with a container command test that uses a timeout in a pod spec

```
apiVersion: v1
kind: Pod
metadata:
  labels:
    test: health-check
  name: my-application
# ...
spec:
  containers:
    - name: goproxy-app 1
      args:
        image: registry.k8s.io/goproxy:0.1 2
      livenessProbe: 3
        exec: 4
          command: 5
            - /bin/bash
            - '-c'
            - timeout 60 /opt/eap/bin/livenessProbe.sh
        periodSeconds: 10 6
        successThreshold: 1 7
        failureThreshold: 3 8
# ...
```

- 1 The container name.
- 2 Specify the container image to deploy.
- 3 The liveness probe.
- 4 The type of probe, here a container command probe.
- 5 The command line to execute inside the container.
- 6 How often in seconds to perform the probe.
- 7 The number of consecutive successes needed to show success after a failure.
- 8 The number of times to try the probe after a failure.

Sample readiness probe and liveness probe with a TCP socket test in a deployment

```
kind: Deployment
apiVersion: apps/v1
metadata:
  labels:
    test: health-check
  name: my-application
spec:
  # ...
  template:
    spec:
      containers:
      - resources: {}
        readinessProbe: ❶
          tcpSocket:
            port: 8080
          timeoutSeconds: 1
          periodSeconds: 10
          successThreshold: 1
          failureThreshold: 3
        terminationMessagePath: /dev/termination-log
        name: ruby-ex
        livenessProbe: ❷
          tcpSocket:
            port: 8080
          initialDelaySeconds: 15
          timeoutSeconds: 1
          periodSeconds: 10
          successThreshold: 1
          failureThreshold: 3
      # ...
```

❶ The readiness probe.

❷ The liveness probe.

11.2. CONFIGURING HEALTH CHECKS USING THE CLI

To configure readiness, liveness, and startup probes, add one or more probes to the specification for the pod that contains the container which you want to perform the health checks



NOTE

If you want to add or edit health checks in an existing pod, you must edit the pod **DeploymentConfig** object or use the web console. You cannot use the CLI to add or edit health checks for an existing pod.

Procedure

To add probes for a container:

1. Create a **Pod** object to add one or more probes:


```

apiVersion: v1
kind: Pod
metadata:
  labels:
    test: health-check
  name: my-application
spec:
  containers:
    - name: my-container ❶
      args:
      image: registry.k8s.io/goproxy:0.1 ❷
      livenessProbe: ❸
        tcpSocket: ❹
          port: 8080 ❺
        initialDelaySeconds: 15 ❻
        periodSeconds: 20 ❼
        timeoutSeconds: 10 ❽
      readinessProbe: ❾
        httpGet: ❿
          host: my-host 11
          scheme: HTTPS 12
          path: /healthz
          port: 8080 13
      startupProbe: 14
        exec: 15
          command: 16
            - cat
            - /tmp/healthy
        failureThreshold: 30 17
        periodSeconds: 20 18
        timeoutSeconds: 10 19

```

- ❶ Specify the container name.
- ❷ Specify the container image to deploy.
- ❸ Optional: Create a Liveness probe.
- ❹ Specify a test to perform, here a TCP Socket test.
- ❺ Specify the port on which the container is listening.
- ❻ Specify the time, in seconds, after the container starts before the probe can be scheduled.
- ❼ Specify the number of seconds to perform the probe. The default is **10**. This value must be greater than **timeoutSeconds**.
- ❽ Specify the number of seconds of inactivity after which the probe is assumed to have failed. The default is **1**. This value must be lower than **periodSeconds**.
- ❾ Optional: Create a Readiness probe.
- ❿ Specify the type of test to perform, here an HTTP test.

- 11 Specify a host IP address. When **host** is not defined, the **PodIP** is used.
- 12 Specify **HTTP** or **HTTPS**. When **scheme** is not defined, the **HTTP** scheme is used.
- 13 Specify the port on which the container is listening.
- 14 Optional: Create a Startup probe.
- 15 Specify the type of test to perform, here an Container Execution probe.
- 16 Specify the commands to execute on the container.
- 17 Specify the number of times to try the probe after a failure.
- 18 Specify the number of seconds to perform the probe. The default is **10**. This value must be greater than **timeoutSeconds**.
- 19 Specify the number of seconds of inactivity after which the probe is assumed to have failed. The default is **1**. This value must be lower than **periodSeconds**.



NOTE

If the **initialDelaySeconds** value is lower than the **periodSeconds** value, the first Readiness probe occurs at some point between the two periods due to an issue with timers.

The **timeoutSeconds** value must be lower than the **periodSeconds** value.

2. Create the **Pod** object:

```
$ oc create -f <file-name>.yaml
```

3. Verify the state of the health check pod:

```
$ oc describe pod my-application
```

Example output

```
Events:
  Type    Reason      Age   From              Message
  ----    -
  Normal  Scheduled   9s    default-scheduler Successfully assigned openshift-logging/liveness-exec to ip-10-0-143-40.ec2.internal
  Normal  Pulling     2s    kubelet, ip-10-0-143-40.ec2.internal pulling image "registry.k8s.io/liveness"
  Normal  Pulled      1s    kubelet, ip-10-0-143-40.ec2.internal Successfully pulled image "registry.k8s.io/liveness"
  Normal  Created     1s    kubelet, ip-10-0-143-40.ec2.internal Created container
  Normal  Started     1s    kubelet, ip-10-0-143-40.ec2.internal Started container
```

The following is the output of a failed probe that restarted a container:

Sample Liveness check output with unhealthy container

```
$ oc describe pod pod1
```

