



OpenShift Container Platform 3.5

Developer Guide

OpenShift Container Platform 3.5 Developer Reference

OpenShift Container Platform 3.5 Developer Guide

OpenShift Container Platform 3.5 Developer Reference

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

These topics help developers set up and configure a workstation to develop and deploy applications in an OpenShift Container Platform cloud environment with a command-line interface (CLI). This guide provides detailed instructions and examples to help developers: Monitor and browse projects with the web console Configure and utilize the CLI Generate configurations using templates Manage builds and webhooks Define and trigger deployments Integrate external services (databases, SaaS endpoints)

Table of Contents

CHAPTER 1. OVERVIEW	13
CHAPTER 2. APPLICATION LIFE CYCLE MANAGEMENT	14
2.1. PLANNING YOUR DEVELOPMENT PROCESS	14
2.1.1. Overview	14
2.1.2. Using OpenShift Container Platform as Your Development Environment	14
2.1.3. Bringing an Application to Deploy on OpenShift Container Platform	15
2.2. CREATING NEW APPLICATIONS	16
2.2.1. Overview	16
2.2.2. Creating an Application Using the CLI	16
2.2.2.1. Creating an Application From Source Code	16
2.2.2.2. Creating an Application From an Image	18
2.2.2.3. Creating an Application From a Template	18
2.2.2.4. Further Modifying Application Creation	19
2.2.2.4.1. Specifying Environment Variables	20
2.2.2.4.2. Specifying Build Environment Variables	20
2.2.2.4.3. Specifying Labels	21
2.2.2.4.4. Viewing the Output Without Creation	21
2.2.2.4.5. Creating Objects With Different Names	21
2.2.2.4.6. Creating Objects in a Different Project	21
2.2.2.4.7. Creating Multiple Objects	22
2.2.2.4.8. Grouping Images and Source in a Single Pod	22
2.2.2.4.9. Searching for Images, Templates, and Other Inputs	22
2.2.3. Creating an Application Using the Web Console	22
2.3. PROMOTING APPLICATIONS ACROSS ENVIRONMENTS	26
2.3.1. Overview	26
2.3.2. Application Components	26
2.3.2.1. API Objects	26
2.3.2.2. Images	28
2.3.2.3. Summary	28
2.3.3. Deployment Environments	29
2.3.3.1. Considerations	29
2.3.3.2. Summary	29
2.3.4. Methods and Tools	29
2.3.4.1. Managing API Objects	30
2.3.4.1.1. Exporting API Object State	30
2.3.4.1.2. Importing API Object State	31
2.3.4.1.2.1. Initial Creation	31
2.3.4.1.2.2. Iterative Modification	31
2.3.4.2. Managing Images and Image Streams	32
2.3.4.2.1. Moving Images	32
2.3.4.2.1.1. When Staging Environments Share a Registry	32
2.3.4.2.1.2. When Staging Environments Use Different Registries	32
2.3.4.2.2. Deploying	32
2.3.4.2.3. Automating Promotion Flows with Jenkins	33
2.3.4.2.4. Promotion Caveats	33
2.3.4.2.4.1. API Object References	33
2.3.4.2.4.2. Image Registry References	34
2.3.4.3. Summary	34
2.3.5. Scenarios and Examples	34
2.3.5.1. Setting up for Promotion	35

2.3.5.2. Repeatable Promotion Process	36
2.3.5.3. Repeatable Promotion Process Using Jenkins	38
CHAPTER 3. AUTHENTICATION	40
3.1. WEB CONSOLE AUTHENTICATION	40
3.2. CLI AUTHENTICATION	40
CHAPTER 4. AUTHORIZATION	42
4.1. OVERVIEW	42
4.2. CHECKING IF USERS CAN CREATE PODS	42
4.3. DETERMINING WHAT YOU CAN DO AS AN AUTHENTICATED USER	42
CHAPTER 5. PROJECTS	44
5.1. OVERVIEW	44
5.2. CREATING A PROJECT	44
5.3. VIEWING PROJECTS	44
5.4. CHECKING PROJECT STATUS	45
5.5. FILTERING BY LABELS	45
5.6. BOOKMARKING PAGE STATES	46
5.7. DELETING A PROJECT	46
CHAPTER 6. MIGRATING APPLICATIONS	48
6.1. OVERVIEW	48
6.2. MIGRATING DATABASE APPLICATIONS	49
6.2.1. Overview	49
6.2.2. Supported Databases	49
6.2.3. MySQL	49
6.2.4. PostgreSQL	51
6.2.5. MongoDB	53
6.3. MIGRATING WEB FRAMEWORK APPLICATIONS	55
6.3.1. Overview	55
6.3.2. Python	55
6.3.3. Ruby	56
6.3.4. PHP	57
6.3.5. Perl	57
6.3.6. Node.js	58
6.3.7. JBoss EAP	59
6.3.8. JBoss WS (Tomcat)	60
6.3.9. JBoss AS (Wildfly 10)	60
6.3.10. Supported JBoss/XPaaS Versions	61
6.4. QUICKSTART EXAMPLES	62
6.4.1. Overview	62
6.4.2. Workflow	62
6.5. CONTINUOUS INTEGRATION AND DEPLOYMENT (CI/CD)	63
6.5.1. Overview	63
6.5.2. Jenkins	63
6.6. WEBHOOKS AND ACTION HOOKS	64
6.6.1. Overview	64
6.6.2. Webhooks	64
6.6.3. Action Hooks	64
6.7. S2I TOOL	65
6.7.1. Overview	65
6.7.2. Creating a Container Image	65
6.8. SUPPORT GUIDE	66

6.8.1. Overview	66
6.8.2. Supported Databases	66
6.8.3. Supported Languages	66
6.8.4. Supported Frameworks	66
6.8.5. Supported Markers	67
6.8.6. Supported Environment Variables	68
CHAPTER 7. APPLICATION TUTORIALS	70
7.1. OVERVIEW	70
7.2. QUICKSTART TEMPLATES	70
7.2.1. Overview	70
7.2.2. Web Framework Quickstart Templates	70
7.3. RUBY ON RAILS	71
7.3.1. Overview	71
7.3.2. Local Workstation Setup	71
7.3.2.1. Setting Up the Database	71
7.3.3. Writing Your Application	72
7.3.3.1. Creating a Welcome Page	73
7.3.3.2. Configuring the Application for OpenShift Container Platform	73
7.3.3.3. Storing Your Application in Git	74
7.3.4. Deploying Your Application to OpenShift Container Platform	75
7.3.4.1. Creating the Database Service	75
7.3.4.2. Creating the Frontend Service	76
7.3.4.3. Creating a Route for Your Application	77
7.4. SETTING UP A NEXUS MIRROR FOR MAVEN	77
7.4.1. Introduction	77
7.4.2. Setting up Nexus	78
7.4.2.1. Using Probes to Check for Success	78
7.4.2.2. Adding Persistence to Nexus	79
7.4.3. Connecting to Nexus	79
7.4.4. Confirming Success	80
7.4.5. Additional Resources	80
CHAPTER 8. BUILDS	81
8.1. HOW BUILDS WORK	81
8.1.1. What Is a Build?	81
8.1.2. What Is a BuildConfig?	81
8.2. BASIC BUILD OPERATIONS	83
8.2.1. Starting a Build	83
8.2.2. Canceling a Build	84
8.2.3. Deleting a BuildConfig	84
8.2.4. Viewing Build Details	84
8.2.5. Accessing Build Logs	85
8.3. BUILD INPUTS	86
8.3.1. How Build Inputs Work	86
8.3.2. Dockerfile Source	87
8.3.3. Image Source	87
8.3.4. Git Source	88
8.3.4.1. Using a Proxy	89
8.3.4.2. Source Clone Secrets	89
8.3.4.2.1. Automatically Adding a Source Clone Secret to a Build Configuration	90
8.3.4.2.2. Manually Adding Source Clone Secrets	91
8.3.4.2.3. .Gitconfig File	92

8.3.4.2.4. .gitconfig File for Secured Git	92
8.3.4.2.5. Basic Authentication	93
8.3.4.2.6. SSH Key Authentication	93
8.3.4.2.7. Trusted Certificate Authorities	94
8.3.4.2.8. Combinations	95
8.3.5. Binary (Local) Source	95
8.3.6. Input Secrets	96
8.3.6.1. Adding Input Secrets	97
8.3.6.2. Source-to-Image Strategy	98
8.3.6.3. Docker Strategy	98
8.3.6.4. Custom Strategy	99
8.3.7. Using External Artifacts	99
8.3.8. Using Docker Credentials for Private Registries	100
8.4. BUILD OUTPUT	101
8.4.1. Build Output Overview	101
8.4.2. Output Image Environment Variables	102
8.4.3. Output Image Labels	102
8.4.4. Output Image Digest	103
8.4.5. Using Docker Credentials for Private Registries	103
8.5. BUILD STRATEGY OPTIONS	103
8.5.1. Source-to-Image Strategy Options	103
8.5.1.1. Force Pull	104
8.5.1.2. Incremental Builds	104
8.5.1.3. Extended Builds	104
8.5.1.3.1. Testing your Application	106
8.5.1.4. Overriding Builder Image Scripts	106
8.5.1.5. Environment Variables	106
8.5.1.5.1. Environment Files	106
8.5.1.5.2. BuildConfig Environment	107
8.5.1.6. Adding Secrets via Web Console	107
8.5.1.6.1. Enabling Pulling and Pushing	107
8.5.1.7. Ignoring Source Files	107
8.5.2. Docker Strategy Options	107
8.5.2.1. FROM Image	108
8.5.2.2. Dockerfile Path	108
8.5.2.3. No Cache	108
8.5.2.4. Force Pull	108
8.5.2.5. Environment Variables	108
8.5.2.6. Adding Secrets via Web Console	109
8.5.2.6.1. Enabling Pulling and Pushing	109
8.5.3. Custom Strategy Options	109
8.5.3.1. FROM Image	109
8.5.3.2. Exposing the Docker Socket	110
8.5.3.3. Secrets	110
8.5.3.3.1. Adding Secrets via Web Console	110
8.5.3.3.2. Enabling Pulling and Pushing	110
8.5.3.4. Force Pull	111
8.5.3.5. Environment Variables	111
8.5.4. Pipeline Strategy Options	111
8.5.4.1. Providing the Jenkinsfile	111
8.6. TRIGGERING BUILDS	112
8.6.1. Build Triggers Overview	112
8.6.2. Webhook Triggers	112

8.6.2.1. GitHub Webhooks	112
8.6.2.2. Generic Webhooks	113
8.6.2.3. Displaying Webhook URLs	115
8.6.3. Image Change Triggers	115
8.6.4. Configuration Change Triggers	116
8.7. BUILD HOOKS	117
8.7.1. Build Hooks Overview	117
8.7.2. Configuring Post Commit Build Hooks	117
8.7.2.1. Using the CLI	118
8.8. BUILD RUN POLICY	119
8.8.1. Build Run Policy Overview	119
8.8.2. Serial Run Policy	119
8.8.3. SerialLatestOnly Run Policy	119
8.8.4. Parallel Run Policy	120
8.9. ADVANCED BUILD OPERATIONS	120
8.9.1. Setting Build Resources	120
8.9.2. Setting Maximum Duration	121
8.9.3. Assigning Builds to Specific Nodes	122
8.9.4. Chaining Builds	122
8.10. BUILD TROUBLESHOOTING	124
8.10.1. Requested Access to Resources Denied	124
CHAPTER 9. DEPLOYMENTS	125
9.1. HOW DEPLOYMENTS WORK	125
9.1.1. What Is a Deployment?	125
9.1.2. Creating a Deployment Configuration	125
9.2. BASIC DEPLOYMENT OPERATIONS	127
9.2.1. Starting a Deployment	127
9.2.2. Viewing a Deployment	127
9.2.3. Rolling Back a Deployment	127
9.2.4. Executing Commands Inside a Container	128
9.2.5. Viewing Deployment Logs	128
9.2.6. Setting Deployment Triggers	129
9.2.6.1. Configuration Change Trigger	129
9.2.6.2. ImageChange Trigger	129
9.2.6.2.1. Using the Command Line	130
9.2.7. Setting Deployment Resources	130
9.2.8. Manual Scaling	131
9.2.9. Assigning Pods to Specific Nodes	132
9.2.10. Running a Pod with a Different Service Account	132
9.2.11. Adding Secrets to Deployment Configurations from the Web Console	133
9.3. DEPLOYMENT STRATEGIES	133
9.3.1. What Are Deployment Strategies?	133
9.3.2. Rolling Strategy	133
9.3.2.1. Canary Deployments	133
9.3.2.2. When to Use a Rolling Deployment	133
9.3.2.3. Rolling Example	135
9.3.3. Recreate Strategy	136
9.3.3.1. When to Use a Recreate Deployment	136
9.3.4. Custom Strategy	137
9.3.5. Lifecycle Hooks	138
9.3.5.1. Pod-based Lifecycle Hook	138
9.3.5.2. Using the Command Line	139

9.4. ADVANCED DEPLOYMENT STRATEGIES	140
9.4.1. Blue-Green Deployment	140
9.4.1.1. When to Use a Blue-Green Deployment	140
9.4.1.2. Blue-Green Deployment Example	140
9.4.1.2.1. Using a Route and Two Services	140
9.4.2. A/B Deployment	141
9.4.2.1. When to Use an A/B Deployment	141
9.4.2.2. A/B Deployment Example	141
9.4.2.2.1. One Service, Multiple Deployment Configurations	141
9.4.3. Proxy Shard / Traffic Splitter	143
9.4.4. N-1 Compatibility	143
9.4.5. Graceful Termination	143
9.5. KUBERNETES DEPLOYMENTS SUPPORT	144
9.5.1. New Object Type: Deployments	144
9.5.2. Kubernetes Deployments vs Deployment Configurations	145
9.5.2.1. Deployment Configuration-Specific Features	145
9.5.2.1.1. Automatic Rollbacks	145
9.5.2.1.2. Triggers	145
9.5.2.1.3. Lifecycle Hooks	145
9.5.2.1.4. Custom Strategies	145
9.5.2.1.5. Canary Deployments	146
9.5.2.1.6. Test Deployments	146
9.5.2.2. Kubernetes Deployment-Specific Features	146
9.5.2.2.1. Rollover	146
9.5.2.2.2. Proportional Scaling	146
9.5.2.2.3. Pausing Mid-rollout	146
CHAPTER 10. TEMPLATES	147
10.1. OVERVIEW	147
10.2. UPLOADING A TEMPLATE	147
10.3. CREATING FROM TEMPLATES USING THE WEB CONSOLE	147
10.4. CREATING FROM TEMPLATES USING THE CLI	149
10.4.1. Labels	150
10.4.2. Parameters	150
10.4.3. Generating a List of Objects	151
10.5. MODIFYING AN UPLOADED TEMPLATE	152
10.6. USING THE INSTANT APP AND QUICKSTART TEMPLATES	152
10.7. WRITING TEMPLATES	153
10.7.1. Description	153
10.7.2. Labels	154
10.7.3. Parameters	155
10.7.4. Object List	157
10.7.5. Other Recommendations	158
10.7.6. Creating a Template from Existing Objects	158
CHAPTER 11. OPENING A REMOTE SHELL TO CONTAINERS	160
11.1. OVERVIEW	160
11.2. START A SECURE SHELL SESSION	160
11.3. SECURE SHELL SESSION HELP	160
CHAPTER 12. SERVICE ACCOUNTS	161
12.1. OVERVIEW	161
12.2. USER NAMES AND GROUPS	161
12.3. DEFAULT SERVICE ACCOUNTS AND ROLES	162

12.4. MANAGING SERVICE ACCOUNTS	162
12.5. ENABLING SERVICE ACCOUNT AUTHENTICATION	163
12.6. MANAGED SERVICE ACCOUNTS	163
12.7. INFRASTRUCTURE SERVICE ACCOUNTS	164
12.8. SERVICE ACCOUNTS AND SECRETS	164
12.9. MANAGING ALLOWED SECRETS	165
12.10. USING A SERVICE ACCOUNT'S CREDENTIALS INSIDE A CONTAINER	166
12.11. USING A SERVICE ACCOUNT'S CREDENTIALS EXTERNALLY	166
CHAPTER 13. MANAGING IMAGES	168
13.1. OVERVIEW	168
13.2. TAGGING IMAGES	168
13.2.1. Adding Tags to Image Streams	168
13.2.2. Recommended Tagging Conventions	169
13.2.3. Removing Tags from Image Streams	170
13.2.4. Referencing Images in Image Streams	170
13.3. IMAGE PULL POLICY	173
13.4. ACCESSING THE INTERNAL REGISTRY	173
13.5. USING IMAGE PULL SECRETS	174
13.5.1. Allowing Pods to Reference Images Across Projects	174
13.5.2. Allowing Pods to Reference Images from Other Secured Registries	174
13.6. IMPORTING TAG AND IMAGE METADATA	175
13.6.1. Importing Images from Insecure Registries	177
13.6.1.1. ImageStream Tag Policies	178
13.6.1.1.1. Insecure Tag Import Policy	178
13.6.1.1.2. Reference Policy	178
13.6.2. Importing Images from Private Registries	179
13.6.3. Adding Trusted Certificates for External Registries	179
13.6.4. Importing Images Across Projects	179
13.6.5. Creating an Image Stream by Manually Pushing an Image	180
13.7. WRITING IMAGE STREAM DEFINITIONS	181
CHAPTER 14. IMAGE SIGNATURES	183
14.1. OVERVIEW	183
14.2. SIGNING IMAGES USING ATOMIC CLI	183
14.3. ACCESSING IMAGE SIGNATURES USING REGISTRY API	183
14.3.1. Writing Image Signatures	184
14.3.2. Reading Image Signatures	184
14.4. VERIFYING IMAGE SIGNATURES USING OPENSIFT CLI	185
CHAPTER 15. QUOTAS AND LIMIT RANGES	186
15.1. OVERVIEW	186
15.2. QUOTAS	186
15.2.1. Viewing Quotas	186
15.2.2. Resources Managed by Quota	189
15.2.3. Quota Scopes	191
15.2.4. Quota Enforcement	192
15.2.5. Requests vs Limits	192
15.3. LIMIT RANGES	192
15.3.1. Viewing Limit Ranges	193
15.3.2. Container Limits	194
15.3.3. Pod Limits	195
15.4. COMPUTE RESOURCES	196
15.4.1. CPU Requests	197

