



# **OpenShift Enterprise 3.2**

## **Architecture**

OpenShift Enterprise 3.2 Architecture Information



# OpenShift Enterprise 3.2 Architecture

---

OpenShift Enterprise 3.2 Architecture Information

## Legal Notice

Copyright © 2018 Red Hat, Inc.

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

Learn the architecture of OpenShift Enterprise 3.2 including the infrastructure and core components. These topics also cover authentication, networking and source code management.

# Table of Contents

<b>CHAPTER 1. OVERVIEW .....</b>	<b>6</b>
1.1. WHAT ARE THE LAYERS?	6
1.2. WHAT IS THE OPENSIFT ENTERPRISE ARCHITECTURE?	6
1.3. HOW IS OPENSIFT ENTERPRISE SECURED?	7
<b>CHAPTER 2. INFRASTRUCTURE COMPONENTS .....</b>	<b>8</b>
2.1. KUBERNETES INFRASTRUCTURE	8
2.1.1. Overview	8
2.1.2. Masters	8
2.1.2.1. High Availability Masters	8
2.1.3. Nodes	9
2.1.3.1. Kubelet	10
2.1.3.2. Service Proxy	10
2.1.3.3. Node Object Definition	10
2.2. IMAGE REGISTRY	11
2.2.1. Overview	11
2.2.2. Integrated OpenShift Enterprise Registry	11
2.2.3. Third Party Registries	11
2.2.3.1. Authentication	11
2.3. WEB CONSOLE	11
2.3.1. Overview	11
2.3.2. Browser Requirements	12
2.3.3. CLI Downloads	12
2.3.4. Project Overviews	13
2.3.5. JVM Console	14
<b>CHAPTER 3. CORE CONCEPTS .....</b>	<b>17</b>
3.1. OVERVIEW	17
3.2. CONTAINERS AND IMAGES	17
3.2.1. Containers	17
3.2.2. Images	17
3.2.3. Container Registries	18
3.3. PODS AND SERVICES	18
3.3.1. Pods	18
3.3.2. Services	21
3.3.2.1. Service externalIPs	21
3.3.2.2. Service ingressIPs	22
3.3.2.3. Service NodePort	23
3.3.2.4. Service Proxy Mode	23
3.3.3. Labels	23
3.4. PROJECTS AND USERS	24
3.4.1. Users	24
3.4.2. Namespaces	25
3.4.3. Projects	25
3.5. BUILDS AND IMAGE STREAMS	26
3.5.1. Builds	26
3.5.1.1. Docker Build	26
3.5.1.2. Source-to-Image (S2I) Build	26
3.5.1.3. Custom Build	27
3.5.2. Image Streams	27
3.5.2.1. Image Stream Mappings	28
3.6. DEPLOYMENTS	31

3.6.1. Replication Controllers	31
3.6.2. Jobs	32
3.6.3. Deployments and Deployment Configurations	33
3.7. ROUTES	34
3.7.1. Overview	34
3.7.2. Routers	34
3.7.2.1. Template Routers	35
3.7.3. Available Router Plug-ins	35
3.7.3.1. HAProxy Template Router	35
3.7.3.2. F5 Router	37
3.7.4. Route Host Names	38
3.7.5. Route Types	39
3.7.6. Path Based Routes	39
3.7.7. Secured Routes	40
3.7.8. Router Sharding	44
3.8. TEMPLATES	45
3.8.1. Overview	45
3.8.2. Parameters	46
<b>CHAPTER 4. ADDITIONAL CONCEPTS .....</b>	<b>48</b>
4.1. NETWORKING	48
4.1.1. OpenShift Enterprise DNS	48
4.1.2. Network Plugins	49
4.1.3. OpenShift Enterprise SDN	49
4.2. OPENSIFT SDN	49
4.2.1. Overview	49
4.2.2. Design on Masters	50
4.2.3. Design on Nodes	50
4.2.4. Packet Flow	51
4.2.5. Network Isolation	51
4.3. AUTHENTICATION	52
4.3.1. Overview	52
4.3.2. Users and Groups	52
4.3.3. API Authentication	52
4.3.4. OAuth	53
4.4. AUTHORIZATION	55
4.4.1. Overview	55
4.4.2. Evaluating Authorization	57
4.4.3. Cluster Policy and Local Policy	58
4.4.4. Roles	58
4.4.4.1. Updating Cluster Roles	60
4.4.5. Security Context Constraints	60
4.4.5.1. SCC Strategies	63
4.4.5.1.1. RunAsUser	63
4.4.5.1.2. SELinuxContext	63
4.4.5.1.3. SupplementalGroups	63
4.4.5.1.4. FSGroup	63
4.4.5.2. Controlling Volumes	64
4.4.5.3. Admission	65
4.4.5.3.1. SCC Prioritization	65
4.4.5.3.2. Understanding Pre-allocated Values and Security Context Constraints	66
4.5. PERSISTENT STORAGE	67
4.5.1. Overview	67

4.5.2. Lifecycle of a Volume and Claim	67
4.5.2.1. Provisioning	68
4.5.2.2. Binding	68
4.5.2.3. Using	68
4.5.2.4. Releasing	68
4.5.2.5. Reclaiming	68
4.5.3. Persistent Volumes	68
4.5.3.1. Types of Persistent Volumes	69
4.5.3.2. Capacity	69
4.5.3.3. Access Modes	69
4.5.3.4. Recycling Policy	70
4.5.3.5. Phase	70
4.5.4. Persistent Volume Claims	71
4.5.4.1. Access Modes	71
4.5.4.2. Resources	71
4.5.4.3. Claims As Volumes	71
4.6. REMOTE COMMANDS	72
4.6.1. Overview	72
4.6.2. Server Operation	72
4.7. PORT FORWARDING	72
4.7.1. Overview	72
4.7.2. Server Operation	72
4.8. SOURCE CONTROL MANAGEMENT	73
4.9. ADMISSION CONTROLLERS	73
4.10. OTHER API OBJECTS	74
4.10.1. LimitRange	74
4.10.2. ResourceQuota	74
4.10.3. Resource	74
4.10.4. Secret	74
4.10.5. PersistentVolume	74
4.10.6. PersistentVolumeClaim	74
4.10.7. OAuth Objects	74
4.10.7.1. OAuthClient	75
4.10.7.2. OAuthClientAuthorization	75
4.10.7.3. OAuthAuthorizeToken	76
4.10.7.4. OAuthAccessToken	77
4.10.8. User Objects	78
4.10.8.1. Identity	78
4.10.8.2. User	78
4.10.8.3. UserIdentityMapping	79
4.10.8.4. Group	79
<b>CHAPTER 5. REVISION HISTORY: ARCHITECTURE .....</b>	<b>81</b>
5.1. TUE MAY 02 2017	81
5.2. MON JAN 09 2017	81
5.3. MON NOV 21 2016	81
5.4. TUE OCT 04 2016	81
5.5. TUE SEP 06 2016	81
5.6. MON AUG 08 2016	82
5.7. WED JUL 27 2016	82
5.8. THU JUL 14 2016	82
5.9. WED JUN 15 2016	82
5.10. FRI JUN 10 2016	82

5.11. MON MAY 30 2016	83
5.12. THU MAY 12 2016	83





# CHAPTER 1. OVERVIEW

OpenShift v3 is a layered system designed to expose underlying Docker-formatted container image and Kubernetes concepts as accurately as possible, with a focus on easy composition of applications by a developer. For example, install Ruby, push code, and add MySQL.

Unlike OpenShift v2, more flexibility of configuration is exposed after creation in all aspects of the model. The concept of an application as a separate object is removed in favor of more flexible composition of "services", allowing two web containers to reuse a database or expose a database directly to the edge of the network.

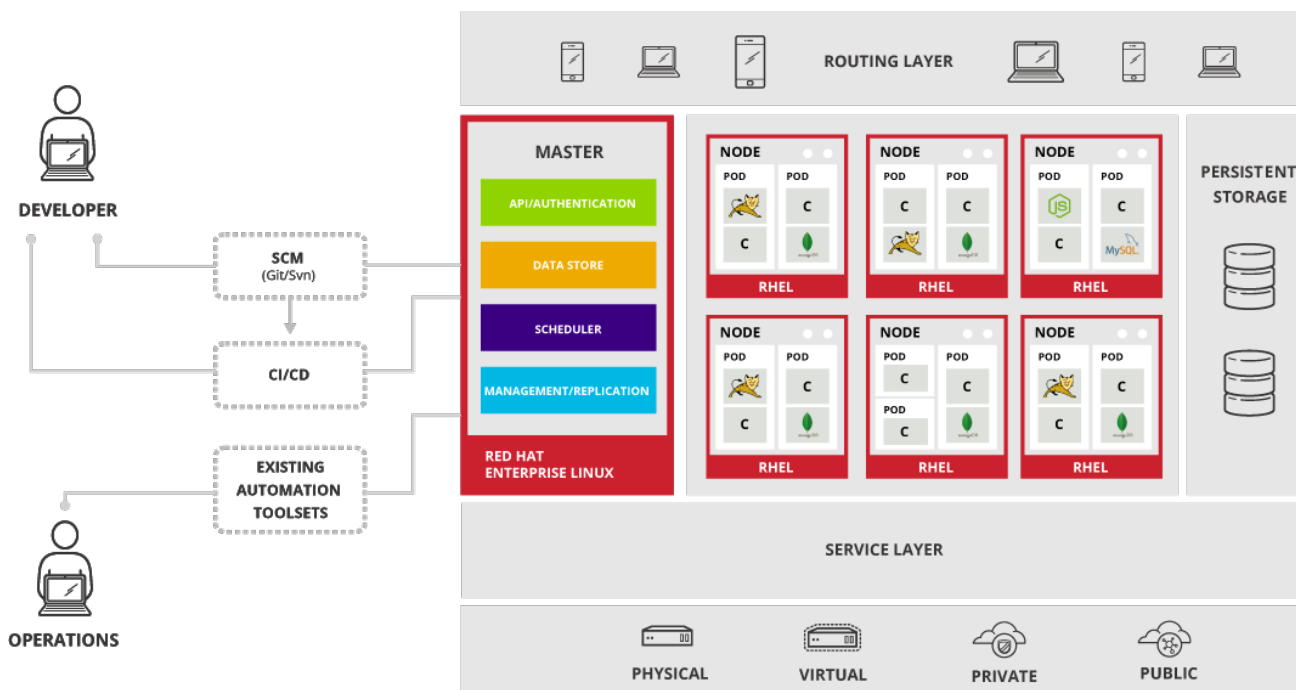
## 1.1. WHAT ARE THE LAYERS?

The Docker service provides the abstraction for packaging and creating Linux-based, lightweight [container images](#). Kubernetes provides the [cluster management](#) and orchestrates containers on multiple hosts.

OpenShift Enterprise adds:

- Source code management, [builds](#), and [deployments](#) for developers
- Managing and promoting [images](#) at scale as they flow through your system
- Application management at scale
- Team and user tracking for organizing a large developer organization

**Figure 1.1. OpenShift Enterprise Architecture Overview**



## 1.2. WHAT IS THE OPENS SHIFT ENTERPRISE ARCHITECTURE?

OpenShift Enterprise has a microservices-based architecture of smaller, decoupled units that work together. It runs on top of a [Kubernetes cluster](#), with data about the objects stored in [etcd](#), a reliable clustered key-value store. Those services are broken down by function:

- [REST APIs](#), which expose each of the [core objects](#).
- Controllers, which read those APIs, apply changes to other objects, and report status or write back to the object.

Users make calls to the REST API to change the state of the system. Controllers use the REST API to read the user's desired state, and then try to bring the other parts of the system into sync. For example, when a user requests a [build](#) they create a "build" object. The build controller sees that a new build has been created, and runs a process on the cluster to perform that build. When the build completes, the controller updates the build object via the REST API and the user sees that their build is complete.

The controller pattern means that much of the functionality in OpenShift Enterprise is extensible. The way that builds are run and launched can be customized independently of how images are managed, or how [deployments](#) happen. The controllers are performing the "business logic" of the system, taking user actions and transforming them into reality. By customizing those controllers or replacing them with your own logic, different behaviors can be implemented. From a system administration perspective, this also means the API can be used to script common administrative actions on a repeating schedule. Those scripts are also controllers that watch for changes and take action. OpenShift Enterprise makes the ability to customize the cluster in this way a first-class behavior.

To make this possible, controllers leverage a reliable stream of changes to the system to sync their view of the system with what users are doing. This event stream pushes changes from etcd to the REST API and then to the controllers as soon as changes occur, so changes can ripple out through the system very quickly and efficiently. However, since failures can occur at any time, the controllers must also be able to get the latest state of the system at startup, and confirm that everything is in the right state. This resynchronization is important, because it means that even if something goes wrong, then the operator can restart the affected components, and the system double checks everything before continuing. The system should eventually converge to the user's intent, since the controllers can always bring the system into sync.

## 1.3. HOW IS OPENSIFT ENTERPRISE SECURED?

The OpenShift Enterprise and Kubernetes APIs [authenticate](#) users who present credentials, and then [authorize](#) them based on their role. Both developers and administrators can be authenticated via a number of means, primarily [OAuth tokens](#) and SSL certificate authorization.

Developers (clients of the system) typically make REST API calls from a [client program](#) like `oc` or to the [web console](#) via their browser, and use OAuth bearer tokens for most communications. Infrastructure components (like nodes) use client certificates generated by the system that contain their identities. Infrastructure components that run in containers use a token associated with their [service account](#) to connect to the API.

Authorization is handled in the OpenShift Enterprise policy engine, which defines actions like "create pod" or "list services" and groups them into roles in a policy document. Roles are bound to users or groups by the user or group identifier. When a user or service account attempts an action, the policy engine checks for one or more of the roles assigned to the user (e.g., cluster administrator or administrator of the current project) before allowing it to continue.

Since every container that runs on the cluster is associated with a service account, it is also possible to associate [secrets](#) to those service accounts and have them automatically delivered into the container. This enables the infrastructure to manage secrets for pulling and pushing images, builds, and the deployment components, and also allows application code to easily leverage those secrets.


## CHAPTER 2. INFRASTRUCTURE COMPONENTS

### 2.1. KUBERNETES INFRASTRUCTURE

#### 2.1.1. Overview

Within OpenShift Enterprise, Kubernetes manages containerized applications across a set of containers or hosts and provides mechanisms for deployment, maintenance, and application-scaling. The Docker service packages, instantiates, and runs containerized applications.

A Kubernetes cluster consists of one or more masters and a set of nodes. You can optionally configure your masters for [high availability](#) (HA) to ensure that the cluster has no single point of failure.



#### NOTE

OpenShift Enterprise uses Kubernetes 1.2 and Docker 1.9.

#### 2.1.2. Masters

The master is the host or hosts that contain the master components, including the API server, controller manager server, and **etcd**. The master manages [nodes](#) in its Kubernetes cluster and schedules [pods](#) to run on nodes.

**Table 2.1. Master Components**

Component	Description
API Server	The Kubernetes API server validates and configures the data for pods, services, and replication controllers. It also assigns pods to nodes and synchronizes pod information with service configuration. Can be run as a standalone process.
<b>etcd</b>	<b>etcd</b> stores the persistent master state while other components watch <b>etcd</b> for changes to bring themselves into the desired state. <b>etcd</b> can be optionally configured for high availability, typically deployed with 2n+1 peer services.
Controller Manager Server	The controller manager server watches <b>etcd</b> for changes to replication controller objects and then uses the API to enforce the desired state. Can be run as a standalone process. Several such processes create a cluster with one active leader at a time.
HAProxy	Optional, used when configuring <a href="#">highly-available masters</a> with the <b>native</b> method to balance load between API master endpoints.  The <a href="#">advanced installation method</a> can configure HAProxy for you with the <b>native</b> method. Alternatively, you can use the <b>native</b> method but pre-configure your own load balancer of choice.


##### 2.1.2.1. High Availability Masters

While in a single master configuration, the availability of running applications remains if the master or any of its services fail. However, failure of master services reduces the ability of the system to respond to application failures or creation of new applications. You can optionally configure your masters for high

availability (HA) to ensure that the cluster has no single point of failure.

To mitigate concerns about availability of the master, two activities are recommended:

1. A [runbook](#) entry should be created for reconstructing the master. A runbook entry is a necessary backstop for any highly-available service. Additional solutions merely control the frequency that the runbook must be consulted. For example, a cold standby of the master host can adequately fulfill SLAs that require no more than minutes of downtime for creation of new applications or recovery of failed application components.
2. Use a high availability solution to configure your masters and ensure that the cluster has no single point of failure. The [advanced installation method](#) provides specific examples using the **native** HA method and configuring HAProxy. You can also take the concepts and apply them towards your existing HA solutions using the **native** method instead of HAProxy.



## NOTE

Moving from a single master cluster to multiple masters after installation is not supported.

When using the **native** HA method with HAProxy, master components have the following availability:

**Table 2.2. Availability Matrix with HAProxy**


Role	Style	Notes
<b>etcd</b>	Active-active	Fully redundant deployment with load balancing
API Server	Active-active	Managed by HAProxy
Controller Manager Server	Active-passive	One instance is elected as a cluster leader at a time
HAProxy	Active-passive	Balances load between API master endpoints

### 2.1.3. Nodes

A node provides the runtime environments for containers. Each node in a Kubernetes cluster has the required services to be managed by the [master](#). Nodes also have the required services to run pods, including the Docker service, a [kubelet](#), and a [service proxy](#).

OpenShift Enterprise creates nodes from a cloud provider, physical systems, or virtual systems. Kubernetes interacts with [node objects](#) that are a representation of those nodes. The master uses the information from node objects to validate nodes with health checks. A node is ignored until it passes the health checks, and the master continues checking nodes until they are valid. The [Kubernetes documentation](#) has more information on node management.

Administrators can [manage nodes](#) in an OpenShift Enterprise instance using the CLI. To define full configuration and security options when launching node servers, use [dedicated node configuration files](#).



## IMPORTANT

The recommended maximum number of nodes is 300.

### 2.1.3.1. Kubelet

Each node has a kubelet that updates the node as specified by a container manifest, which is a YAML file that describes a pod. The kubelet uses a set of manifests to ensure that its containers are started and that they continue to run. A sample manifest can be found in the [Kubernetes documentation](#).

A container manifest can be provided to a kubelet by:

- A file path on the command line that is checked every 20 seconds.
- An HTTP endpoint passed on the command line that is checked every 20 seconds.
- The kubelet watching an **etcd** server, such as `/registry/hosts/${hostname -f}`, and acting on any changes.
- The kubelet listening for HTTP and responding to a simple API to submit a new manifest.

### 2.1.3.2. Service Proxy

Each node also runs a simple network proxy that reflects the services defined in the API on that node. This allows the node to do simple TCP and UDP stream forwarding across a set of back ends.

### 2.1.3.3. Node Object Definition

The following is an example node object definition in Kubernetes:

```
apiVersion: v1 ❶
kind: Node ❷
metadata:
  creationTimestamp: null
  labels: ❸
    kubernetes.io/hostname: node1.example.com
  name: node1.example.com ❹
spec:
  externalID: node1.example.com ❺
status:
  nodeInfo:
    bootID: ""
    containerRuntimeVersion: ""
    kernelVersion: ""
    kubeProxyVersion: ""
    kubeletVersion: ""
    machineID: ""
    osImage: ""
    systemUUID: ""
```

- ❶ **apiVersion** defines the API version to use.
- ❷ **kind** set to **Node** identifies this as a definition for a node object.
- ❸ **metadata.labels** lists any [labels](#) that have been added to the node.
- ❹ **metadata.name** is a required value that defines the name of the node object. This value is shown in the **NAME** column when running the **oc get nodes** command.

- 5** `spec.externalID` defines the fully-qualified domain name where the node can be reached. Defaults to the `metadata.name` value when empty.

The [REST API Reference](#) has more details on these definitions.

## 2.2. IMAGE REGISTRY

### 2.2.1. Overview

OpenShift Enterprise can utilize any server implementing the Docker registry API as a source of images, including the Docker Hub, private registries run by third parties, and the integrated OpenShift Enterprise registry.

### 2.2.2. Integrated OpenShift Enterprise Registry

OpenShift Enterprise provides an integrated container registry that adds the ability to provision new image repositories on the fly. This allows users to automatically have a place for their [builds](#) to push the resulting images.

Whenever a new image is pushed to the integrated registry, the registry notifies OpenShift Enterprise about the new image, passing along all the information about it, such as the namespace, name, and image metadata. Different pieces of OpenShift react to new images, creating new [builds](#) and [deployments](#).

### 2.2.3. Third Party Registries

OpenShift Enterprise can create containers using images from third party registries, but it is unlikely that these registries offer the same image notification support as the integrated OpenShift Enterprise registry. In this situation OpenShift Enterprise will fetch tags from the remote registry upon imagestream creation. Refreshing the fetched tags is as simple as running `oc import-image <stream>`. When new images are detected, the previously-described build and deployment reactions occur.