Example output

```
....

Events:
  Type    Reason          Age          From              Message
  ----    -
  Normal  Scheduled       <unknown>    kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Successfully
assigned aaa/liveness-http to ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj
  Normal  AddedInterface  47s          multus            Add eth0
[10.129.2.11/23]
  Normal  Pulled          46s          kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Successfully pulled image "registry.k8s.io/liveness" in 773.406244ms
  Normal  Pulled          28s          kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Successfully pulled image "registry.k8s.io/liveness" in 233.328564ms
  Normal  Created         10s (x3 over 46s) kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Created container liveness
  Normal  Started         10s (x3 over 46s) kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Started container liveness
  Warning Unhealthy       10s (x6 over 34s) kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-
snzrj Liveness probe failed: HTTP probe failed with statuscode: 500
  Normal  Killing         10s (x2 over 28s) kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Container liveness failed liveness probe, will be restarted
  Normal  Pulling         10s (x3 over 47s) kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Pulling image "registry.k8s.io/liveness"
  Normal  Pulled          10s          kubelet, ci-ln-37hz77b-f76d1-wdpjv-worker-b-snzrj  Successfully pulled image "registry.k8s.io/liveness" in 244.116568ms
```

11.3. MONITORING APPLICATION HEALTH USING THE DEVELOPER PERSPECTIVE

You can use the **Developer** perspective to add three types of health probes to your container to ensure that your application is healthy:

- Use the Readiness probe to check if the container is ready to handle requests.
- Use the Liveness probe to check if the container is running.
- Use the Startup probe to check if the application within the container has started.

You can add health checks either while creating and deploying an application, or after you have deployed an application.

11.4. EDITING HEALTH CHECKS USING THE DEVELOPER PERSPECTIVE

You can use the **Topology** view to edit health checks added to your application, modify them, or add more health checks.

Prerequisites

- You have switched to the **Developer** perspective in the web console.
- You have created and deployed an application on OpenShift Container Platform using the **Developer** perspective.
- You have added health checks to your application.

Procedure

1. In the **Topology** view, right-click your application and select **Edit Health Checks**. Alternatively, in the side panel, click the **Actions** drop-down list and select **Edit Health Checks**.
2. In the **Edit Health Checks** page:
 - To remove a previously added health probe, click the **Remove** icon adjoining it.
 - To edit the parameters of an existing probe:
 - a. Click the **Edit Probe** link next to a previously added probe to see the parameters for the probe.
 - b. Modify the parameters as required, and click the check mark to save your changes.
 - To add a new health probe, in addition to existing health checks, click the add probe links. For example, to add a Liveness probe that checks if your container is running:
 - a. Click **Add Liveness Probe**, to see a form containing the parameters for the probe.
 - b. Edit the probe parameters as required.
3. Click **Save** to save your modifications and add the additional probes to your container. You are redirected to the **Topology** view.
4. In the side panel, verify that the probes have been added by clicking on the deployed pod under the **Pods** section.
5. In the **Pod Details** page, click the listed container in the **Containers** section.
6. In the **Container Details** page, verify that the Liveness probe - **HTTP Get 10.129.4.65:8080/** has been added to the container, in addition to the earlier existing probes.



NOTE

The **Timeout** value must be lower than the **Period** value. The **Timeout** default value is **1**. The **Period** default value is **10**.

- c. Click the check mark at the bottom of the form. The **Liveness Probe Added** message is displayed.

11.5. MONITORING HEALTH CHECK FAILURES USING THE DEVELOPER PERSPECTIVE

In case an application health check fails, you can use the **Topology** view to monitor these health check violations.

Prerequisites

Prerequisites

- You have switched to the **Developer** perspective in the web console.
- You have created and deployed an application on OpenShift Container Platform using the **Developer** perspective.
- You have added health checks to your application.

Procedure

1. In the **Topology** view, click on the application node to see the side panel.
2. Click the **Observe** tab to see the health check failures in the **Events (Warning)** section.
3. Click the down arrow adjoining **Events (Warning)** to see the details of the health check failure.

Additional resources

- For details on switching to the **Developer** perspective in the web console, see [About the Developer perspective](#).
- For details on adding health checks while creating and deploying an application, see **Advanced Options** in the [Creating applications using the Developer perspective](#) section.

CHAPTER 12. EDITING APPLICATIONS

You can edit the configuration and the source code of the application you create using the **Topology** view.

12.1. PREREQUISITES

- You have the appropriate [roles and permissions](#) in a project to create and modify applications in OpenShift Container Platform.
- You have [created and deployed an application on OpenShift Container Platform using the Developer perspective](#).
- You have [logged in to the web console](#) and have switched to [the Developer perspective](#).

12.2. EDITING THE SOURCE CODE OF AN APPLICATION USING THE DEVELOPER PERSPECTIVE

You can use the **Topology** view in the **Developer** perspective to edit the source code of your application.

Procedure

- In the **Topology** view, click the **Edit Source code** icon, displayed at the bottom-right of the deployed application, to access your source code and modify it.



NOTE

This feature is available only when you create applications using the **From Git**, **From Catalog**, and the **From Dockerfile** options.

If the **Eclipse Che** Operator is installed in your cluster, a Che workspace () is created and you are directed to the workspace to edit your source code. If it is not installed, you will be directed to the Git repository () your source code is hosted in.

12.3. EDITING THE APPLICATION CONFIGURATION USING THE DEVELOPER PERSPECTIVE

You can use the **Topology** view in the **Developer** perspective to edit the configuration of your application.



NOTE

Currently, only configurations of applications created by using the **From Git**, **Container Image**, **From Catalog**, or **From Dockerfile** options in the **Add** workflow of the **Developer** perspective can be edited. Configurations of applications created by using the CLI or the **YAML** option from the **Add** workflow cannot be edited.