15.4.2. Viewing Compute Resources	197
15.4.3. CPU Limits	197
15.4.4. Memory Requests	198
15.4.5. Memory Limits	198
15.4.6. Quality of Service Tiers	198
15.4.7. Specifying Compute Resources via CLI	199
15.4.8. Opaque Integer Resources	199
15.5. PROJECT RESOURCE LIMITS	200
CHAPTER 16. GETTING TRAFFIC INTO A CLUSTER	201
16.1. GETTING TRAFFIC INTO A CLUSTER	201
16.2. USING A ROUTER TO GET TRAFFIC INTO THE CLUSTER	201
16.2.1. Overview	201
16.2.2. Administrator Prerequisites	202
16.2.2.1. Defining the Public IP Range	202
16.2.3. Create a Project and Service	203
16.2.4. Expose the Service to Create a Route	203
16.2.5. Configure the Router	204
16.2.6. Configure IP Failover using VIPs	205
16.3. USING A LOAD BALANCER TO GET TRAFFIC INTO THE CLUSTER	205
16.3.1. Overview	205
16.3.2. Administrator Prerequisites	206
16.3.2.1. Defining the Public IP Range	206
16.3.3. Create a Project and Service	207
16.3.4. Expose the Service to Create a Route	207
16.3.5. Create the Load Balancer Service	208
16.3.6. Configuring Networking	210
16.3.7. Configure IP Failover using VIPs	211
16.4. USING A SERVICE EXTERNAL IP TO GET TRAFFIC INTO THE CLUSTER	212
16.4.1. Overview	212
16.4.2. Administrator Prerequisites	212
16.4.2.1. Defining the Public IP Range	213
16.4.3. Create a Project and Service	213
16.4.4. Expose the Service to Create a Route	214
16.4.5. Assigning an IP Address to the Service	215
16.4.6. Configuring Networking	216
16.4.7. Configure IP Failover using VIPs	219
16.5. USING A NODEPORT TO GET TRAFFIC INTO THE CLUSTER	219
16.5.1. Overview	219
16.5.2. Administrator Prerequisites	220
16.5.3. Configuring the Service	220
CHAPTER 17. ROUTES	222
17.1. OVERVIEW	222
17.2. CREATING ROUTES	222
17.3. LOAD BALANCING FOR A/B TESTING	224
CHAPTER 18. INTEGRATING EXTERNAL SERVICES	226
18.1. OVERVIEW	226
18.2. DEFINING A SERVICE FOR AN EXTERNAL DATABASE	226
18.2.1. Step 1: Define a Service	226
18.2.1.1. Using an IP address	226
18.2.1.2. Using an External Domain Name	227
18.2.2. Step 2: Consume a Service	228

18.3. EXTERNAL SAAS PROVIDER	229
18.3.1. Using an IP address and Endpoints	230
18.3.2. Using an External Domain Name	232
CHAPTER 19. SECRETS	233
19.1. USING SECRETS	233
19.1.1. Properties of Secrets	234
19.1.2. Creating Secrets	234
19.1.3. Types of Secrets	235
19.1.4. Updating Secrets	235
19.2. SECRETS IN VOLUMES AND ENVIRONMENT VARIABLES	236
19.3. IMAGE PULL SECRETS	236
19.4. SOURCE CLONE SECRETS	236
19.5. SERVICE SERVING CERTIFICATE SECRETS	236
19.6. RESTRICTIONS	236
19.6.1. Secret Data Keys	237
19.7. EXAMPLES	237
19.8. TROUBLESHOOTING	238
CHAPTER 20. CONFIGMAPS	240
20.1. OVERVIEW	240
20.2. CREATING CONFIGMAPS	240
20.2.1. Creating from Directories	241
20.2.2. Creating from Files	242
20.2.3. Creating from Literal Values	243
20.3. USE CASES: CONSUMING CONFIGMAPS IN PODS	244
20.3.1. Consuming in Environment Variables	244
20.3.2. Setting Command-line Arguments	245
20.3.3. Consuming in Volumes	246
20.4. EXAMPLE: CONFIGURING REDIS	247
20.5. RESTRICTIONS	249
CHAPTER 21. USING DAEMONSETS	250
21.1. OVERVIEW	250
21.2. CREATING DAEMONSETS	250
CHAPTER 22. POD AUTOSCALING	253
22.1. OVERVIEW	253
22.2. REQUIREMENTS FOR USING HORIZONTAL POD AUTOSCALERS	253
22.3. SUPPORTED METRICS	253
22.4. AUTOSCALING	253
22.5. CREATING A HORIZONTAL POD AUTOSCALER	254
22.6. VIEWING A HORIZONTAL POD AUTOSCALER	254
CHAPTER 23. MANAGING VOLUMES	256
23.1. OVERVIEW	256
23.2. GENERAL CLI USAGE	256
23.3. ADDING VOLUMES	257
Examples	258
23.4. UPDATING VOLUMES	258
Examples	259
23.5. REMOVING VOLUMES	259
Examples	259
23.6. LISTING VOLUMES	260

Examples	260
23.7. SPECIFYING A SUB-PATH	260
CHAPTER 24. USING PERSISTENT VOLUMES	262
24.1. OVERVIEW	262
24.2. REQUESTING STORAGE	262
24.3. VOLUME AND CLAIM BINDING	262
24.4. CLAIMS AS VOLUMES IN PODS	263
24.5. VOLUME AND CLAIM PRE-BINDING	263
CHAPTER 25. EXECUTING REMOTE COMMANDS	266
25.1. OVERVIEW	266
25.2. BASIC USAGE	266
25.3. PROTOCOL	266
CHAPTER 26. COPYING FILES TO OR FROM A CONTAINER	268
26.1. OVERVIEW	268
26.2. BASIC USAGE	268
26.3. BACKING UP AND RESTORING DATABASES	268
26.4. REQUIREMENTS	269
26.5. SPECIFYING THE COPY SOURCE	269
26.6. SPECIFYING THE COPY DESTINATION	270
26.7. DELETING FILES AT THE DESTINATION	270
26.8. CONTINUOUS SYNCING ON FILE CHANGE	270
26.9. ADVANCED RSYNC FEATURES	270
CHAPTER 27. PORT FORWARDING	271
27.1. OVERVIEW	271
27.2. BASIC USAGE	271
27.3. PROTOCOL	271
CHAPTER 28. SHARED MEMORY	273
28.1. OVERVIEW	273
28.2. POSIX SHARED MEMORY	273
CHAPTER 29. APPLICATION HEALTH	275
29.1. OVERVIEW	275
29.2. CONTAINER HEALTH CHECKS USING PROBES	275
CHAPTER 30. EVENTS	277
30.1. OVERVIEW	277
30.2. VIEWING EVENTS WITH THE CLI	277
30.3. VIEWING EVENTS IN THE CONSOLE	277
30.4. COMPREHENSIVE LIST OF EVENTS	277
CHAPTER 31. DOWNWARD API	287
31.1. OVERVIEW	287
31.2. SELECTING FIELDS	287
31.3. CONSUMING THE CONTAINER VALUES USING THE DOWNWARD API	287
31.3.1. Using Environment Variables	287
31.3.2. Using the Volume Plug-in	288
31.4. CONSUMING CONTAINER RESOURCES USING THE DOWNWARD API	290
31.4.1. Using Environment Variables	290
31.4.2. Using the Volume Plug-in	291
CHAPTER 32. MANAGING ENVIRONMENT VARIABLES	293

32.1. SETTING AND UNSETTING ENVIRONMENT VARIABLES	293
32.2. LIST ENVIRONMENT VARIABLES	293
32.3. SET ENVIRONMENT VARIABLES	293
32.3.1. Automatically Added Environment Variables	294
32.4. UNSET ENVIRONMENT VARIABLES	294
CHAPTER 33. JOBS	295
33.1. OVERVIEW	295
33.2. CREATING A JOB	295
33.2.1. Known Limitations	296
33.3. SCALING A JOB	296
33.4. SETTING MAXIMUM DURATION	296
CHAPTER 34. CRON JOBS	298
34.1. OVERVIEW	298
34.2. CREATING A CRON JOB	298
34.3. CLEANING UP AFTER A CRON JOB	299
CHAPTER 35. CREATE FROM URL	301
35.1. OVERVIEW	301
35.2. USING AN IMAGE STREAM AND IMAGE TAG	301
35.2.1. Query String Parameters	301
35.2.1.1. Example	302
35.3. USING A TEMPLATE	302
35.3.1. Query String Parameters	302
35.3.1.1. Example	303
CHAPTER 36. REVISION HISTORY: DEVELOPER GUIDE	304
36.1. FRI FEB 23 2018	304
36.2. FRI FEB 16 2018	304
36.3. TUE FEB 06 2018	304
36.4. THU JAN 25 2018	304
36.5. MON JAN 08 2018	305
36.6. FRI DEC 22 2017	305
36.7. MON DEC 11 2017	305
36.8. TUE NOV 21 2017	306
36.9. FRI NOV 10 2017	306
36.10. FRI NOV 03 2017	306
36.11. MON OCT 16 2017	306
36.12. MON SEP 18 2017	306
36.13. WED SEP 06 2017	307
36.14. FRI AUG 25 2017	307
36.15. TUE JUL 18 2017	307
36.16. WED JUL 12 2017	307
36.17. WED JUL 05 2017	308
36.18. TUE JUN 27 2017	308
36.19. MON JUN 19 2017	308
36.20. TUE JUN 13 2017	308
36.21. MON MAY 15 2017	308
36.22. WED APR 12 2017	308

CHAPTER 1. OVERVIEW

This guide is intended for application developers, and provides instructions for setting up and configuring a workstation to develop and deploy applications in an OpenShift Container Platform cloud environment. This includes detailed instructions and examples to help developers:

1. [Create new applications](#)
2. [Monitor and configure projects](#)
3. [Generate configurations using templates](#)
4. [Manage builds, including build strategy options and webhooks](#)
5. [Define deployments, including deployment strategies](#)
6. [Create and manage routes](#)
7. [Create and configure secrets](#)
8. [Integrate external services, such as databases and SaaS endpoints](#)
9. [Check application health using probes](#)

CHAPTER 2. APPLICATION LIFE CYCLE MANAGEMENT


2.1. PLANNING YOUR DEVELOPMENT PROCESS

2.1.1. Overview

OpenShift Container Platform is designed for building and deploying applications. Depending on how much you want to involve OpenShift Container Platform in your development process, you can choose to:

- focus your development within an OpenShift Container Platform project, using it to build an application from scratch then continuously develop and manage its lifecycle, or
- bring an application (e.g., binary, container image, source code) you have already developed in a separate environment and deploy it onto OpenShift Container Platform.

2.1.2. Using OpenShift Container Platform as Your Development Environment



You can begin your application's development from scratch using OpenShift Container Platform directly. Consider the following steps when planning this type of development process:

Initial Planning

- What does your application do?
- What programming language will it be developed in?

Access to OpenShift Container Platform

- OpenShift Container Platform should be installed by this point, either by yourself or an administrator within your organization.

Develop

- Using your editor or IDE of choice, create a basic skeleton of an application. It should be developed enough to tell OpenShift Container Platform [what kind of application it is](#).
- Push the code to your Git repository.

Generate

- [Create a basic application](#) using the `oc new-app` command. OpenShift Container Platform generates build and deployment configurations.

Manage


- Start developing your application code.
- Ensure your application builds successfully.

- Continue to locally develop and polish your code.
- Push your code to a Git repository.
- Is any extra configuration needed? Explore the [Developer Guide](#) for more options.

Verify

- You can verify your application in a number of ways. You can push your changes to your application's Git repository, and use OpenShift Container Platform to rebuild and redeploy your application. Alternatively, you can hot deploy using `rsync` to synchronize your code changes into a running pod.

2.1.3. Bringing an Application to Deploy on OpenShift Container Platform



Another possible application development strategy is to develop locally, then use OpenShift Container Platform to deploy your fully developed application. Use the following process if you plan to have application code already, then want to build and deploy onto an OpenShift Container Platform installation when completed:

Initial Planning

- What does your application do?
- What programming language will it be developed in?

Develop

- Develop your application code using your editor or IDE of choice.
- Build and test your application code locally.
- Push your code to a Git repository.

Access to OpenShift Container Platform

- OpenShift Container Platform should be installed by this point, either by yourself or an administrator within your organization.

Generate

- [Create a basic application](#) using the `oc new-app` command. OpenShift Container Platform generates build and deployment configurations.

Verify

- Ensure that the application that you have built and deployed in the above Generate step is successfully running on OpenShift Container Platform.

Manage

- Continue to develop your application code until you are happy with the results.

- Rebuild your application in OpenShift Container Platform to accept any newly pushed code.
- Is any extra configuration needed? Explore the [Developer Guide](#) for more options.

2.2. CREATING NEW APPLICATIONS

2.2.1. Overview

You can create a new OpenShift Container Platform application from components including source or binary code, images and/or templates by using either the OpenShift CLI or web console.

2.2.2. Creating an Application Using the CLI

2.2.2.1. Creating an Application From Source Code

The **new-app** command allows you to create applications from source code in a local or remote Git repository.

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

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



NOTE

If using a local Git repository, the repository should have a remote named **origin** that points to a URL accessible by the OpenShift Container Platform cluster. If there is no recognised remote, **new-app** will create a [binary build](#).

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

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

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

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

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

OpenShift Container Platform automatically [detects](#) whether the **Docker**, **Pipeline** or **Sourcebuild strategy** should be used, and in the case of **Source** builds, [detects an appropriate language builder image](#).

Build Strategy Detection

If a **Jenkinsfile** exists in the root or specified context directory of the source repository when creating a new application, OpenShift Container Platform generates a [Pipeline build strategy](#). Otherwise, if a

Dockerfile is found, OpenShift Container Platform generates a [Docker build strategy](#). Otherwise, it generates a [Source build strategy](#).

You can override the build strategy by setting the `--strategy` flag to either `docker`, `pipeline` or `source`.

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



NOTE

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

Language Detection

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

Table 2.1. Languages Detected by new-app

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

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

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

For example, to use the `myproject/my-ruby` image stream with the source in a remote repository:

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


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

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

2.2.2.2. Creating an Application From an Image

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

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



NOTE

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

For example, to create an application from the DockerHub MySQL image:

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

To create an application using an image in a private registry, specify the full Docker image specification:

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



NOTE

If the registry containing the image is not [secured with SSL](#), cluster administrators must ensure that the Docker daemon on the OpenShift Container Platform node hosts is run with the **--insecure-registry** flag pointing to that registry. You must also tell **new-app** that the image comes from an insecure registry with the **--insecure-registry** flag.

You can create an application from an existing [image stream](#) and optional [image stream tag](#):

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

2.2.2.3. Creating an Application From a Template

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

To create an application from a stored template:

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

To directly use a template in your local file system, without first storing it in OpenShift Container Platform, use the `-f | --file` argument:

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

Template Parameters

When creating an application based on a [template](#), use the `-p | --param` argument to set parameter values defined by the template:

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

You can store your parameters in a file, then use that file with `--param-file` when instantiating a template. If you want to read the parameters from standard input, use `--param-file=-`:

```
$ cat helloworld.params
ADMIN_USERNAME=admin
ADMIN_PASSWORD=myspassword
$ oc new-app ruby-helloworld-sample --param-file=helloworld.params
$ cat helloworld.params | oc new-app ruby-helloworld-sample --param-file=-
```

2.2.2.4. Further Modifying Application Creation

The `new-app` command generates OpenShift Container Platform objects that will build, deploy, and run the application being created. Normally, these objects are created in the current project using names derived from the input source repositories or the input images. However, `new-app` allows you to modify this behavior.

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

Table 2.2. new-app Output Objects

Object	Description
BuildConfig	A BuildConfig is created for each source repository specified in the command line. The BuildConfig specifies the strategy to use, the source location, and the build output location.
ImageStreams	For BuildConfig , two ImageStreams are usually created. One represents the input image. With Source builds, this is the builder image. With Docker builds, this is the FROM image. The second one represents the output image. If a container image was specified as input to <code>new-app</code> , then an image stream is created for that image as well.

Object	Description
DeploymentConfig	A DeploymentConfig is created either to deploy the output of a build, or a specified image. The new-app command creates emptyDir volumes for all Docker volumes that are specified in containers included in the resulting DeploymentConfig .
Service	The new-app command attempts to detect exposed ports in input images. It uses the lowest numeric exposed port to generate a service that exposes that port. In order to expose a different port, after new-app has completed, simply use the oc expose command to generate additional services.
Other	Other objects may be generated when instantiating templates , according to the template.

2.2.2.4.1. Specifying Environment Variables

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

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


The variables can also be read from file using the **--env-file** argument:

```
$ cat postgresql.env
POSTGRESQL_USER=user
POSTGRESQL_DATABASE=db
POSTGRESQL_PASSWORD=password
$ oc new-app openshift/postgresql-92-centos7 --env-file=postgresql.env
```

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

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

See [Managing Environment Variables](#) for more information.



NOTE

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

2.2.2.4.2. Specifying Build Environment Variables

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


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

The variables can also be read from a file using the **--build-env-file** argument:

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

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

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

2.2.2.4.3. Specifying Labels

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

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

2.2.2.4.4. Viewing the Output Without Creation

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

To output **new-app** artifacts to a file, edit them, then create them:

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

2.2.2.4.5. Creating Objects With Different Names

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

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

2.2.2.4.6. Creating Objects in a Different Project

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

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

2.2.2.4.7. Creating Multiple Objects

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

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

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



NOTE

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

2.2.2.4.8. Grouping Images and Source in a Single Pod

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

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

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

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


2.2.2.4.9. Searching for Images, Templates, and Other Inputs

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

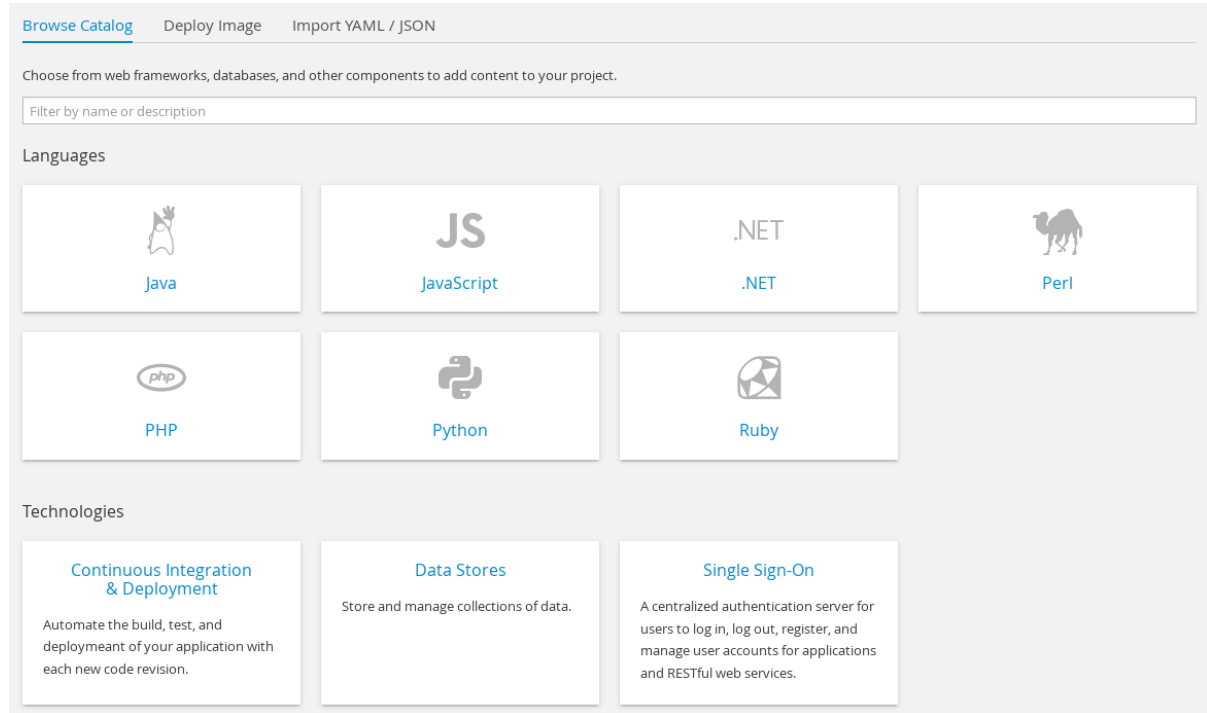
```
$ oc new-app --search php
```

2.2.3. Creating an Application Using the Web Console

1. While in the desired project, click **Add to Project**:



2. Select either a builder image from the list of images in your project, or from the global library:




NOTE

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

```
kind: "ImageStream"
apiVersion: "v1"
metadata:
  name: "ruby"
  creationTimestamp: null
spec:
  dockerImageRepository:
    "registry.access.redhat.com/openshift3/ruby-20-rhel7"
  tags:
    -
      name: "2.0"
      annotations:
        description: "Build and run Ruby 2.0 applications"
        iconClass: "icon-ruby"
        tags: "builder,ruby" 1
        supports: "ruby:2.0,ruby"
        version: "2.0"
```

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

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



*** Name** 2

nodejs-ex

Identifies the resources created for this application.