#### 2.2.3.1. Authentication

OpenShift Enterprise can communicate with registries to access private image repositories using credentials supplied by the user. This allows OpenShift to push and pull images to and from private repositories. The [Authentication](#) topic has more information.

## 2.3. WEB CONSOLE

### 2.3.1. Overview

The OpenShift Enterprise web console is a user interface accessible from a web browser. Developers can use the web console to visualize, browse, and manage the contents of [projects](#).



#### NOTE

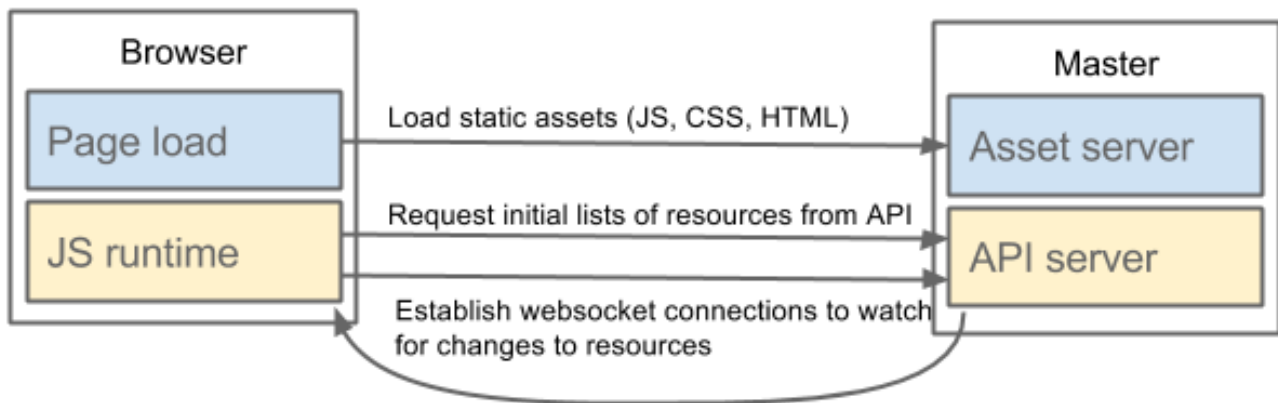
JavaScript must be enabled to use the web console. For the best experience, use a web browser that supports [WebSockets](#).

The web console is started as part of the [master](#). All static assets required to run the web console are

served from the **openshift** binary. Administrators can also [customize the web console](#) using extensions, which let you run scripts and load custom stylesheets when the web console loads. You can change the look and feel of nearly any aspect of the user interface in this way.

When you access the web console from a browser, it first loads all required static assets. It then makes requests to the OpenShift Enterprise APIs using the values defined from the **openshift start** option **--public-master**, or from the related [master configuration file](#) parameter **masterPublicURL**. The web console uses WebSockets to maintain a persistent connection with the API server and receive updated information as soon as it is available.

**Figure 2.1. Web Console Request Architecture**



The configured host names and IP addresses for the web console are whitelisted to access the API server safely even when the browser would consider the requests to be [cross-origin](#). To access the API server from a web application using a different host name, you must whitelist that host name by specifying the **--cors-allowed-origins** option on **openshift start** or from the related [master configuration file](#) parameter **corsAllowedOrigins**.

### 2.3.2. Browser Requirements

Review the [tested integrations](#) for OpenShift Enterprise. The following browser versions and operating systems can be used to access the web console.


**Table 2.3. Browser Requirements**

Browser (Latest Stable)	Operating System
Firefox	Fedora 23, Windows 8
Internet Explorer	Windows 8
Chrome	Fedora 23, Windows 8, and MacOSX
Safari	MacOSX, iPad 2, iPhone 4

### 2.3.3. CLI Downloads

You can download and unpack the CLI from the [About page on the web console](#) for use on Linux, MacOSX, and Windows clients if your cluster administrator has enabled it :





# OpenShift by Red Hat®

## About

OpenShift is Red Hat's Platform-as-a-Service (PaaS) that allows developers to quickly develop, host, and scale applications in a cloud environment.

## Version

**OpenShift Master:** v  
**Kubernetes Master:** v

## Command Line Tools

With the OpenShift command line interface (CLI), you can create applications and manage OpenShift projects from a terminal. You can download the `oc` client tool using the links below. For more information about downloading and installing it, please refer to the [Get Started with the CLI](#) documentation.

**Download `oc` :**  
[Latest Release](#)

After downloading and installing it, you can start by logging in using this current session token:

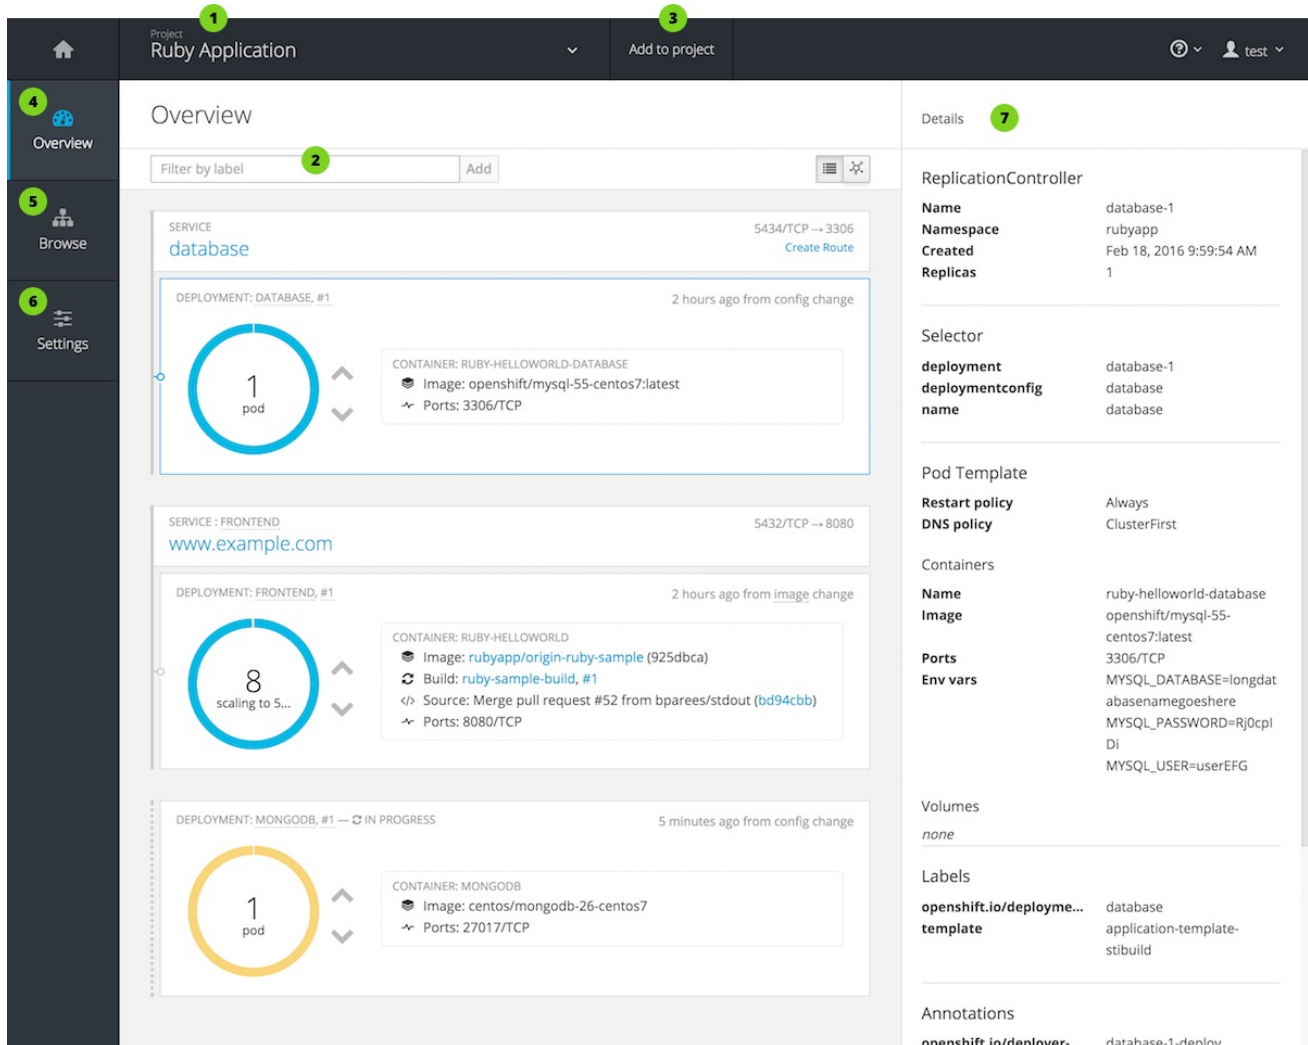
```
oc login https://:8443 --token=...
```

For other information about the command line tools, check the [CLI Reference](#) and [Basic CLI Operations](#).

### 2.3.4. Project Overviews

After [logging in](#), the web console provides developers with an overview for the currently selected [project](#):

**Figure 2.2. Web Console Project Overview**



The screenshot shows the OpenShift Web Console interface for a project named 'Ruby Application'. The sidebar on the left has a 'Project' dropdown set to 'Ruby Application' and an 'Add to project' button. The main content area is divided into two sections: 'Overview' and 'Details'.

**Overview Section:**

- Filter by label:** A search bar with an 'Add' button.
- Service: database** (5434/TCP → 3306):
  - Deployment: DATABASE, #1** (2 hours ago from config change): 1 pod.
  - Container: RUBY-HELLOWORLD-DATABASE**: Image: openshift/mysql-55-centos7:latest, Ports: 3306/TCP.
- Service: FRONTEND** (5432/TCP → 8080):
  - Deployment: FRONTEND, #1** (2 hours ago from image change): 8 pods, scaling to 5...
  - Container: RUBY-HELLOWORLD**: Image: rubyapp/origin-ruby-sample (925dbca), Build: ruby-sample-build, #1, Source: Merge pull request #52 from bpares/stdout (bd94cbb), Ports: 8080/TCP.
- Deployment: MONGODB, #1** (5 minutes ago from config change): 1 pod, IN PROGRESS.
- Container: MONGODB**: Image: centos/mongodb-26-centos7, Ports: 27017/TCP.

**Details Section (ReplicationController):**

- Name:** database-1
- Namespace:** rubyapp
- Created:** Feb 18, 2016 9:59:54 AM
- Replicas:** 1
- Selector:**
  - deployment:** database-1
  - deploymentconfig:** database
  - name:** database
- Pod Template:**
  - Restart policy:** Always
  - DNS policy:** ClusterFirst
- Containers:**
  - Name:** ruby-helloworld-database
  - Image:** openshift/mysql-55-centos7:latest
  - Ports:** 3306/TCP
  - Env vars:** MYSQL\_DATABASE=longdat, abasenamegoeshere, MYSQL\_PASSWORD=Rj0cplDi, MYSQL\_USER=userEFG
- Volumes:** none
- Labels:**
  - openshift.io/deployme...:** database
  - template:** application-template-stibuild
- Annotations:**
  - openshift.io/deployer-...:** database-1-deploy

The project selector allows you to [switch between projects](#) you have access to.

[Filter](#) the contents of a project page by using the [labels](#) of a resource.


Create new applications [using a source repository](#) or [using a template](#).

The **Overview** tab (currently selected) visualizes the contents of your project with a high-level view of each component.

The **Browse** tab explores the different [objects types](#) within your project: [Builds](#), [Deployments](#), [Image Streams](#), [Pods](#), and [Services](#).

The **Settings** tab provides general information about your project, as well as the [quota and limits](#) that are set on your project.

When you click on one of your objects in the **Overview** page, the **Details** pane displays detailed information about that object. In this example, the **cakephp-mysql-example** deployment is selected, and the **Details** pane is displaying details on the related [replication controller](#).



#### NOTE

[Cockpit](#) is automatically installed and enabled in OpenShift Enterprise 3.1 and later to help you monitor your development environment. [Red Hat Enterprise Linux Atomic Host: Getting Started with Cockpit](#) provides more information on using Cockpit.

### 2.3.5. JVM Console

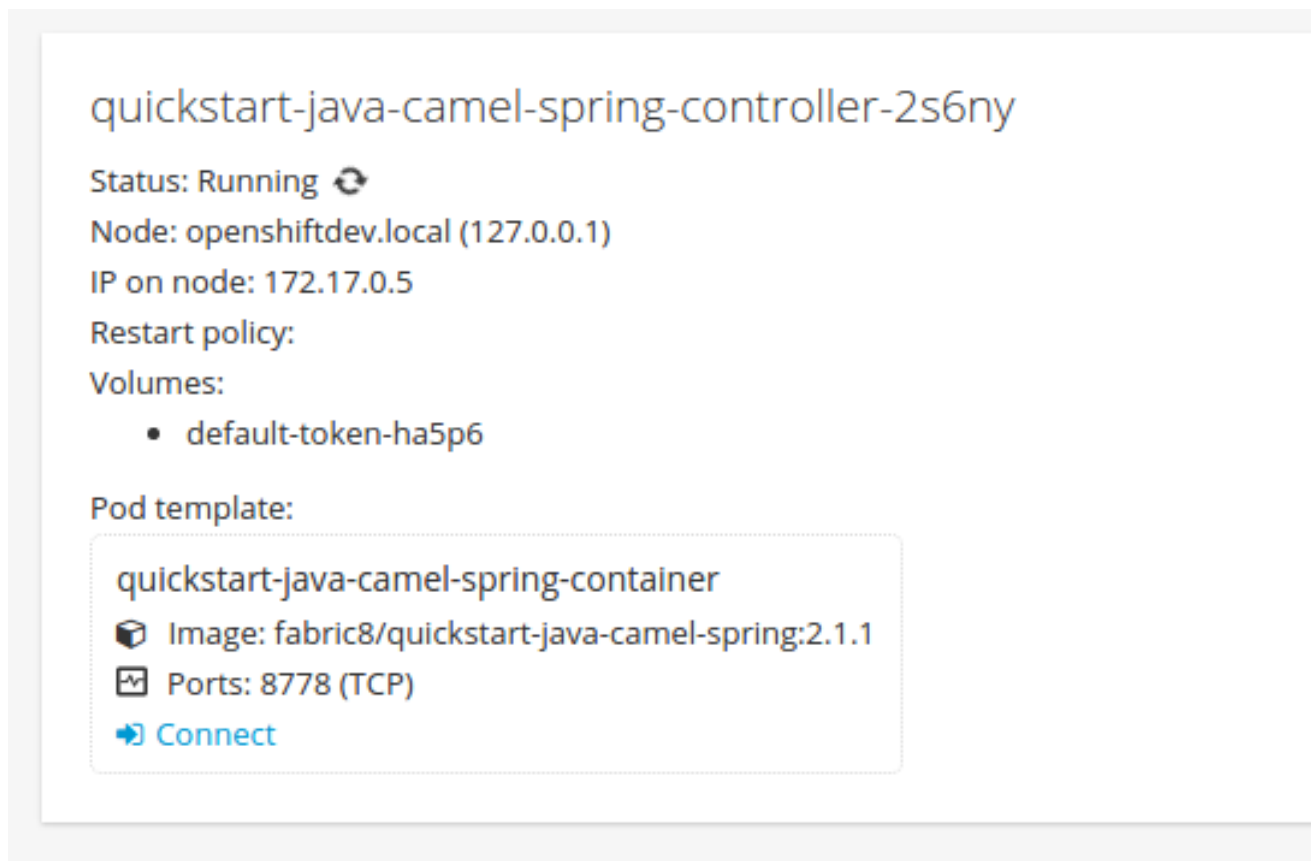


#### NOTE

This feature is currently in Technology Preview and not intended for production use.

For pods based on Java images, the web console also exposes access to a [hawt.io](#)-based JVM console for viewing and managing any relevant integration components. A **Connect** link is displayed in the pod's details on the *Browse* → *Pods* page, provided the container has a port named **jolokia**.

Figure 2.3. Pod with a Link to the JVM Console



After connecting to the JVM console, different pages are displayed depending on which components are relevant to the connected pod.

Figure 2.4. JVM Console

Connected to quickstart-java-camel-spring-container

[← Back](#)

JMX Threads Camel

Total: 9 Runnable: 3 Timed waiting: 2 Waiting: 4

Filter... ✕

ID	State	Name	Waited Time	Blocked Time	Native	Suspended
15		Thread-5	1 hour			
14		Camel (camel-1) thread #0 - file://src/data	1 hour			
9		Jolokia Agent Cleanup Thread				
8		Thread-3		279 ms	(in native)	
6		server-timer	1 hour			
4		Signal Dispatcher				
3		Finalizer	1 hour			
2		Reference Handler	1 hour	10 ms		
1		main				

The following pages are available:

Page	Description
JMX	View and manage JMX domains and mbeans.
Threads	View and monitor the state of threads.
ActiveMQ	View and manage Apache ActiveMQ brokers.
Camel	View and and manage Apache Camel routes and dependencies.

## CHAPTER 3. CORE CONCEPTS

### 3.1. OVERVIEW

The following topics provide high-level, architectural information on core concepts and objects you will encounter when using OpenShift Enterprise. Many of these objects come from Kubernetes, which is extended by OpenShift Enterprise to provide a more feature-rich development lifecycle platform.

- [Containers and images](#) are the building blocks for deploying your applications.
- [Pods and services](#) allow for containers to communicate with each other and proxy connections.
- [Projects and users](#) provide the space and means for communities to organize and manage their content together.
- [Builds and image streams](#) allow you to build working images and react to new images.
- [Deployments](#) add expanded support for the software development and deployment lifecycle.
- [Routes](#) announce your service to the world.
- [Templates](#) allow for many objects to be created at once based on customized parameters.

### 3.2. CONTAINERS AND IMAGES

#### 3.2.1. Containers

The basic units of OpenShift Enterprise applications are called *containers*. [Linux container technologies](#) are lightweight mechanisms for isolating running processes so that they are limited to interacting with only their designated resources.

Many application instances can be running in containers on a single host without visibility into each others' processes, files, network, and so on. Typically, each container provides a single service (often called a "micro-service"), such as a web server or a database, though containers can be used for arbitrary workloads.

The Linux kernel has been incorporating capabilities for container technologies for years. More recently the Docker project has developed a convenient management interface for Linux containers on a host. OpenShift Enterprise and Kubernetes add the ability to orchestrate Docker-formatted containers across multi-host installations.

Though you do not directly interact with the Docker CLI or service when using OpenShift Enterprise, understanding their capabilities and terminology is important for understanding their role in OpenShift Enterprise and how your applications function inside of containers. The **docker** RPM is available as part of RHEL 7, as well as CentOS and Fedora, so you can experiment with it separately from OpenShift Enterprise. Refer to the article [Get Started with Docker Formatted Container Images on Red Hat Systems](#) for a guided introduction.

#### 3.2.2. Images

Containers in OpenShift Enterprise are based on Docker-formatted container *images*. An image is a binary that includes all of the requirements for running a single container, as well as metadata describing its needs and capabilities.

You can think of it as a packaging technology. Containers only have access to resources defined in the image unless you give the container additional access when creating it. By deploying the same image in multiple containers across multiple hosts and load balancing between them, OpenShift Enterprise can provide redundancy and horizontal scaling for a service packaged into an image.

You can use the Docker CLI directly to build images, but OpenShift Enterprise also supplies builder images that assist with creating new images by adding your code or configuration to existing images.