Prerequisites

Ensure that you have created an application using the **From Git**, **Container Image**, **From Catalog**, or **From Dockerfile** options in the **Add** workflow.

Procedure

1. After you have created an application and it is displayed in the **Topology** view, right-click the application to see the edit options available.

Figure 12.1. Edit application



2. Click **Edit application-name** to see the **Add** workflow you used to create the application. The form is pre-populated with the values you had added while creating the application.
3. Edit the necessary values for the application.



NOTE

You cannot edit the **Name** field in the **General** section, the CI/CD pipelines, or the **Create a route to the application** field in the **Advanced Options** section.

- Click **Save** to restart the build and deploy a new image.

Figure 12.2. Edit and redeploy application

The screenshot displays the OpenShift console interface for the 'nodejs-ex-git' application. The left sidebar shows a diagram of the application architecture, including a 'nodejs-ex-git' service and a 'django-ex-git' build. The main panel shows the 'Resources' tab with details for Pods, Builds, and Services.

nodejs-ex-git [Actions]

Details **Resources** Monitoring

Pods

Pod Name	Status	View logs
nodejs-ex-git-57fd9cc6d8-snzsf	Running	View logs

Builds

Build Name	Status	View logs
nodejs-ex-git	Start Build	
Build #2	is complete (a few seconds ago)	View logs
Build #1	is complete (5 hours ago)	View logs

Services

Service Name	Port
nodejs-ex-git	Service port: 8080-tcp → Pod Port: 8080

CHAPTER 13. PRUNING OBJECTS TO RECLAIM RESOURCES

Over time, API objects created in OpenShift Container Platform can accumulate in the cluster's etcd data store through normal user operations, such as when building and deploying applications.

Cluster administrators can periodically prune older versions of objects from the cluster that are no longer required. For example, by pruning images you can delete older images and layers that are no longer in use, but are still taking up disk space.

13.1. BASIC PRUNING OPERATIONS

The CLI groups prune operations under a common parent command:

```
$ oc adm prune <object_type> <options>
```

This specifies:

- The **<object_type>** to perform the action on, such as **groups, builds, deployments, or images**.
- The **<options>** supported to prune that object type.

13.2. PRUNING GROUPS

To prune groups records from an external provider, administrators can run the following command:

```
$ oc adm prune groups \
  --sync-config=path/to/sync/config [<options>]
```

Table 13.1. **oc adm prune groups** flags

Options	Description
--confirm	Indicate that pruning should occur, instead of performing a dry-run.
--blacklist	Path to the group blacklist file.
--whitelist	Path to the group whitelist file.
--sync-config	Path to the synchronization configuration file.

Procedure

1. To see the groups that the prune command deletes, run the following command:

```
$ oc adm prune groups --sync-config=ldap-sync-config.yaml
```

2. To perform the prune operation, add the **--confirm** flag:

```
$ oc adm prune groups --sync-config=ldap-sync-config.yaml --confirm
```

13.3. PRUNING DEPLOYMENT RESOURCES

You can prune resources associated with deployments that are no longer required by the system, due to age and status.

The following command prunes replication controllers associated with **DeploymentConfig** objects:

```
$ oc adm prune deployments [<options>]
```



NOTE

To also prune replica sets associated with **Deployment** objects, use the **--replica-sets** flag. This flag is currently a Technology Preview feature.

Table 13.2. **oc adm prune deployments** flags

Option	Description
--confirm	Indicate that pruning should occur, instead of performing a dry-run.
--keep-complete=<N>	Per the DeploymentConfig object, keep the last N replication controllers that have a status of Complete and replica count of zero. The default is 5 .
--keep-failed=<N>	Per the DeploymentConfig object, keep the last N replication controllers that have a status of Failed and replica count of zero. The default is 1 .
--keep-younger-than=<duration>	Do not prune any replication controller that is younger than <duration> relative to the current time. Valid units of measurement include nanoseconds (ns), microseconds (us), milliseconds (ms), seconds (s), minutes (m), and hours (h). The default is 60m .
--orphans	Prune all replication controllers that no longer have a DeploymentConfig object, has status of Complete or Failed , and has a replica count of zero.
--replica-sets=true false	If true, replica sets are included in the pruning process. The default is false.  IMPORTANT This flag is a Technology Preview feature.

Procedure

1. To see what a pruning operation would delete, run the following command:

```
$ oc adm prune deployments --orphans --keep-complete=5 --keep-failed=1 \
--keep-younger-than=60m
```

- To actually perform the prune operation, add the **--confirm** flag:

```
$ oc adm prune deployments --orphans --keep-complete=5 --keep-failed=1 \
--keep-younger-than=60m --confirm
```

13.4. PRUNING BUILDS

To prune builds that are no longer required by the system due to age and status, administrators can run the following command:

```
$ oc adm prune builds [<options>]
```

Table 13.3. **oc adm prune builds** flags

Option	Description
--confirm	Indicate that pruning should occur, instead of performing a dry-run.
--orphans	Prune all builds whose build configuration no longer exists, status is complete, failed, error, or canceled.
--keep-complete=<N>	Per build configuration, keep the last N builds whose status is complete. The default is 5 .
--keep-failed=<N>	Per build configuration, keep the last N builds whose status is failed, error, or canceled. The default is 1 .
--keep-younger-than=<duration>	Do not prune any object that is younger than <duration> relative to the current time. The default is 60m .