*** Git Repository URL** 3

https://github.com/openshift/nodejs-ex.git

Sample repository for nodejs: https://github.com/openshift/nodejs-ex.git [Try it ↗](#)**Git Reference**

master

Optional branch, tag, or commit.

Context Dir

/

Optional subdirectory for the application source code, used as the context directory for the build.

Routing 4[? About Routing](#)☒ Create a route to the application**Hostname**

www.example.com

Public hostname for the route. If not specified, a hostname is generated. The hostname can't be changed after the route is created.

Path

/

Path that the router watches to route traffic to the service.

Target Port

8080/TCP

Target port for traffic.

[Show options for secured routes](#)**Build Configuration** 5[? About Build Configuration](#)☒ Configure a webhook build trigger [?](#)☒ Automatically build a new image when the builder image changes [?](#)☒ Launch the first build when the build configuration is created**Environment Variables (Build and Runtime)** [?](#)

name

value

✕

[Add environment variable](#)**Deployment Configuration** 6[? About Deployment Configuration](#)**Autodeploy when**☒ New image is available☒ Deployment configuration changes**Environment Variables (Runtime only)** [?](#)[Show image environment variables](#)

name

value

✕

[Add environment variable](#)**Scaling** 7[? About Scaling](#)**Strategy**

Manual

Scale replicas manually or automatically based on CPU usage.

[Learn more ↗](#)**Replicas**

1

The number of instances of your image.

Resource Limits [About Resource Limits](#)

CPU

Request

millicores

The amount of CPU the container requests.

Limit

millicores

The amount of CPU the container is limited to use.

Memory

Request

MiB

The amount of memory the container requests.

Limit

MiB

The amount of Memory the container is limited to use.

Labels [About labels](#)

The following labels are being added automatically. If you want to override them, you can do so below.

app nodejs-ex

Each label is applied to each created resource.

Name Value ×

[Add label](#)

[Create](#) [Cancel](#)

The builder image name and description.

The application name used for the generated OpenShift Container Platform objects.

The Git repository URL, reference, and context directory for your source code.

Routing configuration section for making this application publicly accessible.

Build configuration section for customizing [build triggers](#).

Deployment configuration section for customizing [deployment triggers](#) and image environment variables.

Replica [scaling](#) section for configuring the number of running instances of the application.

The [labels](#) to assign to all items generated for the application. You can add and edit labels for all objects here.



NOTE

To see all of the configuration options, click the "Show advanced build and deployment options" link.

2.3. PROMOTING APPLICATIONS ACROSS ENVIRONMENTS

2.3.1. Overview

Application promotion means moving an application through various runtime environments, typically with an increasing level of maturity. For example, an application might start out in a development environment, then be promoted to a stage environment for further testing, before finally being promoted into a production environment. As changes are introduced in the application, again the changes will start in development and be promoted through stage and production.

The "application" today is more than just the source code written in Java, Perl, Python, etc. It is more now than the static web content, the integration scripts, or the associated configuration for the language specific runtimes for the application. It is more than the application specific archives consumed by those language specific runtimes.

In the context of OpenShift Container Platform and its combined foundation of Kubernetes and Docker, additional application artifacts include:

- *Docker container images* with their rich set of metadata and associated tooling.
- *Environment variables* that are injected into containers for application use.
- *API objects* (also known as resource definitions; see [Core Concepts](#)) of OpenShift Container Platform, which:
 - are injected into containers for application use.
 - dictate how OpenShift Container Platform manages containers and pods.

In examining how to promote applications in OpenShift Container Platform, this topic will:

- Elaborate on these new artifacts introduced to the application definition.
- Describe how you can demarcate the different environments for your application promotion pipeline.
- Discuss methodologies and tools for managing these new artifacts.
- Provide examples that apply the various concepts, constructs, methodologies, and tools to application promotion.

2.3.2. Application Components


2.3.2.1. API Objects

With regard to OpenShift Container Platform and Kubernetes resource definitions (the items newly introduced to the application inventory), there are a couple of key design points for these API objects that are relevant to revisit when considering the topic of application promotion.

First, as highlighted throughout OpenShift Container Platform documentation, every API object can be expressed via either JSON or YAML, making it easy to manage these resource definitions via traditional source control and scripting.

Also, the API objects are designed such that there are portions of the object which specify the desired state of the system, and other portions which reflect the status or current state of the system. This can be thought of as inputs and outputs. The input portions, when expressed in JSON or YAML, in particular

are items that fit naturally as source control managed (SCM) artifacts.



NOTE

Remember, the input or specification portions of the API objects can be totally static or dynamic in the sense that [variable substitution via template processing](#) is possible on instantiation.

The result of these points with respect to API objects is that with their expression as JSON or YAML files, you can treat the configuration of the application as code.

Conceivably, almost any of the API objects may be considered an application artifact by your organization. Listed below are the objects most commonly associated with deploying and managing an application:

BuildConfigs

This is a special case resource in the context of application promotion. While a **BuildConfig** is certainly a part of the application, especially from a developer's perspective, typically the **BuildConfig** is not promoted through the pipeline. It produces the **Image** that is promoted (along with other items) through the pipeline.

Templates

In terms of application promotion, **Templates** can serve as the starting point for setting up resources in a given staging environment, especially with the parameterization capabilities. Additional post-instantiation modifications are very conceivable though when applications move through a promotion pipeline. See [Scenarios and Examples](#) for more on this.

Routes

These are the most typical resources that differ stage to stage in the application promotion pipeline, as tests against different stages of an application access that application via its **Route**. Also, remember that you have options with regard to manual specification or auto-generation of host names, as well as the HTTP-level security of the **Route**.

Services

If reasons exist to avoid **Routers** and **Routes** at given application promotion stages (perhaps for simplicity's sake for individual developers at early stages), an application can be accessed via the **Cluster** IP address and port. If used, some management of the address and port between stages could be warranted.

Endpoints

Certain application-level services (e.g., database instances in many enterprises) may not be managed by OpenShift Container Platform. If so, then creating those **Endpoints** yourself, along with the necessary modifications to the associated **Service** (omitting the selector field on the **Service**) are activities that are either duplicated or shared between stages (based on how you delineate your environment).

Secrets

The sensitive information encapsulated by **Secrets** are shared between staging environments when the corresponding entity (either a **Service** managed by OpenShift Container Platform or an external service managed outside of OpenShift Container Platform) the information pertains to is shared. If there are different versions of the said entity in different stages of your application promotion pipeline, it may be necessary to maintain a distinct **Secret** in each stage of the pipeline or to make modifications to it as it traverses through the pipeline. Also, take care that if you are storing the **Secret** as JSON or YAML in an SCM, some form of encryption to protect the sensitive information may be warranted.

DeploymentConfigs

This object is the primary resource for defining and scoping the environment for a given application promotion pipeline stage; it controls how your application starts up. While there are aspects of it that will be common across all the different stage, undoubtedly there will be modifications to this object as it progresses through your application promotion pipeline to reflect differences in the environments for each stage, or changes in behavior of the system to facilitate testing of the different scenarios your application must support.

ImageStreams, ImageStreamTags, and ImageStreamImage

Detailed in the [Images](#) and [Image Streams](#) sections, these objects are central to the OpenShift Container Platform additions around managing container images.

ServiceAccounts and RoleBindings

Management of permissions to other API objects within OpenShift Container Platform, as well as the external services, are intrinsic to managing your application. Similar to **Secrets**, the **ServiceAccounts** and **RoleBindings** can objects vary in how they are shared between the different stages of your application promotion pipeline based on your needs to share or isolate those different environments.

PersistentVolumeClaims

Relevant to stateful services like databases, how much these are shared between your different application promotion stages directly correlates to how your organization shares or isolates the copies of your application data.

ConfigMaps

A useful decoupling of **Pod** configuration from the **Pod** itself (think of an environment variable style configuration), these can either be shared by the various staging environments when consistent **Pod** behavior is desired. They can also be modified between stages to alter **Pod** behavior (usually as different aspects of the application are vetted at different stages).

2.3.2.2. Images

As noted earlier, container images are now artifacts of your application. In fact, of the new applications artifacts, images and the management of images are the key pieces with respect to application promotion. In some cases, an image might encapsulate the entirety of your application, and the application promotion flow consists solely of managing the image.

Images are not typically managed in a SCM system, just as application binaries were not in previous systems. However, just as with binaries, installable artifacts and corresponding repositories (that is, RPMs, RPM repositories, Nexus, etc.) arose with similar semantics to SCMs, similar constructs and terminology around image management that are similar to SCMs have arisen:

- Image registry == SCM server
- Image repository == SCM repository

As images reside in registries, application promotion is concerned with ensuring the appropriate image exists in a registry that can be accessed from the environment that needs to run the application represented by that image.

Rather than reference images directly, application definitions typically abstract the reference into an image stream. This means the image stream will be another API object that makes up the application components. For more details on image streams, see [Core Concepts](#).

2.3.2.3. Summary

Now that the application artifacts of note, images and API objects, have been detailed in the context of

application promotion within OpenShift Container Platform, the notion of *where* you run your application in the various stages of your promotion pipeline is next the point of discussion.

2.3.3. Deployment Environments

A deployment environment, in this context, describes a distinct space for an application to run during a particular stage of a CI/CD pipeline. Typical environments include **development**, **test**, **stage**, and **production**, for example. The boundaries of an environment can be defined in different ways, such as:

- Via labels and unique naming within a single project.
- Via distinct projects within a cluster.
- Via distinct clusters.

And it is conceivable that your organization leverages all three.

2.3.3.1. Considerations

Typically, you will consider the following heuristics in how you structure the deployment environments:

- How much resource sharing the various stages of your promotion flow allow
- How much isolation the various stages of your promotion flow require
- How centrally located (or geographically dispersed) the various stages of your promotion flow are

Also, some important reminders on how OpenShift Container Platform clusters and projects relate to image registries:

- Multiple project in the same cluster can access the same image streams.
- Multiple clusters can access the same external registries.
- Clusters can only share a registry if the OpenShift Container Platform internal image registry is exposed via a route.

2.3.3.2. Summary

After deployment environments are defined, promotion flows with delineation of stages within a pipeline can be implemented. The methods and tools for constructing those promotion flow implementations are the next point of discussion.

2.3.4. Methods and Tools

Fundamentally, application promotion is a process of moving the aforementioned application components from one environment to another. The following subsections outline tools that can be used to move the various components by hand, before advancing to discuss holistic solutions for automating application promotion.



NOTE

There are a number of insertion points available during both the build and deployment processes. They are defined within **BuildConfig** and **DeploymentConfig** API objects. These hooks allow for the invocation of custom scripts which can interact with deployed components such as databases, and with the OpenShift Container Platform cluster itself.

Therefore, it is possible to use these hooks to perform component management operations that effectively move applications between environments, for example by performing an image tag operation from within a hook. However, the various hook points are best suited to managing an application's lifecycle within a given environment (for example, using them to perform database schema migrations when a new version of the application is deployed), rather than to move application components between environments.

2.3.4.1. Managing API Objects

Resources, as defined in one environment, will be exported as JSON or YAML file content in preparation for importing it into a new environment. Therefore, the expression of API objects as JSON or YAML serves as the unit of work as you promote API objects through your application pipeline. The **oc** CLI is used to export and import this content.

TIP

While not required for promotion flows with OpenShift Container Platform, with the JSON or YAML stored in files, you can consider storing and retrieving the content from a SCM system. This allows you to leverage the versioning related capabilities of the SCM, including the creation of branches, and the assignment of and query on various labels or tags associated to versions.

2.3.4.1.1. Exporting API Object State

API object specifications should be captured with **oc export**. This operation removes environment specific data from the object definitions (e.g., current namespace or assigned IP addresses), allowing them to be recreated in different environments (unlike **oc get** operations, which output an unfiltered state of the object).

Use of **oc label**, which allows for adding, modifying, or removing labels on API objects, can prove useful as you organize the set of object collected for promotion flows, because labels allow for selection and management of groups of pods in a single operation. This makes it easier to export the correct set of objects and, because the labels will carry forward when the objects are created in a new environment, they also make for easier management of the application components in each environment.



NOTE

API objects often contain references such as a **DeploymentConfig** that references a **Secret**. When moving an API object from one environment to another, you must ensure that such references are also moved to the new environment.

Similarly, API objects such as a **DeploymentConfig** often contain references to **ImageStreams** that reference an external registry. When moving an API object from one environment to another, you must ensure such references are resolvable within the new environment, meaning that the reference must be resolvable and the **ImageStream** must reference an accessible registry in the new environment. See [Moving Images](#) and [Promotion Caveats](#) for more detail.

2.3.4.1.2. Importing API Object State

2.3.4.1.2.1. Initial Creation

The first time an application is being introduced into a new environment, it is sufficient to take the JSON or YAML expressing the specifications of your API objects and run **oc create** to create them in the appropriate environment. When using **oc create**, keep the **--save-config** option in mind. Saving configuration elements on the object in its annotation list facilitates the later use of **oc apply** to modify the object.

2.3.4.1.2.2. Iterative Modification

After the various staging environments are initially established, as promotion cycles commence and the application moves from stage to stage, the updates to your application can include modification of the API objects that are part of the application. Changes in these API objects are conceivable since they represent the configuration for the OpenShift Container Platform system. Motivations for such changes include:

- Accounting for environmental differences between staging environments.
- Verifying various scenarios your application supports.

Transfer of the API objects to the next stage's environment is accomplished via use of the **oc CLI**. While a rich set of **oc** commands which modify API objects exist, this topic focuses on **oc apply**, which computes and applies differences between objects.

Specifically, you can view **oc apply** as a three-way merge that takes in files or stdin as the input along with an existing object definition. It performs a three-way merge between:

1. the input into the command,
2. the current version of the object, and
3. the most recent user specified object definition stored as an annotation in the current object.

The existing object is then updated with the result.

If further customization of the API objects is necessary, as in the case when the objects are not expected to be identical between the source and target environments, **oc** commands such as **oc set** can be used to modify the object after applying the latest object definitions from the upstream environment.

Some specific usages are cited in [Scenarios and Examples](#).

2.3.4.2. Managing Images and Image Streams

Images in OpenShift Container Platform are managed via a series of API objects as well. However, managing images are so central to application promotion that discussion of the tools and API objects most directly tied to images warrant separate discussion. Both manual and automated forms exist to assist you in managing image promotion (the propagation of images through your pipeline).

2.3.4.2.1. Moving Images



NOTE

For all the detailed caveats around managing images, refer to the [Managing Images](#) topic.

2.3.4.2.1.1. When Staging Environments Share a Registry

When your staging environments share the same OpenShift Container Platform registry, for example if they are all on the same OpenShift Container Platform cluster, there are two operations that are the basic means of *moving* your images between the stages of your application promotion pipeline:

1. First, analogous to `docker tag` and `git tag`, the `oc tag` command allows you to update an OpenShift Container Platform image stream with a reference to a specific image. It also allows you to copy references to specific versions of an image from one image stream to another, even across different projects in a cluster.
2. Second, the `oc import -image` serves as a bridge between external registries and image streams. It imports the metadata for a given image from the registry and stores it into the image stream as an [image stream tag](#). Various `BuildConfigs` and `DeploymentConfigs` in your project can reference those specific images.

2.3.4.2.1.2. When Staging Environments Use Different Registries

More advanced usage occurs when your staging environments leverage different OpenShift Container Platform registries. [Accessing the Internal Registry](#) spells out the steps in detail, but in summary you can:

1. Use the `docker` command in conjunction with obtaining the OpenShift Container Platform access token to supply into your `docker login` command.
2. After being logged into the OpenShift Container Platform registry, use `docker pull`, `docker tag` and `docker push` to transfer the image.
3. After the image is available in the registry of the next environment of your pipeline, use `oc tag` as needed to populate any image streams.

2.3.4.2.2. Deploying

Whether changing the underlying application image or the API objects that configure the application, a deployment is typically necessary to pick up the promoted changes. If the images for your application change (for example, due to an `oc tag` operation or a `docker push` as part of promoting an image from an upstream environment), `ImageChangeTriggers` on your `DeploymentConfig` can trigger

the new deployment. Similarly, if the **DeploymentConfig** API object itself is being changed, a **ConfigChangeTrigger** can initiate a deployment when the API object is updated by the promotion step (for example, `oc apply`).

Otherwise, the `oc` commands that facilitate manual deployment include:

- **oc rollout**: The new approach to manage deployments, including pause and resume semantics and richer features around managing history.
- **oc rollback**: Allows for reversion to a previous deployment; in the promotion scenario, if testing of a new version encounters issues, confirming it still works with the previous version could be warranted.

2.3.4.2.3. Automating Promotion Flows with Jenkins

After you understand the components of your application that need to be moved between environments when promoting it and the steps required to move the components, you can start to orchestrate and automate the workflow. OpenShift Container Platform provides a Jenkins image and plug-ins to help with this process.

The OpenShift Container Platform Jenkins image is detailed in [Using Images](#), including the set of OpenShift Container Platform-centric plug-ins that facilitate the integration of Jenkins, and Jenkins Pipelines. Also, the [Pipeline build strategy](#) facilitates the integration between Jenkins Pipelines and OpenShift Container Platform. All of these focus on enabling various aspects of CI/CD, including application promotion.

When moving beyond manual execution of application promotion steps, the Jenkins-related features provided by OpenShift Container Platform should be kept in mind:

- OpenShift Container Platform provides a Jenkins image that is heavily customized to greatly ease deployment in an OpenShift Container Platform cluster.
- The Jenkins image contains the OpenShift Pipeline plug-in, which provides building blocks for implementing promotion workflows. These building blocks include the triggering of Jenkins jobs as image streams change, as well as the triggering of builds and deployments within those jobs.
- **BuildConfigs** employing the OpenShift Container Platform Jenkins Pipeline build strategy enable execution of Jenkinsfile-based Jenkins Pipeline jobs. Pipeline jobs are the strategic direction within Jenkins for complex promotion flows and can leverage the steps provided by the OpenShift Pipeline Plug-in.

2.3.4.2.4. Promotion Caveats

2.3.4.2.4.1. API Object References

API objects can reference other objects. A common use for this is to have a **DeploymentConfig** that references an image stream, but other reference relationships may also exist.

When copying an API object from one environment to another, it is critical that all references can still be resolved in the target environment. There are a few reference scenarios to consider:

- The reference is "local" to the project. In this case, the referenced object resides in the same project as the object that references it. Typically the correct thing to do is to ensure that you copy the referenced object into the target environment in the same project as the object referencing it.