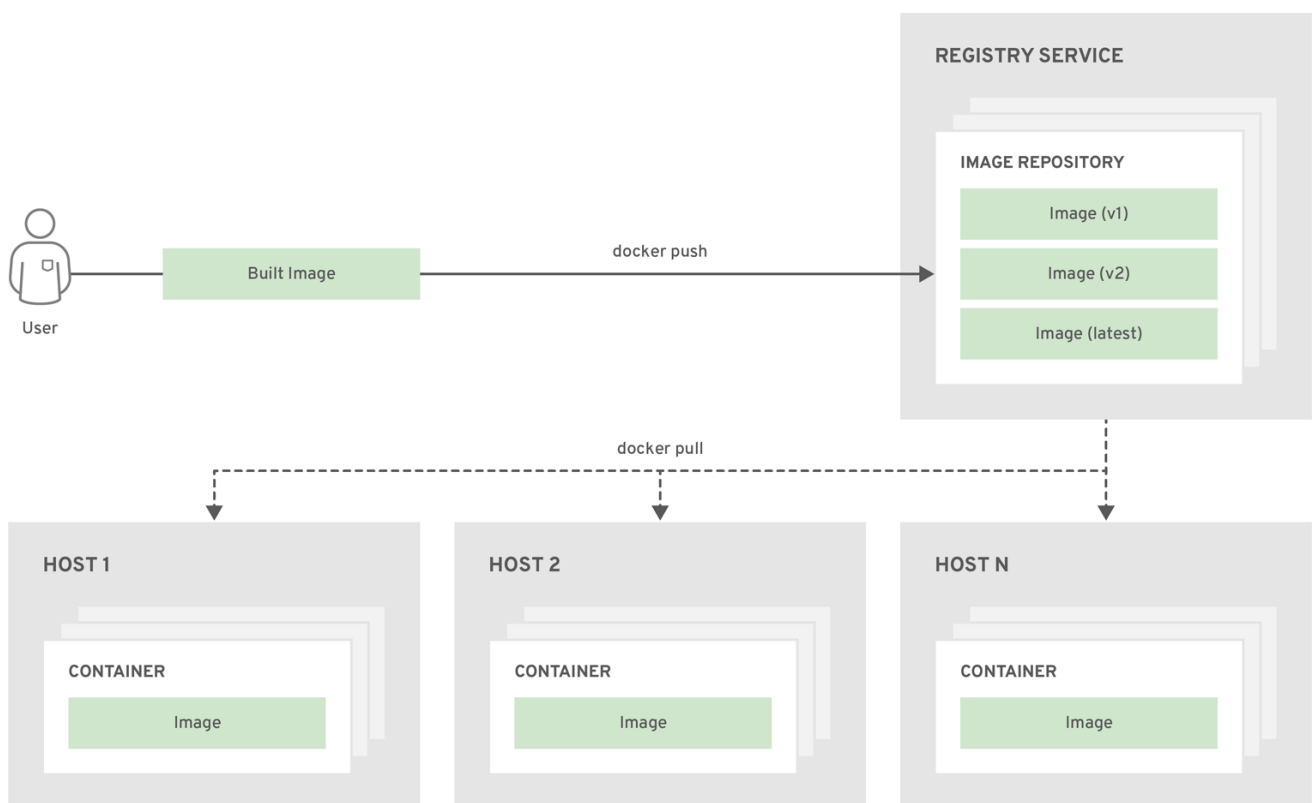
Because applications develop over time, a single image name can actually refer to many different versions of the "same" image. Each different image is referred to uniquely by its hash (a long hexadecimal number e.g. **fd44297e2ddb050ec4f...**) which is usually shortened to 12 characters (e.g. **fd44297e2ddb**).

Rather than version numbers, the Docker service allows applying tags (such as **v1**, **v2.1**, **GA**, or the default **latest**) in addition to the image name to further specify the image desired, so you may see the same image referred to as **centos** (implying the **latest** tag), **centos:centos7**, or **fd44297e2ddb**.

### 3.2.3. Container Registries

A container registry is a service for storing and retrieving Docker-formatted container images. A registry contains a collection of one or more image repositories. Each image repository contains one or more tagged images. Docker provides its own registry, the [Docker Hub](#), and you can also use private or third-party registries. Red Hat provides a registry at **registry.access.redhat.com** for subscribers. OpenShift Enterprise can also supply its own internal registry for managing custom container images.

The relationship between containers, images, and registries is depicted in the following diagram:



OPENSHIFT\_415489\_0218

## 3.3. PODS AND SERVICES

### 3.3.1. Pods

OpenShift Enterprise leverages the Kubernetes concept of a pod, which is one or more [containers](#) deployed together on one host, and the smallest compute unit that can be defined, deployed, and managed.

Pods are the rough equivalent of OpenShift Enterprise v2 gears, with containers the rough equivalent of v2 cartridge instances. Each pod is allocated its own internal IP address, therefore owning its entire port space, and containers within pods can share their local storage and networking.

Pods have a lifecycle; they are defined, then they are assigned to run on a node, then they run until their container(s) exit or they are removed for some other reason. Pods, depending on policy and exit code, may be removed after exiting, or may be retained to enable access to the logs of their containers.

OpenShift Enterprise treats pods as largely immutable; changes cannot be made to a pod definition while it is running. OpenShift Enterprise implements changes by terminating an existing pod and recreating it with modified configuration, base image(s), or both. Pods are also treated as expendable, and do not maintain state when recreated. Therefore manage pods with higher-level [controllers](#), rather than directly by users.

The following example definition of a pod provides a long-running service, which is actually a part of the OpenShift Enterprise infrastructure: the private Docker registry. It demonstrates many features of pods, most of which are discussed in other topics and thus only briefly mentioned here.

### Pod object definition example

```

apiVersion: v1
kind: Pod
metadata:
  annotations: { ... }
  labels:
    deployment: docker-registry-1
    deploymentconfig: docker-registry
    docker-registry: default
  generateName: docker-registry-1-
spec:
  containers:
    - env:
        - name: OPENSIFT_CA_DATA
          value: ...
        - name: OPENSIFT_CERT_DATA
          value: ...
        - name: OPENSIFT_INSECURE
          value: "false"
        - name: OPENSIFT_KEY_DATA
          value: ...
        - name: OPENSIFT_MASTER
          value: https://master.example.com:8443
      image: openshift/origin-docker-registry:v0.6.2
      imagePullPolicy: IfNotPresent
      name: registry
      ports:
        - containerPort: 5000
          protocol: TCP
      resources: {}
      securityContext: { ... }
      volumeMounts:

```

```

- mountPath: /registry
  name: registry-storage
- mountPath: /var/run/secrets/kubernetes.io/serviceaccount
  name: default-token-br6yz
  readOnly: true
dnsPolicy: ClusterFirst
imagePullSecrets:
- name: default-dockercfg-at06w
restartPolicy: Always
serviceAccount: default
volumes:
- emptyDir: {}
  name: registry-storage
- name: default-token-br6yz
  secret:
    secretName: default-token-br6yz

```

- 1 Pods can be "tagged" with one or more [labels](#), which can then be used to select and manage groups of pods in a single operation. The labels are stored in key-value format in the **metadata** hash. One label in this example is **docker-registry=default**.
- 2 Pods must have a unique name within their [namespace](#). A pod definition may specify the basis of a name with the **generateName** attribute and random characters will be added automatically to generate a unique name.
- 3 **containers** specifies an array of container definitions; in this case (as with most), just one.
- 4 Environment variables can be specified to pass necessary values to each container.
- 5 Each container in the pod is instantiated from its own [Docker-formatted container image](#).
- 6 The container can bind to ports which will be made available on the pod's IP.
- 7 OpenShift Enterprise defines a [security context](#) for containers which specifies whether they are allowed to run as privileged containers, run as a user of their choice, and more. The default context is very restrictive but administrators can modify this as needed.
- 8 The container specifies where external storage volumes should be mounted within the container. In this case, there is a volume for storing the registry's data, and one for access to credentials the registry needs for making requests against the OpenShift Enterprise API.
- 9 Pods making requests against the OpenShift Enterprise API is a common enough pattern that there is a **serviceAccount** field for specifying which [service account](#) user the pod should authenticate as when making the requests. This enables fine-grained access control for custom infrastructure components.
- 10 The pod defines storage volumes that are available to its container(s) to use. In this case, it provides an ephemeral volume for the registry storage and a **secret** volume containing the service account credentials.





## NOTE

This pod definition does not include attributes that are filled by OpenShift Enterprise automatically after the pod is created and its lifecycle begins. The [Kubernetes API documentation](#) has complete details of the pod REST API object attributes, and the [Kubernetes pod documentation](#) has details about the functionality and purpose of pods.

### 3.3.2. Services

A Kubernetes **service** serves as an internal load balancer. It identifies a set of replicated [pods](#) in order to proxy the connections it receives to them. Backing pods can be added to or removed from a service arbitrarily while the service remains consistently available, enabling anything that depends on the service to refer to it at a consistent internal address.

Services are assigned an IP address and port pair that, when accessed, proxy to an appropriate backing pod. A service uses a label selector to find all the containers running that provide a certain network service on a certain port.

Like pods, services are REST objects. The following example shows the definition of a service for the pod defined above:

#### Service object definition example

```
apiVersion: v1
kind: Service
metadata:
  name: docker-registry ❶
spec:
  selector: ❷
    docker-registry: default
  portalIP: 172.30.136.123 ❸
  ports:
    - nodePort: 0
      port: 5000 ❹
      protocol: TCP
      targetPort: 5000 ❺
```

- ❶ The service name **docker-registry** is also used to construct an environment variable with the service IP that is inserted into other pods in the same namespace. The maximum name length is 63 characters.
- ❷ The label selector identifies all pods with the **docker-registry=default** label attached as its backing pods.
- ❸ Virtual IP of the service, allocated automatically at creation from a pool of internal IPs.
- ❹ Port the service listens on.
- ❺ Port on the backing pods to which the service forwards connections.

The [Kubernetes documentation](#) has more information on services.

#### 3.3.2.1. Service externalIPs

In addition to the cluster's internal IP addresses, the user can configure IP addresses that are external to the cluster. The administrator is responsible for ensuring that traffic arrives at a node with this IP.

The externalIPs must be selected by the admin from the **ExternalIPNetworkCIDRs** range configured in **master-config.yaml** file. When **master-config.yaml** is changed, the master service must be restarted.

### Example 3.1. Sample ExternalIPNetworkCIDR /etc/origin/master/master-config.yaml

```
networkConfig:
  ExternalIPNetworkCIDR: 172.47.0.0/24
```

### Example 3.2. Service externalIPs Definition (JSON)

```
{
  "kind": "Service",
  "apiVersion": "v1",
  "metadata": {
    "name": "my-service"
  },
  "spec": {
    "selector": {
      "app": "MyApp"
    },
    "ports": [
      {
        "name": "http",
        "protocol": "TCP",
        "port": 80,
        "targetPort": 9376
      }
    ],
    "externalIPs" : [
      "80.11.12.10"
    ]
  }
}
```

- 1** List of External IP addresses on which the **port** is exposed. In addition to the internal IP addresses)

### 3.3.2.2. Service ingressIPs

In non-cloud clusters, externalIP addresses can be automatically assigned from a pool of addresses. This eliminates the need for the administrator manually assigning them.

The pool is configured in **/etc/origin/master/master-config.yaml** file. After changing this file, restart the master service.

The **ingressIPNetworkCIDR** is set to **172.29.0.0/16** by default. If the cluster environment is not already using this private range, use the default range or set a custom range.

**NOTE**

If you are using [high availability](#), then this range must be less than 256 addresses.

**Example 3.3. Sample ingressIPNetworkCIDR /etc/origin/master/master-config.yaml**

```
networkConfig:
  ingressIPNetworkCIDR: 172.29.0.0/16
```

**3.3.2.3. Service NodePort**

Setting the service **type=NodePort** will allocate a port from a flag-configured range (default: 30000-32767), and each node will proxy that port (the same port number on every node) into your service.

The selected port will be reported in the service configuration, under **spec.ports[\*].nodePort**.

To specify a custom port just place the port number in the nodePort field. The custom port number must be in the configured range for nodePorts. When '**master-config.yaml**' is changed the master service must be restarted.

**Example 3.4. Sample servicesNodePortRange /etc/origin/master/master-config.yaml**

```
kubernetesMasterConfig:
  servicesNodePortRange: ""
```

The service will be visible as both the **<NodeIP>:spec.ports[].nodePort** and **spec.clusterIp:spec.ports[].port**

**NOTE**

Setting a nodePort is a privileged operation.

**3.3.2.4. Service Proxy Mode**

OpenShift Enterprise has two different implementations of the service-routing infrastructure. The default implementation is entirely **iptables**-based, and uses probabilistic **iptables** rewriting rules to distribute incoming service connections between the endpoint pods. The older implementation uses a user space process to accept incoming connections and then proxy traffic between the client and one of the endpoint pods.

The **iptables**-based implementation is much more efficient, but it requires that all endpoints are always able to accept connections; the user space implementation is slower, but can try multiple endpoints in turn until it finds one that works. If you have good [readiness checks](#) (or generally reliable nodes and pods), then the **iptables**-based service proxy is the best choice. Otherwise, you can enable the user space-based proxy when installing, or after deploying the cluster by editing the node configuration file.

**3.3.3. Labels**

Labels are used to organize, group, or select API objects. For example, [pods](#) are "tagged" with labels, and then [services](#) use label selectors to identify the pods they proxy to. This makes it possible for

services to reference groups of pods, even treating pods with potentially different containers as related entities.

Most objects can include labels in their metadata. Labels can be used to group arbitrarily-related objects; for example, all of the [pods](#), [services](#), [replication controllers](#), and [deployment configurations](#) of a particular application can be grouped.

Labels are simple key-value pairs, as in the following example:

### Key-value pairs example

```
labels:
  key1: value1
  key2: value2
```

Consider:

- A pod consisting of an **nginx** container, with the label **role=webserver**.
- A pod consisting of an **Apache httpd** container, with the same label **role=webserver**.

A service or replication controller that is defined to use pods with the **role=webserver** label treats both of these pods as part of the same group.

The [Kubernetes labels documentation](#) has more information about labels.

## 3.4. PROJECTS AND USERS

### 3.4.1. Users

Interaction with OpenShift Enterprise is associated with a user. An OpenShift Enterprise user object represents an actor which may be granted permissions in the system by [adding roles to them or to their groups](#).

Several types of users can exist:

<b>Regular users</b>	This is the way most interactive OpenShift Enterprise users will be represented. Regular users are created automatically in the system upon first login, or can be created via the API. Regular users are represented with the <b>User</b> object. Examples: <b>joe</b> <b>alice</b>
<b>System users</b>	Many of these are created automatically when the infrastructure is defined, mainly for the purpose of enabling the infrastructure to interact with the API securely. They include a cluster administrator (with access to everything), a per-node user, users for use by routers and registries, and various others. Finally, there is an <b>anonymous</b> system user that is used by default for unauthenticated requests. Examples: <b>system:admin</b> <b>system:openshift-registry</b> <b>system:node:node1.example.com</b>
<b>Service accounts</b>	These are special system users associated with projects; some are created automatically when the project is first created, while project administrators can create more for the purpose of defining access to the contents of each <a href="#">project</a> . Service accounts are represented with the <b>ServiceAccount</b> object. Examples: <b>system:serviceaccount:default:deployer</b> <b>system:serviceaccount:foo:builder</b>

Every user must [authenticate](#) in some way in order to access OpenShift Enterprise. API requests with no authentication or invalid authentication are authenticated as requests by the **anonymous** system user. Once authenticated, policy determines what the user is [authorized](#) to do.

### 3.4.2. Namespaces

A Kubernetes namespace provides a mechanism to scope resources in a cluster. In OpenShift Enterprise, a [project](#) is a Kubernetes namespace with additional annotations.

Namespaces provide a unique scope for:

- Named resources to avoid basic naming collisions.
- Delegated management authority to trusted users.
- The ability to limit community resource consumption.

Most objects in the system are scoped by namespace, but some are excepted and have no namespace, including nodes and users.

The [Kubernetes documentation](#) has more information on namespaces.

### 3.4.3. Projects

A project is a Kubernetes namespace with additional annotations, and is the central vehicle by which access to resources for regular users is managed. A project allows a community of users to organize and manage their content in isolation from other communities. Users must be given access to projects by administrators, or if allowed to create projects, automatically have access to their own projects.

Projects can have a separate **name**, **displayName**, and **description**.

- The mandatory **name** is a unique identifier for the project and is most visible when using the CLI tools or API. The maximum name length is 63 characters.
- The optional **displayName** is how the project is displayed in the web console (defaults to **name**).
- The optional **description** can be a more detailed description of the project and is also visible in the web console.

Each project scopes its own set of:

<b>Objects</b>	Pods, services, replication controllers, etc.
<b>Policies</b>	Rules for which users can or cannot perform actions on objects.
<b>Constraints</b>	Quotas for each kind of object that can be limited.
<b>Service accounts</b>	Service accounts act automatically with designated access to objects in the project.

Cluster administrators can [create projects](#) and [delegate administrative rights](#) for the project to any member of the user community. Cluster administrators can also allow developers to create [their own projects](#).

Developers and administrators can [interact with projects](#) using [the CLI](#) or the [web console](#).

## 3.5. BUILDS AND IMAGE STREAMS

### 3.5.1. Builds

A [build](#) is the process of transforming input parameters into a resulting object. Most often, the process is used to transform input parameters or source code into a runnable image. A [BuildConfig](#) object is the definition of the entire build process.

OpenShift Enterprise leverages Kubernetes by creating Docker-formatted containers from build images and pushing them to a [container registry](#).

Build objects share common characteristics: inputs for a build, the need to complete a build process, logging the build process, publishing resources from successful builds, and publishing the final status of the build. Builds take advantage of resource restrictions, specifying limitations on resources such as CPU usage, memory usage, and build or pod execution time.

The OpenShift Enterprise build system provides extensible support for *build strategies* that are based on selectable types specified in the build API. There are three build strategies available:

- [Docker build](#)
- [Source-to-Image \(S2I\) build](#)
- [Custom build](#)

By default, Docker builds and S2I builds are supported.

The resulting object of a build depends on the builder used to create it. For Docker and S2I builds, the resulting objects are runnable images. For Custom builds, the resulting objects are whatever the builder image author has specified.

For a list of build commands, see the [Developer's Guide](#).

For more information on how OpenShift Enterprise leverages Docker for builds, see the [upstream documentation](#).

#### 3.5.1.1. Docker Build

The Docker build strategy invokes the [docker build](#) command, and it therefore expects a repository with a **Dockerfile** and all required artifacts in it to produce a runnable image.

#### 3.5.1.2. Source-to-Image (S2I) Build

[Source-to-Image \(S2I\)](#) is a tool for building reproducible, Docker-formatted container images. It produces ready-to-run images by injecting application source into a container image and assembling a new image. The new image incorporates the base image (the builder) and built source and is ready to use with the **docker run** command. S2I supports incremental builds, which re-use previously downloaded dependencies, previously built artifacts, etc.

The advantages of S2I include the following:

Image flexibility	S2I scripts can be written to inject application code into almost any existing Docker-formatted container image, taking advantage of the existing ecosystem. Note that, currently, S2I relies on <b>tar</b> to inject application source, so the image needs to be able to process tarred content.
Speed	With S2I, the assemble process can perform a large number of complex operations without creating a new layer at each step, resulting in a fast process. In addition, S2I scripts can be written to re-use artifacts stored in a previous version of the application image, rather than having to download or build them each time the build is run.
Patchability	S2I allows you to rebuild the application consistently if an underlying image needs a patch due to a security issue.
Operational efficiency	By restricting build operations instead of allowing arbitrary actions, as a <b>Dockerfile</b> would allow, the PaaS operator can avoid accidental or intentional abuses of the build system.
Operational security	Building an arbitrary <b>Dockerfile</b> exposes the host system to root privilege escalation. This can be exploited by a malicious user because the entire Docker build process is run as a user with Docker privileges. S2I restricts the operations performed as a root user and can run the scripts as a non-root user.
User efficiency	S2I prevents developers from performing arbitrary <b>yum install</b> type operations, which could slow down development iteration, during their application build.
Ecosystem	S2I encourages a shared ecosystem of images where you can leverage best practices for your applications.
Reproducibility	Produced images can include all inputs including specific versions of build tools and dependencies. This ensures that the image can be reproduced precisely.

### 3.5.1.3. Custom Build

The Custom build strategy allows developers to define a specific builder image responsible for the entire build process. Using your own builder image allows you to customize your build process.

A [Custom builder image](#) is a plain Docker-formatted container image embedded with build process logic, for example for building RPMs or base images. The **openshift/origin-custom-docker-builder** image is available on the [Docker Hub](#) registry as an example implementation of a Custom builder image.

### 3.5.2. Image Streams

An *image stream* comprises any number of [Docker-formatted container images](#) identified by tags. It presents a single virtual view of related images, similar to an image repository, and may contain images from any of the following:

1. Its own image repository in OpenShift Enterprise's [integrated registry](#)
2. Other image streams
3. Image repositories from external registries

Image streams can be used to automatically perform an action when new images are created. Builds and deployments can watch an image stream to receive notifications when new images are added and react by performing a build or deployment, respectively.

For example, if a deployment is using a certain image and a new version of that image is created, a deployment could be automatically performed.



## NOTE

See the [Developer Guide](#) for details on managing images and image streams.

### Example 3.5. Image Stream Object Definition

```
apiVersion: v1
kind: ImageStream
metadata:
  annotations:
    openshift.io/generated-by: OpenShiftNewApp
  creationTimestamp: 2016-01-29T13:33:49Z
  generation: 1
  labels:
    app: ruby-sample-build
    template: application-template-stibuild
  name: origin-ruby-sample
  namespace: test
  resourceVersion: "633"
  selflink: /oapi/v1/namespaces/test/imagestreams/origin-ruby-sample
  uid: ee2b9405-c68c-11e5-8a99-525400f25e34
spec: {}
status:
  dockerImageRepository: 172.30.56.218:5000/test/origin-ruby-sample
  tags:
    - items:
        - created: 2016-01-29T13:40:11Z
          dockerImageReference: 172.30.56.218:5000/test/origin-ruby-sample@sha256:47463d94eb5c049b2d23b03a9530bf944f8f967a0fe79147dd6b9135bf7dd13d
          generation: 1
          image:
            sha256:47463d94eb5c049b2d23b03a9530bf944f8f967a0fe79147dd6b9135bf7dd13d
            tag: latest
```

### 3.5.2.1. Image Stream Mappings

When the [integrated registry](#) receives a new image, it creates and sends an **ImageStreamMapping** to OpenShift Enterprise, providing the image's namespace (i.e., its project), name, tag, and image metadata.