Procedure

- To see what a pruning operation would delete, run the following command:

```
$ oc adm prune builds --orphans --keep-complete=5 --keep-failed=1 \
--keep-younger-than=60m
```

- To actually perform the prune operation, add the **--confirm** flag:

```
$ oc adm prune builds --orphans --keep-complete=5 --keep-failed=1 \
--keep-younger-than=60m --confirm
```



NOTE

Developers can enable automatic build pruning by modifying their build configuration.

Additional resources

- [Performing advanced builds → Pruning builds](#)

13.5. AUTOMATICALLY PRUNING IMAGES

Images from the OpenShift image registry that are no longer required by the system because of age, status, or exceed limits being automatically pruned. As a cluster administrator, can configure or suspend the pruning Custom Resource (CR).

Prerequisites

- You have access to an OpenShift Container Platform cluster using an account with cluster administrator permissions.
- Install the **oc** CLI.



IMPORTANT

The behavior of the Image Registry Operator for managing the pruner is independent to the **managementState** specified on the **ClusterOperator** object of the Image Registry Operator. If the Image Registry Operator is not in the **Managed** state, the image pruner can still be configured and managed by the Pruning Custom Resource.

However, the **managementState** of the Image Registry Operator alters the behavior of the deployed image pruner job:

- **Managed:** the **--prune-registry** flag for the image pruner is set to **true**.
- **Removed:** the **--prune-registry** flag for the image pruner is set to **false**, meaning it only prunes image metadata in etcd.

Procedure

- Verify that the object named **imagepruners.imageregistry.operator.openshift.io/cluster** contains the following **spec** and **status** fields:

```
spec:
  schedule: 0 0 * * *
  suspend: false
  keepTagRevisions: 3
  keepYoungerThanDuration: 60m
  keepYoungerThan: 3600000000000
  resources: {}
  affinity: {}
  nodeSelector: {}
  tolerations: []
  successfulJobsHistoryLimit: 3
  failedJobsHistoryLimit: 3
status:
  observedGeneration: 2
  conditions:
  - type: Available
    status: "True"
    lastTransitionTime: 2019-10-09T03:13:45
```

```

reason: Ready
message: "Periodic image pruner has been created."
- type: Scheduled
  status: "True"
  lastTransitionTime: 2019-10-09T03:13:45
reason: Scheduled
message: "Image pruner job has been scheduled."
- type: Failed
  status: "False"
  lastTransitionTime: 2019-10-09T03:13:45
reason: Succeeded
message: "Most recent image pruning job succeeded."

```

- **schedule:** **CronJob** formatted schedule. This is an optional field, default is daily at midnight.
- **suspend:** If set to **true**, the **CronJob** running pruning is suspended. This is an optional field, default is **false**. The initial value on new clusters is **false**.
- **keepTagRevisions:** The number of revisions per tag to keep. This is an optional field, default is **3**. The initial value is **3**.
- **keepYoungerThanDuration:** Retain images younger than this duration. This is an optional field. If a value is not specified, either **keepYoungerThan** or the default value **60m** (60 minutes) is used.
- **keepYoungerThan:** Deprecated. The same as **keepYoungerThanDuration**, but the duration is specified as an integer in nanoseconds. This is an optional field. When **keepYoungerThanDuration** is set, this field is ignored.
- **resources:** Standard pod resource requests and limits. This is an optional field.
- **affinity:** Standard pod affinity. This is an optional field.
- **nodeSelector:** Standard pod node selector. This is an optional field.
- **tolerations:** Standard pod tolerations. This is an optional field.
- **successfulJobsHistoryLimit:** The maximum number of successful jobs to retain. Must be greater than or equal to **1** to ensure metrics are reported. This is an optional field, default is **3**. The initial value is **3**.
- **failedJobsHistoryLimit:** The maximum number of failed jobs to retain. Must be greater than or equal **1** to ensure metrics are reported. This is an optional field, default is **3**. The initial value is **3**.
- **observedGeneration:** The generation observed by the Operator.
- **conditions:** The standard condition objects with the following types:
 - **Available:** Indicates if the pruning job has been created. Reasons can be **Ready** or **Error**.
 - **Scheduled:** Indicates if the next pruning job has been scheduled. Reasons can be **Scheduled**, **Suspended**, or **Error**.
 - **Failed:** Indicates if the most recent pruning job failed.

13.6. MANUALLY PRUNING IMAGES

The pruning custom resource enables automatic image pruning for the images from the OpenShift image registry. Administrators can manually prune images with the **oc adm prune images** **<image_prune_option>** command. For example:

```
$ oc adm prune images <image_prune_option> 1
```

- 1 For more information about available pruning options, see "Manual image pruning command options".

This command removes images that are no longer required by the system.

Depending on your needs, you can prune images based on their age and tag history, or prune images that cause a project to exceed its defined storage limits.

13.6.1. Considerations when pruning images

Consider the following information before manually pruning images:

- Pruning with the **--namespace** flag does not remove images. It only removes image streams, because images are cluster-scoped resources. Limiting pruning to a particular namespace makes it impossible to calculate current usage.
- By default, the integrated registry caches metadata of blobs to reduce the number of requests to storage, and to increase request-processing speed. Pruning does not update the integrated registry cache. Images that still contain pruned layers after pruning will be broken because the pruned layers that have metadata in the cache will not be pushed. Therefore, you must redeploy the registry to clear the cache after pruning:

```
$ oc rollout restart deployment/image-registry -n openshift-image-registry
```

- If the integrated registry uses a Redis cache, you must clean the database manually.
- If redeploying the registry after pruning is not an option, then you must permanently disable the cache.
- **oc adm prune images** operations require a route for your registry. Registry routes are not created by default.

13.6.2. Limitations when pruning images

The following limitations apply when pruning an image:

- Pruning images from external registries is unsupported.
- When an image is pruned, all references to the image are removed from all image streams that contain the image in **status.tags**.
- Image layers that are no longer referenced by any images are removed.