- The reference is to an object in another project. This is typical when an image stream in a shared project is used by multiple application projects (see [Managing Images](#)). In this case, when copying the referencing object to the new environment, you must update the reference as needed so it can be resolved in the target environment. That may mean:
 - Changing the project the reference points to, if the shared project has a different name in the target environment.
 - Moving the referenced object from the shared project into the local project in the target environment and updating the reference to point to the local project when moving the primary object into the target environment.
 - Some other combination of copying the referenced object into the target environment and updating references to it.

In general, the guidance is to consider objects referenced by the objects being copied to a new environment and ensure the references are resolvable in the target environment. If not, take appropriate action to fix the references and make the referenced objects available in the target environment.

2.3.4.2.4.2. Image Registry References

Image streams point to image repositories to indicate the source of the image they represent. When an image stream is moved from one environment to another, it is important to consider whether the registry and repository reference should also change:

- If different image registries are used to assert isolation between a test environment and a production environment.
- If different image repositories are used to separate test and production-ready images.

If either of these are the case, the image stream must be modified when it is copied from the source environment to the target environment so that it resolves to the correct image. This is in addition to performing the steps described in [Scenarios and Examples](#) to copy the image from one registry and repository to another.

2.3.4.3. Summary

At this point, the following have been defined:

- New application artifacts that make up a deployed application.
- Correlation of application promotion activities to tools and concepts provided by OpenShift Container Platform.
- Integration between OpenShift Container Platform and the CI/CD pipeline engine Jenkins.

Putting together examples of application promotion flows within OpenShift Container Platform is the final step for this topic.

2.3.5. Scenarios and Examples

Having defined the new application artifact components introduced by the Docker, Kubernetes, and OpenShift Container Platform ecosystems, this section covers how to promote those components between environments using the mechanisms and tools provided by OpenShift Container Platform.

Of the components making up an application, the image is the primary artifact of note. Taking that premise and extending it to application promotion, the core, fundamental application promotion pattern is image promotion, where the unit of work is the image. The vast majority of application promotion scenarios entails management and propagation of the image through the promotion pipeline.

Simpler scenarios solely deal with managing and propagating the image through the pipeline. As the promotion scenarios broaden in scope, the other application artifacts, most notably the API objects, are included in the inventory of items managed and propagated through the pipeline.

This topic lays out some specific examples around promoting images as well as API objects, using both manual and automated approaches. But first, note the following on setting up the environment(s) for your application promotion pipeline.

2.3.5.1. Setting up for Promotion

After you have completed development of the initial revision of your application, the next logical step is to package up the contents of the application so that you can transfer to the subsequent staging environments of your promotion pipeline.

1. First, group all the API objects you view as transferable and apply a common `label` to them:

```
labels:
  promotion-group: <application_name>
```

As previously described, the `oc label` command facilitates the management of labels with your various API objects.


TIP

If you initially define your API objects in a OpenShift Container Platform template, you can easily ensure all related objects have the common label you will use to query on when exporting in preparation for a promotion.

2. You can leverage that label on subsequent queries. For example, consider the following set of `oc` command invocations that would then achieve the transfer of your application's API objects:

```
$ oc login <source_environment>
$ oc project <source_project>
$ oc export dc,is,svc,route,secret,sa -l promotion-group=
<application_name> -o yaml > export.yaml
$ oc login <target_environment>
$ oc new-project <target_project> ❶
$ oc create -f export.yaml
```

- ❶ Alternatively, `oc project <target_project>` if it already exists.

**NOTE**

On the `oc export` command, whether or not you include the `is` type for image streams depends on how you choose to manage images, image streams, and registries across the different environments in your pipeline. The caveats around this are discussed below. See also the [Managing Images](#) topic.

3. You must also get any tokens necessary to operate against each registry used in the different staging environments in your promotion pipeline. For each environment:

- a. Log in to the environment:

```
$ oc login <each_environment_with_a_unique_registry>
```

- b. Get the access token with:

```
$ oc whoami -t
```

- c. Copy and paste the token value for later use.

2.3.5.2. Repeatable Promotion Process

After the initial setup of the different staging environments for your pipeline, a set of repeatable steps to validate each iteration of your application through the promotion pipeline can commence. These basic steps are taken each time the image or API objects in the source environment are changed:

Move updated images → Move updated API objects → Apply environment specific customizations

1. Typically, the first step is promoting any updates to the image(s) associated with your application to the next stage in the pipeline. As noted above, the key differentiator in promoting images is whether the OpenShift Container Platform registry is shared or not between staging environments.

- a. If the registry is shared, simply leverage `oc tag`:

```
$ oc tag <project_for_stage_N>/<imagestream_name_for_stage_N>:
<tag_for_stage_N>
<project_for_stage_N+1>/<imagestream_name_for_stage_N+1>:
<tag_for_stage_N+1>
```

- b. If the registry is not shared, you can leverage the access tokens for each of your promotion pipeline registries as you log into both the source and destination registries, pulling, tagging, and pushing your application images accordingly:

- i. Log in to the source environment registry:

```
$ docker login -u <username> -e <any_email_address> -p
<token_value> <src_env_registry_ip>:<port>
```

- ii. Pull your application's image:

```
$ docker pull <src_env_registry_ip>:<port>/<namespace>/<image
name>:<tag>
```


- iii. Tag your application's image to the destination registry's location, updating namespace, name, and tag as needed to conform to the destination staging environment:

```
$ docker tag <src_env_registry_ip>:<port>/<namespace>/<image
name>:<tag> <dest_env_registry_ip>:<port>/<namespace>/<image
name>:<tag>
```

- iv. Log into the destination staging environment registry:

```
$ docker login -u <username> -e <any_email_address> -p
<token_value> <dest_env_registry_ip>:<port>
```

- v. Push the image to its destination:

```
$ docker push <dest_env_registry_ip>:<port>/<namespace>/<image
name>:<tag>
```

TIP

To automatically import new versions of an image from an external registry, the **oc tag** command has a **--scheduled** option. If used, the image the **ImageStreamTag** references will be periodically pulled from the registry hosting the image.

2. Next, there are the cases where the evolution of your application necessitates fundamental changes to your API objects or additions and deletions from the set of API objects that make up the application. When such evolution in your application's API objects occurs, the OpenShift Container Platform CLI provides a broad range of options to transfer to changes from one staging environment to the next.

- a. Start in the same fashion as you did when you initially set up your promotion pipeline:

```
$ oc login <source_environment>
$ oc project <source_project>
$ oc export dc,is,svc,route,secret,sa -l template=
<application_template_name> -o yaml > export.yaml
$ oc login <target_environment>
$ oc <target_project>
```

- b. Rather than simply creating the resources in the new environment, update them. You can do this a few different ways:

- i. The more conservative approach is to leverage **oc apply** and merge the new changes to each API object in the target environment. In doing so, you can **--dry-run=true** option and examine the resulting objects prior to actually changing the objects:

```
$ oc apply -f export.yaml --dry-run=true
```

If satisfied, actually run the **apply** command:

```
$ oc apply -f export.yaml
```

The **apply** command optionally takes additional arguments that help with more complicated scenarios. See **oc apply --help** for more details.

- ii. Alternatively, the simpler but more aggressive approach is to leverage **oc replace**. There is no dry run with this update and replace. In the most basic form, this involves executing:

```
$ oc replace -f export.yaml
```

As with **apply**, **replace** optionally takes additional arguments for more sophisticated behavior. See **oc replace --help** for more details.

3. The previous steps automatically handle new API objects that were introduced, but if API objects were deleted from the source environment, they must be manually deleted from the target environment using **oc delete**.
4. Tuning of the environment variables cited on any of the API objects may be necessary as the desired values for those may differ between staging environments. For this, use **oc set env**:

```
$ oc set env <api_object_type>/<api_object_ID> <env_var_name>=  
<env_var_value>
```

5. Finally, trigger a new deployment of the updated application using the **oc rollout** command or one of the other mechanisms discussed in the [Deployments](#) section above.

2.3.5.3. Repeatable Promotion Process Using Jenkins

The [OpenShift Sample](#) job defined in the [Jenkins Docker Image](#) for OpenShift Container Platform is an example of image promotion within OpenShift Container Platform within the constructs of Jenkins. Setup for this sample is located in the [OpenShift Origin source repository](#).

This sample includes:

- Use of **Jenkins as the CI/CD engine**.
- Use of the **OpenShift Pipeline plug-in for Jenkins**. This plug-in provides a subset of the functionality provided by the **oc** CLI for OpenShift Container Platform packaged as Jenkins Freestyle and DSL Job steps. Note that the **oc** binary is also included in the Jenkins Docker Image for OpenShift Container Platform, and can also be used to interact with OpenShift Container Platform in Jenkins jobs.
- The OpenShift Container Platform-provided **templates for Jenkins**. There is a template for both ephemeral and persistent storage.
- A **sample application**: defined in the [OpenShift Origin source repository](#), this application leverages **ImageStreams**, **ImageChangeTriggers**, **ImageStreamTags**, **BuildConfigs**, and separate **DeploymentConfigs** and **Services** corresponding to different stages in the promotion pipeline.

The following examines the various pieces of the OpenShift Sample job in more detail:

1. [The first step](#) is the equivalent of an **oc scale dc frontend --replicas=0** call. This step is intended to bring down any previous versions of the application image that may be running.
2. [The second step](#) is the equivalent of an **oc start-build frontend** call.

3. **The third step** is the equivalent of an `oc rollout latest dc/frontend` call.
4. **The fourth step** is the "test" for this sample. It ensures that the associated service for this application is in fact accessible from a network perspective. Under the covers, a socket connection is attempted against the IP address and port associated with the OpenShift Container Platform service. Of course, additional tests can be added (if not via OpenShift Pipeline plug-in steps, then via use of the Jenkins Shell step to leverage OS-level commands and scripts to test your application).
5. **The fifth step** commences under that assumption that the testing of your application passed and hence intends to mark the image as "ready". In this step, a new `prod` tag is created for the application image off of the `latest` image. With the `frontend DeploymentConfig` having an `ImageChangeTrigger` defined for that tag, the corresponding "production" deployment is launched.
6. **The sixth and last step** is a verification step, where the plug-in confirms that OpenShift Container Platform launched the desired number of replicas for the "production" deployment.

CHAPTER 3. AUTHENTICATION

3.1. WEB CONSOLE AUTHENTICATION

When accessing the [web console](#) from a browser at `<master_public_addr>:8443`, you are automatically redirected to a login page.

Review the [browser versions and operating systems](#) that can be used to access the web console.

You can provide your login credentials on this page to obtain a token to make API calls. After logging in, you can navigate your projects using the [web console](#).

3.2. CLI AUTHENTICATION

You can authenticate from the command line using the CLI command `oc login`. You can [get started with the CLI](#) by running this command without any options:

```
$ oc login
```

The command's interactive flow helps you establish a session to an OpenShift Container Platform server with the provided credentials. If any information required to successfully log in to an OpenShift Container Platform server is not provided, the command prompts for user input as required. The [configuration](#) is automatically saved and is then used for every subsequent command.

All configuration options for the `oc login` command, listed in the `oc login --help` command output, are optional. The following example shows usage with some common options:

```
$ oc login [-u=<username>] \
  [-p=<password>] \
  [-s=<server>] \
  [-n=<project>] \
  [--certificate-authority=</path/to/file.crt>|--insecure-skip-tls-verify]
```


The following table describes these common options:

Table 3.1. Common CLI Configuration Options

Option	Syntax	Description
-s, --server	<pre>\$ oc login -s= <server></pre>	Specifies the host name of the OpenShift Container Platform server. If a server is provided through this flag, the command does not ask for it interactively. This flag can also be used if you already have a CLI configuration file and want to log in and switch to another server.
-u, --username and -p, --password	<pre>\$ oc login -u= <username> -p= <password></pre>	Allows you to specify the credentials to log in to the OpenShift Container Platform server. If user name or password are provided through these flags, the command does not ask for it interactively. These flags can also be used if you already have a configuration file with a session token established and want to log in and switch to another user name.

Option	Syntax	Description
-n, --namespace	<pre>\$ oc login -u= <username> -p= <password> -n= <project></pre>	A global CLI option which, when used with oc login , allows you to specify the project to switch to when logging in as a given user.
--certificate-authority	<pre>\$ oc login -- certificate- authority= <path/to/file. crt></pre>	Correctly and securely authenticates with an OpenShift Container Platform server that uses HTTPS. The path to a certificate authority file must be provided.
--insecure-skip-tls-verify	<pre>\$ oc login -- insecure-skip- tls-verify</pre>	Allows interaction with an HTTPS server bypassing the server certificate checks; however, note that it is not secure. If you try to oc login to a HTTPS server that does not provide a valid certificate, and this or the --certificate-authority flags were not provided, oc login will prompt for user input to confirm (y/N kind of input) about connecting insecurely.

CLI configuration files allow you to easily [manage multiple CLI profiles](#).



NOTE

If you have access to administrator credentials but are no longer logged in as the [default system user system:admin](#), you can log back in as this user at any time as long as the credentials are still present in your [CLI configuration file](#). The following command logs in and switches to the **default** project:

```
$ oc login -u system:admin -n default
```

CHAPTER 4. AUTHORIZATION

4.1. OVERVIEW

This topic contains [authorization tasks](#) for application developers and their capabilities, as dictated by the cluster administrator.

4.2. CHECKING IF USERS CAN CREATE PODS

Using the `scc-review` and `scc-subject-review` options, you can see if an individual user, or a user under a specific service account, can create or update a pod.

Using the `scc-review` option, you can check if a service account can create or update a pod. The command outputs the security context constraints that admit the resource.

For example, to check if a user with the `system:serviceaccount:projectname:default` service account can create a pod:

```
$ oc policy scc-review -z system:serviceaccount:projectname:default -f my_resource.yaml
```

You can also use the `scc-subject-review` option to check whether a specific user can create or update a pod:

```
$ oc policy scc-subject-review -u <username> -f my_resource.yaml
```

To check if a user belonging to a specific group can create a pod in a specific file:

```
$ oc policy scc-subject-review -u <username> -g <groupname> -f my_resource.yaml
```

4.3. DETERMINING WHAT YOU CAN DO AS AN AUTHENTICATED USER

From within your OpenShift Container Platform project, you can determine what [verbs](#) you can perform against all namespace-scoped resources (including third-party resources).

The `can-i` command option tests scopes in terms of the user and role.

```
$ oc policy can-i --list --loglevel=8
```

The output helps you to determine what API request to make to gather the information.

To receive information back in a user-readable format, run:

```
$ oc policy can-i --list
```

The output provides a full list.

To determine if you can perform specific verbs, run:

```
$ oc policy can-i <verb> <resource>
```

[User scopes](#) can provide more information about a given scope. For example:

```
$ oc policy can-i <verb> <resource> --scopes=user:info
```

CHAPTER 5. PROJECTS

5.1. OVERVIEW

A [project](#) allows a community of users to organize and manage their content in isolation from other communities.

5.2. CREATING A PROJECT

If [allowed](#) by your cluster administrator, you can create a new project using [the CLI](#) or the [web console](#).

To create a new project using the CLI:

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

For example:

```
$ oc new-project hello-openshift \
  --description="This is an example project to demonstrate OpenShift v3" \
  --display-name="Hello OpenShift"
```



NOTE

The number of projects you are allowed to create [may be limited by the system administrator](#). Once your limit is reached, you may need to delete an existing project in order to create a new one.

5.3. VIEWING PROJECTS

When viewing projects, you are restricted to seeing only the projects you have access to view based on the [authorization policy](#).

To view a list of projects:

```
$ oc get projects
```

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



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

You can also use the [web console](#) to view and change between projects. After [authenticating](#) and logging in, you are presented with a list of projects that you have access to:

Projects

New Project

Rails Example

A project admin can add you to a role on a project by running the command `oc policy add-role-to-user <role> admin -n <projectname>`

If you use [the CLI](#) to [create a new project](#), you can then refresh the page in the browser to see the new project.

Selecting a project brings you to the [project overview](#) for that project.

5.4. CHECKING PROJECT STATUS

The `oc status` command provides a high-level overview of the current project, with its components and their relationships. This command takes no argument:

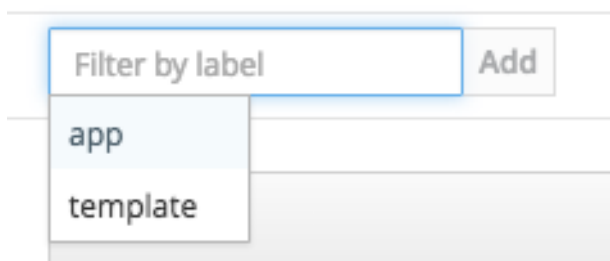
```
$ oc status
```

5.5. FILTERING BY LABELS

You can filter the contents of a project page in the [web console](#) by using the [labels](#) of a resource. You can pick from a suggested label name and values, or type in your own. Multiple filters can be added. When multiple filters are applied, resources must match all of the filters to remain visible.

To filter by labels:

1. Select a label type:



2. Select one of the following:

exists	Verify that the label name exists, but ignore its value.
in	Verify that the label name exists and is equal to one of the selected values.

not in	Verify that the label name does not exist, or is not equal to any of the selected values.
--------	---

exists

in ...

not in ...

Name	Last D
postgresql	#

a. If you selected **in** or **not in**, select a set of values then select **Filter**:

Name	Last D
------	--------

- After adding filters, you can stop filtering by selecting **Clear all filters** or by clicking individual filters to remove them:

5.6. BOOKMARKING PAGE STATES

The OpenShift Container Platform [web console](#) now bookmarks page states, which is helpful in saving label filters and other settings.

When you do something to change the page's state, like switching between tabs, the URL in the browser's navigation bar is automatically updated.

5.7. DELETING A PROJECT

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

To delete a project using the CLI:

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

CHAPTER 6. MIGRATING APPLICATIONS

6.1. OVERVIEW

This topic covers the migration procedure of OpenShift version 2 (v2) applications to OpenShift version 3 (v3).



NOTE

This topic uses some terminology that is specific to OpenShift v2. [Comparing OpenShift Enterprise 2 and OpenShift Enterprise 3](#) provides insight on the differences between the two versions and the language used.

To migrate OpenShift v2 applications to OpenShift Container Platform v3, all cartridges in the v2 application must be recorded as each v2 cartridge is equivalent with a corresponding image or template in OpenShift Container Platform v3 and they must be migrated individually. For each cartridge, all dependencies or required packages also must be recorded, as they must be included in the v3 images.

The general migration procedure is:

1. Back up the v2 application.
 - Web cartridge: The source code can be backed up to a Git repository such as by pushing to a repository on GitHub.
 - Database cartridge: The database can be backed up using a dump command (`mongodump`, `mysqldump`, `pg_dump`) to back up the database.
 - Web and database cartridges: `rhc` client tool provides snapshot ability to back up multiple cartridges:

```
$ rhc snapshot save <app_name>
```

The snapshot is a tar file that can be unzipped, and its content is application source code and the database dump.

2. If the application has a database cartridge, create a v3 database application, sync the database dump to the pod of the new v3 database application, then restore the v2 database in the v3 database application with database restore commands.
3. For a web framework application, edit the application source code to make it v3 compatible. Then, add any dependencies or packages required in appropriate files in the Git repository. Convert v2 environment variables to corresponding v3 environment variables.
4. Create a v3 application from source (your Git repository) or from a quickstart with your Git URL. Also, add the database service parameters to the new application to link the database application to the web application.
5. In v2, there is an integrated Git environment and your applications automatically rebuild and restart whenever a change is pushed to your v2 Git repository. In v3, in order to have a build automatically triggered by source code changes pushed to your public Git repository, you must set up a [webhook](#) after the initial build in v3 is completed.