This information is used to create a new image (if it does not already exist) and to tag the image into the image stream. OpenShift Enterprise stores complete metadata about each image, such as commands, entrypoint, and environment variables. Images in OpenShift Enterprise are immutable and the maximum name length is 63 characters.



**NOTE**

See the [Developer Guide](#) for details on manually tagging images.

The following **ImageStreamMapping** example results in an image being tagged as **test/origin-ruby-sample:latest**:

**Example 3.6. Image Stream Mapping Object Definition**

```

apiVersion: v1
kind: ImageStreamMapping
metadata:
  creationTimestamp: null
  name: origin-ruby-sample
  namespace: test
tag: latest
image:
  dockerImageLayers:
    - name:
sha256:5f70bf18a086007016e948b04aed3b82103a36bea41755b6cddfaf10ace3c6ef
  size: 0
    - name:
sha256:ee1dd2cb6df21971f4af6de0f1d7782b81fb63156801cfde2bb47b4247c23c29
  size: 196634330
    - name:
sha256:5f70bf18a086007016e948b04aed3b82103a36bea41755b6cddfaf10ace3c6ef
  size: 0
    - name:
sha256:5f70bf18a086007016e948b04aed3b82103a36bea41755b6cddfaf10ace3c6ef
  size: 0
    - name:
sha256:ca062656bff07f18bff46be00f40cfbb069687ec124ac0aa038fd676cfaea092
  size: 177723024
    - name:
sha256:63d529c59c92843c395befd065de516ee9ed4995549f8218eac6ff088bfa6b6e
  size: 55679776
    - name:
sha256:92114219a04977b5563d7dff71ec4caa3a37a15b266ce42ee8f43dba9798c966
  size: 11939149
  dockerImageMetadata:
    Architecture: amd64
    Config:
      Cmd:
        - /usr/libexec/s2i/run
      Entrypoint:
        - container-entrypoint
      Env:
        - RACK_ENV=production
        - OPENSIFT_BUILD_NAMESPACE=test
        - OPENSIFT_BUILD_SOURCE=https://github.com/openshift/ruby-hello-
world.git
        - EXAMPLE=sample-app
        - OPENSIFT_BUILD_NAME=ruby-sample-build-1
        - PATH=/opt/app-root/src/bin:/opt/app-
root/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin

```

```

- STI_SCRIPTS_URL=image:///usr/libexec/s2i
- STI_SCRIPTS_PATH=/usr/libexec/s2i
- HOME=/opt/app-root/src
- BASH_ENV=/opt/app-root/etc/scl_enable
- ENV=/opt/app-root/etc/scl_enable
- PROMPT_COMMAND=. /opt/app-root/etc/scl_enable
- RUBY_VERSION=2.2
ExposedPorts:
  8080/tcp: {}
Labels:
  build-date: 2015-12-23
  io.k8s.description: Platform for building and running Ruby 2.2
  applications
  io.k8s.display-name: 172.30.56.218:5000/test/origin-ruby-
  sample:latest
  io.openshift.build.commit.author: Ben Parees
  <bparees@users.noreply.github.com>
  io.openshift.build.commit.date: Wed Jan 20 10:14:27 2016 -0500
  io.openshift.build.commit.id:
  00cad9c392d39d5ef9117cbc8a31db0889eedd442
  io.openshift.build.commit.message: 'Merge pull request #51 from
  php-coder/fix_url_and_sti'
  io.openshift.build.commit.ref: master
  io.openshift.build.image: centos/ruby-22-
  centos7@sha256:3a335d7d8a452970c5b4054ad7118ff134b3a6b50a2bb6d0c07c746e8
  986b28e
  io.openshift.build.source-location:
  https://github.com/openshift/ruby-hello-world.git
  io.openshift.builder-base-version: 8d95148
  io.openshift.builder-version:
  8847438ba06307f86ac877465eadc835201241df
  io.openshift.expose-services: 8080:http
  io.openshift.s2i.scripts-url: image:///usr/libexec/s2i
  io.openshift.tags: builder,ruby,ruby22
  io.s2i.scripts-url: image:///usr/libexec/s2i
  license: GPLv2
  name: CentOS Base Image
  vendor: CentOS
User: "1001"
WorkingDir: /opt/app-root/src
Container:
  86e9a4a3c760271671ab913616c51c9f3cea846ca524bf07c04a6f6c9e103a76
ContainerConfig:
  AttachStdout: true
  Cmd:
  - /bin/sh
  - -c
  - tar -C /tmp -xf - && /usr/libexec/s2i/assemble
  Entrypoint:
  - container-entrypoint
  Env:
  - RACK_ENV=production
  - OPENSIFT_BUILD_NAME=ruby-sample-build-1
  - OPENSIFT_BUILD_NAMESPACE=test
  - OPENSIFT_BUILD_SOURCE=https://github.com/openshift/ruby-hello-
  world.git

```

```

- EXAMPLE=sample-app
- PATH=/opt/app-root/src/bin:/opt/app-
root/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin
- STI_SCRIPTS_URL=image:///usr/libexec/s2i
- STI_SCRIPTS_PATH=/usr/libexec/s2i
- HOME=/opt/app-root/src
- BASH_ENV=/opt/app-root/etc/scl_enable
- ENV=/opt/app-root/etc/scl_enable
- PROMPT_COMMAND=. /opt/app-root/etc/scl_enable
- RUBY_VERSION=2.2
ExposedPorts:
  8080/tcp: {}
Hostname: ruby-sample-build-1-build
Image: centos/ruby-22-
centos7@sha256:3a335d7d8a452970c5b4054ad7118ff134b3a6b50a2bb6d0c07c746e8
986b28e
  OpenStdin: true
  StdinOnce: true
  User: "1001"
  WorkingDir: /opt/app-root/src
Created: 2016-01-29T13:40:00Z
DockerVersion: 1.8.2.fc21
Id: 9d7fd5e2d15495802028c569d544329f4286dcd1c9c085ff5699218dbaa69b43
Parent:
57b08d979c86f4500dc8cad639c9518744c8dd39447c055a3517dc9c18d6fccd
Size: 441976279
apiVersion: "1.0"
kind: DockerImage
dockerImageMetadataVersion: "1.0"
dockerImageReference: 172.30.56.218:5000/test/origin-ruby-
sample@sha256:47463d94eb5c049b2d23b03a9530bf944f8f967a0fe79147dd6b9135bf
7dd13d

```

## 3.6. DEPLOYMENTS

### 3.6.1. Replication Controllers

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 acts to instantiate 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.

A replication controller configuration consists of:

1. The number of replicas desired (which can be adjusted at runtime).
2. A pod definition to use when creating a replicated pod.
3. 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 would require its replica count to be adjusted by an external auto-scaler.

A replication controller is a core Kubernetes object called **ReplicationController**.

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

```
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.

### 3.6.2. Jobs

A job is similar to a replication controller, in that its purpose is to create pods for specified reasons. The difference is that replication controllers are designed for pods that will be continuously running, whereas jobs are for one-time pods. A job tracks any successful completions and when the specified amount of completions have been reached, the job itself is completed.

The following example computes  $\pi$  to 2000 places, prints it out, then completes:

```
apiVersion: extensions/v1
kind: Job
metadata:
  name: pi
spec:
  selector:
    matchLabels:
      app: pi
```

```

template:
  metadata:
    name: pi
    labels:
      app: pi
  spec:
    containers:
    - name: pi
      image: perl
      command: ["perl", "-Mbignum=bpi", "-wle", "print bpi(2000)"]
      restartPolicy: Never

```

See the [Jobs](#) topic for more information on how to use jobs.

### 3.6.3. Deployments and Deployment Configurations

Building on replication controllers, OpenShift Enterprise adds expanded support for the software development and deployment lifecycle with the concept of deployments. In the simplest case, a deployment just creates a new replication controller and lets it start up pods. However, OpenShift Enterprise deployments 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 OpenShift Enterprise **DeploymentConfiguration** object defines the following details of a deployment:

1. The elements of a **ReplicationController** definition.
2. Triggers for creating a new deployment automatically.
3. The strategy for transitioning between deployments.
4. Life cycle 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 in order 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.

For detailed instructions on how to create and interact with deployments, refer to [Deployments](#).

Here is an example **DeploymentConfiguration** definition with some omissions and callouts:

```

apiVersion: 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 2
  strategy:
    type: Rolling 3

```

1. A **ConfigChange** trigger causes a new deployment to be created any time the replication controller template changes.
2. An **ImageChange** trigger causes a new deployment to be created each time a new version of the backing image is available in the named image stream.
3. The default **Rolling** strategy makes a downtime-free transition between deployments.

## 3.7. ROUTES

### 3.7.1. Overview


An OpenShift Enterprise route exposes a [service](#) at a host name, like *www.example.com*, so that external clients can reach it by name.

DNS resolution for a host name is handled separately from routing. Your administrator may have configured a [DNS wildcard entry](#) that will resolve to the OpenShift Enterprise node that is running the OpenShift Enterprise router. If you are using a different host name you may need to modify its DNS records independently to resolve to the node that is running the router.

Each route consists of a name (limited to 63 characters), a service selector, and an optional security configuration.

### 3.7.2. Routers

An OpenShift Enterprise administrator can deploy *routers* to nodes in an OpenShift Enterprise cluster, which enable routes [created by developers](#) to be used by external clients. The routing layer in OpenShift Enterprise is pluggable, and two available [router plug-ins](#) are provided and supported by default.



#### NOTE

See the [Installation and Configuration](#) guide for information on deploying a router.

A router uses the service selector to find the [service](#) and the endpoints backing the service. When both router and service provide load balancing, OpenShift Enterprise uses the router load balancing. A router detects relevant changes in the IP addresses of its services, and adapts its configuration accordingly. This is useful for custom routers to communicate modifications of API objects to an external routing solution.

The path of a request starts with the DNS resolution of a host name to one or more routers. The suggested method is to define a cloud domain with a wildcard DNS entry pointing to one or more virtual IP (VIP) addresses backed by multiple router instances. Routes using names and addresses outside the cloud domain require configuration of individual DNS entries.

When there are fewer VIP addresses than routers, the routers corresponding to the number of addresses

are *active* and the rest are *passive*. A passive router is also known as a *hot-standby* router. For example, with two VIP addresses and three routers, you have an "active-active-passive" configuration. See [High Availability](#) for more information on router VIP configuration.

Routes can be [sharded](#) among the set of routers. Administrators can set up sharding on a cluster-wide basis and users can set up sharding for the namespace in their project. Sharding allows the operator to define multiple router groups. Each router in the group serves only a subset of traffic.

OpenShift Enterprise routers provide external host name mapping and load balancing of [service](#) end points over protocols that pass distinguishing information directly to the router; the host name must be present in the protocol in order for the router to determine where to send it.

Router plug-ins assume they can bind to host ports 80 (HTTP) and 443 (HTTPS), by default. This means that routers must be placed on nodes where those ports are not otherwise in use. Alternatively, a router can be configured to listen on other ports by setting the **ROUTER\_SERVICE\_HTTP\_PORT** and **ROUTER\_SERVICE\_HTTPS\_PORT** environment variables.

Because a router binds to ports on the host node, only one router listening on those ports can be on each node if the router uses host networking (the default). Cluster networking is configured such that all routers can access all pods in the cluster.

Routers support the following protocols:

- HTTP
- HTTPS (with SNI)
- WebSockets
- TLS with SNI



#### NOTE

WebSocket traffic uses the same route conventions and supports the same TLS termination types as other traffic.

### 3.7.2.1. Template Routers

A *template router* is a type of router that provides certain infrastructure information to the underlying router implementation, such as:

- A wrapper that watches endpoints and routes.
- Endpoint and route data, which is saved into a consumable form.
- Passing the internal state to a configurable template and executing the template.
- Calling a reload script.

### 3.7.3. Available Router Plug-ins

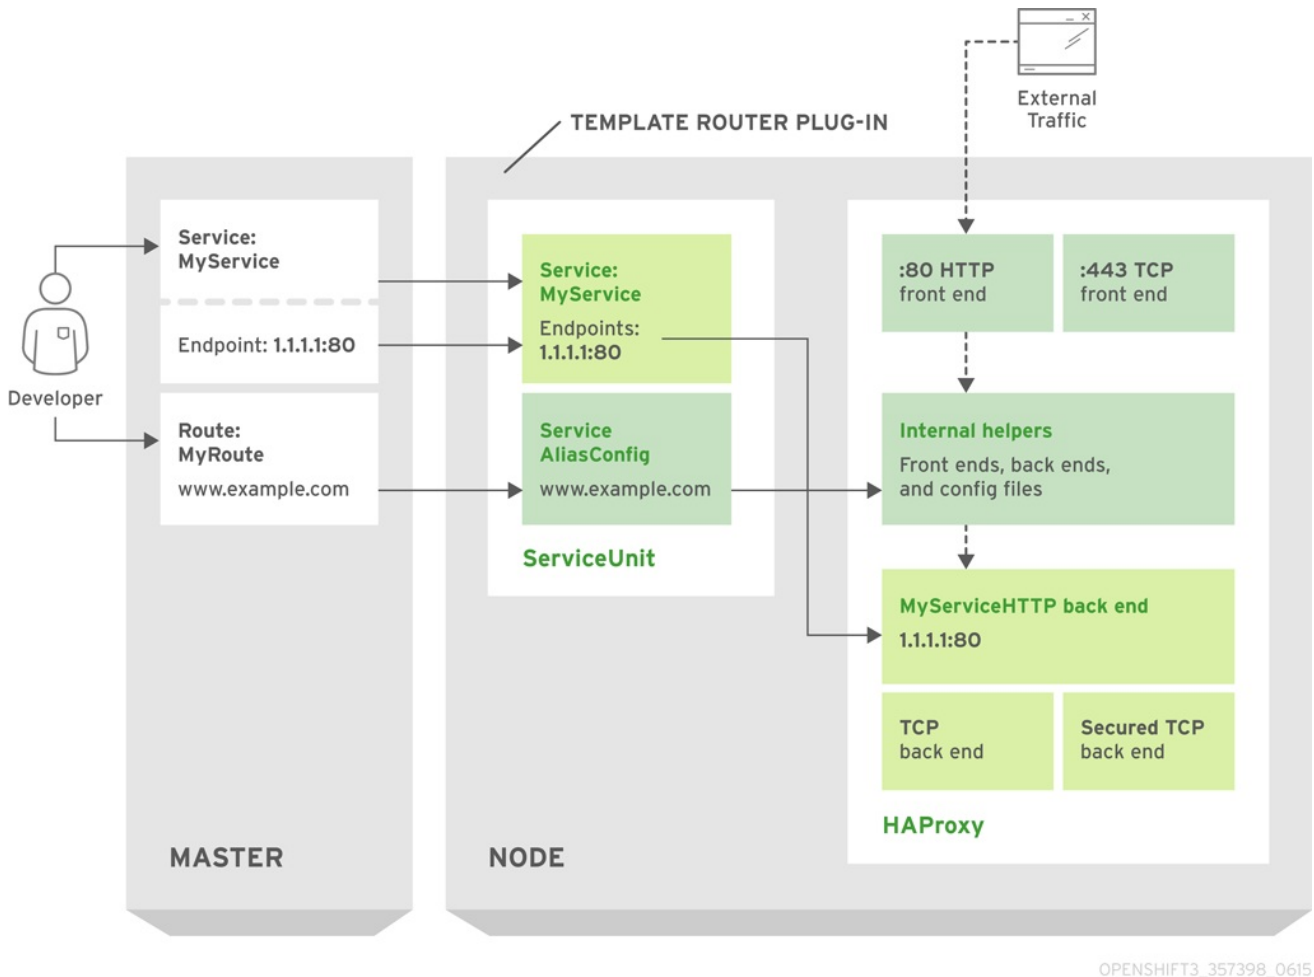
The following router plug-ins are provided and supported in OpenShift Enterprise. Instructions on deploying these routers are available in [Deploying a Router](#).

#### 3.7.3.1. HAProxy Template Router

The HAProxy template router implementation is the reference implementation for a template router plug-in. It uses the **openshift3/ose-haproxy-router** repository to run an HAProxy instance alongside the template router plug-in.

The following diagram illustrates how data flows from the master through the plug-in and finally into an HAProxy configuration:

**Figure 3.1. HAProxy Router Data Flow**



## Sticky Sessions

Implementing sticky sessions is up to the underlying router configuration. The default HAProxy template implements sticky sessions using the **balance source** directive which balances based on the source IP. In addition, the template router plug-in provides the service name and namespace to the underlying implementation. This can be used for more advanced configuration such as implementing stick-tables that synchronize between a set of peers.

Specific configuration for this router implementation is stored in the **haproxy-config.template** file located in the **/var/lib/haproxy/conf** directory of the router container.



## NOTE

The **balance source** directive does not distinguish between external client IP addresses; because of the NAT configuration, the originating IP address (HAProxy remote) is the same. Unless the HAProxy router is running with **hostNetwork: true**, all external clients will be routed to a single pod.



### 3.7.3.2. F5 Router



#### NOTE

The F5 router plug-in is available starting in OpenShift Enterprise 3.0.2.

The F5 router plug-in integrates with an existing **F5 BIG-IP®** system in your environment. **F5 BIG-IP®** version 11.4 or newer is required in order to have the F5 iControl REST API. The F5 router supports [unsecured](#), [edge terminated](#), [re-encryption terminated](#), and [passthrough terminated](#) routes matching on HTTP vhost and request path.

The F5 router has feature parity with the [HAProxy template router](#), and has additional features over the **F5 BIG-IP®** support in OpenShift Enterprise 2. Compared with the **routing-daemon** used in earlier versions, the F5 router additionally supports:

- path-based routing (using policy rules),
- re-encryption (implemented using client and server SSL profiles), and
- passthrough of encrypted connections (implemented using an iRule that parses the SNI protocol and uses a data group that is maintained by the F5 router for the servername lookup).



#### NOTE

Passthrough routes are a special case: path-based routing is technically impossible with passthrough routes because **F5 BIG-IP®** itself does not see the HTTP request, so it cannot examine the path. The same restriction applies to the template router; it is a technical limitation of passthrough encryption, not a technical limitation of OpenShift Enterprise.

### Routing Traffic to Pods Through the SDN

Because **F5 BIG-IP®** is external to the [OpenShift Enterprise SDN](#), a cluster administrator must create a peer-to-peer tunnel between **F5 BIG-IP®** and a host that is on the SDN, typically an OpenShift Enterprise node host. This [ramp node](#) can be configured as [unschedulable](#) for pods so that it will not be doing anything except act as a gateway for the **F5 BIG-IP®** host. It is also possible to configure multiple such hosts and use the OpenShift Enterprise **ipfailover** feature for redundancy; the **F5 BIG-IP®** host would then need to be configured to use the **ipfailover** VIP for its tunnel's remote endpoint.

### F5 Integration Details

The operation of the F5 router is similar to that of the OpenShift Enterprise **routing-daemon** used in earlier versions. Both use REST API calls to:

- create and delete pools,
- add endpoints to and delete them from those pools, and
- configure policy rules to route to pools based on vhost.

Both also use **scp** and **ssh** commands to upload custom TLS/SSL certificates to **F5 BIG-IP®**.

The F5 router configures pools and policy rules on virtual servers as follows:

- When a user creates or deletes a route on OpenShift Enterprise, the router creates a pool to **F5**

**BIG-IP®** for the route (if no pool already exists) and adds a rule to, or deletes a rule from, the policy of the appropriate vserver: the HTTP vserver for non-TLS routes, or the HTTPS vserver for edge or re-encrypt routes. In the case of edge and re-encrypt routes, the router also uploads and configures the TLS certificate and key. The router supports host- and path-based routes.



## NOTE

Passthrough routes are a special case: to support those, it is necessary to write an iRule that parses the SNI ClientHello handshake record and looks up the servername in an F5 data-group. The router creates this iRule, associates the iRule with the vserver, and updates the F5 data-group as passthrough routes are created and deleted. Other than this implementation detail, passthrough routes work the same way as other routes.

- When a user creates a service on OpenShift Enterprise, the router adds a pool to **F5 BIG-IP®** (if no pool already exists). As endpoints on that service are created and deleted, the router adds and removes corresponding pool members.
- When a user deletes the route and all endpoints associated with a particular pool, the router deletes that pool.

### 3.7.4. Route Host Names

In order for services to be exposed externally, an OpenShift Enterprise route allows you to associate a service with an externally-reachable host name. This edge host name is then used to route traffic to the service.

When two routes claim the same host, the oldest route wins. If additional routes with different path fields are defined in the same namespace, those paths will be added. If multiple routes with the same path are used, the oldest takes priority.

#### Example 3.7. A Route with a Specified Host:

```
apiVersion: v1
kind: Route
metadata:
  name: host-route
spec:
  host: www.example.com 1
  to:
    kind: Service
    name: service-name
```

1 1 1 Specifies the externally-reachable host name used to expose a service.

#### Example 3.8. A Route Without a Host:

```
apiVersion: v1
kind: Route
metadata:
  name: no-route-hostname
spec:
```

```
to:
  kind: Service
  name: service-name
```

If a host name is not provided as part of the route definition, then OpenShift Enterprise automatically generates one for you. The generated host name is of the form:

```
<route-name>[-<namespace>].<suffix>
```

The following example shows the OpenShift Enterprise-generated host name for the above configuration of a route without a host added to a namespace **mynamespace**:

### Example 3.9. Generated Host Name

```
no-route-hostname-mynamespace.router.default.svc.cluster.local 1
```

- 1** The generated host name suffix is the default routing subdomain **router.default.svc.cluster.local**.

A cluster administrator can also [customize the suffix used as the default routing subdomain](#) for their environment.

## 3.7.5. Route Types

Routes can be either secured or unsecured. Secure routes provide the ability to use several types of TLS termination to serve certificates to the client. Routers support [edge](#), [passthrough](#), and [re-encryption](#) termination.

### Example 3.10. Unsecured Route Object YAML Definition

```
apiVersion: v1
kind: Route
metadata:
  name: route-unsecured
spec:
  host: www.example.com
  to:
    kind: Service
    name: service-name
```

Unsecured routes are simplest to configure, as they require no key or certificates, but secured routes offer security for connections to remain private.

A secured route is one that specifies the TLS termination of the route. The available types of termination are [described below](#).

## 3.7.6. Path Based Routes

Path based routes specify a path component that can be compared against a URL (which requires that

the traffic for the route be HTTP based) such that multiple routes can be served using the same hostname, each with a different path. Routers should match routes based on the most specific path to the least; however, this depends on the router implementation. The following table shows example routes and their accessibility:

**Table 3.1. Route Availability**

Route	When Compared to	Accessible
<i>www.example.com/test</i>	<i>www.example.com/test</i>	Yes
	<i>www.example.com</i>	No
<i>www.example.com/test</i> and <i>www.example.com</i>	<i>www.example.com/test</i>	Yes
	<i>www.example.com</i>	Yes
<i>www.example.com</i>	<i>www.example.com/test</i>	Yes (Matched by the host, not the route)
	<i>www.example.com</i>	Yes

#### Example 3.11. An Unsecured Route with a Path:

```

apiVersion: v1
kind: Route
metadata:
  name: route-unsecured
spec:
  host: www.example.com
  path: "/test" 1
  to:
    kind: Service
    name: service-name

```

1 The path is the only added attribute for a path-based route.



#### NOTE

Path-based routing is not available when using passthrough TLS, as the router does not terminate TLS in that case and cannot read the contents of the request.

### 3.7.7. Secured Routes

Secured routes specify the TLS termination of the route and, optionally, provide a key and certificate(s).



## NOTE

TLS termination in OpenShift Enterprise relies on [SNI](#) for serving custom certificates. Any non-SNI traffic received on port 443 is handled with TLS termination and a default certificate (which may not match the requested hostname, resulting in validation errors).

Secured routes can use any of the following three types of secure TLS termination.

### Edge Termination

With edge termination, TLS termination occurs at the router, prior to proxying traffic to its destination. TLS certificates are served by the front end of the router, so they must be configured into the route, otherwise the [router's default certificate](#) will be used for TLS termination.

#### Example 3.12. A Secured Route Using Edge Termination

```
apiVersion: v1
kind: Route
metadata:
  name: route-edge-secured 1
spec:
  host: www.example.com
  to:
    kind: Service
    name: service-name 2
  tls:
    termination: edge 3
    key: |- 4
      -----BEGIN PRIVATE KEY-----
      [...]
      -----END PRIVATE KEY-----
    certificate: |- 5
      -----BEGIN CERTIFICATE-----
      [...]
      -----END CERTIFICATE-----
    caCertificate: |- 6
      -----BEGIN CERTIFICATE-----
      [...]
      -----END CERTIFICATE-----
```

- 1 2 The name of the object, which is limited to 63 characters.
- 3 The **termination** field is **edge** for edge termination.
- 4 The **key** field is the contents of the PEM format key file.
- 5 The **certificate** field is the contents of the PEM format certificate file.
- 6 An optional CA certificate may be required to establish a certificate chain for validation.

Because TLS is terminated at the router, connections from the router to the endpoints over the internal network are not encrypted.

Edge-terminated routes can specify an **insecureEdgeTerminationPolicy** that enables traffic on insecure schemes (**HTTP**) to be disabled, allowed or redirected. The allowed values for **insecureEdgeTerminationPolicy** are: **None** or empty (for disabled), **Allow** or **Redirect**. The default **insecureEdgeTerminationPolicy** is to disable traffic on the insecure scheme. A common use case is to allow content to be served via a secure scheme but serve the assets (example images, stylesheets and javascript) via the insecure scheme.

### Example 3.13. A Secured Route Using Edge Termination Allowing HTTP Traffic

```
apiVersion: v1
kind: Route
metadata:
  name: route-edge-secured-allow-insecure ❶
spec:
  host: www.example.com
  to:
    kind: Service
    name: service-name ❷
  tls:
    termination: edge ❸
    insecureEdgeTerminationPolicy: Allow ❹
    [ ... ]
```

- ❶ ❷ The name of the object, which is limited to 63 characters.
- ❸ The **termination** field is **edge** for edge termination.
- ❹ The insecure policy to allow requests sent on an insecure scheme **HTTP**.

### Example 3.14. A Secured Route Using Edge Termination Redirecting HTTP Traffic to HTTPS

```
apiVersion: v1
kind: Route
metadata:
  name: route-edge-secured-redirect-insecure ❶
spec:
  host: www.example.com
  to:
    kind: Service
    name: service-name ❷
  tls:
    termination: edge ❸
    insecureEdgeTerminationPolicy: Redirect ❹
    [ ... ]
```

- ❶ ❷ The name of the object, which is limited to 63 characters.
- ❸ The **termination** field is **edge** for edge termination.
- ❹ The insecure policy to redirect requests sent on an insecure scheme **HTTP** to a secure scheme **HTTPS**.

## Passthrough Termination

With passthrough termination, encrypted traffic is sent straight to the destination without the router providing TLS termination. Therefore no key or certificate is required.

### Example 3.15. A Secured Route Using Passthrough Termination

```
apiVersion: v1
kind: Route
metadata:
  name: route-passthrough-secured ❶
spec:
  host: www.example.com
  to:
    kind: Service
    name: service-name ❷
  tls:
    termination: passthrough ❸
```

- ❶ ❷ The name of the object, which is limited to 63 characters.
- ❸ The **termination** field is set to **passthrough**. No other encryption fields are needed.

The destination pod is responsible for serving certificates for the traffic at the endpoint. This is currently the only method that can support requiring client certificates (also known as two-way authentication).

## Re-encryption Termination

Re-encryption is a variation on edge termination where the router terminates TLS with a certificate, then re-encrypts its connection to the endpoint which may have a different certificate. Therefore the full path of the connection is encrypted, even over the internal network. The router uses health checks to determine the authenticity of the host.

### Example 3.16. A Secured Route Using Re-Encrypt Termination

```
apiVersion: v1
kind: Route
metadata:
  name: route-pt-secured ❶
spec:
  host: www.example.com
  to:
    kind: Service
    name: service-name ❷
  tls:
    termination: reencrypt ❸
    key: [as in edge termination]
    certificate: [as in edge termination]
    caCertificate: [as in edge termination]
    destinationCACertificate: |- ❹
```

```
-----BEGIN CERTIFICATE-----
[...]
```

- 1 2 The name of the object, which is limited to 63 characters.
- 3 The **termination** field is set to **reencrypt**. Other fields are as in edge termination.
- 4 The **destinationCACertificate** field specifies a CA certificate to validate the endpoint certificate, securing the connection from the router to the destination. This field is required, but only for re-encryption.

### 3.7.8. Router Sharding

In OpenShift Enterprise, each route can have any number of [labels](#) in its **metadata** field. A router uses *selectors* (also known as a *selection expression*) to select a subset of routes from the entire pool of routes to serve. A selection expression can also involve labels on the route's namespace. The selected routes form a *router shard*. You can [create](#) and [modify](#) router shards independently from the routes, themselves.

This design supports *traditional* sharding as well as *overlapped* sharding. In traditional sharding, the selection results in no overlapping sets and a route belongs to exactly one shard. In overlapped sharding, the selection results in overlapping sets and a route can belong to many different shards. For example, a single route may belong to a **SLA=high** shard (but not **SLA=medium** or **SLA=low** shards), as well as a **geo=west** shard (but not a **geo=east** shard).

Another example of overlapped sharding is a set of routers that select based on namespace of the route:

Router	Selection	Namespaces
router-1	<b>A* — J*</b>	<b>A*, B*, C*, D*, E*, F*, G*, H*, I*, J*</b>
router-2	<b>K* — T*</b>	<b>K*, L*, M*, N*, O*, P*, Q*, R*, S*, T*</b>
router-3	<b>Q* — Z*</b>	<b>Q*, R*, S*, T*, U*, V*, W*, X*, Y*, Z*</b>

Both **router-2** and **router-3** serve routes that are in the namespaces **Q\*, R\*, S\*, T\***. To change this example from overlapped to traditional sharding, we could change the selection of **router-2** to **K\* — P\***, which would eliminate the overlap.

When routers are sharded, a given route is bound to zero or more routers in the group. The route binding ensures uniqueness of the route across the shard. Uniqueness allows secure and non-secure versions of the same route to exist within a single shard. This implies that routes now have a visible life cycle that moves from created to bound to active.

In the sharded environment the first route to hit the shard reserves the right to exist there indefinitely, even across restarts.

During a green/blue deployment a route may be selected in multiple routers. An OpenShift Enterprise application administrator may wish to bleed traffic from one version of the application to another and then turn off the old version.



Sharding can be done by the administrator at a cluster level and by the user at a project/namespace level. When namespace labels are used, the service account for the router must have **cluster-reader** permission to permit the router to access the labels in the namespace.



## NOTE

For two or more routes that claim the same host name, the resolution order is based on the age of the route and the oldest route would win the claim to that host. In the case of sharded routers, routes are selected based on their labels matching the router's selection criteria. There is no consistent way to determine when labels are added to a route. So if an older route claiming an existing host name is "re-labelled" to match the router's selection criteria, it will replace the existing route based on the above mentioned resolution order (oldest route wins).

## 3.8. TEMPLATES

### 3.8.1. Overview

A template describes a set of [objects](#) that can be parameterized and processed to produce a list of objects for creation by OpenShift Enterprise. The objects to create can include anything that users have permission to create within a project, for example [services](#), [build configurations](#), and [deployment configurations](#). A template may also define a set of [labels](#) to apply to every object defined in the template.

#### Example 3.17. A Simple Template Object Definition (YAML)

```
apiVersion: v1
kind: Template
metadata:
  name: redis-template ①
  annotations:
    description: "Description" ②
    iconClass: "icon-redis" ③
    tags: "database,nosql" ④
  objects: ⑤
  - apiVersion: v1
    kind: Pod
    metadata:
      name: redis-master
    spec:
      containers:
        - env:
            - name: REDIS_PASSWORD
              value: ${REDIS_PASSWORD} ⑥
          image: dockerfile/redis
          name: master
          ports:
            - containerPort: 6379
              protocol: TCP
      parameters: ⑦
      - description: Password used for Redis authentication
        from: '[A-Z0-9]{8}' ⑧
        generate: expression
```

```

    name: REDIS_PASSWORD
  labels:
    redis: master

```

- 1 The name of the template
- 2 Optional description for the template
- 3 The icon that will be shown in the UI for this template; the name of a CSS class defined in [the web console source](#) (search content for "openshift-logos-icon").
- 4 A list of arbitrary tags that this template will have in the UI
- 5 A list of objects the template will create (in this case, a single pod)
- 6 Parameter value that will be substituted during processing
- 7 A list of parameters for the template
- 8 An expression used to generate a random password if not specified
- 9 A list of labels to apply to all objects on create

A template describes a set of related object definitions to be created together, as well as a set of parameters for those objects. For example, an application might consist of a frontend web application backed by a database; each consists of a service object and deployment configuration object, and they share a set of credentials (parameters) for the frontend to authenticate to the backend. The template can be processed, either specifying parameters or allowing them to be automatically generated (for example, a unique DB password), in order to instantiate the list of objects in the template as a cohesive application.

Templates can be processed from a definition in a file or from an existing OpenShift Enterprise API object. Cluster administrators can [define standard templates in the API](#) that are available for all users to process, while users can [define their own templates](#) within their own projects.

Administrators and developers can [interact with templates](#) using the CLI and web console.

### 3.8.2. Parameters

Templates allow you to define parameters which take on a value. That value is then substituted wherever the parameter is referenced. References can be defined in any text field in the objects list field.

Each parameter describes a variable and the variable value which can be referenced in any text field in the **objects** list field. During processing, the value can be set explicitly or it can be generated by OpenShift Enterprise.

An explicit value can be set as the parameter default using the **value** field:

```

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

```

The **generate** field can be set to 'expression' to specify generated values. The **from** field should specify the pattern for generating the value using a pseudo regular expression syntax:

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

In the example above, processing will generate 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**, 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}**.

## CHAPTER 4. ADDITIONAL CONCEPTS

### 4.1. NETWORKING

Kubernetes ensures that pods are able to network with each other, and allocates each pod an IP address from an internal network. This ensures all containers within the pod behave as if they were on the same host. Giving each pod its own IP address means that pods can be treated like physical hosts or virtual machines in terms of port allocation, networking, naming, service discovery, load balancing, application configuration, and migration.

Creating links between pods is unnecessary. However, it is not recommended that you have a pod talk to another directly by using the IP address. Instead, we recommend that you create a [service](#), then interact with the service.

#### 4.1.1. OpenShift Enterprise DNS

If you are running multiple [services](#), such as frontend and backend services for use with multiple pods, in order for the frontend pods to communicate with the backend services, environment variables are created for user names, service IP, and more. If the service is deleted and recreated, a new IP address can be assigned to the service, and requires the frontend pods to be recreated in order to pick up the updated values for the service IP environment variable. Additionally, the backend service has to be created before any of the frontend pods to ensure that the service IP is generated properly and that it can be provided to the frontend pods as an environment variable.

For this reason, OpenShift Enterprise has a built-in DNS so that the services can be reached by the service DNS as well as the service IP/port. OpenShift Enterprise supports split DNS by running [SkyDNS](#) on the master that answers DNS queries for services. The master listens to port 53 by default.

When the node starts, the following message indicates the Kubelet is correctly resolved to the master:

```
0308 19:51:03.118430      4484 node.go:197] Started Kubelet for node
openshiftdev.local, server at 0.0.0.0:10250
I0308 19:51:03.118459      4484 node.go:199]   Kubelet is setting 10.0.2.15
as a
DNS nameserver for domain "local"
```

If the second message does not appear, the Kubernetes service may not be available.

On a node host, each container's nameserver has the master name added to the front, and the default search domain for the container will be `<pod_namespace>.cluster.local`. The container will then direct any nameserver queries to the master before any other nameservers on the node, which is the default behavior for Docker-formatted containers. The master will answer queries on the `.cluster.local` domain that have the following form:

**Table 4.1. DNS Example Names**

Object Type	Example
Default	<code>&lt;pod_namespace&gt;.cluster.local</code>
Services	<code>&lt;service&gt;.&lt;pod_namespace&gt;.svc.cluster.local</code>
Endpoints	<code>&lt;name&gt;.&lt;namespace&gt;.endpoints.cluster.local</code>

This prevents having to restart frontend pods in order to pick up new services, which creates a new IP for the service. This also removes the need to use environment variables, as pods can use the service DNS. Also, as the DNS does not change, you can reference database services as **db.local** in config files. Wildcard lookups are also supported, as any lookups resolve to the service IP, and removes the need to create the backend service before any of the frontend pods, since the service name (and hence DNS) is established upfront.

This DNS structure also covers headless services, where a portal IP is not assigned to the service and the kube-proxy does not load-balance or provide routing for its endpoints. Service DNS can still be used and responds with multiple A records, one for each pod of the service, allowing the client to round-robin between each pod.

### 4.1.2. Network Plugins

OpenShift Enterprise supports the same plugin model as Kubernetes for networking pods. The following network plugins are currently supported by OpenShift Enterprise.

### 4.1.3. OpenShift Enterprise SDN

OpenShift Enterprise deploys a software-defined networking (SDN) approach for connecting pods in an OpenShift Enterprise cluster. The OpenShift Enterprise SDN connects all pods across all node hosts, providing a unified cluster network.

OpenShift Enterprise SDN is automatically installed and configured as part of the Ansible-based installation procedure. Further administration should not be required; however, [further details on the design and operation of OpenShift Enterprise SDN](#) are provided for those who are curious or need to troubleshoot problems.

## 4.2. OPENSIFT SDN

### 4.2.1. Overview

OpenShift Enterprise uses a software-defined networking (SDN) approach to provide a unified cluster network that enables communication between pods across the OpenShift Enterprise cluster. This pod network is established and maintained by the OpenShift Enterprise SDN, which configures an overlay network using Open vSwitch (OVS).

OpenShift Enterprise SDN provides two SDN plug-ins for configuring the pod network:

- The **ovs-subnet** plug-in is the original plug-in which provides a "flat" pod network where every pod can communicate with every other pod and service.
- The **ovs-multitenant** plug-in provides OpenShift Enterprise project level isolation for pods and services. Each project receives a unique Virtual Network ID (VNID) that identifies traffic from pods assigned to the project. Pods from different projects cannot send packets to or receive packets from pods and services of a different project. However, projects which receive VNID 0 are more privileged in that they are allowed to communicate with all other pods, and all other pods can communicate with them. In OpenShift Enterprise clusters, the **default** project has VNID 0. This facilitates certain services like the load balancer, etc. to communicate with all other pods in the cluster and vice versa.

Following is a detailed discussion of the design and operation of OpenShift Enterprise SDN, which may be useful for troubleshooting.

**NOTE**

Information on configuring the SDN on masters and nodes is available in [Configuring the SDN](#).

### 4.2.2. Design on Masters

On an OpenShift Enterprise master, OpenShift Enterprise SDN maintains a registry of nodes, stored in **etcd**. When the system administrator registers a node, OpenShift Enterprise SDN allocates an unused subnet from the cluster network and stores this subnet in the registry. When a node is deleted, OpenShift Enterprise SDN deletes the subnet from the registry and considers the subnet available to be allocated again.

In the default configuration, the cluster network is the **10.128.0.0/14** network (i.e. **10.128.0.0 - 10.131.255.255**), and nodes are allocated **/23** subnets (i.e., **10.128.0.0/23**, **10.128.2.0/23**, **10.128.4.0/23**, and so on). This means that the cluster network has 512 subnets available to assign to nodes, and a given node is allocated 510 addresses that it can assign to the containers running on it. The size and address range of the cluster network are configurable, as is the host subnet size.

Note that OpenShift Enterprise SDN on a master does not configure the local (master) host to have access to any cluster network. Consequently, a master host does not have access to pods via the cluster network, unless it is also running as a node.

When using the **ovs-multitenant** plug-in, the OpenShift Enterprise SDN master also watches for the creation and deletion of projects, and assigns VXLAN VNIDs to them, which will be used later by the nodes to isolate traffic correctly.

### 4.2.3. Design on Nodes

On a node, OpenShift Enterprise SDN first registers the local host with the SDN master in the aforementioned registry so that the master allocates a subnet to the node.

Next, OpenShift Enterprise SDN creates and configures six network devices:

- **br0**, the OVS bridge device that pod containers will be attached to. OpenShift Enterprise SDN also configures a set of non-subnet-specific flow rules on this bridge. The **ovs-multitenant** plug-in does this immediately.
- **lbr0**, a Linux bridge device, which is configured as the Docker service's bridge and given the cluster subnet gateway address (eg, 10.128.x.1/23).
- **tun0**, an OVS internal port (port 2 on **br0**). This *also* gets assigned the cluster subnet gateway address, and is used for external network access. OpenShift Enterprise SDN configures **netfilter** and routing rules to enable access from the cluster subnet to the external network via NAT.
- **vlinuxbr** and **vovsbr**, two Linux peer virtual Ethernet interfaces. **vlinuxbr** is added to **lbr0** and **vovsbr** is added to **br0** (port 9 with the **ovs-subnet** plug-in and port 3 with the **ovs-multitenant** plug-in) to provide connectivity for containers created directly with the Docker service outside of OpenShift Enterprise.
- **vxlan0**, the OVS VXLAN device (port 1 on **br0**), which provides access to containers on remote nodes.

Each time a pod is started on the host, OpenShift Enterprise SDN:

1. moves the host side of the pod's veth interface pair from the **lbr0** bridge (where the Docker service placed it when starting the container) to the OVS bridge **br0**.
2. adds OpenFlow rules to the OVS database to route traffic addressed to the new pod to the correct OVS port.
3. in the case of the **ovs-multitenant** plug-in, adds OpenFlow rules to tag traffic coming from the pod with the pod's VNID, and to allow traffic into the pod if the traffic's VNID matches the pod's VNID (or is the privileged VNID 0). Non-matching traffic is filtered out by a generic rule.

The pod is allocated an IP address in the cluster subnet by the Docker service itself because the Docker service is told to use the **lbr0** bridge, which OpenShift Enterprise SDN has assigned the cluster gateway address (eg. 10.128.x.1/23). Note that the **tun0** is also assigned the cluster gateway IP address because it is the default gateway for all traffic destined for external networks, but these two interfaces do not conflict because the **lbr0** interface is only used for IPAM and no OpenShift Enterprise SDN pods are connected to it.

OpenShift Enterprise SDN nodes also watch for subnet updates from the SDN master. When a new subnet is added, the node adds OpenFlow rules on **br0** so that packets with a destination IP address in the remote subnet go to **vxlان0** (port 1 on **br0**) and thus out onto the network. The **ovs-subnet** plug-in sends all packets across the VXLAN with VNID 0, but the **ovs-multitenant** plug-in uses the appropriate VNID for the source container.

#### 4.2.4. Packet Flow

Suppose you have two containers, A and B, where the peer virtual Ethernet device for container A's **eth0** is named **vethA** and the peer for container B's **eth0** is named **vethB**.



#### NOTE

If the Docker service's use of peer virtual Ethernet devices is not already familiar to you, review [Docker's advanced networking documentation](#).

Now suppose first that container A is on the local host and container B is also on the local host. Then the flow of packets from container A to container B is as follows:

**eth0** (in A's netns) → **vethA** → **br0** → **vethB** → **eth0** (in B's netns)

Next, suppose instead that container A is on the local host and container B is on a remote host on the cluster network. Then the flow of packets from container A to container B is as follows:

**eth0** (in A's netns) → **vethA** → **br0** → **vxlان0** → network<sup>[1]</sup> → **vxlان0** → **br0** → **vethB** → **eth0** (in B's netns)

Finally, if container A connects to an external host, the traffic looks like:

**eth0** (in A's netns) → **vethA** → **br0** → **tun0** → (NAT) → **eth0** (physical device) → Internet

Almost all packet delivery decisions are performed with OpenFlow rules in the OVS bridge **br0**, which simplifies the plug-in network architecture and provides flexible routing. In the case of the **ovs-multitenant** plug-in, this also provides enforceable [network isolation](#).

#### 4.2.5. Network Isolation

You can use the **ovs-multitenant** plug-in to achieve network isolation. When a packet exits a pod

assigned to a non-default project, the OVS bridge **br0** tags that packet with the project's assigned VNID. If the packet is directed to another IP address in the node's cluster subnet, the OVS bridge only allows the packet to be delivered to the destination pod if the VNIDs match.

If a packet is received from another node via the VXLAN tunnel, the Tunnel ID is used as the VNID, and the OVS bridge only allows the packet to be delivered to a local pod if the tunnel ID matches the destination pod's VNID.

Packets destined for other cluster subnets are tagged with their VNID and delivered to the VXLAN tunnel with a tunnel destination address of the node owning the cluster subnet.

As described before, VNID 0 is privileged in that traffic with any VNID is allowed to enter any pod assigned VNID 0, and traffic with VNID 0 is allowed to enter any pod. Only the **default** OpenShift Enterprise project is assigned VNID 0; all other projects are assigned unique, isolation-enabled VNIDs. Cluster administrators can optionally [control the pod network](#) for the project using the administrator CLI.

## 4.3. AUTHENTICATION

### 4.3.1. Overview

The authentication layer identifies the user associated with requests to the OpenShift Enterprise API. The authorization layer then uses information about the requesting user to determine if the request should be allowed.

As an administrator, you can [configure authentication](#) using a [master configuration file](#).

### 4.3.2. Users and Groups

A *user* in OpenShift Enterprise is an entity that can make requests to the OpenShift Enterprise API. Typically, this represents the account of a developer or administrator that is interacting with OpenShift Enterprise.

A user can be assigned to one or more *groups*, each of which represent a certain set of users. Groups are useful when [managing authorization policies](#) to grant permissions to multiple users at once, for example allowing access to [objects](#) within a [project](#), versus granting them to users individually.

In addition to explicitly defined groups, there are also system groups, or *virtual groups*, that are automatically provisioned by OpenShift. These can be seen when [viewing cluster bindings](#).

In the default set of virtual groups, note the following in particular:

Virtual Group	Description
<b>system:authenticated</b>	Automatically associated with all authenticated users.
<b>system:authenticated:oauth</b>	Automatically associated with all users authenticated with an OAuth access token.
<b>system:unauthenticated</b>	Automatically associated with all unauthenticated users.

### 4.3.3. API Authentication



Requests to the OpenShift Enterprise API are authenticated using the following methods:

### OAuth Access Tokens

- Obtained from the OpenShift Enterprise OAuth server using the `<master>/oauth/authorize` and `<master>/oauth/token` endpoints.
- Sent as an **Authorization: Bearer...** header or an **access\_token=...** query parameter

### X.509 Client Certificates

- Requires a HTTPS connection to the API server.
- Verified by the API server against a trusted certificate authority bundle.
- The API server creates and distributes certificates to controllers to authenticate themselves.

Any request with an invalid access token or an invalid certificate is rejected by the authentication layer with a 401 error.

If no access token or certificate is presented, the authentication layer assigns the **system:anonymous** virtual user and the **system:unauthenticated** virtual group to the request. This allows the authorization layer to determine which requests, if any, an anonymous user is allowed to make.

See the [REST API Overview](#) for more information and examples.

### 4.3.4. OAuth

The OpenShift Enterprise master includes a built-in OAuth server. Users obtain OAuth access tokens to authenticate themselves to the API.

When a person requests a new OAuth token, the OAuth server uses the configured [identity provider](#) to determine the identity of the person making the request.

It then determines what user that identity maps to, creates an access token for that user, and returns the token for use.

### OAuth Clients

Every request for an OAuth token must specify the OAuth client that will receive and use the token. The following OAuth clients are automatically created when starting the OpenShift Enterprise API:

OAuth Client	Usage
<b>openshift-web-console</b>	Requests tokens for the web console.
<b>openshift-browser-client</b>	Requests tokens at <code>&lt;master&gt;/oauth/token/request</code> with a user-agent that can handle interactive logins.
<b>openshift-challenging-client</b>	Requests tokens with a user-agent that can handle <b>WWW-Authenticate</b> challenges.

To register additional clients:

```
$ oc create -f <(echo '{
  "kind": "OAuthClient",
  "apiVersion": "v1",
  "metadata": {
    "name": "demo" 1
  },
  "secret": "...", 2
  "redirectURIs": [
    "http://www.example.com/" 3
  ]
}' )
```

- 1 The **name** of the OAuth client is used as the **client\_id** parameter when making requests to **<master>/oauth/authorize** and **<master>/oauth/token**.
- 2 The **secret** is used as the **client\_secret** parameter when making requests to **<master>/oauth/token**.
- 3 The **redirect\_uri** parameter specified in requests to **<master>/oauth/authorize** and **<master>/oauth/token** must be equal to (or prefixed by) one of the URIs in **redirectURIs**.

## Integrations

All requests for OAuth tokens involve a request to **<master>/oauth/authorize**. Most authentication integrations place an authenticating proxy in front of this endpoint, or configure OpenShift Enterprise to validate credentials against a backing [identity provider](#). Requests to **<master>/oauth/authorize** can come from user-agents that cannot display interactive login pages, such as the CLI. Therefore, OpenShift Enterprise supports authenticating using a **WWW-Authenticate** challenge in addition to interactive login flows.

If an authenticating proxy is placed in front of the **<master>/oauth/authorize** endpoint, it should send unauthenticated, non-browser user-agents **WWW-Authenticate** challenges, rather than displaying an interactive login page or redirecting to an interactive login flow.



### NOTE

To prevent cross-site request forgery (CSRF) attacks against browser clients, Basic authentication challenges should only be sent if a **X-CSRF-Token** header is present on the request. Clients that expect to receive Basic **WWW-Authenticate** challenges should set this header to a non-empty value.

If the authenticating proxy cannot support **WWW-Authenticate** challenges, or if OpenShift Enterprise is configured to use an identity provider that does not support **WWW-Authenticate** challenges, users can visit **<master>/oauth/token/request** using a browser to obtain an access token manually.

## Obtaining OAuth Tokens

The OAuth server supports standard [authorization code grant](#) and the [implicit grant](#) OAuth authorization flows.

When requesting an OAuth token using the implicit grant flow (**response\_type=token**) with a `client_id` configured to request WWW-Authenticate challenges (like **openshift-challenging-client**), these are the possible server responses from `/oauth/authorize`, and how they should be handled:

Status	Content	Client response
302	<b>Location</b> header containing an <b>access_token</b> parameter in the URL fragment ( <a href="#">RFC 4.2.2</a> )	Use the <b>access_token</b> value as the OAuth token
302	<b>Location</b> header containing an <b>error</b> query parameter ( <a href="#">RFC 4.1.2.1</a> )	Fail, optionally surfacing the <b>error</b> (and optional <b>error_description</b> ) query values to the user
302	Other <b>Location</b> header	Follow the redirect, and process the result using these rules
401	<b>WWW-Authenticate</b> header present	Respond to challenge if type is recognized (e.g. <b>Basic</b> , <b>Negotiate</b> , etc), resubmit request, and process the result using these rules
401	<b>WWW-Authenticate</b> header missing	No challenge authentication is possible. Fail and show response body (which might contain links or details on alternate methods to obtain an OAuth token)
Other	Other	Fail, optionally surfacing response body to the user

## 4.4. AUTHORIZATION

### 4.4.1. Overview

Authorization policies determine whether a user is allowed to perform a given [action](#) within a project. This allows platform administrators to use the [cluster policy](#) to control who has various access levels to the OpenShift Enterprise platform itself and all projects. It also allows developers to use [local policy](#) to control who has access to their [projects](#). Note that authorization is a separate step from [authentication](#), which is more about determining the identity of who is taking the action.

Authorization is managed using:

<b>Rules</b>	Sets of permitted <a href="#">verbs</a> on a set of <a href="#">objects</a> . For example, whether something can <b>create</b> pods.
<b>Roles</b>	Collections of rules. <a href="#">Users and groups</a> can be associated with, or <i>bound</i> to, multiple <a href="#">roles</a> at the same time.
<b>Bindings</b>	Associations between users and/or groups with a <a href="#">role</a> .

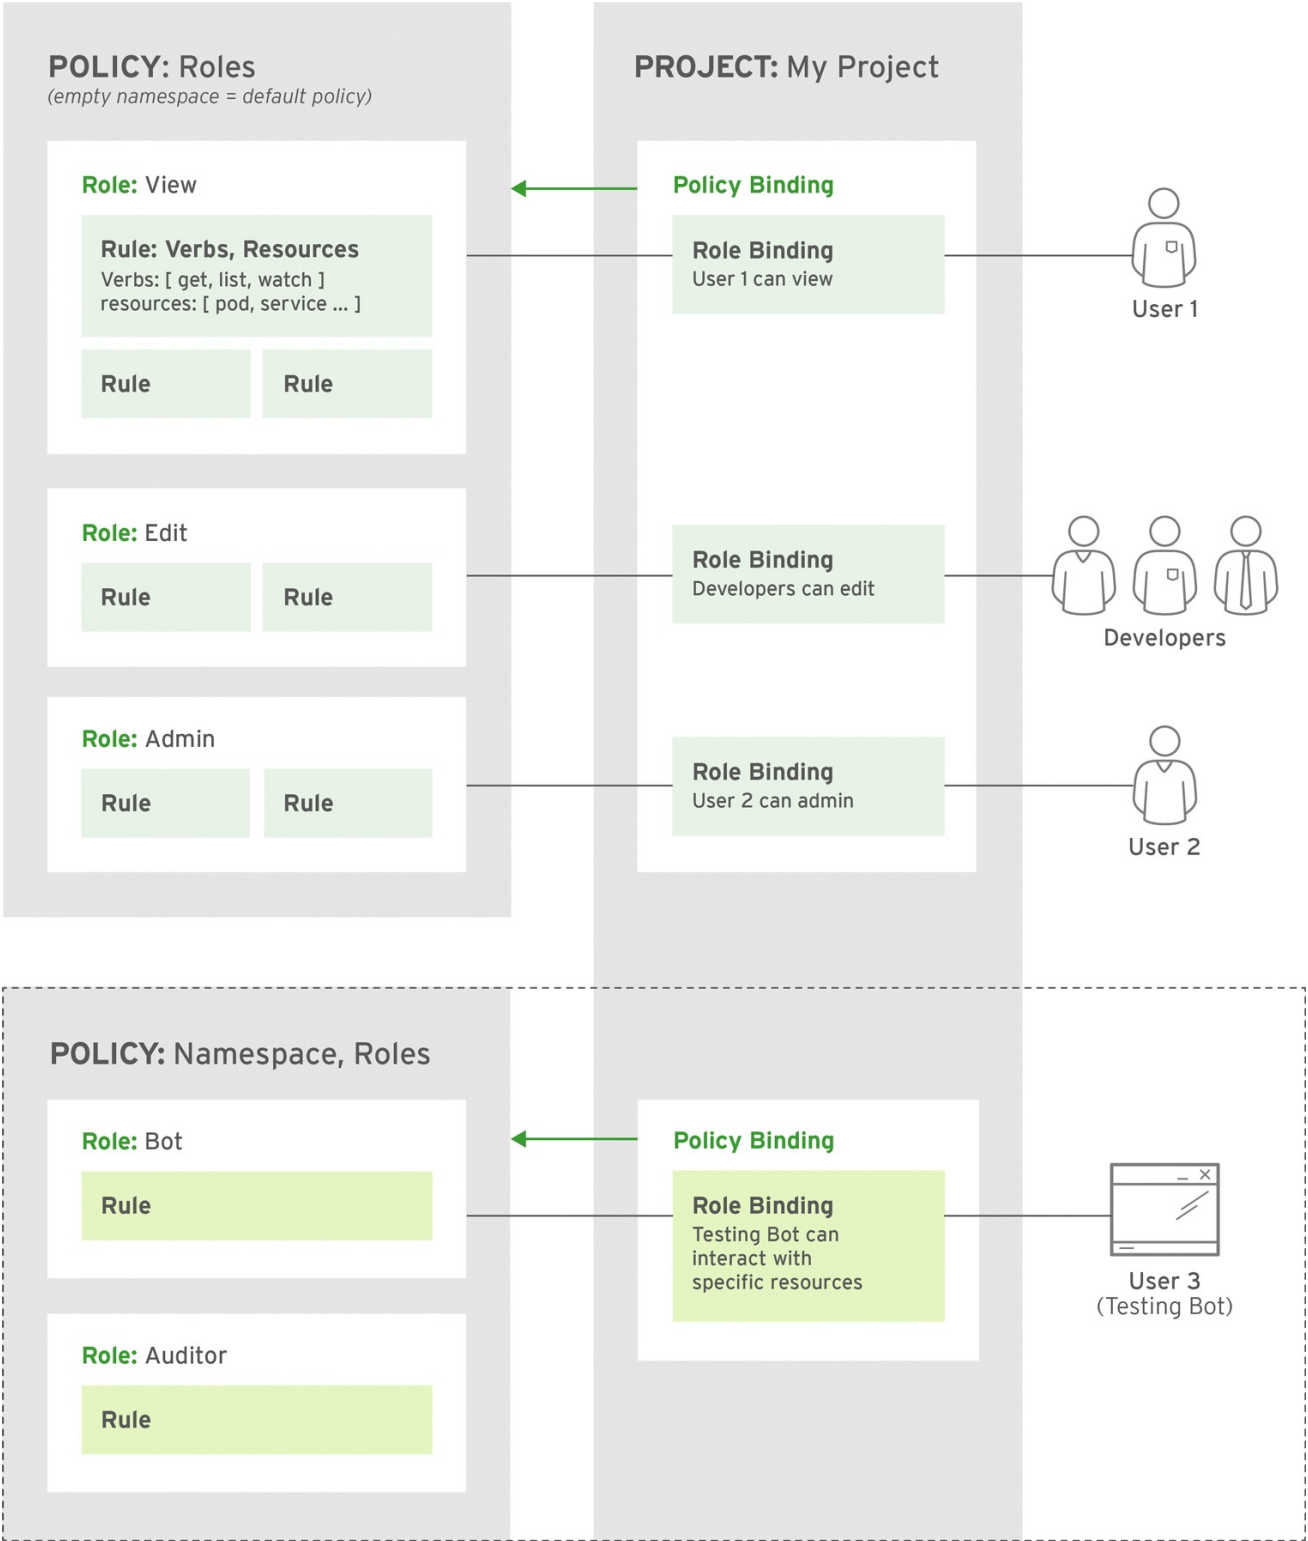
Cluster administrators can visualize rules, roles, and bindings [using the CLI](#). For example, consider the following excerpt from viewing a policy, showing rule sets for the **admin** and **basic-user** [default roles](#):

```
admin    Verbs      Resources                                Resource Names Extension
[create delete get list update watch] [projects
resourcegroup:exposedkube resourcegroup:exposedopenshift
resourcegroup:granter secrets]      []
[get list watch] [resourcegroup:allkube resourcegroup:allkube-status
resourcegroup:allopenshift-status resourcegroup:policy] []
basic-user Verbs      Resources                                Resource Names Extension
[get]       [users]                [~]
[list]      [projectrequests]      []
[list]      [projects]              []
[create]    [subjectaccessreviews] []
IsPersonalSubjectAccessReview
```

The following excerpt from viewing policy bindings shows the above roles bound to various users and groups:

```
RoleBinding[admins]:
  Role: admin
  Users: [alice system:admin]
  Groups: []
RoleBinding[basic-user]:
  Role: basic-user
  Users: [joe]
  Groups: [devel]
```

The relationships between the the policy roles, policy bindings, users, and developers are illustrated below.



OPENSIFT\_393387\_0316

### 4.4.2. Evaluating Authorization

Several factors are combined to make the decision when OpenShift Enterprise evaluates authorization:

Identity	In the context of authorization, both the user name and list of groups the user belongs to.
----------	---

<b>Action</b>	The action being performed. In most cases, this consists of:	
	Project	The <a href="#">project</a> being accessed.
	Verb	Can be <b>get</b> , <b>list</b> , <b>create</b> , <b>update</b> , <b>delete</b> , <b>deletecollection</b> or <b>watch</b> .
	Resource Name	The API endpoint being accessed.
<b>Bindings</b>	The full list of <a href="#">bindings</a> .	

OpenShift Enterprise evaluates authorizations using the following steps:

1. The identity and the project-scoped action is used to find all bindings that apply to the user or their groups.
2. Bindings are used to locate all the roles that apply.
3. Roles are used to find all the rules that apply.
4. The action is checked against each rule to find a match.
5. If no matching rule is found, the action is then denied by default.

#### 4.4.3. Cluster Policy and Local Policy

There are two levels of authorization policy:

<b>Cluster policy</b>	<a href="#">Roles</a> and bindings that are applicable across all projects. Roles that exist in the cluster policy are considered <i>cluster roles</i> . Cluster bindings can only reference cluster roles.
<b>Local policy</b>	<a href="#">Roles</a> and bindings that are scoped to a given project. Roles that exist only in a local policy are considered <i>local roles</i> . Local bindings can reference both cluster and local roles.

This two-level hierarchy allows re-usability over multiple projects through the cluster policy while allowing customization inside of individual projects through local policies.

During evaluation, both the cluster bindings and the local bindings are used. For example:

1. Cluster-wide "allow" rules are checked.
2. Locally-bound "allow" rules are checked.
3. Deny by default.

#### 4.4.4. Roles

Roles are collections of policy [rules](#), which are sets of permitted verbs that can be performed on a set of resources. OpenShift Enterprise includes a set of default roles that can be added to users and groups in the [cluster policy](#) or in a [local policy](#).

Default Role	Description
<b>admin</b>	A project manager. If used in a <a href="#">local binding</a> , an <b>admin</b> user will have rights to view any resource in the project and modify any resource in the project except for role creation and quota. If the <b>cluster-admin</b> wants to allow an <b>admin</b> to modify roles, the <b>cluster-admin</b> must create a project-scoped <b>Policy</b> object using JSON.
<b>basic-user</b>	A user that can get basic information about projects and users.
<b>cluster-admin</b>	A super-user that can perform any action in any project. When granted to a user within a local policy, they have full control over quota and roles and every action on every resource in the project.
<b>cluster-status</b>	A user that can get basic cluster status information.
<b>edit</b>	A user that can modify most objects in a project, but does not have the power to view or modify roles or bindings.
<b>self-provisioner</b>	A user that can create their own projects.
<b>view</b>	A user who cannot make any modifications, but can see most objects in a project. They cannot view or modify roles or bindings.

## TIP

Remember that [users and groups](#) can be associated with, or *bound* to, multiple roles at the same time.

Cluster administrators can visualize these roles, including a matrix of the verbs and resources each are associated using the CLI to [view the cluster roles](#). Additional **system:** roles are listed as well, which are used for various OpenShift Enterprise system and component operations.

By default in a local policy, only the binding for the **admin** role is immediately listed when using the CLI to [view local bindings](#). However, if other default roles are added to users and groups within a local policy, they become listed in the CLI output, as well.

If you find that these roles do not suit you, a **cluster-admin** user can create a **policyBinding** object named **<projectname>:default** with the CLI using a JSON file. This allows the project **admin** to bind users to roles that are defined only in the **<projectname>** local policy.



## IMPORTANT

The **cluster-** role assigned by the project administrator is limited in a project. It is not the same **cluster-** role granted by the **cluster-admin** or **system:admin**.

Cluster roles are [roles](#) defined at the cluster level, but can be bound either at the cluster level or at the project level.

[Learn how to create a local role for a project.](#)

#### 4.4.4.1. Updating Cluster Roles

After any [OpenShift Enterprise cluster upgrade](#), the recommended default roles may have been updated. See [Updating Policy Definitions](#) for instructions on getting to the new recommendations using:

```
$ oadm policy reconcile-cluster-roles
```

#### 4.4.5. Security Context Constraints

In addition to [authorization policies](#) that control what a user can do, OpenShift Enterprise provides *security context constraints* (SCC) that control the actions that a [pod](#) can perform and what it has the ability to access. Administrators can [manage SCCs](#) using the CLI.

SCCs are also very useful for [managing access to persistent storage](#).

SCCs are objects that define a set of conditions that a pod must run with in order to be accepted into the system. They allow an administrator to control the following:

1. Running of [privileged containers](#).
2. Capabilities a container can request to be added.
3. Use of host directories as volumes.
4. The SELinux context of the container.
5. The user ID.
6. The use of host namespaces and networking.
7. Allocating an **FSGroup** that owns the pod's volumes
8. Configuring allowable supplemental groups
9. Requiring the use of a read only root file system
10. Controlling the usage of volume types

Seven SCCs are added to the cluster by default, and are viewable by cluster administrators using the CLI:

```
$ oc get scc
NAME          PRIV          CAPS          SELINUX          RUNASUSER
FSGROUP       SUPGROUP     PRIORITY     READONLYROOTFS  VOLUMES
anyuid        false        []           MustRunAs        RunAsAny
RunAsAny      RunAsAny     10           false            [configMap
downwardAPI   emptyDir    persistentVolumeClaim secret]
hostaccess    false        []           MustRunAs        MustRunAsRange
MustRunAs     RunAsAny     <none>       false            [configMap
downwardAPI   emptyDir    hostPath    persistentVolumeClaim secret]
hostmount-anyuid false        []           MustRunAs        RunAsAny
RunAsAny      RunAsAny     <none>       false            [configMap
downwardAPI   emptyDir    hostPath    persistentVolumeClaim secret]
hostnetwork   false        []           MustRunAs        MustRunAsRange
```



MustRunAs	MustRunAs	<none>	false	[configMap
downwardAPI	emptyDir	persistentVolumeClaim	secret]	
nonroot	false	[]	MustRunAs	MustRunAsNonRoot
RunAsAny	RunAsAny	<none>	false	[configMap
downwardAPI	emptyDir	persistentVolumeClaim	secret]	
privileged	true	[]	RunAsAny	RunAsAny
RunAsAny	RunAsAny	<none>	false	[*]
restricted	false	[]	MustRunAs	MustRunAsRange
MustRunAs	RunAsAny	<none>	false	[configMap
downwardAPI	emptyDir	persistentVolumeClaim	secret]	

The definition for each SCC is also viewable by cluster administrators using the CLI. For example, for the privileged SCC:

```
# oc export scc/privileged

allowHostDirVolumePlugin: true
allowHostIPC: true
allowHostNetwork: true
allowHostPID: true
allowHostPorts: true
allowPrivilegedContainer: true
allowedCapabilities: null
apiVersion: v1
defaultAddCapabilities: null
fsGroup: ①
  type: RunAsAny
groups: ②
- system:cluster-admins
- system:nodes
kind: SecurityContextConstraints
metadata:
  annotations:
    kubernetes.io/description: 'privileged allows access to all privileged
and host
  features and the ability to run as any user, any group, any fsGroup,
and with
  any SELinux context. WARNING: this is the most relaxed SCC and
should be used
  only for cluster administration. Grant with caution.'
  creationTimestamp: null
  name: privileged
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities: null
runAsUser: ③
  type: RunAsAny
seLinuxContext: ④
  type: RunAsAny
supplementalGroups: ⑤
  type: RunAsAny
users: ⑥
- system:serviceaccount:default:registry
- system:serviceaccount:default:router
```

```
- system:serviceaccount:openshift-infra:build-controller
volumes:
- '*'
```

- 1 The **FSGroup** strategy which dictates the allowable values for the Security Context
- 2 The groups that have access to this SCC
- 3 The run as user strategy type which dictates the allowable values for the Security Context
- 4 The SELinux context strategy type which dictates the allowable values for the Security Context
- 5 The supplemental groups strategy which dictates the allowable supplemental groups for the Security Context
- 6 The users who have access to this SCC


The **users** and **groups** fields on the SCC control which SCCs can be used. By default, cluster administrators, nodes, and the build controller are granted access to the privileged SCC. All authenticated users are granted access to the restricted SCC.

The privileged SCC:

- allows privileged pods.
- allows host directories to be mounted as volumes.
- allows a pod to run as any user.
- allows a pod to run with any MCS label.
- allows a pod to use the host's IPC namespace.
- allows a pod to use the host's PID namespace.
- allows a pod to use any FSGroup.
- allows a pod to use any supplemental group.

The restricted SCC:

- ensures pods cannot run as privileged.
- ensures pods cannot use host directory volumes.
- requires that a pod run as a user in a pre-allocated range of UIDs.
- requires that a pod run with a pre-allocated MCS label.
- allows a pod to use any FSGroup.
- allows a pod to use any supplemental group.

**NOTE**

For more information about each SCC, see the **kubernetes.io/description** annotation available on the SCC.

SCCs are comprised of settings and strategies that control the security features a pod has access to. These settings fall into three categories:

<b>Controlled by a boolean</b>	Fields of this type default to the most restrictive value. For example, <b>AllowPrivilegedContainer</b> is always set to <b>false</b> if unspecified.
<b>Controlled by an allowable set</b>	Fields of this type are checked against the set to ensure their value is allowed.
<b>Controlled by a strategy</b>	Items that have a strategy to generate a value provide: <ul style="list-style-type: none"> <li>• A mechanism to generate the value, and</li> <li>• A mechanism to ensure that a specified value falls into the set of allowable values.</li> </ul>

#### 4.4.5.1. SCC Strategies

##### 4.4.5.1.1. RunAsUser

1. **MustRunAs** - Requires a **runAsUser** to be configured. Uses the configured **runAsUser** as the default. Validates against the configured **runAsUser**.
2. **MustRunAsRange** - Requires minimum and maximum values to be defined if not using pre-allocated values. Uses the minimum as the default. Validates against the entire allowable range.
3. **MustRunAsNonRoot** - Requires that the pod be submitted with a non-zero **runAsUser** or have the **USER** directive defined in the image. No default provided.
4. **RunAsAny** - No default provided. Allows any **runAsUser** to be specified.

##### 4.4.5.1.2. SELinuxContext

1. **MustRunAs** - Requires **seLinuxOptions** to be configured if not using pre-allocated values. Uses **seLinuxOptions** as the default. Validates against **seLinuxOptions**.
2. **RunAsAny** - No default provided. Allows any **seLinuxOptions** to be specified.

##### 4.4.5.1.3. SupplementalGroups

1. **MustRunAs** - Requires at least one range to be specified if not using pre-allocated values. Uses the minimum value of the first range as the default. Validates against all ranges.
2. **RunAsAny** - No default provided. Allows any **supplementalGroups** to be specified.

##### 4.4.5.1.4. FSGroup


1. **MustRunAs** - Requires at least one range to be specified if not using pre-allocated values. Uses the minimum value of the first range as the default. Validates against the first ID in the first range.
2. **RunAsAny** - No default provided. Allows any **fsGroup** ID to be specified.

#### 4.4.5.2. Controlling Volumes

The usage of specific volume types can be controlled by setting the **volumes** field of the SCC. The allowable values of this field correspond to the volume sources that are defined when creating a volume:

- **azureFile**
- **flocker**
- **flexVolume**
- **hostPath**
- **emptyDir**
- **gcePersistentDisk**
- **awsElasticBlockStore**
- **gitRepo**
- **secret**
- **nfs**
- **iscsi**
- **glusterfs**
- **persistentVolumeClaim**
- **rbd**
- **cinder**
- **cephFS**
- **downwardAPI**
- **fc**
- **configMap**
- **\***

The recommended minimum set of allowed volumes for new SCCs are **configMap**, **downwardAPI**, **emptyDir**, **persistentVolumeClaim**, and **secret**.



#### NOTE

\* is a special value to allow the use of all volume types.

**NOTE**

For backwards compatibility, the usage of **allowHostDirVolumePlugin** overrides settings in the **volumes** field. For example, if **allowHostDirVolumePlugin** is set to false but allowed in the **volumes** field, then the **hostPath** value will be removed from **volumes**.

**4.4.5.3. Admission**

*Admission control* with SCCs allows for control over the creation of resources based on the capabilities granted to a user.

In terms of the SCCs, this means that an admission controller can inspect the user information made available in the context to retrieve an appropriate set of SCCs. Doing so ensures the pod is authorized to make requests about its operating environment or to generate a set of constraints to apply to the pod.

The set of SCCs that admission uses to authorize a pod are determined by the user identity and groups that the user belongs to. Additionally, if the pod specifies a service account, the set of allowable SCCs includes any constraints accessible to the service account.

Admission uses the following approach to create the final security context for the pod:

1. Retrieve all SCCs available for use.
2. Generate field values for security context settings that were not specified on the request.
3. Validate the final settings against the available constraints.

If a matching set of constraints is found, then the pod is accepted. If the request cannot be matched to an SCC, the pod is rejected.

A pod must validate every field against the SCC. The following are examples for just two of the fields that must be validated:

**NOTE**

These examples are in the context of a strategy using the preallocated values.

**A FSGroup SCC Strategy of MustRunAs**

If the pod defines a **fsGroup** ID, then that ID must equal the default **FSGroup** ID. Otherwise, the pod is not validated by that SCC and the next SCC is evaluated. If the **FSGroup** strategy is **RunAsAny** and the pod omits a **fsGroup** ID, then the pod matches the SCC based on **FSGroup** (though other strategies may not validate and thus cause the pod to fail).

**A SupplementalGroups SCC Strategy of MustRunAs**

If the pod specification defines one or more **SupplementalGroups** IDs, then the pod's IDs must equal one of the IDs in the namespace's **openshift.io/sa.scc.supplemental-groups** annotation. Otherwise, the pod is not validated by that SCC and the next SCC is evaluated. If the **SupplementalGroups** setting is **RunAsAny** and the pod specification omits a **SupplementalGroups** ID, then the pod matches the SCC based on **SupplementalGroups** (though other strategies may not validate and thus cause the pod to fail).

**4.4.5.3.1. SCC Prioritization**

SCCs have a priority field that affects the ordering when attempting to validate a request by the admission controller. A higher priority SCC is moved to the front of the set when sorting. When the complete set of available SCCs are determined they are ordered by:

1. Highest priority first, nil is considered a 0 priority
2. If priorities are equal, the SCCs will be sorted from most restrictive to least restrictive
3. If both priorities and restrictions are equal the SCCs will be sorted by name

By default, the anyuid SCC granted to cluster administrators is given priority in their SCC set. This allows cluster administrators to run pods as any user by without specifying a **RunAsUser** on the pod's **SecurityContext**. The administrator may still specify a **RunAsUser** if they wish.

#### 4.4.5.3.2. Understanding Pre-allocated Values and Security Context Constraints

The admission controller is aware of certain conditions in the security context constraints that trigger it to look up pre-allocated values from a namespace and populate the security context constraint before processing the pod. Each SCC strategy is evaluated independently of other strategies, with the pre-allocated values (where allowed) for each policy aggregated with pod specification values to make the final values for the various IDs defined in the running pod.

The following SCCs cause the admission controller to look for pre-allocated values when no ranges are defined in the pod specification:

1. A **RunAsUser** strategy of **MustRunAsRange** with no minimum or maximum set. Admission looks for the **openshift.io/sa.scc.uid-range** annotation to populate range fields.
2. An **SELinuxContext** strategy of **MustRunAs** with no level set. Admission looks for the **openshift.io/sa.scc.mcs** annotation to populate the level.
3. A **FSGroup** strategy of **MustRunAs**. Admission looks for the **openshift.io/sa.scc.supplemental-groups** annotation.
4. A **SupplementalGroups** strategy of **MustRunAs**. Admission looks for the **openshift.io/sa.scc.supplemental-groups** annotation.

During the generation phase, the security context provider will default any values that are not specifically set in the pod. Defaulting is based on the strategy being used:

1. **RunAsAny** and **MustRunAsNonRoot** strategies do not provide default values. Thus, if the pod needs a field defined (for example, a group ID), this field must be defined inside the pod specification.
2. **MustRunAs** (single value) strategies provide a default value which is always used. As an example, for group IDs: even if the pod specification defines its own ID value, the namespace's default field will also appear in the pod's groups.
3. **MustRunAsRange** and **MustRunAs** (range-based) strategies provide the minimum value of the range. As with a single value **MustRunAs** strategy, the namespace's default value will appear in the running pod. If a range-based strategy is configurable with multiple ranges, it will provide the minimum value of the first configured range.

**NOTE**

**FSGroup** and **SupplementalGroups** strategies fall back to the **openshift.io/sa.scc.uid-range** annotation if the **openshift.io/sa.scc.supplemental-groups** annotation does not exist on the namespace. If neither exist, the SCC will fail to create.

**NOTE**

By default, the annotation-based **FSGroup** strategy configures itself with a single range based on the minimum value for the annotation. For example, if your annotation reads **1/3**, the **FSGroup** strategy will configure itself with a minimum and maximum of **1**. If you want to allow more groups to be accepted for the **FSGroup** field, you can configure a custom SCC that does not use the annotation.

**NOTE**

The **openshift.io/sa.scc.supplemental-groups** annotation accepts a comma delimited list of blocks in the format of **<start>/<length>** or **<start>-<end>**. The **openshift.io/sa.scc.uid-range** annotation accepts only a single block.

## 4.5. PERSISTENT STORAGE

### 4.5.1. Overview

Managing storage is a distinct problem from managing compute resources. OpenShift Enterprise leverages the Kubernetes persistent volume (PV) framework to allow administrators to provision persistent storage for a cluster. Using persistent volume claims (PVCs), developers can request PV resources without having specific knowledge of the underlying storage infrastructure.

PVCs are specific to a [project](#) and are created and used by developers as a means to use a PV. PV resources on their own are not scoped to any single project; they can be shared across the entire OpenShift Enterprise cluster and claimed from any project. After a PV has been [bound](#) to a PVC, however, that PV cannot then be bound to additional PVCs. This has the effect of scoping a bound PV to a single [namespace](#) (that of the binding project).

PVs are defined by a **PersistentVolume** API object, which represents a piece of existing networked storage in the cluster that has been provisioned by an administrator. It is a resource in the cluster just like a node is a cluster resource. PVs are volume plug-ins like **Volumes**, but have a lifecycle independent of any individual [pod](#) that uses the PV. PV objects capture the details of the implementation of the storage, be that NFS, iSCSI, or a cloud-provider-specific storage system.

**IMPORTANT**

High-availability of storage in the infrastructure is left to the underlying storage provider.

PVCs are defined by a **PersistentVolumeClaim** API object, which represents a request for storage by a developer. It is similar to a pod in that pods consume node resources and PVCs consume PV resources. For example, pods can request specific levels of resources (e.g., CPU and memory), while PVCs can request specific [storage capacity](#) and [access modes](#) (e.g., they can be mounted once read/write or many times read-only).

### 4.5.2. Lifecycle of a Volume and Claim

PVs are resources in the cluster. PVCs are requests for those resources and also act as claim checks to the resource. The interaction between PVs and PVCs have the following lifecycle.

#### 4.5.2.1. Provisioning

A cluster administrator creates some number of PVs. They carry the details of the real storage that is available for use by cluster users. They exist in the API and are available for consumption.

#### 4.5.2.2. Binding

A user creates a **PersistentVolumeClaim** with a specific amount of storage requested and with certain access modes. A control loop in the master watches for new PVCs, finds a matching PV (if possible), and binds them together. The user will always get at least what they asked for, but the volume may be in excess of what was requested.

Claims remain unbound indefinitely if a matching volume does not exist. Claims are bound as matching volumes become available. For example, a cluster provisioned with many 50Gi volumes would not match a PVC requesting 100Gi. The PVC can be bound when a 100Gi PV is added to the cluster.

#### 4.5.2.3. Using

Pods use claims as volumes. The cluster inspects the claim to find the bound volume and mounts that volume for a pod. For those volumes that support multiple access modes, the user specifies which mode is desired when using their claim as a volume in a pod.

Once a user has a claim and that claim is bound, the bound PV belongs to the user for as long as they need it. Users schedule pods and access their claimed PVs by including a **persistentVolumeClaim** in their pod's volumes block. See [below](#) for syntax details.

#### 4.5.2.4. Releasing

When a user is done with a volume, they can delete the PVC object from the API which allows reclamation of the resource. The volume is considered "released" when the claim is deleted, but it is not yet available for another claim. The previous claimant's data remains on the volume which must be handled according to policy.

#### 4.5.2.5. Reclaiming

The reclaim policy of a **PersistentVolume** tells the cluster what to do with the volume after it is released. Currently, volumes can either be *retained* or *recycled*.

Retention allows for manual reclamation of the resource. For those volume plug-ins that support it, recycling performs a basic scrub on the volume (e.g., `rm -rf /<volume>/*`) and makes it available again for a new claim.

### 4.5.3. Persistent Volumes

Each PV contains a **spec** and **status**, which is the specification and status of the volume.

#### Example 4.1. Persistent Volume Object Definition

```
apiVersion: v1
kind: PersistentVolume
metadata:
```



```

    name: pv0003
  spec:
    capacity:
      storage: 5Gi
    accessModes:
      - ReadWriteOnce
    persistentVolumeReclaimPolicy: Recycle
    nfs:
      path: /tmp
      server: 172.17.0.2

```

#### 4.5.3.1. Types of Persistent Volumes

OpenShift Enterprise supports the following **PersistentVolume** plug-ins:

- [NFS](#)
- [HostPath](#)
- [GlusterFS](#)
- [Ceph RBD](#)
- [OpenStack Cinder](#)
- [AWS Elastic Block Store \(EBS\)](#)
- [GCE Persistent Disk](#)
- [iSCSI](#)
- [Fibre Channel](#)

#### 4.5.3.2. Capacity

Generally, a PV will have a specific storage capacity. This is set using the PV's **capacity** attribute. See the [Kubernetes Resource Model](#) to understand the units expected by **capacity**.

Currently, storage capacity is the only resource that can be set or requested. Future attributes may include IOPS, throughput, etc.

#### 4.5.3.3. Access Modes

A **PersistentVolume** can be mounted on a host in any way supported by the resource provider. Providers will have different capabilities and each PV's access modes are set to the specific modes supported by that particular volume. For example, NFS can support multiple read/write clients, but a specific NFS PV might be exported on the server as read-only. Each PV gets its own set of access modes describing that specific PV's capabilities.

Claims are matched to volumes with similar access modes. The only two matching criteria are access modes and size. A claim's access modes represent a request. Therefore, the user may be granted more, but never less. For example, if a claim requests RWO, but the only volume available was an NFS PV (RWO+ROX+RWX), the claim would match NFS because it supports RWO.

Direct matches are always attempted first. The volume's modes must match or contain more modes than you requested. The size must be greater than or equal to what is expected. If two types of volumes (NFS and iSCSI, for example) both have the same set of access modes, then either of them will match a claim with those modes. There is no ordering between types of volumes and no way to choose one type over another.

All volumes with the same modes are grouped, then sorted by size (smallest to largest). The binder gets the group with matching modes and iterates over each (in size order) until one size matches.

The access modes are:

Access Mode	CLI Abbreviation	Description
ReadWriteOnce	<b>RWO</b>	The volume can be mounted as read-write by a single node.
ReadOnlyMany	<b>ROX</b>	The volume can be mounted read-only by many nodes.
ReadWriteMany	<b>RWX</b>	The volume can be mounted as read-write by many nodes.



### IMPORTANT

A volume's **AccessModes** are descriptors of the volume's capabilities. They are not enforced constraints. The storage provider is responsible for runtime errors resulting from invalid use of the resource.

For example, a GCE Persistent Disk has **AccessModes ReadWriteOnce** and **ReadOnlyMany**. The user must mark their claims as **read-only** if they want to take advantage of the volume's ability for ROX. Errors in the provider show up at runtime as mount errors.

#### 4.5.3.4. Recycling Policy

The current recycling policies are:

Recycling Policy	Description
Retain	Manual reclamation
Recycle	Basic scrub (e.g, <b>rm -rf /&lt;volume&gt;/*</b> )



### NOTE

Currently, only NFS and HostPath support the 'Recycle' recycling policy.

#### 4.5.3.5. Phase

A volumes can be found in one of the following phases:

Phase	Description
Available	A free resource that is not yet bound to a claim.
Bound	The volume is bound to a claim.
Released	The claim has been deleted, but the resource is not yet reclaimed by the cluster.
Failed	The volume has failed its automatic reclamation.

The CLI shows the name of the PVC bound to the PV.

#### 4.5.4. Persistent Volume Claims

Each PVC contains a **spec** and **status**, which is the specification and status of the claim.

##### Example 4.2. Persistent Volume Claim Object Definition

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: myclaim
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 8Gi
```

##### 4.5.4.1. Access Modes

Claims use the same conventions as volumes when requesting storage with specific access modes.