13.6.3. Image prune conditions

OpenShift Container Platform supports two methodologies for pruning images:

1. Pruning by age and tag
2. Pruning by size limit

These methodologies are mutually exclusive. You must choose whether to prune by age and tag, or by size limit. Regardless of the method that you choose, the image pruner checks to ensure that images in use are not removed.

An image is only pruned if it meets the primary condition **and** is not actively referenced by a system component.

13.6.3.1. Pruning an image by age and tag

Pruning an image by age and tag is the default pruning strategy. It identifies images for removal by using the **--keep-younger-than** and **--keep-tag-revisions** flags. To prune an image by age and tag, the image must be older than the **--keep-younger-than** threshold, not one of the most recent tag revisions, and cannot be in use by an active workload.

For an image to be pruned by age and tag, **all** of the following conditions must be met:

1. The image is managed by OpenShift Container Platform or has the **openshift.io/image.managed** annotation.
2. The image is older than the time specified by the **--keep-younger-than** flag.
3. The image is not one of the most recent images for its tag, as specified by the **--keep-tag-revisions** flag.
4. The image is **not** currently referenced by any of the following active or recent API objects:
 - Pods or image streams created more recently than the **--keep-younger-than** duration.
 - Running or pending pods
 - Deployments, replication controllers, replica sets, or stateful sets.
 - Builds, build configurations, jobs, or cronjobs.

An image is only removed if it is old, not a recent tag revision, and is confirmed to have no active references by system components.

13.6.3.2. Pruning an image by size limit

Pruning an image by size limit uses the **--prune-over-size-limit** flag. This method is used to bring a project back under its defined image storage limit.



NOTE

The **--prune-over-size-limit** flag cannot be combined with the **--keep-tag-revisions** flag nor the **--keep-younger-than** flags. Doing so returns information that this operation is not allowed.

For an image to be pruned using this method, all of the following conditions must be true:

1. The image is part of a project that is currently exceeding its smallest defined size limit.
2. The image is selected by the pruner as a candidate for deletion to reduce the total size.
3. The image is not currently referenced by any of the following active API objects:
 - Pods that are in a **running** or **pending** state.
 - Deployments, replication controllers, replica sets, or stateful sets.
 - Builds, build configurations, jobs, or cronjobs.

With this method, the primary trigger is the project's size, but the safety check to ensure that the image is not actively in use is still performed.

13.6.4. Running image prune operations

Use the following procedure to run an image prune operation

Prerequisites

- You must be logged into the CLI with an access token.
- You must have the **system:image-pruner** cluster role or greater (for example, **cluster-admin**).
- The image registry must be exposed.
- You have reviewed the "Considerations when manually pruning images" section of this document.

Procedure

1. Optional: To preview which images would be pruned, enter the following command. This command prints a list of the images, image streams, and pods that would be removed. Note that nothing is deleted until you add the **--confirm** flag.

```
$ oc adm prune images <image_prune_option_one> <image_prune_option_two> 1
```

- 1 For more information about available pruning options, see "Manual image pruning command options".

2. Review the output to confirm the list of images, image streams, and pods to be removed.
3. Run the **oc adm prune images** command with the appropriate options for your cluster. Add the **--confirm** flag to confirm deletion. For example:


```
$ oc adm prune images <image_prune_option_one> <image_prune_option_two> --confirm
```

13.6.5. Using secure or insecure connections

The secure connection is the preferred and recommended approach. It is done over HTTPS protocol with a mandatory certificate verification. The **prune** command always attempts to use it if possible. If it is not possible, in some cases it can fall-back to insecure connection, which is dangerous. In this case, either certificate verification is skipped or plain HTTP protocol is used.

The fall-back to insecure connection is allowed in the following cases unless **--certificate-authority** is specified:

1. The **prune** command is run with the **--force-insecure** option.
2. The provided **registry-url** is prefixed with the **http://** scheme.
3. The provided **registry-url** is a local-link address or **localhost**.
4. The configuration of the current user allows for an insecure connection. This can be caused by the user either logging in using **--insecure-skip-tls-verify** or choosing the insecure connection when prompted.



IMPORTANT

If the registry is secured by a certificate authority different from the one used by OpenShift Container Platform, it must be specified using the **--certificate-authority** flag. Otherwise, the **prune** command fails with an error.

13.6.6. Image pruning CLI options

The following table describes the options you can use with the **oc adm prune images** **<image_prune_option>** command.

Table 13.4. Manual image pruning command options

Option	Description
--all	Include images that were not pushed to the registry, but have been mirrored by pullthrough. This is on by default. To limit the pruning to images that were pushed to the integrated registry, pass --all=false .
--certificate-authority	The path to a certificate authority file to use when communicating with the OpenShift Container Platform-managed registries. Defaults to the certificate authority data from the current user's configuration file. If provided, a secure connection is initiated.
--confirm	Indicate that pruning should occur, instead of performing a test-run. This requires a valid route to the integrated container image registry. If this command is run outside of the cluster network, the route must be provided using --registry-url .
--force-insecure	Use caution with this option. Allow an insecure connection to the container registry that is hosted via HTTP or has an invalid HTTPS certificate.
--keep-tag-revisions=<N>	For each imagestream, keep up to at most N image revisions per tag (default 3).