6.2. MIGRATING DATABASE APPLICATIONS

6.2.1. Overview

This topic reviews how to migrate MySQL, PostgreSQL, and MongoDB database applications from OpenShift version 2 (v2) to OpenShift version 3 (v3).

6.2.2. Supported Databases

v2	v3
MongoDB: 2.4	MongoDB: 2.4, 2.6
MySQL: 5.5	MySQL: 5.5, 5.6
PostgreSQL: 9.2	PostgreSQL: 9.2, 9.4

6.2.3. MySQL

1. Export all databases to a dump file and copy it to a local machine (into the current directory):

```
$ rhc ssh <v2_application_name>
$ mysqldump --skip-lock-tables -h $OPENSHIFT_MYSQL_DB_HOST -P
${OPENSHIFT_MYSQL_DB_PORT:-3306} -u ${OPENSHIFT_MYSQL_DB_USERNAME:-
'admin'} \
  --password="$OPENSHIFT_MYSQL_DB_PASSWORD" --all-databases > ~/app-
root/data/all.sql
$ exit
```

2. Download **dbdump** to your local machine:

```
$ mkdir mysqldumpdir
$ rhc scp -a <v2_application_name> download mysqldumpdir app-
root/data/all.sql
```

3. Create a v3 **mysql-persistent** pod from template:

```
$ oc new-app mysql-persistent -p \
  MYSQL_USER=<your_v2_mysql_username> -p \
  MYSQL_PASSWORD=<your_v2_mysql_password> -p MYSQL_DATABASE=
<your_v2_database_name>
```

4. Check to see if the pod is ready to use:

```
$ oc get pods
```

5. When the pod is up and running, copy database archive files to your v3 MySQL pod:

```
$ oc rsync /local/mysqldumpdir <mysql_pod_name>:/var/lib/mysql/data
```

6. Restore the database in the v3 running pod:

```
$ oc rsh <mysql_pod>
$ cd /var/lib/mysql/data/mysqldumpdir
```

In v3, to restore databases you need to access MySQL as **root** user.

In v2, the **\$OPENSIFT_MYSQL_DB_USERNAME** had full privileges on all databases. In v3, you must grant privileges to **\$MYSQL_USER** for each database.

```
$ mysql -u root
$ source all.sql
```

Grant all privileges on **<dbname>** to **<your_v2_username>@localhost**, then flush privileges.

7. Remove the dump directory from the pod:

```
$ cd ../; rm -rf /var/lib/mysql/data/mysqldumpdir
```

Supported MySQL Environment Variables

v2	v3
OPENSIFT_MYSQL_DB_HOST	[service_name]_SERVICE_HOST
OPENSIFT_MYSQL_DB_PORT	[service_name]_SERVICE_PORT
OPENSIFT_MYSQL_DB_USERNAME	MYSQL_USER
OPENSIFT_MYSQL_DB_PASSWORD	MYSQL_PASSWORD
OPENSIFT_MYSQL_DB_URL	
OPENSIFT_MYSQL_DB_LOG_DIR	
OPENSIFT_MYSQL_VERSION	
OPENSIFT_MYSQL_DIR	
OPENSIFT_MYSQL_DB_SOCKET	
OPENSIFT_MYSQL_IDENT	
OPENSIFT_MYSQL_AIO	MYSQL_AIO
OPENSIFT_MYSQL_MAX_ALLOWED_PACKET	MYSQL_MAX_ALLOWED_PACKET
OPENSIFT_MYSQL_TABLE_OPEN_CACHE	MYSQL_TABLE_OPEN_CACHE

v2	v3
OPENSIFT_MYSQL_SORT_BUFFER_SIZE	MYSQL_SORT_BUFFER_SIZE
OPENSIFT_MYSQL_LOWER_CASE_TABLE_NAMES	MYSQL_LOWER_CASE_TABLE_NAMES
OPENSIFT_MYSQL_MAX_CONNECTIONS	MYSQL_MAX_CONNECTIONS
OPENSIFT_MYSQL_FT_MIN_WORD_LEN	MYSQL_FT_MIN_WORD_LEN
OPENSIFT_MYSQL_FT_MAX_WORD_LEN	MYSQL_FT_MAX_WORD_LEN
OPENSIFT_MYSQL_DEFAULT_STORAGE_ENGINE	
OPENSIFT_MYSQL_TIMEZONE	
	MYSQL_DATABASE
	MYSQL_ROOT_PASSWORD
	MYSQL_MASTER_USER
	MYSQL_MASTER_PASSWORD

6.2.4. PostgreSQL

1. Back up the v2 PostgreSQL database from the gear:

```
$ rhc ssh -a <v2-application_name>
$ mkdir ~/app-root/data/tmp
$ pg_dump <database_name> | gzip > ~/app-root/data/tmp/<database_name>.gz
```

2. Extract the backup file back to your local machine:

```
$ rhc scp -a <v2_application_name> download <local_dest> app-root/data/tmp/<db-name>.gz
$ gzip -d <database-name>.gz
```



NOTE

Save the backup file to a separate folder for step 4.

3. Create the PostgreSQL service using the v2 application database name, user name and password to create the new service:

```
$ oc new-app postgresql-persistent -p POSTGRESQL_DATABASE=dbname -p
POSTGRESQL_PASSWORD=password -p POSTGRESQL_USER=username
```

4. Check to see if the pod is ready to use:

```
$ oc get pods
```

5. When the pod is up and running, sync the backup directory to pod:

```
$ oc rsync /local/path/to/dir
<postgresql_pod_name>:/var/lib/pgsql/data
```

6. Remotely access the pod:

```
$ oc rsh <pod_name>
```

7. Restore the database:

```
psql dbname < /var/lib/pgsql/data/<database_backup_file>
```

8. Remove all backup files that are no longer needed:

```
$ rm /var/lib/pgsql/data/<database-backup-file>
```

Supported PostgreSQL Environment Variables

v2	v3
OPENSIFT_POSTGRESQL_DB_HOST	[service_name]_SERVICE_HOST
OPENSIFT_POSTGRESQL_DB_PORT	[service_name]_SERVICE_PORT
OPENSIFT_POSTGRESQL_DB_USERNAME	POSTGRESQL_USER
OPENSIFT_POSTGRESQL_DB_PASSWORD	POSTGRESQL_PASSWORD
OPENSIFT_POSTGRESQL_DB_LOG_DIR	
OPENSIFT_POSTGRESQL_DB_PID	
OPENSIFT_POSTGRESQL_DB_SOCKET_DIR	
OPENSIFT_POSTGRESQL_DB_URL	
OPENSIFT_POSTGRESQL_VERSION	
OPENSIFT_POSTGRESQL_SHARED_BUFFERS	

v2	v3
OPENSIFT_POSTGRESQL_MAX_CONNECTIONS	
OPENSIFT_POSTGRESQL_MAX_PREPARED_TRANSACTIONS	
OPENSIFT_POSTGRESQL_DATESTYLE	
OPENSIFT_POSTGRESQL_LOCALE	
OPENSIFT_POSTGRESQL_CONFIG	
OPENSIFT_POSTGRESQL_SSL_ENABLED	
	POSTGRESQL_DATABASE
	POSTGRESQL_ADMIN_PASSWORD

6.2.5. MongoDB



NOTE

- For OpenShift v3: MongoDB shell version 3.2.6
- For OpenShift v2: MongoDB shell version 2.4.9

1. Remotely access the v2 application via the **ssh** command:

```
$ rhc ssh <v2_application_name>
```

2. Run **mongodump**, specifying a single database with **-d <database_name> -c <collections>**. Without those options, dump all databases. Each database is dumped in its own directory:

```
$ mongodump -h $OPENSIFT_MONGODB_DB_HOST -o app-root/repo/mydbdump
-u 'admin' -p $OPENSIFT_MONGODB_DB_PASSWORD
$ cd app-root/repo/mydbdump/<database_name>; tar -cvzf dbname.tar.gz
$ exit
```

3. Download **dbdump** to a local machine in the **mongodump** directory:

```
$ mkdir mongodump
$ rhc scp -a <v2 appname> download mongodump \
app-root/repo/mydbdump/<dbname>/dbname.tar.gz
```

4. Start a MongoDB pod in v3. Because the latest image (3.2.6) does not include **mongo-tools**, to use **mongorestore** or **mongoimport** commands you need to edit the default **mongodb-persistent** template to specify the image tag that contains the **mongo-tools**,

“mongodb:2.4”. For that reason, the following **oc export** command and edit are necessary:

```
$ oc export template mongodb-persistent -n openshift -o json >
mongodb-24persistent.json
```

Edit L80 of **mongodb-24persistent.json**; replace **mongodb:latest** with **mongodb:2.4**.

```
$ oc new-app --template=mongodb-persistent -n <project-name-that-
template-was-created-in> \
  MONGODB_USER=user_from_v2_app -p \
  MONGODB_PASSWORD=password_from_v2_db -p \
  MONGODB_DATABASE=v2_dbname -p \
  MONGODB_ADMIN_PASSWORD=password_from_v2_db
$ oc get pods
```

- When the mongodb pod is up and running, copy the database archive files to the v3 MongoDB pod:

```
$ oc rsync local/path/to/mongodump
<mongodb_pod_name>:/var/lib/mongodb/data
$ oc rsh <mongodb_pod>
```

- In the MongoDB pod, complete the following for each database you want to restore:

```
$ cd /var/lib/mongodb/data/mongodump
$ tar -xzf dbname.tar.gz
$ mongorestore -u $MONGODB_USER -p $MONGODB_PASSWORD -d dbname -v
/var/lib/mongodb/data/mongodump
```

- Check if the database is restored:

```
$ mongo admin -u $MONGODB_USER -p $MONGODB_ADMIN_PASSWORD
$ use dbname
$ show collections
$ exit
```

- Remove the **mongodump** directory from the pod:

```
$ rm -rf /var/lib/mongodb/data/mongodump
```

Supported MongoDB Environment Variables

v2	v3
OPENSIFT_MONGODB_DB_HOST	[service_name]_SERVICE_HOST
OPENSIFT_MONGODB_DB_PORT	[service_name]_SERVICE_PORT
OPENSIFT_MONGODB_DB_USERNAME	MONGODB_USER
OPENSIFT_MONGODB_DB_PASSWORD	MONGODB_PASSWORD

v2	v3
OPENSIFT_MONGODB_DB_URL	
OPENSIFT_MONGODB_DB_LOG_DIR	
	MONGODB_DATABASE
	MONGODB_ADMIN_PASSWORD
	MONGODB_NOPREALLOC
	MONGODB_SMALLFILES
	MONGODB_QUIET
	MONGODB_REPLICA_NAME
	MONGODB_KEYFILE_VALUE

6.3. MIGRATING WEB FRAMEWORK APPLICATIONS

6.3.1. Overview

This topic reviews how to migrate Python, Ruby, PHP, Perl, Node.js, JBoss EAP, JBoss WS (Tomcat), and Wildfly 10 (JBoss AS) web framework applications from OpenShift version 2 (v2) to OpenShift version 3 (v3).

6.3.2. Python

1. Set up a new GitHub repository and add it as a remote branch to the current, local v2 Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>.git
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. Ensure that all important files such as *setup.py*, *wsgi.py*, *requirements.txt*, and *etc* are pushed to new repository.

- Ensure all required packages for your application are included in *requirements.txt*.

4. Use the `oc` command to launch a new Python application from the builder image and source code:

```
$ oc new-app --strategy=source
```

```
python:3.3~https://github.com/<github-id>/<repo-name> --name=<app-name> -e
<ENV_VAR_NAME>=<env_var_value>
```

Supported Python Versions

v2	v3
Python: 2.6, 2.7, 3.3	Supported Container Images
Django	Django-psql-example (quickstart)

6.3.3. Ruby

1. Set up a new GitHub repository and add it as a remote branch to the current, local v2 Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>.git
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. If you do not have a Gemfile and are running a simple rack application, copy this Gemfile into the root of your source:

```
https://github.com/openshift/ruby-ex/blob/master/Gemfile
```



NOTE

The latest version of the **rack** gem that supports Ruby 2.0 is 1.6.4, so the Gemfile needs to be modified to `gem 'rack', "1.6.4"`.

For Ruby 2.2 or later, use the **rack** gem 2.0 or later.

4. Use the **oc** command to launch a new Ruby application from the builder image and source code:

```
$ oc new-app --strategy=source
ruby:2.0~https://github.com/<github-id>/<repo-name>.git
```

Supported Ruby Versions

v2	v3
Ruby: 1.8, 1.9, 2.0	Supported Container Images
Ruby on Rails: 3, 4	Rails-postgresql-example (quickstart)

v2	v3
Sinatra	

6.3.4. PHP

1. Set up a new GitHub repository and add it as a remote branch to the current, local v2 Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. Use the `oc` command to launch a new PHP application from the builder image and source code:

```
$ oc new-app https://github.com/<github-id>/<repo-name>.git
--name=<app-name> -e <ENV_VAR_NAME>=<env_var_value>
```

Supported PHP Versions

v2	v3
PHP: 5.3, 5.4	Supported Container Images
PHP 5.4 with Zend Server 6.1	
CodeIgniter 2	
HHVM	
Laravel 5.0	
	cakephp-mysql-example (quickstart)

6.3.5. Perl

1. Set up a new GitHub repository and add it as a remote branch to the current, local v2 Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. Edit the local Git repository and push changes upstream to make it v3 compatible:

- a. In v2, CPAN modules reside in `.openshift/cpan.txt`. In v3, the s2i builder looks for a file named `cpanfile` in the root directory of the source.

```
$ cd <local-git-repository>
$ mv .openshift/cpan.txt cpanfile
```

Edit `cpanfile`, as it has a slightly different format:

format of <code>cpanfile</code>	format of <code>cpan.txt</code>
requires 'cpan::mod';	cpan::mod
requires 'Dancer';	Dancer
requires 'YAML';	YAML

- b. Remove `.openshift` directory



NOTE

In v3, `action_hooks` and `cron` tasks are not supported in the same way. See [Action Hooks](#) for more information.

4. Use the `oc` command to launch a new Perl application from the builder image and source code:

```
$ oc new-app https://github.com/<github-id>/<repo-name>.git
```

Supported Perl Versions

v2	v3
Perl: 5.10	Supported Container Images
	Dancer-mysql-example (quickstart)

6.3.6. Node.js

1. Set up a new GitHub repository and add it as a remote branch to the current, local Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. Edit the local Git repository and push changes upstream to make it v3 compatible:

- a. Remove the **.openshift** directory.



NOTE

In v3, **action_hooks** and **cron** tasks are not supported in the same way. See [Action Hooks](#) for more information.

- b. Edit **server.js**.

- L116 server.js: 'self.app = express();'
- L25 server.js: self.ipaddress = '0.0.0.0';
- L26 server.js: self.port = 8080;



NOTE

Lines(L) are from the base V2 cartridge **server.js**.

4. Use the **oc** command to launch a new Node.js application from the builder image and source code:

```
$ oc new-app https://github.com/<github-id>/<repo-name>.git
--name=<app-name> -e <ENV_VAR_NAME>=<env_var_value>
```

Supported Node.js Versions

v2	v3
Node.js 0.10	Supported Container Images
	Nodejs-mongodb-example. This quickstart template only supports Node.js version 6.

6.3.7. JBoss EAP

1. Set up a new GitHub repository and add it as a remote branch to the current, local Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. If the repository includes pre-built **.war** files, they need to reside in the **deployments** directory off the root directory of the repository.
4. Create the new application using the JBoss EAP 6 builder image (jboss-eap64-openshift) and

the source code repository from GitHub:

```
$ oc new-app --strategy=source jboss-eap64-openshift~https://github.com/<github-id>/<repo-name>.git
```

6.3.8. JBoss WS (Tomcat)

1. Set up a new GitHub repository and add it as a remote branch to the current, local Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. If the repository includes pre-built **.war** files, they need to reside in the **deployments** directory off the root directory of the repository.
4. Create the new application using the JBoss Web Server 3 (Tomcat 7) builder image (jboss-webserver30-tomcat7) and the source code repository from GitHub:

```
$ oc new-app --strategy=source  
jboss-webserver30-tomcat7-openshift~https://github.com/<github-id>/<repo-name>.git  
--name=<app-name> -e <ENV_VAR_NAME>=<env_var_value>
```

6.3.9. JBoss AS (Wildfly 10)

1. Set up a new GitHub repository and add it as a remote branch to the current, local Git repository:

```
$ git remote add <remote-name> https://github.com/<github-id>/<repo-name>
```

2. Push the local v2 source code to the new repository:

```
$ git push -u <remote-name> master
```

3. Edit the local Git repository and push the changes upstream to make it v3 compatible:

- a. Remove **.openshift** directory.




NOTE

In v3, **action_hooks** and **cron** tasks are not supported in the same way. See [Action Hooks](#) for more information.

- b. Add the **deployments** directory to the root of the source repository. Move the **.war** files to 'deployments' directory.

4. Use the `oc` command to launch a new Wildfly application from the builder image and source code:

```
$ oc new-app https://github.com/<github-id>/<repo-name>.git
--image-stream="openshift/wildfly:10.0" --name=<app-name> -e
<ENV_VAR_NAME>=<env_var_value>
```



NOTE

The argument `--name` is optional to specify the name of your application. The argument `-e` is optional to add environment variables that are needed for build and deployment processes, such as `OPENSIFT_PYTHON_DIR`.

6.3.10. Supported JBoss/XPaas Versions

v2	v3
JBoss App Server 7	
Tomcat 6 (JBoss EWS 1.0)	Supported Container Images
Tomcat 7 (JBoss EWS 2.0)	Supported Container Images
Vert.x 2.1	
WildFly App Server 10	
WildFly App Server 8.2.1.Final	
WildFly App Server 9	
CapeDwarf	
JBoss Data Virtualization 6	Supported Container Images
JBoss Enterprise App Platform (EAP) 6	Supported Container Images
JBoss Unified Push Server 1.0.0.Beta1, Beta2	
JBoss BPM Suite	Supported Container Images
JBoss BRMS	Supported Container Images
	jboss-eap70-openshift: 1.3-Beta
	eap64-https-s2i
	eap64-mongodb-persistent-s2i

v2	v3
	eap64-mysql-persistent-s2i
	eap64-psql-persistent-s2i

6.4. QUICKSTART EXAMPLES

6.4.1. Overview

Although there is no clear-cut migration path for v2 quickstart to v3 quickstart, the following quickstarts are currently available in v3. If you have an application with a database, rather than using `oc new-app` to create your application, then `oc new-app` again to start a separate database service and linking the two with common environment variables, you can use one of the following to instantiate the linked application and database at once, from your GitHub repository containing your source code. You can list all available templates with `oc get templates -n openshift`:

- CakePHP MySQL <https://github.com/openshift/cakephp-ex>
 - template: cakephp-mysql-example
- Node.js MongoDB <https://github.com/openshift/nodejs-ex>
 - template: nodejs-mongodb-example
- Django PostgreSQL <https://github.com/openshift/django-ex>
 - template: django-psql-example
- Dancer MySQL <https://github.com/openshift/dancer-ex>
 - template: dancer-mysql-example
- Rails PostgreSQL <https://github.com/openshift/rails-ex>
 - template: rails-postgresql-example

6.4.2. Workflow

Run a `git clone` of one of the above template URLs locally. Add and commit your application source code and push a GitHub repository, then start a v3 quickstart application from one of the templates listed above:

1. Create a GitHub repository for your application.
2. Clone a quickstart template and add your GitHub repository as a remote:

```
$ git clone <one-of-the-template-URLs-listed-above>
$ cd <your local git repository>
$ git remote add upstream <https://github.com/<git-id>/<quickstart-repo>.git>
$ git push -u upstream master
```

3. Commit and push your source code to GitHub:

```
$ cd <your local repository>
$ git commit -am "added code for my app"
$ git push origin master
```

4. Create a new application in v3:

```
$ oc new-app --template=<template> \
-p SOURCE_REPOSITORY_URL=<https://github.com/<git-id>/<quickstart_repo>.git> \
-p DATABASE_USER=<your_db_user> \
-p DATABASE_NAME=<your_db_name> \
-p DATABASE_PASSWORD=<your_db_password> \
-p DATABASE_ADMIN_PASSWORD=<your_db_admin_password> ❶
```

❶ Only applicable for MongoDB.