##### 4.5.4.2. Resources

Claims, like pods, can request specific quantities of a resource. In this case, the request is for storage. The same [resource model](#) applies to both volumes and claims.

##### 4.5.4.3. Claims As Volumes

Pods access storage by using the claim as a volume. Claims must exist in the same namespace as the pod using the claim. The cluster finds the claim in the pod's namespace and uses it to get the **PersistentVolume** backing the claim. The volume is then mounted to the host and into the pod:

```
kind: Pod
apiVersion: v1
metadata:
  name: mypod
spec:
```

```
containers:
  - name: myfrontend
    image: dockerfile/nginx
    volumeMounts:
      - mountPath: "/var/www/html"
        name: mypd
volumes:
  - name: mypd
    persistentVolumeClaim:
      claimName: myclaim
```

## 4.6. REMOTE COMMANDS

### 4.6.1. Overview

OpenShift Enterprise takes advantage of a feature built into Kubernetes to support executing commands in containers. This is implemented using HTTP along with a multiplexed streaming protocol such as [SPDY](#) or [HTTP/2](#).

Developers can [use the CLI](#) to execute remote commands in containers.

### 4.6.2. Server Operation

The Kubelet handles remote execution requests from clients. Upon receiving a request, it upgrades the response, evaluates the request headers to determine what streams (**stdin**, **stdout**, and/or **stderr**) to expect to receive, and waits for the client to create the streams.

After the Kubelet has received all the streams, it executes the command in the container, copying between the streams and the command's **stdin**, **stdout**, and **stderr**, as appropriate. When the command terminates, the Kubelet closes the upgraded connection, as well as the underlying one.

Architecturally, there are options for running a command in a container. The supported implementation currently in OpenShift Enterprise invokes **nsenter** directly on the node host to enter the container's namespaces prior to executing the command. However, custom implementations could include using **docker exec**, or running a "helper" container that then runs **nsenter** so that **nsenter** is not a required binary that must be installed on the host.

## 4.7. PORT FORWARDING

### 4.7.1. Overview

OpenShift Enterprise takes advantage of a feature built into Kubernetes to support port forwarding to pods. This is implemented using HTTP along with a multiplexed streaming protocol such as [SPDY](#) or [HTTP/2](#).

Developers can [use the CLI](#) to port forward to a pod. The CLI listens on each local port specified by the user, forwarding via the [described protocol](#).

### 4.7.2. Server Operation

The Kubelet handles port forward requests from clients. Upon receiving a request, it upgrades the response and waits for the client to create port forwarding streams. When it receives a new stream, it copies data between the stream and the pod's port.

Architecturally, there are options for forwarding to a pod's port. The supported implementation currently in OpenShift Enterprise invokes **nsenter** directly on the node host to enter the pod's network namespace, then invokes **socat** to copy data between the stream and the pod's port. However, a custom implementation could include running a "helper" pod that then runs **nsenter** and **socat**, so that those binaries are not required to be installed on the host.

## 4.8. SOURCE CONTROL MANAGEMENT

OpenShift Enterprise takes advantage of preexisting source control management (SCM) systems hosted either internally (such as an in-house Git server) or externally (for example, on [GitHub](#), [Bitbucket](#), etc.). Currently, OpenShift Enterprise only supports [Git](#) solutions.

SCM integration is tightly coupled with [builds](#), the two points being:

- Creating a **BuildConfig** using a repository, which allows building your application inside of OpenShift Enterprise. You can create a **BuildConfig** [manually](#) or let OpenShift create it [automatically](#) by inspecting your repository.
- [Triggering a build](#) upon repository changes.

## 4.9. ADMISSION CONTROLLERS

Admission control plug-ins intercept requests to the master API prior to persistence of a resource, but after the request is authenticated and authorized.

Each admission control plug-in is run in sequence before a request is accepted into the cluster. If any plug-in in the sequence rejects the request, the entire request is rejected immediately, and an error is returned to the end-user.

Admission control plug-ins may modify the incoming object in some cases to apply system configured defaults. In addition, admission control plug-ins may modify related resources as part of request processing to do things such as incrementing quota usage.



### WARNING

The OpenShift Enterprise master has a default list of plug-ins that are enabled by default for each type of resource (Kubernetes and OpenShift Enterprise). These are required for the proper functioning of the master. Modifying these lists is not recommended unless you strictly know what you are doing. Future versions of the product may use a different set of plug-ins and may change their ordering. If you do override the default list of plug-ins in the master configuration file, you are responsible for updating it to reflect requirements of newer versions of the OpenShift Enterprise master.

Cluster administrators can configure some admission control plug-ins to control certain behavior, such as:

- [Limiting Number of Self-Provisioned Projects Per User](#)
- [Configuring Global Build Defaults and Overrides](#)

- [Controlling Pod Placement](#)

## 4.10. OTHER API OBJECTS

### 4.10.1. LimitRange

A limit range provides a mechanism to enforce min/max limits placed on resources in a Kubernetes [namespace](#).

By adding a limit range to your namespace, you can enforce the minimum and maximum amount of CPU and Memory consumed by an individual pod or container.

See the [Kubernetes documentation](#) for more information.

### 4.10.2. ResourceQuota

Kubernetes can limit both the number of objects created in a [namespace](#), and the total amount of resources requested across objects in a namespace. This facilitates sharing of a single Kubernetes cluster by several teams, each in a namespace, as a mechanism of preventing one team from starving another team of cluster resources.

See [Cluster Administration](#) and [Kubernetes documentation](#) for more information on **ResourceQuota**.

### 4.10.3. Resource

A Kubernetes **Resource** is something that can be requested by, allocated to, or consumed by a pod or container. Examples include memory (RAM), CPU, disk-time, and network bandwidth.

See the [Developer Guide](#) and [Kubernetes documentation](#) for more information.

### 4.10.4. Secret

[Secrets](#) are storage for sensitive information, such as keys, passwords, and certificates. They are accessible by the intended pod(s), but held separately from their definitions.

### 4.10.5. PersistentVolume

A [persistent volume](#) is an object (**PersistentVolume**) in the infrastructure provisioned by the cluster administrator. Persistent volumes provide durable storage for stateful applications.

See the [Kubernetes documentation](#) for more information.

### 4.10.6. PersistentVolumeClaim

A **PersistentVolumeClaim** object is a [request for storage by a pod author](#). Kubernetes matches the claim against the pool of available volumes and binds them together. The claim is then used as a volume by a pod. Kubernetes makes sure the volume is available on the same node as the pod that requires it.

See the [Kubernetes documentation](#) for more information.

### 4.10.7. OAuth Objects

#### 4.10.7.1. OAuthClient

An **OAuthClient** represents an OAuth client, as described in [RFC 6749, section 2](#).

The following **OAuthClient** objects are automatically created:

<b>openshift -web- console</b>	Client used to request tokens for the web console
<b>openshift -browser- client</b>	Client used to request tokens at /oauth/token/request with a user-agent that can handle interactive logins
<b>openshift - challengi ng-client</b>	Client used to request tokens with a user-agent that can handle WWW-Authenticate challenges

#### Example 4.3. OAuthClient Object Definition

```
kind: "OAuthClient"
apiVersion: "v1"
metadata:
  name: "openshift-web-console" ❶
  selflink: "/osapi/v1/oauthClients/openshift-web-console"
  resourceVersion: "1"
  creationTimestamp: "2015-01-01T01:01:01Z"
respondWithChallenges: false ❷
secret: "45e27750-a8aa-11e4-b2ea-3c970e4b7ffe" ❸
redirectURIs:
  - "https://localhost:8443" ❹
```

- ❶ The **name** is used as the **client\_id** parameter in OAuth requests.
- ❷ When **respondWithChallenges** is set to **true**, unauthenticated requests to **/oauth/authorize** will result in **WWW-Authenticate** challenges, if supported by the configured authentication methods.
- ❸ The value in the **secret** parameter is used as the **client\_secret** parameter in an authorization code flow.
- ❹ One or more absolute URLs can be placed in the **redirectURIs** section. The **redirect\_uri** parameter sent with authorization requests must be prefixed by one of the specified **redirectURIs**.

#### 4.10.7.2. OAuthClientAuthorization

An **OAuthClientAuthorization** represents an approval by a **User** for a particular **OAuthClient** to be given an **OAuthAccessToken** with particular scopes.

Creation of **OAuthClientAuthorization** objects is done during an authorization request to the **OAuth** server.

#### Example 4.4. OAuthClientAuthorization Object Definition

```
---
kind: "OAuthClientAuthorization"
apiVersion: "v1"
metadata:
  name: "bob:openshift-web-console"
  resourceVersion: "1"
  creationTimestamp: "2015-01-01T01:01:01-00:00"
clientName: "openshift-web-console"
userName: "bob"
userID: "9311ac33-0fde-11e5-97a1-3c970e4b7ffe"
scopes: []
----
```

#### 4.10.7.3. OAuthAuthorizeToken

An **OAuthAuthorizeToken** represents an **OAuth** authorization code, as described in [RFC 6749, section 1.3.1](#).

An **OAuthAuthorizeToken** is created by a request to the `/oauth/authorize` endpoint, as described in [RFC 6749, section 4.1.1](#).

An **OAuthAuthorizeToken** can then be used to obtain an **OAuthAccessToken** with a request to the `/oauth/token` endpoint, as described in [RFC 6749, section 4.1.3](#).

#### Example 4.5. OAuthAuthorizeToken Object Definition

```
kind: "OAuthAuthorizeToken"
apiVersion: "v1"
metadata:
  name: "MDAwYjM5YjMtMzM1MC00NDY4LTkxODItOTA2OTE2YzE0M2Fj" 1
  resourceVersion: "1"
  creationTimestamp: "2015-01-01T01:01:01-00:00"
clientName: "openshift-web-console" 2
expiresIn: 300 3
scopes: []
redirectURI: "https://localhost:8443/console/oauth" 4
userName: "bob" 5
userID: "9311ac33-0fde-11e5-97a1-3c970e4b7ffe" 6
```

1 **name** represents the token name, used as an authorization code to exchange for an **OAuthAccessToken**.

2 The **clientName** value is the **OAuthClient** that requested this token.

3 The **expiresIn** value is the expiration in seconds from the **creationTimestamp**.

4



The **redirectURI** value is the location where the user was redirected to during the authorization flow that resulted in this token.

- 5 **userName** represents the name of the User this token allows obtaining an OAuthAccessToken for.
- 6 **userID** represents the UID of the User this token allows obtaining an OAuthAccessToken for.

#### 4.10.7.4. OAuthAccessToken

An **OAuthAccessToken** represents an **OAuth** access token, as described in [RFC 6749, section 1.4](#).

An **OAuthAccessToken** is created by a request to the **/oauth/token** endpoint, as described in [RFC 6749, section 4.1.3](#).

Access tokens are used as bearer tokens to authenticate to the API.

##### Example 4.6. OAuthAccessToken Object Definition

```
kind: "OAuthAccessToken"
apiVersion: "v1"
metadata:
  name: "ODli0GE5ZmMtYzczYi00Nzk1LTg4MGEtNzQyZmUxZmUwY2Vh" 1
  resourceVersion: "1"
  creationTimestamp: "2015-01-01T01:01:02-00:00"
clientName: "openshift-web-console" 2
expiresIn: 86400 3
scopes: []
redirectURI: "https://localhost:8443/console/oauth" 4
userName: "bob" 5
userID: "9311ac33-0fde-11e5-97a1-3c970e4b7ffe" 6
authorizeToken: "MDAwYjM5YjMtMzM1MC00NDY4LTkxODItOTA2OTEyZjE0M2Fj" 7
```

- 1 **name** is the token name, which is used as a bearer token to authenticate to the API.
- 2 The **clientName** value is the OAuthClient that requested this token.
- 3 The **expiresIn** value is the expiration in seconds from the creationTimestamp.
- 4 The **redirectURI** is where the user was redirected to during the authorization flow that resulted in this token.
- 5 **userName** represents the User this token allows authentication as.
- 6 **userID** represents the User this token allows authentication as.
- 7 **authorizeToken** is the name of the OAuthAuthorizationToken used to obtain this token, if any.

## 4.10.8. User Objects

### 4.10.8.1. Identity

When a user logs into OpenShift Enterprise, they do so using a configured [identity provider](#). This determines the user's identity, and provides that information to OpenShift Enterprise.

OpenShift Enterprise then looks for a **UserIdentityMapping** for that **Identity**:

- If the **Identity** already exists, but is not mapped to a **User**, login fails.
- If the **Identity** already exists, and is mapped to a **User**, the user is given an **OAuthAccessToken** for the mapped **User**.
- If the **Identity** does not exist, an **Identity**, **User**, and **UserIdentityMapping** are created, and the user is given an **OAuthAccessToken** for the mapped **User**.

#### Example 4.7. Identity Object Definition

```
kind: "Identity"
apiVersion: "v1"
metadata:
  name: "anypassword:bob" ①
  uid: "9316ebad-0fde-11e5-97a1-3c970e4b7ffe"
  resourceVersion: "1"
  creationTimestamp: "2015-01-01T01:01:01-00:00"
providerName: "anypassword" ②
providerUserName: "bob" ③
user:
  name: "bob" ④
  uid: "9311ac33-0fde-11e5-97a1-3c970e4b7ffe" ⑤
```

- ① The identity name must be in the form `providerName:providerUserName`.
- ② **providerName** is the name of the identity provider.
- ③ **providerUserName** is the name that uniquely represents this identity in the scope of the identity provider.
- ④ The **name** in the **user** parameter is the name of the user this identity maps to.
- ⑤ The **uid** represents the UID of the user this identity maps to.

### 4.10.8.2. User

A **User** represents an actor in the system. Users are granted permissions by [adding roles to users or to their groups](#).

User objects are created automatically on first login, or can be created via the API.

#### Example 4.8. User Object Definition

-

```

kind: "User"
apiVersion: "v1"
metadata:
  name: "bob" ❶
  uid: "9311ac33-0fde-11e5-97a1-3c970e4b7ffe"
  resourceVersion: "1"
  creationTimestamp: "2015-01-01T01:01:01-00:00"
identities:
  - "anypassword:bob" ❷
fullName: "Bob User" ❸

```

<1> `name` is the user name used when adding roles to a user.  
 <2> The values in `identities` are Identity objects that map to this user. May be `null` or empty for users that cannot log in.  
 <3> The `fullName` value is an optional display name of user.

### 4.10.8.3. UserIdentityMapping

A **UserIdentityMapping** maps an **Identity** to a **User**.

Creating, updating, or deleting a **UserIdentityMapping** modifies the corresponding fields in the **Identity** and **User** objects.

An **Identity** can only map to a single **User**, so logging in as a particular identity unambiguously determines the **User**.

A **User** can have multiple identities mapped to it. This allows multiple login methods to identify the same **User**.

#### Example 4.9. UserIdentityMapping Object Definition

```

kind: "UserIdentityMapping"
apiVersion: "v1"
metadata:
  name: "anypassword:bob" ❶
  uid: "9316ebad-0fde-11e5-97a1-3c970e4b7ffe"
  resourceVersion: "1"
identity:
  name: "anypassword:bob"
  uid: "9316ebad-0fde-11e5-97a1-3c970e4b7ffe"
user:
  name: "bob"
  uid: "9311ac33-0fde-11e5-97a1-3c970e4b7ffe"

```

<1> `\*UserIdentityMapping\*` name matches the mapped `\*Identity\*` name

### 4.10.8.4. Group

A **Group** represents a list of users in the system. Groups are granted permissions by [adding roles to users or to their groups](#).

**Example 4.10. Group Object Definition**

```
kind: "Group"
apiVersion: "v1"
metadata:
  name: "developers" 1
  creationTimestamp: "2015-01-01T01:01:01-00:00"
users:
  - "bob" 2
```

**1 1 1** **name** is the group name used when adding roles to a group.

**2 2** The values in **users** are the names of User objects that are members of this group.

---

[1] After this point, device names refer to devices on container B's host.

## CHAPTER 5. REVISION HISTORY: ARCHITECTURE

### 5.1. TUE MAY 02 2017

Affected Topic	Description of Change
<a href="#">Additional Concepts</a> → <a href="#">F5 BIG-IP® Router Plug-in</a>	Removed information about integration with HAProxy router.

### 5.2. MON JAN 09 2017

Affected Topic	Description of Change
<a href="#">Additional Concepts</a> → <a href="#">Authorization</a>	Added clarifying details about cluster roles.

### 5.3. MON NOV 21 2016

Affected Topic	Description of Change
<a href="#">Additional Concepts</a> → <a href="#">Authorization</a>	Updated to reflect that seven, not six, <a href="#">Security Context Constraints</a> are added to the cluster by default.
<a href="#">Additional Concepts</a> → <a href="#">Persistent Storage</a>	Removed statement indicating that HostPath is for testing only.

### 5.4. TUE OCT 04 2016

Affected Topic	Description of Change
<a href="#">Core Concepts</a> → <a href="#">Routes</a>	Added <a href="#">Router Configuration Parameters</a> .

### 5.5. TUE SEP 06 2016

Affected Topic	Description of Change
<a href="#">Core Concepts</a> → <a href="#">Routes</a>	Fixed table markup in Table 3.1 in <a href="#">HAProxy Template Router</a> section.
	Added a Note box on the implication of older routes with a sharded setup.

## 5.6. MON AUG 08 2016

Affected Topic	Description of Change
<a href="#">Core Concepts → Routes</a>	Improved explanation of <i>traditional</i> vs <i>overlapped</i> <a href="#">router sharding</a> , including an example of namespace-based sharding.

## 5.7. WED JUL 27 2016

Affected Topic	Description of Change
<a href="#">Additional Concepts → Admission Controllers</a>	Added a Warning box indicating that OpenShift Enterprise master has a default list of plug-ins that are enabled by default for each type of resource.

## 5.8. THU JUL 14 2016

Affected Topic	Description of Change
<a href="#">Core Concepts → Pods and Services</a>	Minor update for clarity in the <a href="#">Pods</a> section.
	Fixed wording in the <a href="#">Pods</a> section.
	Added information about the <a href="#">service proxy mode</a> and how to configure it.

## 5.9. WED JUN 15 2016

Affected Topic	Description of Change
<a href="#">Additional Concepts → Authentication</a>	Removed the Impersonation section, which does not apply to OpenShift Enterprise 3.2.

## 5.10. FRI JUN 10 2016

Affected Topic	Description of Change
<a href="#">Infrastructure Components → Web Console</a>	Added a Note introducing Cockpit to the <a href="#">Project Overviews</a> section.
<a href="#">Core Concepts → Builds and Image Streams</a>	Added <b>Reproducibility</b> to the list of S2I advantages in the <a href="#">Source-to-Image (S2I) Build</a> section.
<a href="#">Additional Concepts → Authorization</a>	Added a Note box indicating that <a href="#">Roles</a> need to be assigned to administer the setup with an external user.

Affected Topic	Description of Change
----------------	-----------------------

## 5.11. MON MAY 30 2016

Affected Topic	Description of Change
<a href="#">Additional Concepts → Authentication</a>	Added the Impersonation section.
<a href="#">Core Concepts → Routes</a>	Added the <a href="#">Router Sharding</a> section.

## 5.12. THU MAY 12 2016

OpenShift Enterprise 3.2 initial release.

Affected Topic	Description of Change
<a href="#">Core Concepts → Builds and Image Streams</a>	Moved many task-based sections to the new <a href="#">Managing Images</a> topic in the Developer Guide.
<a href="#">Infrastructure Components → Web Console</a>	Updated the Web Console <a href="#">Project Overview</a> screenshot.
	Added a <a href="#">CLI Downloads</a> section, including a screenshot of the new <b>About</b> page.
<a href="#">Infrastructure Components → Kubernetes Infrastructure</a>	Added a note that OSE 3.2 uses Kubernetes 1.2 and Docker 1.9.
	Removed support for the Pacemaker HA method.
<a href="#">Additional Concepts → Authentication</a>	Added the <b>system:authenticated:oauth</b> virtual group to the <a href="#">Users and Groups</a> section.
<a href="#">Additional Concepts → Admission Controllers</a>	New topic discussing admission controllers and the configurable plug-ins.