Option	Description
--keep-younger-than=<duration>	Do not prune any image that is younger than <duration> relative to the current time. Alternately, do not prune any image that is referenced by any other object that is younger than <duration> relative to the current time (default 60m).
--prune-over-size-limit	Prune each image that exceeds the smallest limit defined in the same project. This flag cannot be combined with --keep-tag-revisions nor --keep-younger-than .
--registry-url	The address to use when contacting the registry. The command attempts to use a cluster-internal URL determined from managed images and image streams. In case it fails (the registry cannot be resolved or reached), an alternative route that works needs to be provided using this flag. The registry hostname can be prefixed by https:// or http:// , which enforces particular connection protocol.
--prune-registry	In conjunction with the conditions stipulated by the other options, this option controls whether the data in the registry corresponding to the OpenShift Container Platform image API object is pruned. By default, image pruning processes both the image API objects and corresponding data in the registry. This option is useful when you are only concerned with removing etcd content, to reduce the number of image objects but are not concerned with cleaning up registry storage, or if you intend to do that separately by hard pruning the registry during an appropriate maintenance window for the registry.

13.6.6.1. Additional information about the **--prune-registry** flag

You can separate the removal of OpenShift Container Platform image API objects from the removal of image data in the registry by passing in the **--prune-registry=false** flag. For example, the following command prunes only the API objects, leaving the registry storage untouched:

```
$ oc adm prune images --keep-tag-revisions=3 --keep-younger-than=60m --confirm --prune-registry=false
```

Then, you can perform a hard prune of the registry to remove the associated image data. This approach can narrow the timing window for race conditions compared to pruning both in a single command.

However, timing windows are not completely eliminated. For example, a pod might still be created that references an image while that image is being identified for pruning. You should track any API objects created during pruning to ensure that they do not reference deleted content.

Re-running the pruning without the **--prune-registry** option, or with **--prune-registry=true**, does not remove the associated registry storage for images previously pruned with **--prune-registry=false**. Those images can only be removed from registry storage by performing a hard prune of the registry. For more information, see "Hard pruning the registry".

13.6.7. Image pruning problems

13.6.7.1. Images not being pruned

If your images keep accumulating and the **prune** command removes just a small portion of what you expect, ensure that you understand the image prune conditions that must apply for an image to be considered a candidate for pruning.

Ensure that images you want removed occur at higher positions in each tag history than your chosen tag revisions threshold. For example, consider an old and obsolete image named **sha256:abz**. By running the following command in your namespace, where the image is tagged, the image is tagged three times in a single image stream named **myapp**:

```
$ oc get is -n <namespace> -o go-template='{{range $isi, $is := .items}}{{range $ti, $tag := $is.status.tags}}\n{{range $ii, $item := $tag.items}}{{if eq $item.image "sha256:<hash>"}}{{$is.metadata.name}}:{{tag.tag}} at position {{$ii}} out of {{len $tag.items}}\n\n{{end}}{{end}}{{end}}'
```

Example output

```
myapp:v2 at position 4 out of 5
myapp:v2.1 at position 2 out of 2
myapp:v2.1-may-2016 at position 0 out of 1
```

When default options are used, the image is never pruned because it occurs at position **0** in a history of **myapp:v2.1-may-2016** tag. For an image to be considered for pruning, the administrator must either:

- Specify **--keep-tag-revisions=0** with the **oc adm prune images** command.



WARNING

This action removes all the tags from all the namespaces with underlying images, unless they are younger or they are referenced by objects younger than the specified threshold.

- Delete all the **istags** where the position is below the revision threshold, which means **myapp:v2.1** and **myapp:v2.1-may-2016**.
- Move the image further in the history, either by running new builds pushing to the same **istag**, or by tagging other image. This is not always desirable for old release tags.

Tags having a date or time of a particular image's build in their names should be avoided, unless the image must be preserved for an undefined amount of time. Such tags tend to have just one image in their history, which prevents them from ever being pruned.

13.6.7.2. Using a secure connection against insecure registry

If you see a message similar to the following in the output of the **oc adm prune images** command, then your registry is not secured and the **oc adm prune images** client attempts to use a secure connection:

```
error: error communicating with registry: Get https://172.30.30.30:5000/healthz: http: server gave HTTP response to HTTPS client
```

- The recommended solution is to secure the registry. Otherwise, you can force the client to use an insecure connection by appending **--force-insecure** to the command; however, this is not recommended.

13.6.7.3. Using an insecure connection against a secured registry

If you see one of the following errors in the output of the **oc adm prune images** command, it means that your registry is secured using a certificate signed by a certificate authority other than the one used by **oc adm prune images** client for connection verification:

```
error: error communicating with registry: Get http://172.30.30.30:5000/healthz: malformed HTTP response "\x15\x03\x01\x00\x02\x02"
error: error communicating with registry: [Get https://172.30.30.30:5000/healthz: x509: certificate signed by unknown authority, Get http://172.30.30.30:5000/healthz: malformed HTTP response "\x15\x03\x01\x00\x02\x02"]
```

By default, the certificate authority data stored in the user's configuration files is used; the same is true for communication with the master API.

Use the **--certificate-authority** option to provide the right certificate authority for the container image registry server.

13.6.7.4. Using the wrong certificate authority

The following error means that the certificate authority used to sign the certificate of the secured container image registry is different from the authority used by the client:

```
error: error communicating with registry: Get https://172.30.30.30:5000/: x509: certificate signed by unknown authority
```

Make sure to provide the right one with the flag **--certificate-authority**.

As a workaround, the **--force-insecure** flag can be added instead. However, this is not recommended.

Additional resources

- [Accessing the registry](#)
- [Exposing the registry](#)
- See [Image Registry Operator in OpenShift Container Platform](#) for information on how to create a registry route.

13.7. HARD PRUNING THE REGISTRY

The OpenShift Container Registry can accumulate blobs that are not referenced by the OpenShift Container Platform cluster's etcd. The basic pruning images procedure, therefore, is unable to operate on them. These are called *orphaned blobs*.