You should now have 2 pods running, a web framework pod, and a database pod. The web framework pod environment should match the database pod environment. You can list the environment variables with `oc set env pod/<pod_name> --list`:

- **DATABASE_NAME** is now **<DB_SERVICE>_DATABASE**
- **DATABASE_USER** is now **<DB_SERVICE>_USER**
- **DATABASE_PASSWORD** is now **<DB_SERVICE>_PASSWORD**
- **DATABASE_ADMIN_PASSWORD** is now **MONGODB_ADMIN_PASSWORD** (only applicable for MongoDB)
If no **SOURCE_REPOSITORY_URL** is specified, the template will use the template URL (<https://github.com/openshift/<quickstart>-ex>) listed above as the source repository, and a **hello-welcome** application will be started.

5. If you are migrating a database, export databases to a dump file and restore the database in the new v3 database pod. Refer to the steps outlined in [Database Applications](#), skipping the `oc new-app` step as the database pod is already up and running.

6.5. CONTINUOUS INTEGRATION AND DEPLOYMENT (CI/CD)

6.5.1. Overview

This topic reviews the differences in continuous integration and deployment (CI/CD) applications between OpenShift version 2 (v2) and OpenShift version 3 (v3) and how to migrate these applications into the v3 environment.

6.5.2. Jenkins

The Jenkins applications in OpenShift version 2 (v2) and OpenShift version 3 (v3) are configured differently due to fundamental differences in architecture. For example, in v2, the application uses an integrated Git repository that is hosted in the gear to store the source code. In v3, the source code is located in a public or private Git repository that is hosted outside of the pod.

Furthermore, in OpenShift v3, Jenkins jobs can not only be triggered by source code changes, but also by changes in ImageStream, which are changes on the images that are used to build the application

along with its source code. As a result, it is highly recommended that you migrate the Jenkins application manually by creating a new Jenkins application in v3, and then re-creating jobs with the configurations that are suitable to OpenShift v3 environment.

Consult these resources for more information on how to create a Jenkins application, configure jobs, and use Jenkins plug-ins properly:

- <https://github.com/openshift/origin/blob/master/examples/jenkins/README.md>
- <https://github.com/openshift/jenkins-plugin/blob/master/README.md>
- <https://github.com/openshift/origin/blob/master/examples/sample-app/README.md>

6.6. WEBHOOKS AND ACTION HOOKS

6.6.1. Overview

This topic reviews the differences in webhooks and action hooks between OpenShift version 2 (v2) and OpenShift version 3 (v3) and how to migrate these applications into the v3 environment.

6.6.2. Webhooks

1. After creating a **BuildConfig** from a GitHub repository, run:

```
$ oc describe bc/<name-of-your-BuildConfig>
```

This will output a webhook GitHub URL that looks like:

```
<https://api.starter-us-east-1.openshift.com:443/oapi/v1/namespaces/nsname/buildconfigs/bcname/webhooks/secret/github>.
```

2. Cut and paste this URL into GitHub, from the GitHub web console.
3. In your GitHub repository, select **Add Webhook** from **Settings** → **Webhooks & Services**.
4. Paste the URL output (similar to above) into the **Payload URL** field.
5. Set the **Content Type** to **application/json**.
6. Click **Add webhook**.

You should see a message from GitHub stating that your webhook was successfully configured.

Now, whenever you push a change to your GitHub repository, a new build will automatically start, and upon a successful build a new deployment will start.



NOTE

If you delete or recreate your application, you will have to update the **Payload URL** field in GitHub with the new **BuildConfig** webhook url.

6.6.3. Action Hooks

In OpenShift version 2 (v2), there are `build`, `deploy`, `post_deploy`, and `pre_build` scripts or `action_hooks` that are located in the `.openshift/action_hooks` directory. While there is no one-to-one mapping of function for these in v3, the [S2I tool](#) in v3 does have the option of adding [customizable scripts](#), either in a designated URL or in the `.s2i/bin` directory of your source repository.

OpenShift version 3 (v3) also offers a [post-build hook](#) for running basic testing of an image after it is built and before it is pushed to the registry. [Deployment hooks](#) are configured in the deployment configuration.

In v2, `action_hooks` are commonly used to set up environment variables. In v2, any environment variables should be passed with:

```
$ oc new-app <source-url> -e ENV_VAR=env_var
```

or:

```
$ oc new-app <template-name> -p ENV_VAR=env_var
```

Also, environment variables can be added or changed using:

```
$ oc set env dc/<name-of-dc>
ENV_VAR1=env_var1 ENV_VAR2=env_var2'
```

6.7. S2I TOOL

6.7.1. Overview

The [Source-to-Image \(S2I\) tool](#) injects application source code into a container image and the final product is a new and ready-to-run container image that incorporates the builder image and built source code. The S2I tool can be installed on your local machine without OpenShift Container Platform from [the repository](#).

The S2I tool is a very powerful tool to test and verify your application and images locally before using them on OpenShift Container Platform.

6.7.2. Creating a Container Image

1. Identify the builder image that is needed for the application. Red Hat offers multiple builder images for different languages including [Python](#), [Ruby](#), [Perl](#), [PHP](#), and [Node.js](#). Other images are available from [the community space](#).
2. S2I can build images from source code in a local file system or from a Git repository. To build a new container image from the builder image and the source code:

```
$ s2i build <source-location> <builder-image-name> <output-image-name>
```



NOTE

`<source-location>` can either be a Git repository URL or a directory to source code in a local file system.

3. Test the built image with the Docker daemon:

```
$ docker run -d --name <new-name> -p <port-number>:<port-number>  
<output-image-name>  
$ curl localhost:<port-number>
```

4. Push the new image to the [OpenShift registry](#).

5. Create a new application from the image in the OpenShift registry using the `oc` command:

```
$ oc new-app <image-name>
```

6.8. SUPPORT GUIDE

6.8.1. Overview

This topic reviews supported languages, frameworks, databases, and markers for OpenShift version 2 (v2) and OpenShift version 3 (v3).

6.8.2. Supported Databases

See the [Supported Databases](#) section of the Database Applications topic.

6.8.3. Supported Languages

- [PHP](#)
- [Python](#)
- [Perl](#)
- [Node.js](#)
- [Ruby](#)
- [JBoss/xPaaS](#)

6.8.4. Supported Frameworks

Table 6.1. Supported Frameworks

v2	v3
Jenkins Server	jenkins-persistent
Drupal 7	
Ghost 0.7.5	
WordPress 4	

v2	v3
Ceylon	
Go	
MEAN	

6.8.5. Supported Markers

Table 6.2. Python

v2	v3
pip_install	If your repository contains <i>requirements.txt</i> , then pip is invoked by default. Otherwise, pip is not used.

Table 6.3. Ruby

v2	v3
disable_asset_compilation	This can be done by setting DISABLE_ASSET_COMPILATION environment variable to true on the buildconfig strategy definition.

Table 6.4. Perl

v2	v3
enable_cpan_tests	This can be done by setting ENABLE_CPAN_TEST environment variable to true on the build configuration .

Table 6.5. PHP

v2	v3
use_composer	composer is always used if the source repository includes a <i>composer.json</i> in the root directory.

Table 6.6. Node.js

v2	v3
NODEJS_VERSION	N/A

v2	v3
use_npm	npm is always used to start the application, unless DEV_MODE is set to true , in which case nodemon is used instead.

Table 6.7. JBoss EAP, JBoss WS, WildFly

v2	v3
enable_debugging	This option is controlled via the ENABLE_JPDA environment variable set on the deployment configuration by setting it to any non-empty value.
skip_maven_build	If <i>pom.xml</i> is present, maven will be run.
java7	N/A
java8	JavaEE is using JDK8.

Table 6.8. Jenkins

v2	v3
enable_debugging	N/A

Table 6.9. All

v2	v3
force_clean_build	There is a similar concept in v3, as noCache field in buildconfig forces the container build to rerun each layer. In the S2I build, the incremental flag is false by default, which indicates a clean build .
hot_deploy	Ruby , Python , Perl , PHP , Node.js
enable_public_server_status	N/A
disable_auto_scaling	Autoscaling is off by default and it can be turn on via pod auto-scaling .

6.8.6. Supported Environment Variables

- [MySQL](#)
- [MongoDB](#)

- [PostgreSQL](#)

CHAPTER 7. APPLICATION TUTORIALS

7.1. OVERVIEW

This topic group includes information on how to get your application up and running in OpenShift Container Platform and covers different languages and their frameworks.

7.2. QUICKSTART TEMPLATES

7.2.1. Overview

A quickstart is a basic example of an application running on OpenShift Container Platform. Quickstarts come in a variety of languages and frameworks, and are defined in a [template](#), which is constructed from a set of services, build configurations, and deployment configurations. This template references the necessary images and source repositories to build and deploy the application.

To explore a quickstart, create an application from a template. Your administrator may have already installed these templates in your OpenShift Container Platform cluster, in which case you can simply select it from the web console. See the [template](#) documentation for more information on how to upload, create from, and modify a template.

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

7.2.2. Web Framework Quickstart Templates

These quickstarts provide a basic application of the indicated framework and language:

- CakePHP: a PHP web framework (includes a MySQL database)
 - [Template definition](#)
 - [Source repository](#)
- Django: a Python web framework (includes a PostgreSQL database)
 - [Template definition](#)
 - [Source repository](#)
- NodeJS: a NodeJS web application (includes a MongoDB database)
 - [Template definition](#)
 - [Source repository](#)

- Rails: a Ruby web framework (includes a PostgreSQL database)
 - [Template definition](#)
 - [Source repository](#)

7.3. RUBY ON RAILS

7.3.1. Overview

Ruby on Rails is a popular web framework written in [Ruby](#). This guide covers using Rails 4 on OpenShift Container Platform.



WARNING

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

For this guide you will need:

- Basic Ruby/Rails knowledge
- Locally installed version of Ruby 2.0.0+, Rubygems, Bundler
- Basic Git knowledge
- Running instance of OpenShift Container Platform v3

7.3.2. Local Workstation Setup

First make sure that an instance of OpenShift Container Platform is running and is available. For more info on how to get OpenShift Container Platform up and running check the [installation methods](#). Also make sure that your [oc CLI client is installed](#) and the command is accessible from your command shell, so you can use it to [log in](#) using your email address and password.

7.3.2.1. Setting Up the Database

Rails applications are almost always used with a database. For the local development we chose the PostgreSQL database. To install it type:

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

Next you need to initialize the database with:

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

This command will create the `/var/lib/pgsql/data` directory, in which the data will be stored.

Start the database by typing:

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

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

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

Note that the user we created has no password.

7.3.3. Writing Your Application

If you are starting your Rails application from scratch, you need to install the Rails gem first.

```
$ gem install rails
Successfully installed rails-4.2.0
1 gem installed
```

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

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

Then change into your new application directory.

```
$ cd rails-app
```

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

```
gem 'pg'
```

To generate a new **Gemfile.lock** with all your dependencies run:

```
$ bundle install
```

In addition to using the **postgresql** database with the **pg** gem, you'll also need to ensure the **config/database.yml** is using the **postgresql** adapter.

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

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

Create your application's development and test databases by using this **rake** command:

```
$ rake db:create
```

This will create **development** and **test** database in your PostgreSQL server.

7.3.3.1. Creating a Welcome Page

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

In order to have a custom welcome page we need to do following steps:

- Create a **controller** with an index action
- Create a **view** page for the **welcome** controller **index** action
- Create a **route** that will serve applications root page with the created **controller** and **view**

Rails offers a generator that will do all this necessary steps for you.

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

All the necessary files have been created, now we just need to edit line 2 in **config/routes.rb** file to look like:

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

Run the rails server to verify the page is available.

```
$ rails server
```

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

7.3.3.2. Configuring the Application for OpenShift Container Platform

In order to have your application communicating with the PostgreSQL database service that will be running in OpenShift Container Platform, you will need to edit the **default** section in your **config/database.yml** to use **environment variables**, which you will define later, upon the database service creation.

The **default** section in your edited **config/database.yml** together with pre-defined variables should look like:

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

default: &default
  adapter: postgresql
  encoding: unicode
  # For details on connection pooling, see rails configuration guide
  # http://guides.rubyonrails.org/configuring.html#database-pooling
```

```
pool: <%= ENV["POSTGRESQL_MAX_CONNECTIONS"] || 5 %>
username: <%= user %>
password: <%= password %>
host: <%= ENV["#{db_service}_SERVICE_HOST"] %>
port: <%= ENV["#{db_service}_SERVICE_PORT"] %>
database: <%= ENV["POSTGRESQL_DATABASE"] %>
```

For an example of how the final file should look, see [Ruby on Rails example application config/database.yml](#).

7.3.3.3. Storing Your Application in Git

OpenShift Container Platform requires [git](#), if you don't have it installed you will need to install it.

Building an application in OpenShift Container Platform usually requires that the source code be stored in a [git](#) repository, so you will need to install [git](#) if you do not already have it.

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

```
$ ls -1
app
bin
config
config.ru
db
Gemfile
Gemfile.lock
lib
log
public
Rakefile
README.rdoc
test
tmp
vendor
```

Now run these commands in your Rails app directory to initialize and commit your code to git:

```
$ git init
$ git add .
$ git commit -m "initial commit"
```

Once your application is committed you need to push it to a remote repository. For this you would need a [GitHub account](#), in which you [create a new repository](#).

Set the remote that points to your [git](#) repository:

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

After that, push your application to your remote git repository.

```
$ git push
```

7.3.4. Deploying Your Application to OpenShift Container Platform

To deploy your Ruby on Rails application, create a new [Project](#) for the application:

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

After creating the the `rails-app` [project](#), you will be automatically switched to the new project namespace.

Deploying your application in OpenShift Container Platform involves three steps:

- Creating a database [service](#) from OpenShift Container Platform's [PostgreSQL image](#)
- Creating a frontend [service](#) from OpenShift Container Platform's [Ruby 2.0 builder image](#) and your Ruby on Rails source code, which we wire with the database service
- Creating a route for your application.

7.3.4.1. Creating the Database Service

Your Rails application expects a running database [service](#). For this service use [PostgreSQL database image](#).

To create the database [service](#) you will use the `oc new-app` command. To this command you will need to pass some necessary [environment variables](#) which will be used inside the database container. These [environment variables](#) are required to set the username, password, and name of the database. You can change the values of these [environment variables](#) to anything you would like. The variables we are going to be setting are as follows:

- `POSTGRESQL_DATABASE`
- `POSTGRESQL_USER`
- `POSTGRESQL_PASSWORD`

Setting these variables ensures:

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

For example:

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

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

```
-e POSTGRESQL_ADMIN_PASSWORD=admin_pw
```

To watch the progress of this command:

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

7.3.4.2. Creating the Frontend Service

To bring your application to OpenShift Container Platform, you need to specify a repository in which your application lives, using once again the `oc new-app` command, in which you will need to specify database related [environment variables](#) we setup in the [Creating the Database Service](#):

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

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

You can verify the environment variables have been added by viewing the JSON document of the **rails-app** [DeploymentConfig](#):

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

You should see the following section:

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

To check the build process:

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

Once the build is complete, you can look at the running pods in OpenShift Container Platform.

```
$ oc get pods
```

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

Before your application will be functional, you need to initialize the database by running the database migration script. There are two ways you can do this:

- Manually from the running frontend container:

First you need to exec into frontend container with `rsh` command:

```
$ oc rsh <FRONTEND_POD_ID>
```

Run the migration from inside the container:

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

If you are running your Rails application in a `development` or `test` environment you don't have to specify the `RAILS_ENV` environment variable.

- By adding pre-deployment [lifecycle hooks](#) in your template. For example check the [hooks example](#) in our [Rails example](#) application.

7.3.4.3. Creating a Route for Your Application

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

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



WARNING

It's the user's responsibility to ensure the hostname they specify resolves into the IP address of the router. For more information, check the OpenShift Container Platform documentation on:

- [Routes](#)
- [Configuring a Highly-available Routing Service](#)

7.4. SETTING UP A NEXUS MIRROR FOR MAVEN

7.4.1. Introduction

While developing your application with Java and Maven, you will most likely be building many times. In order to shorten the build times of your pods, Maven dependencies can be cached in a local Nexus repository. This tutorial will guide you through creating a Nexus repository on your cluster.

This tutorial assumes that you are working with a project that is already set up for use with Maven. If you are interested in using Maven with your Java project, it is highly recommended that you look at [their guide](#).

In addition, be sure to check your application's image for Maven mirror capabilities. Many images that use Maven have a `MAVEN_MIRROR_URL` environment variable that you can use to simplify this process. If it does not have this capability, read [the Nexus documentation](#) to configure your build properly.

Furthermore, make sure that you give each pod enough resources to function. You may have to [edit the pod template](#) in the Nexus deployment configuration to request more resources.

7.4.2. Setting up Nexus

1. Download and deploy the official Nexus container image:

```
oc new-app sonatype/nexus
```

2. Create a route by exposing the newly created Nexus service:

```
oc expose svc/nexus
```

3. Use `oc get routes` to find the pod's new external address.

```
oc get routes
```

The output should resemble:

NAME	HOST/PORT	PATH
SERVICES	PORT	TERMINATION
nexus	nexus-myproject.192.168.1.173.xip.io	nexus
8081-tcp		

4. Confirm that Nexus is running by navigating your browser to the URL under **HOST/PORT**. To sign in to Nexus, the default administrator username is `admin`, and the password is `admin123`.



NOTE

Nexus comes pre-configured for the Central Repository, but you may need others for your application. For many Red Hat images, it is recommended to [add the jboss-ga repository](#) at [Maven repository](#).

7.4.2.1. Using Probes to Check for Success

This is a good time to set up [readiness and liveness probes](#). These will periodically check to see that Nexus is running properly.

```
$ oc set probe dc/nexus \
  --liveness \
  --failure-threshold 3 \
  --initial-delay-seconds 30 \
  -- echo ok
$ oc set probe dc/nexus \
  --readiness \
  --failure-threshold 3 \
  --initial-delay-seconds 30 \
  --get-url=http://:8081/nexus/content/groups/public
```

7.4.2.2. Adding Persistence to Nexus



NOTE

If you do not want persistent storage, continue to [Connecting to Nexus](#). However, your cached dependencies and any configuration customization will be lost if the pod is restarted for any reason.

Create a persistent volume claim (PVC) for Nexus, so that the cached dependencies are not lost when the pod running the server terminates. PVCs require available persistent volumes (PV) in the cluster. If there are no PVs available and you do not have administrator access on your cluster, ask your system administrator to create a Read/Write Persistent Volume for you. Otherwise, see [Persistent Storage in OpenShift Container Platform](#) for instructions on creating a persistent volume.

Add a PVC to the Nexus deployment configuration.