Orphaned blobs can occur from the following scenarios:

- Manually deleting an image with **oc delete image <sha256:image-id>** command, which only removes the image from etcd, but not from the registry's storage.
- Pushing to the registry initiated by daemon failures, which causes some blobs to get uploaded, but the image manifest (which is uploaded as the very last component) does not. All unique image blobs become orphans.
- OpenShift Container Platform refusing an image because of quota restrictions.
- The standard image pruner deleting an image manifest, but is interrupted before it deletes the related blobs.
- A bug in the registry pruner, which fails to remove the intended blobs, causing the image objects referencing them to be removed and the blobs becoming orphans.

Hard pruning the registry, a separate procedure from basic image pruning, allows cluster administrators to remove orphaned blobs. You should hard prune if you are running out of storage space in your OpenShift Container Registry and believe you have orphaned blobs.

This should be an infrequent operation and is necessary only when you have evidence that significant numbers of new orphans have been created. Otherwise, you can perform standard image pruning at regular intervals, for example, once a day (depending on the number of images being created).

Procedure

To hard prune orphaned blobs from the registry:

1. **Log in.**

Log in to the cluster with the CLI as **kubeadmin** or another privileged user that has access to the **openshift-image-registry** namespace.

2. **Run a basic image prune**

Basic image pruning removes additional images that are no longer needed. The hard prune does not remove images on its own. It only removes blobs stored in the registry storage. Therefore, you should run this just before the hard prune.

3. **Switch the registry to read-only mode.**

If the registry is not running in read-only mode, any pushes happening at the same time as the prune will either:

- fail and cause new orphans, or
- succeed although the images cannot be pulled (because some of the referenced blobs were deleted).

Pushes will not succeed until the registry is switched back to read-write mode. Therefore, the hard prune must be carefully scheduled.

To switch the registry to read-only mode:

- a. In **configs.imageregistry.operator.openshift.io/cluster**, set **spec.readOnly** to **true**:

–

```
$ oc patch configs.imageregistry.operator.openshift.io/cluster -p '{"spec":
{"readOnly":true}}' --type=merge
```

4. Add the **system:image-pruner** role.

The service account used to run the registry instances requires additional permissions to list some resources.

- a. Get the service account name:

```
$ service_account=$(oc get -n openshift-image-registry \
-o jsonpath='{.spec.template.spec.serviceAccountName}' deploy/image-registry)
```

- b. Add the **system:image-pruner** cluster role to the service account:

```
$ oc adm policy add-cluster-role-to-user \
system:image-pruner -z \
${service_account} -n openshift-image-registry
```

5. Optional: Run the pruner in dry-run mode.

To see how many blobs would be removed, run the hard pruner in dry-run mode. No changes are actually made. The following example references an image registry pod called **image-registry-3-vhndw**:

```
$ oc -n openshift-image-registry exec pod/image-registry-3-vhndw -- /bin/sh -c
'/usr/bin/dockerregistry -prune=check'
```

Alternatively, to get the exact paths for the prune candidates, increase the logging level:

```
$ oc -n openshift-image-registry exec pod/image-registry-3-vhndw -- /bin/sh -c
'REGISTRY_LOG_LEVEL=info /usr/bin/dockerregistry -prune=check'
```

Example output

```
time="2017-06-22T11:50:25.066156047Z" level=info msg="start prune (dry-run mode)"
distribution_version="v2.4.1+unknown" kubernetes_version=v1.6.1+${Format:%h$}
openshift_version=unknown
time="2017-06-22T11:50:25.092257421Z" level=info msg="Would delete blob:
sha256:00043a2a5e384f6b59ab17e2c3d3a3d0a7de01b2cabeb606243e468acc663fa5"
go.version=go1.7.5 instance.id=b097121c-a864-4e0c-ad6c-cc25f8fdf5a6
time="2017-06-22T11:50:25.092395621Z" level=info msg="Would delete blob:
sha256:0022d49612807cb348cab562c072ef34d756adfe0100a61952cbcb87ee6578a"
go.version=go1.7.5 instance.id=b097121c-a864-4e0c-ad6c-cc25f8fdf5a6
time="2017-06-22T11:50:25.092492183Z" level=info msg="Would delete blob:
sha256:0029dd4228961086707e53b881e25eba0564fa80033fbbb2e27847a28d16a37c"
go.version=go1.7.5 instance.id=b097121c-a864-4e0c-ad6c-cc25f8fdf5a6
time="2017-06-22T11:50:26.673946639Z" level=info msg="Would delete blob:
sha256:ff7664dfc213d6cc60fd5c5f5bb00a7bf4a687e18e1df12d349a1d07b2cf7663"
go.version=go1.7.5 instance.id=b097121c-a864-4e0c-ad6c-cc25f8fdf5a6
time="2017-06-22T11:50:26.674024531Z" level=info msg="Would delete blob:
sha256:ff7a933178ccd931f4b5f40f9f19a65be5eeec207e4fad2a5baf28afbef57e"
go.version=go1.7.5 instance.id=b097121c-a864-4e0c-ad6c-cc25f8fdf5a6
time="2017-06-22T11:50:26.674675469Z" level=info msg="Would delete blob:
sha256:ff9b8956794b426cc80bb49a604a0b24a1553aae96b930c6919a6675db3d5e06"
```

```
go.version=go1.7.5 instance.id=b097121c-a864-4e0c-ad6c-cc25f8fdf5a6
...
Would delete 13374 blobs
Would free up 2.835 GiB of disk space
Use -prune=delete to actually delete the data
```

6. Run the hard prune.

Execute the following command inside one running instance of a **image-registry** pod to run the hard prune. The following example references an image registry pod called **image-registry-3-vhndw**:

```
$ oc -n openshift-image-registry exec pod/image-registry-3-vhndw -- /bin/sh -c
'/usr/bin/dockerregistry -prune=delete'
```

Example output

```
Deleted 13374 blobs
Freed up 2.835 GiB of disk space
```

7. Switch the registry back to read-write mode.

After the prune is finished, the registry can be switched back to read-write mode. In **configs.imageregistry.operator.openshift.io/cluster**, set **spec.readOnly** to **false**:

```
$ oc patch configs.imageregistry.operator.openshift.io/cluster -p '{"spec":{"readOnly":false}}' -
-type=merge
```

13.8. PRUNING CRON JOBS

Cron jobs can perform pruning of successful jobs, but might not properly handle failed jobs. Therefore, the cluster administrator should perform regular cleanup of jobs manually. They should also restrict the access to cron jobs to a small group of trusted users and set appropriate quota to prevent the cron job from creating too many jobs and pods.

Additional resources

- [Running tasks in pods using jobs](#)
- [Resource quotas across multiple projects](#)
- [Using RBAC to define and apply permissions](#)

CHAPTER 14. IDLING APPLICATIONS

Cluster administrators can idle applications to reduce resource consumption. This is useful when the cluster is deployed on a public cloud where cost is related to resource consumption.

If any scalable resources are not in use, OpenShift Container Platform discovers and idles them by scaling their replicas to **0**. The next time network traffic is directed to the resources, the resources are unidled by scaling up the replicas, and normal operation continues.

Applications are made of services, as well as other scalable resources, such as deployment configs. The action of idling an application involves idling all associated resources.

14.1. IDLING APPLICATIONS

Idling an application involves finding the scalable resources (deployment configurations, replication controllers, and others) associated with a service. Idling an application finds the service and marks it as idled, scaling down the resources to zero replicas.

You can use the **oc idle** command to idle a single service, or use the **--resource-names-file** option to idle multiple services.

14.1.1. Idling a single service

Procedure

1. To idle a single service, run:

```
$ oc idle <service>
```

14.1.2. Idling multiple services

Idling multiple services is helpful if an application spans across a set of services within a project, or when idling multiple services in conjunction with a script to idle multiple applications in bulk within the same project.

Procedure

1. Create a file containing a list of the services, each on their own line.
2. Idle the services using the **--resource-names-file** option:

```
$ oc idle --resource-names-file <filename>
```



NOTE

The **idle** command is limited to a single project. For idling applications across a cluster, run the **idle** command for each project individually.

14.2. UNIDLING APPLICATIONS

Application services become active again when they receive network traffic and are scaled back up their previous state. This includes both traffic to the services and traffic passing through routes.

Applications can also be manually unidled by scaling up the resources.

Procedure

1. To scale up a DeploymentConfig, run:

```
$ oc scale --replicas=1 dc <dc_name>
```



NOTE

Automatic unidling by a router is currently only supported by the default HAProxy router.

CHAPTER 15. DELETING APPLICATIONS

You can delete applications created in your project.

15.1. DELETING APPLICATIONS USING THE DEVELOPER PERSPECTIVE

You can delete an application and all of its associated components using the **Topology** view in the **Developer** perspective:

1. Click the application you want to delete to see the side panel with the resource details of the application.
2. Click the **Actions** drop-down menu displayed on the upper right of the panel, and select **Delete Application** to see a confirmation dialog box.
3. Enter the name of the application and click **Delete** to delete it.

You can also right-click the application you want to delete and click **Delete Application** to delete it.

CHAPTER 16. USING THE RED HAT MARKETPLACE

The [Red Hat Marketplace](#) is an open cloud marketplace that makes it easy to discover and access certified software for container-based environments that run on public clouds and on-premise.

16.1. RED HAT MARKETPLACE FEATURES

Cluster administrators can use [the Red Hat Marketplace](#) to manage software on OpenShift Container Platform, give developers self-service access to deploy application instances, and correlate application usage against a quota.

16.1.1. Connect OpenShift Container Platform clusters to the Marketplace

Cluster administrators can install a common set of applications on OpenShift Container Platform clusters that connect to the Marketplace. They can also use the Marketplace to track cluster usage against subscriptions or quotas. Users that they add by using the Marketplace have their product usage tracked and billed to their organization.

During the [cluster connection process](#), a Marketplace Operator is installed that updates the image registry secret, manages the catalog, and reports application usage.

16.1.2. Install applications

Cluster administrators can [install Marketplace applications](#) from within OperatorHub in OpenShift Container Platform, or from the [Marketplace web application](#).


You can access installed applications from the web console by clicking **Operators > Installed Operators**

16.1.3. Deploy applications from different perspectives

Developers can access newly installed capabilities and deploy Marketplace applications from the web console.

For example, after a database Operator is installed, a developer can create an instance from the catalog within their project. Database usage is aggregated and reported to the cluster administrator.

Cluster administrators can access Operator installation and application usage information. They can also launch application instances by browsing custom resource definitions (CRDs) in the **Installed Operators** list.



IMPORTANT

Starting with OpenShift Container Platform 4.19, the perspectives in the web console have unified. The **Developer** perspective is no longer enabled by default.

All users can interact with all OpenShift Container Platform web console features. However, if you are not the cluster owner, you might need to request permission to access certain features from the cluster owner.

You can still enable the **Developer** perspective. On the **Getting Started** pane in the web console, you can take a tour of the console, find information on setting up your cluster, view a quick start for enabling the **Developer** perspective, and follow links to explore new features and capabilities.