```
$ oc volumes dc/nexus --add \
  --name 'nexus-volume-1' \
  --type 'pvc' \
  --mount-path '/sonatype-work/' \
  --claim-name 'nexus-pv' \
  --claim-size '1G' \
  --overwrite
```

This removes the previous **emptyDir** volume for the deployment config and adds a claim for one gigabyte of persistent storage mounted at **/sonatype-work**, which is where the dependencies will be stored. Due to the change in configuration, the Nexus pod will be redeployed automatically.

To verify that Nexus is running, refresh the Nexus page in your browser. You can monitor the deployment's progress using:

```
$ oc get pods -w
```

7.4.3. Connecting to Nexus

The next steps demonstrate defining a build that uses the new Nexus repository. The rest of the tutorial uses [this example repository](#) with **wildfly-100-centos7** as a builder, but these changes should work for any project.

The [example builder image](#) supports **MAVEN_MIRROR_URL** as part of its environment, so we can use this to point our builder image to our Nexus repository. If your image does not support consuming an environment variable to configure a Maven mirror, you may need to modify the builder image to provide the correct Maven settings to point to the Nexus mirror.

```
$ oc new-build openshift/wildfly-100-
centos7:latest~https://github.com/openshift/jee-ex.git \
  -e MAVEN_MIRROR_URL='http://nexus.
<Nexus_Project>:8081/nexus/content/groups/public'
$ oc logs build/jee-ex-1 --follow
```

Replace **<Nexus_Project>** with the project name of the Nexus repository. If it is in the same project as the application that is using it, you can remove the **<Nexus_Project>**.. [Learn more about DNS resolution in OpenShift Container Platform](#).

7.4.4. Confirming Success

In your web browser, navigate to `http://<NexusIP>:8081/nexus/content/groups/public` to confirm that it has stored your application's dependencies. You can also check the build logs to see if Maven is using the Nexus mirror. If successful, you should see output referencing the URL `http://nexus:8081`.

7.4.5. Additional Resources

- [Managing Volumes in OpenShift Container Platform](#)
- [Improving Build Time of Java Builds on OpenShift Container Platform](#)
- [Nexus Repository Documentation](#)

CHAPTER 8. BUILDS

8.1. HOW BUILDS WORK

8.1.1. What Is a Build?

A *build* in OpenShift Container Platform is the process of transforming input parameters into a resulting object. Most often, builds are used to transform source code into a runnable container image.

A *build configuration*, or **BuildConfig**, is characterized by a *build strategy* and one or more sources. The strategy determines the aforementioned process, while the sources provide its input.

The build strategies are:

- Source-to-Image (S2I) ([description](#), [options](#))
- Pipeline ([description](#), [options](#))
- Docker ([description](#), [options](#))
- Custom ([description](#), [options](#))

And there are six types of sources that can be given as *build input*:

- [Git](#)
- [Dockerfile](#)
- [Binary](#)
- [Image](#)
- [Input secrets](#)
- [External artifacts](#)

It is up to each build strategy to consider or ignore a certain type of source, as well as to determine how it is to be used. Binary and Git are mutually exclusive source types. Dockerfile and Image can be used by themselves, with each other, or together with either Git or Binary. The Binary source type is unique from the other options in [how it is specified to the system](#).

8.1.2. What Is a BuildConfig?

A build configuration describes a single build definition and a set of [triggers](#) for when a new build should be created. Build configurations are defined by a **BuildConfig**, which is a REST object that can be used in a POST to the API server to create a new instance.

Depending on how you choose to create your application using OpenShift Container Platform, a **BuildConfig** is typically generated automatically for you if you use the web console or CLI, and it can be edited at any time. Understanding the parts that make up a **BuildConfig** and their available options can help if you choose to manually tweak your configuration later.

The following example **BuildConfig** results in a new build every time a container image tag or the source code changes:

BuildConfig Object Definition

```
kind: "BuildConfig"
apiVersion: "v1"
metadata:
  name: "ruby-sample-build" ❶
spec:
  runPolicy: "Serial" ❷
  triggers: ❸
  -
    type: "GitHub"
    github:
      secret: "secret101"
  - type: "Generic"
    generic:
      secret: "secret101"
  -
    type: "ImageChange"
  source: ❹
    git:
      uri: "https://github.com/openshift/ruby-hello-world"
  strategy: ❺
    sourceStrategy:
      from:
        kind: "ImageStreamTag"
        name: "ruby-20-centos7:latest"
  output: ❻
    to:
      kind: "ImageStreamTag"
      name: "origin-ruby-sample:latest"
  postCommit: ❼
    script: "bundle exec rake test"
```

❶ This specification will create a new **BuildConfig** named **ruby-sample-build**.

❷ The **runPolicy** field controls whether builds created from this build configuration can be run simultaneously. The default value is **Serial**, which means new builds will run sequentially, not simultaneously.

❸ You can specify a list of **triggers**, which cause a new build to be created.

❹ The **source** section defines the source of the build. The source type determines the primary source of input, and can be either **Git**, to point to a code repository location, **Dockerfile**, to build from an inline Dockerfile, or **Binary**, to accept binary payloads. It is possible to have multiple sources at once, refer to the documentation for each source type for details.

❺ The **strategy** section describes the build strategy used to execute the build. You can specify a **Source**, **Docker**, or **Custom** strategy here. This above example uses the **ruby-20-centos7** container image that Source-To-Image will use for the application build.

❻ After the container image is successfully built, it will be pushed into the repository described in the **output** section.

❼ The **postCommit** section defines an optional **build hook**.

8.2. BASIC BUILD OPERATIONS

8.2.1. Starting a Build

Manually start a new build from an existing build configuration in your current project using the following command:

```
$ oc start-build <buildconfig_name>
```

Re-run a build using the **--from-build** flag:

```
$ oc start-build --from-build=<build_name>
```

Specify the **--follow** flag to stream the build's logs in stdout:

```
$ oc start-build <buildconfig_name> --follow
```

Specify the **--env** flag to set any desired environment variable for the build:

```
$ oc start-build <buildconfig_name> --env=<key>=<value>
```

Rather than relying on a Git source pull or a Dockerfile for a build, you can also start a build by directly pushing your source, which could be the contents of a Git or SVN working directory, a set of prebuilt binary artifacts you want to deploy, or a single file. This can be done by specifying one of the following options for the **start-build** command:

Option	Description
--from-dir=<directory>	Specifies a directory that will be archived and used as a binary input for the build.
--from-file=<file>	Specifies a single file that will be the only file in the build source. The file is placed in the root of an empty directory with the same file name as the original file provided.
--from-repo=<local_source_repo>	Specifies a path to a local repository to use as the binary input for a build. Add the --commit option to control which branch, tag, or commit is used for the build.

When passing any of these options directly to the build, the contents are streamed to the build and override the current build source settings.



NOTE

Builds triggered from binary input will not preserve the source on the server, so rebuilds triggered by base image changes will use the source specified in the build configuration.

For example, the following command sends the contents of a local Git repository as an archive from the tag **v2** and starts a build:

```
$ oc start-build hello-world --from-repo=../hello-world --commit=v2
```

8.2.2. Canceling a Build

Manually cancel a build using the web console, or with the following CLI command:

```
$ oc cancel-build <build_name>
```

Cancel multiple builds at the same time:

```
$ oc cancel-build <build1_name> <build2_name> <build3_name>
```

Cancel all builds created from the build configuration:

```
$ oc cancel-build bc/<buildconfig_name>
```

Cancel all builds in a given state (for example, **new** or **pending**), ignoring the builds in other states:

```
$ oc cancel-build bc/<buildconfig_name> --state=<state>
```

8.2.3. Deleting a BuildConfig

Delete a **BuildConfig** using the following command:

```
$ oc delete bc <BuildConfigName>
```

This will also delete all builds that were instantiated from this **BuildConfig**. Specify the **--cascade=false** flag if you do not want to delete the builds:

```
$ oc delete --cascade=false bc <BuildConfigName>
```

8.2.4. Viewing Build Details

You can view build details with the web console or by using the **oc describe** CLI command:

```
$ oc describe build <build_name>
```

This displays information such as:

- The build source
- The build strategy
- The output destination
- Digest of the image in the destination registry
- How the build was created

If the build uses the **Docker** or **Source** strategy, the **oc describe** output also includes information about the source revision used for the build, including the commit ID, author, committer, and message.

8.2.5. Accessing Build Logs

You can access build logs using the web console or the CLI.

To stream the logs using the build directly:

```
$ oc logs -f build/<build_name>
```

To stream the logs of the latest build for a build configuration:

```
$ oc logs -f bc/<buildconfig_name>
```

To return the logs of a given version build for a build configuration:

```
$ oc logs --version=<number> bc/<buildconfig_name>
```

Log Verbosity

To enable more verbose output, pass the **BUILD_LOGLEVEL** environment variable as part of the **sourceStrategy** or **dockerStrategy** in a **BuildConfig**:

```
sourceStrategy:
  ...
  env:
    - name: "BUILD_LOGLEVEL"
      value: "2" 1
```

1 Adjust this value to the desired log level.



NOTE

A platform administrator can set the default build verbosity for the entire OpenShift Container Platform instance by configuring **env/BUILD_LOGLEVEL** for the **BuildDefaults** admission controller. This default can be overridden by specifying **BUILD_LOGLEVEL** in a given **BuildConfig**. You can specify a higher priority override on the command line for non-binary builds by passing **--build-loglevel** to **oc start-build**.

Available log levels for Source builds are as follows:

Level 0	Produces output from containers running the <i>assemble</i> script and all encountered errors. This is the default.
Level 1	Produces basic information about the executed process.
Level 2	Produces very detailed information about the executed process.
Level 3	Produces very detailed information about the executed process, and a listing of the archive contents.

Level 4	Currently produces the same information as level 3.
Level 5	Produces everything mentioned on previous levels and additionally provides docker push messages.

8.3. BUILD INPUTS

8.3.1. How Build Inputs Work

A *build input* provides source content for builds to operate on. There are several ways to provide source in OpenShift Container Platform. In order of precedence:

- [Inline Dockerfile definitions](#)
- [Content extracted from existing images](#)
- [Git repositories](#)
- [Binary \(Local\) inputs](#)
- [Input secrets](#)
- [External artifacts](#)

Different inputs can be combined into a single build. As the inline Dockerfile takes precedence, it can overwrite any other file named **Dockerfile** provided by another input. Binary (local) input and Git repositories are mutually exclusive inputs.

Input secrets are useful for when you do not want certain resources or credentials used during a build to be available in the final application image produced by the build, or want to consume a value that is defined in a **Secret** resource. External artifacts can be used to pull in additional files that are not available as one of the other build input types.

Whenever a build is run:

1. A working directory is constructed and all input content is placed in the working directory. For example, the input Git repository is cloned into the working directory, and files specified from input images are copied into the working directory using the target path.
2. The build process changes directories into the **contextDir**, if one is defined.
3. The inline Dockerfile, if any, is written to the current directory.
4. The content from the current directory is provided to the build process for reference by the Dockerfile, custom builder logic, or **assemble** script. This means any input content that resides outside the **contextDir** will be ignored by the build.

The following example of a source definition includes multiple input types and an explanation of how they are combined. For more details on how each input type is defined, see the specific sections for each input type.

```
source:
  git:
```

```

uri: https://github.com/openshift/ruby-hello-world.git ❶
images:
- from:
  kind: ImageStreamTag
  name: myinputimage:latest
  namespace: mynamespace
  paths:
  - destinationDir: app/dir/injected/dir ❷
    sourcePath: /usr/lib/somefile.jar
contextDir: "app/dir" ❸
dockerfile: "FROM centos:7\nRUN yum install -y httpd" ❹

```

- ❶ The repository to be cloned into the working directory for the build.
- ❷ `/usr/lib/somefile.jar` from `myinputimage` will be stored in `<workingdir>/app/dir/injected/dir`.
- ❸ The working directory for the build will become `<original_workingdir>/app/dir`.
- ❹ A Dockerfile with this content will be created in `<original_workingdir>/app/dir`, overwriting any existing file with that name.

8.3.2. Dockerfile Source

When a `dockerfile` value is supplied, the content of this field will be written to disk as a file named **Dockerfile**. This is done after other input sources are processed, so if the input source repository contains a **Dockerfile** in the root directory, it will be overwritten with this content.

The typical use for this field is to provide a **Dockerfile** to a [Docker strategy](#) build.

The source definition is part of the `spec` section in the **BuildConfig**:

```

source:
  dockerfile: "FROM centos:7\nRUN yum install -y httpd" ❶

```

- ❶ The `dockerfile` field contains an inline Dockerfile that will be built.

8.3.3. Image Source

Additional files can be provided to the build process via images. Input images are referenced in the same way the `From` and `To` image targets are defined. This means both container images and [image stream tags](#) can be referenced. In conjunction with the image, you must provide one or more path pairs to indicate the path of the files or directories to copy the image and the destination to place them in the build context.

The source path can be any absolute path within the image specified. The destination must be a relative directory path. At build time, the image will be loaded and the indicated files and directories will be copied into the context directory of the build process. This is the same directory into which the source repository content (if any) is cloned. If the source path ends in `/`, then the content of the directory will be copied, but the directory itself will not be created at the destination.


Image inputs are specified in the `source` definition of the **BuildConfig**:

```

source:
  git:
    uri: https://github.com/openshift/ruby-hello-world.git
  images: ❶
  - from: ❷
    kind: ImageStreamTag
    name: myinputimage:latest
    namespace: mynamespace
  paths: ❸
  - destinationDir: injected/dir ❹
    sourcePath: /usr/lib/somefile.jar ❺
  - from:
    kind: ImageStreamTag
    name: myotherinputimage:latest
    namespace: myothernamespace
  pullSecret: mysecret ❻
  paths:
  - destinationDir: injected/dir
    sourcePath: /usr/lib/somefile.jar

```

- ❶ An array of one or more input images and files.
- ❷ A reference to the image containing the files to be copied.
- ❸ An array of source/destination paths.
- ❹ The directory relative to the build root where the build process can access the file.
- ❺ The location of the file to be copied out of the referenced image.
- ❻ An optional secret provided if credentials are needed to access the input image.

**NOTE**

This feature is not supported for builds using the [Custom Strategy](#).

8.3.4. Git Source

When specified, source code will be fetched from the location supplied.

If an inline Dockerfile is supplied, it will overwrite the *Dockerfile* (if any) in the `contextDir` of the Git repository.

The source definition is part of the `spec` section in the **BuildConfig**:

```

source:
  git: ❶
    uri: "https://github.com/openshift/ruby-hello-world"
    ref: "master"
  contextDir: "app/dir" ❷
  dockerfile: "FROM openshift/ruby-22-centos7\nUSER example" ❸

```


- 1 The `git` field contains the URI to the remote Git repository of the source code. Optionally, specify the `ref` field to check out a specific Git reference. A valid `ref` can be a SHA1 tag or a branch name.
- 2 The `contextDir` field allows you to override the default location inside the source code repository where the build looks for the application source code. If your application exists inside a sub-directory, you can override the default location (the root folder) using this field.
- 3 If the optional `dockerfile` field is provided, it should be a string containing a Dockerfile that overwrites any Dockerfile that may exist in the source repository.

If the `ref` field denotes a pull request, the system will use a `git fetch` operation and then checkout `FETCH_HEAD`.

When no `ref` value is provided, OpenShift Container Platform performs a shallow clone (`--depth=1`). In this case, only the files associated with the most recent commit on the default branch (typically `master`) are downloaded. This results in repositories downloading faster, but without the full commit history. To perform a full `git clone` of the default branch of a specified repository, set `ref` to the name of the default branch (for example `master`).

8.3.4.1. Using a Proxy

If your Git repository can only be accessed using a proxy, you can define the proxy to use in the `source` section of the `BuildConfig`. You can configure both a HTTP and HTTPS proxy to use. Both fields are optional. Domains for which no proxying should be performed can also be specified via the `NoProxy` field.



NOTE

Your source URI must use the HTTP or HTTPS protocol for this to work.

```
source:
  git:
    uri: "https://github.com/openshift/ruby-hello-world"
    httpProxy: http://proxy.example.com
    httpsProxy: https://proxy.example.com
    noProxy: somedomain.com, otherdomain.com
```

Cluster administrators can also [configure a global proxy for Git cloning using Ansible](#) .



NOTE

For Pipeline strategy builds, given the current restrictions with the Git plug-in for Jenkins, any Git operations through the Git plug-in will not leverage the HTTP or HTTPS proxy defined in the `BuildConfig`. The Git plug-in only will use the proxy configured in the Jenkins UI at the Plugin Manager panel. This proxy will then be used for all git interactions within Jenkins, across all jobs. You can find instructions on how to configure proxies through the Jenkins UI at [JenkinsBehindProxy](#).

8.3.4.2. Source Clone Secrets

Builder pods require access to any Git repositories defined as source for a build. Source clone secrets are used to provide the builder pod with access it would not normally have access to, such as private repositories or repositories with self-signed or untrusted SSL certificates.

The following source clone secret configurations are supported.

- [.gitconfig File](#)
- [Basic Authentication](#)
- [SSH Key Authentication](#)
- [Trusted Certificate Authorities](#)



NOTE

You can also use [combinations](#) of these configurations to meet your specific needs.

Builds are run with the **builder** service account, which must have access to any source clone secrets used. Access is granted with the following command:

```
$ oc secrets link builder mysecret
```



NOTE

Limiting secrets to only the service accounts that reference them is disabled by default. This means that if `serviceAccountConfig.limitSecretReferences` is set to `false` (the default setting) in the master configuration file, linking secrets to a service is not required.

8.3.4.2.1. Automatically Adding a Source Clone Secret to a Build Configuration

When a **BuildConfig** is created, OpenShift Container Platform can automatically populate its source clone secret reference. This behaviour allows the resulting **Builds** to automatically use the credentials stored in the referenced **Secret** to authenticate to a remote Git repository, without requiring further configuration.

To use this functionality, a **Secret** containing the Git repository credentials must exist in the namespace in which the **BuildConfig** will later be created. This **Secret** must additionally include one or more annotations prefixed with `build.openshift.io/source-secret-match-uri-`. The value of each of these annotations is a URI pattern, defined as follows. When a **BuildConfig** is created without a source clone secret reference and its Git source URI matches a URI pattern in a **Secret** annotation, OpenShift Container Platform will automatically insert a reference to that **Secret** in the **BuildConfig**.

A URI pattern must consist of:

- a valid scheme (`*://`, `git://`, `http://`, `https://` or `ssh://`).
- a host (`*` or a valid hostname or IP address optionally preceded by `*.`).
- a path (`/*` or `/` followed by any characters optionally including `*` characters).

In all of the above, a `*` character is interpreted as a wildcard.

**NOTE**

URI patterns only match Git source URIs which are conformant to RFC3986. For example, <https://github.com/openshift/origin.git>. They do not match the alternate SSH style that Git also uses. For example, [git@github.com:openshift/origin.git](https://github.com/openshift/origin.git).

It is not valid to attempt to express a URI pattern in the alternate style, or to include a username/password component in a URI pattern.

If multiple **Secrets** match the Git URI of a particular **BuildConfig**, OpenShift Container Platform will select the secret with the longest match. This allows for basic overriding, as in the following example.

The following fragment shows two partial source clone secrets, the first matching any server in the domain `mycorp.com` accessed by HTTPS, and the second overriding access to servers `mydev1.mycorp.com` and `mydev2.mycorp.com`:

```
kind: Secret
apiVersion: v1
metadata:
  name: matches-all-corporate-servers-https-only
  annotations:
    build.openshift.io/source-secret-match-uri-1: https://*.mycorp.com/*
data:
  ...

kind: Secret
apiVersion: v1
metadata:
  name: override-for-my-dev-servers-https-only
  annotations:
    build.openshift.io/source-secret-match-uri-1:
https://mydev1.mycorp.com/*
    build.openshift.io/source-secret-match-uri-2:
https://mydev2.mycorp.com/*
data:
  ...
```

Add a `build.openshift.io/source-secret-match-uri-` annotation to a pre-existing secret using:

```
$ oc annotate secret mysecret \
  'build.openshift.io/source-secret-match-uri-1=https://*.mycorp.com/*'
```

8.3.4.2.2. Manually Adding Source Clone Secrets

Source clone secrets can be added manually to a build configuration by adding a `sourceSecret` field to the `source` section inside the **BuildConfig** and setting it to the name of the `secret` that you created (`basicsecret`, in this example).

```
apiVersion: "v1"
kind: "BuildConfig"
metadata:
  name: "sample-build"
spec:
```

```

output:
  to:
    kind: "ImageStreamTag"
    name: "sample-image:latest"
source:
  git:
    uri: "https://github.com/user/app.git"
    sourceSecret:
      name: "basicsecret"
strategy:
  sourceStrategy:
    from:
      kind: "ImageStreamTag"
      name: "python-33-centos7:latest"

```

**NOTE**

You can also use the `oc set build-secret` command to set the source clone secret on an existing build configuration:

```
$ oc set build-secret --source bc/sample-build basicsecret
```

[Defining Secrets in the BuildConfig](#) provides more information on this topic.

8.3.4.2.3. .Gitconfig File

If the cloning of your application is dependent on a `.gitconfig` file, then you can create a secret that contains it, and then add it to the builder service account, and then your `BuildConfig`.

To create a secret from a `.gitconfig` file:

```
$ oc secrets new mysecret .gitconfig=path/to/.gitconfig
```

**NOTE**

SSL verification can be turned off if `sslVerify=false` is set for the `http` section in your `.gitconfig` file:

```
[http]
    sslVerify=false
```

8.3.4.2.4. .gitconfig File for Secured Git

If your Git server is secured with 2-way SSL and user name with password you must add the certificate files to your source build and add references to the certificate files in the `.gitconfig` file:

1. Add the `client.crt`, `cacert.crt`, and `client.key` files to the `/var/run/secrets/openshift.io/source/` folder in the [application source code](#).
2. In the `.gitconfig` file for the server, add the `[http]` section shown in the following example:

```
# cat .gitconfig
```



```
[user]
    name = <name>
    email = <email>

[http]
    sslVerify = false
    sslCert = /var/run/secrets/openshift.io/source/client.crt
    sslKey = /var/run/secrets/openshift.io/source/client.key
    sslCaInfo = /var/run/secrets/openshift.io/source/cacert.crt
```

3. Create the secret:

```
$ oc secrets new <secret_name> \
--from-literal=username=<user_name> \ 1
--from-literal=password=<password> \ 2
--from-file=.gitconfig=.gitconfig \
--from-
file=client.crt=/var/run/secrets/openshift.io/source/client.crt \
--from-
file=cacert.crt=/var/run/secrets/openshift.io/source/cacert.crt \
--from-
file=client.key=/var/run/secrets/openshift.io/source/client.key
```

- 1 The user's Git user name.
- 2 The password for this user.



IMPORTANT

To avoid having to enter your password again, be sure to specify the S2I image in your builds. However, if you cannot clone the repository, you still need to specify your user name and password to promote the build.

8.3.4.2.5. Basic Authentication

Basic authentication requires either a combination of `--username` and `--password`, or a `token` to authenticate against the SCM server.

Create the `secret` first before using the user name and password to access the private repository:

```
$ oc secrets new-basicauth <secret_name> \
--username=<user_name> \
--password=<password>
```

To create a basic authentication secret with a token:

```
$ oc secrets new-basicauth <secret_name> \
--password=<token>
```

8.3.4.2.6. SSH Key Authentication

SSH key based authentication requires a private SSH key.

The repository keys are usually located in the `$HOME/.ssh/` directory, and are named `id_dsa.pub`, `id_ecdsa.pub`, `id_ed25519.pub`, or `id_rsa.pub` by default. Generate SSH key credentials with the following command:

```
$ ssh-keygen -t rsa -C "your_email@example.com"
```



NOTE

Creating a passphrase for the SSH key prevents OpenShift Container Platform from building. When prompted for a passphrase, leave it blank.

Two files are created: the public key and a corresponding private key (one of `id_dsa`, `id_ecdsa`, `id_ed25519`, or `id_rsa`). With both of these in place, consult your source control management (SCM) system's manual on how to upload the public key. The private key is used to access your private repository.

Before using the SSH key to access the private repository, create the secret first:

```
$ oc secrets new-sshauth sshsecret \
  --ssh-privatekey=$HOME/.ssh/id_rsa
```

8.3.4.2.7. Trusted Certificate Authorities

The set of TLS certificate authorities that are trusted during a `git clone` operation are built into the OpenShift Container Platform infrastructure images. If your Git server uses a self-signed certificate or one signed by an authority not trusted by the image, you can create a secret that contains the certificate or disable TLS verification.

If you create a secret for the CA certificate, OpenShift Container Platform uses it to access your Git server during the `git clone` operation. Using this method is significantly more secure than disabling Git's SSL verification, which accepts any TLS certificate that is presented.

Complete one of the following processes:

- Create a secret with a CA certificate file (recommended).
 - a. If your CA uses Intermediate Certificate Authorities, combine the certificates for all CAs in a `ca.crt` file. Run the following command:

```
$ cat intermediateCA.crt intermediateCA.crt rootCA.crt > ca.crt
```

- b. Create the secret:

```
$ oc create secret generic mycert --from-file=ca.crt=
</path/to/file> 1
```

1 You must use the key name `ca.crt`.

- Disable Git TLS verification.

Set the `GIT_SSL_NO_VERIFY` environment variable to `true` in the appropriate strategy section of your build configuration. You can use the `oc set env` command to manage `BuildConfig` environment variables.

8.3.4.2.8. Combinations

Below are several examples of how you can combine the above methods for creating source clone secrets for your specific needs.

- a. To create an SSH-based authentication secret with a *.gitconfig* file:

```
$ oc secrets new-sshauth sshsecret \
  --ssh-privatekey=$HOME/.ssh/id_rsa \
  --gitconfig=</path/to/file>
```

- b. To create a secret that combines a *.gitconfig* file and CA certificate:

```
$ oc secrets new mysecret \
  ca.crt=path/to/certificate \
  .gitconfig=path/to/.gitconfig
```

- c. To create a basic authentication secret with a CA certificate file:

```
$ oc secrets new-basicauth <secret_name> \
  --username=<user_name> \
  --password=<password> \
  --ca-cert=</path/to/file>
```

- d. To create a basic authentication secret with a *.gitconfig* file:

```
$ oc secrets new-basicauth <secret_name> \
  --username=<user_name> \
  --password=<password> \
  --gitconfig=</path/to/file>
```

- e. To create a basic authentication secret with a *.gitconfig* file and CA certificate file:

```
$ oc secrets new-basicauth <secret_name> \
  --username=<user_name> \
  --password=<password> \
  --gitconfig=</path/to/file> \
  --ca-cert=</path/to/file>
```

8.3.5. Binary (Local) Source

Streaming content from a local file system to the builder is called a **Binary** type build. The corresponding value of `BuildConfig.spec.source.type` is **Binary** for such builds.

This source type is unique in that it is leveraged solely based on your use of the `oc start-build`.



NOTE

Binary type builds require content to be streamed from the local file system, so automatically triggering a binary type build (e.g. via an image change trigger) is not possible, because the binary files cannot be provided. Similarly, you cannot launch binary type builds from the web console.

To utilize binary builds, invoke `oc start-build` with one of these options:

- **--from-file**: The contents of the file you specify are sent as a binary stream to the builder. You can also specify a URL to a file. Then, the builder stores the data in a file with the same name at the top of the build context.
- **--from-dir** and **--from-repo**: The contents are archived and sent as a binary stream to the builder. Then, the builder extracts the contents of the archive within the build context directory. With **--from-dir**, you can also specify a URL to an archive, which will be extracted.
- **--from-archive**: The archive you specify is sent to the builder, where it is extracted within the build context directory. This option behaves the same as **--from-dir**; an archive is created on your host first, whenever the argument to these options is a directory.

In each of the above cases:

- If your **BuildConfig** already has a **Binary** source type defined, it will effectively be ignored and replaced by what the client sends.
- If your **BuildConfig** has a **Git** source type defined, it is dynamically disabled, since **Binary** and **Git** are mutually exclusive, and the data in the binary stream provided to the builder takes precedence.

Instead of a file name, you can pass a URL with HTTP or HTTPS schema to **--from-file** and **--from-archive**. When using **--from-file** with a URL, the name of the file in the builder image is determined by the **Content-Disposition** header sent by the web server, or the last component of the URL path if the header is not present. No form of authentication is supported and it is not possible to use custom TLS certificate or disable certificate validation.

When using `oc new-build --binary=true`, the command ensures that the restrictions associated with binary builds are enforced. The resulting **BuildConfig** will have a source type of **Binary**, meaning that the only valid way to run a build for this **BuildConfig** is to use `oc start-build` with one of the **--from** options to provide the requisite binary data.

The **dockerfile** and **contextDir** [source options](#) have special meaning with binary builds.

dockerfile can be used with any binary build source. If **dockerfile** is used and the binary stream is an archive, its contents serve as a replacement Dockerfile to any Dockerfile in the archive. If **dockerfile** is used with the **--from-file** argument, and the file argument is named **dockerfile**, the value from **dockerfile** replaces the value from the binary stream.

In the case of the binary stream encapsulating extracted archive content, the value of the **contextDir** field is interpreted as a subdirectory within the archive, and, if valid, the builder changes into that subdirectory before executing the build.

8.3.6. Input Secrets

In some scenarios, build operations require credentials to access dependent resources, but it is undesirable for those credentials to be available in the final application image produced by the build. You can define *input secrets* for this purpose.

For example, when building a Node.js application, you can set up your private mirror for Node.js modules. In order to download modules from that private mirror, you have to supply a custom **.npmrc** file for the build that contains a URL, user name, and password. For security reasons, you do not want to expose your credentials in the application image.

This example describes Node.js, but you can use the same approach for adding SSL certificates into the `/etc/ssl/certs` directory, API keys or tokens, license files, and more.

8.3.6.1. Adding Input Secrets

To add an input secret to an existing **BuildConfig**:

1. Create the secret, if it does not exist:

```
$ oc secrets new secret-npmrc .npmrc=~/.npmrc
```

This creates a new secret named **secret-npmrc**, which contains the base64 encoded content of the `~/.npmrc` file.

2. Add the secret to the **source** section in the existing **BuildConfig**:

```
source:
  git:
    uri: https://github.com/openshift/nodejs-ex.git
  secrets:
    - secret:
        name: secret-npmrc
```

To include the secret in a new **BuildConfig**, run the following command:

```
$ oc new-build \
  openshift/nodejs-010-centos7~https://github.com/openshift/nodejs-
  ex.git \
  --build-secret secret-npmrc
```

During the build, the `.npmrc` file is copied into the directory where the source code is located. In OpenShift Container Platform S2I builder images, this is the image working directory, which is set using the **WORKDIR** instruction in the **Dockerfile**. If you want to specify another directory, add a **destinationDir** to the secret definition:

```
source:
  git:
    uri: https://github.com/openshift/nodejs-ex.git
  secrets:
    - secret:
        name: secret-npmrc
        destinationDir: /etc
```

You can also specify the destination directory when creating a new **BuildConfig**:

```
$ oc new-build \
  openshift/nodejs-010-centos7~https://github.com/openshift/nodejs-
  ex.git \
  --build-secret "secret-npmrc:/etc"
```

In both cases, the `.npmrc` file is added to the `/etc` directory of the build environment. Note that for a [Docker strategy](#) the destination directory must be a relative path.

8.3.6.2. Source-to-Image Strategy

When using a **Source** strategy, all defined input secrets are copied to their respective **destinationDir**. If you left **destinationDir** empty, then the secrets are placed in the working directory of the builder image.

The same rule is used when a **destinationDir** is a relative path; the secrets are placed in the paths that are relative to the image's working directory. The **destinationDir** must exist or an error will occur. No directory paths are created during the copy process.



NOTE

Currently, any files with these secrets are world-writable (have **0666** permissions) and will be truncated to size zero after executing the **assemble** script. This means that the secret files will exist in the resulting image, but they will be empty for security reasons.

8.3.6.3. Docker Strategy

When using a **Docker** strategy, you can add all defined input secrets into your container image using the [ADD](#) and [COPY instructions](#) in your **Dockerfile**.

If you do not specify the **destinationDir** for a secret, then the files will be copied into the same directory in which the **Dockerfile** is located. If you specify a relative path as **destinationDir**, then the secrets will be copied into that directory, relative to your **Dockerfile** location. This makes the secret files available to the Docker build operation as part of the context directory used during the build.

Example 8.1. Example of a Dockerfile referencing secret data

```
FROM centos/ruby-22-centos7

USER root
ADD ./secret-dir /secrets
COPY ./secret2 /

# Create a shell script that will output secrets when the image is run
RUN echo '#!/bin/sh' > /secret_report.sh
RUN echo '(test -f /secrets/secret1 && echo -n "secret1=" && cat /secrets/secret1)' >> /secret_report.sh
RUN echo '(test -f /secret2 && echo -n "relative-secret2=" && cat /secret2)' >> /secret_report.sh
RUN chmod 755 /secret_report.sh

CMD ["/bin/sh", "-c", "/secret_report.sh"]
```



NOTE

Users should normally remove their input secrets from the final application image so that the secrets are not present in the container running from that image. However, the secrets will still exist in the image itself in the layer where they were added. This removal should be part of the **Dockerfile** itself.

8.3.6.4. Custom Strategy

When using a **Custom** strategy, all the defined input secrets are available inside the builder container in the `/var/run/secrets/openshift.io/build` directory. The custom build image is responsible for using these secrets appropriately. The **Custom** strategy also allows secrets to be defined as described in [Custom Strategy Options](#).

There is no technical difference between existing strategy secrets and the input secrets. However, your builder image might distinguish between them and use them differently, based on your build use case.

The input secrets are always mounted into the `/var/run/secrets/openshift.io/build` directory or your builder can parse the `$BUILD` environment variable, which includes the full build object.

8.3.7. Using External Artifacts

It is not recommended to store binary files in a source repository. Therefore, you may find it necessary to define a build which pulls additional files (such as Java `.jar` dependencies) during the build process. How this is done depends on the build strategy you are using.

For a **Source** build strategy, you must put appropriate shell commands into the **assemble** script:

.s2i/bin/assemble File

```
#!/bin/sh
APP_VERSION=1.0
wget http://repository.example.com/app/app-$APP_VERSION.jar -O app.jar
```

.s2i/bin/run File

```
#!/bin/sh
exec java -jar app.jar
```



NOTE

For more information on how to control which **assemble** and **run** script is used by a Source build, see [Overriding Builder Image Scripts](#).

For a **Docker** build strategy, you must modify the **Dockerfile** and invoke shell commands with the [RUN instruction](#):

Excerpt of Dockerfile

```
FROM jboss/base-jdk:8

ENV APP_VERSION 1.0
RUN wget http://repository.example.com/app/app-$APP_VERSION.jar -O app.jar

EXPOSE 8080
CMD [ "java", "-jar", "app.jar" ]
```

In practice, you may want to use an environment variable for the file location so that the specific file to be downloaded can be customized using an environment variable defined on the **BuildConfig**, rather than updating the *Dockerfile* or *assemble* script.

You can choose between different methods of defining environment variables:

- [Using the `.s2i/environment` file](#) (only for a Source build strategy)
- [Setting in **BuildConfig**](#)
- [Providing explicitly using `oc start-build --env`](#) (only for builds that are triggered manually)

8.3.8. Using Docker Credentials for Private Registries

You can supply builds with a `.docker/config.json` file with valid credentials for private Docker registries. This allows you to push the output image into a private Docker registry or pull a builder image from the private Docker registry that requires authentication.



NOTE

For the OpenShift Container Platform Docker registry, this is not required because secrets are generated automatically for you by OpenShift Container Platform.

The `.docker/config.json` file is found in your home directory by default and has the following format:

```
auths:
  https://index.docker.io/v1/: 1
    auth: "YWRfbGZhcGU6R2labnRib21ifTE=" 2
    email: "user@example.com" 3
```

- 1 URL of the registry.
- 2 Encrypted password.
- 3 Email address for the login.

You can define multiple Docker registry entries in this file. Alternatively, you can also add authentication entries to this file by running the `docker login` command. The file will be created if it does not exist.

Kubernetes provides [Secret](#) objects, which can be used to store configuration and passwords.

1. Create the secret from your local `.docker/config.json` file:

```
$ oc secrets new dockerhub ~/.docker/config.json
```

This generates a JSON specification of the secret named **dockerhub** and creates the object.

2. Once the secret is created, add it to the builder service account. Each build is run with the **builder** role, so you must give it access your secret with the following command:

```
$ oc secrets link builder dockerhub
```


3. Add a `pushSecret` field into the `output` section of the `BuildConfig` and set it to the name of the secret that you created, which in the above example is `dockerhub`:

```
spec:
  output:
    to:
      kind: "DockerImage"
      name: "private.registry.com/org/private-image:latest"
    pushSecret:
      name: "dockerhub"
```

You can also use the `oc set build-secret` command to set the push secret on the build configuration:

```
$ oc set build-secret --push bc/sample-build dockerhub
```

4. Pull the builder container image from a private Docker registry by specifying the `pullSecret` field, which is part of the build strategy definition:

```
strategy:
  sourceStrategy:
    from:
      kind: "DockerImage"
      name: "docker.io/user/private_repository"
    pullSecret:
      name: "dockerhub"
```

You can also use the `oc set build-secret` command to set the pull secret on the build configuration:

```
$ oc set build-secret --pull bc/sample-build dockerhub
```



NOTE

This example uses `pullSecret` in a Source build, but it is also applicable in Docker and Custom builds.

8.4. BUILD OUTPUT

8.4.1. Build Output Overview

Builds that use the **Docker** or **Source** strategy result in the creation of a new container image. The image is then pushed to the container image registry specified in the `output` section of the **Build** specification.

If the output kind is `ImageStreamTag`, then the image will be pushed to the integrated OpenShift Container Platform registry and tagged in the specified image stream. If the output is of type `DockerImage`, then the name of the output reference will be used as a Docker push specification. The specification may contain a registry or will default to DockerHub if no registry is specified. If the output section of the build specification is empty, then the image will not be pushed at the end of the build.

Output to an `ImageStreamTag`

```
spec:
  output:
    to:
      kind: "ImageStreamTag"
      name: "sample-image:latest"
```

Output to a Docker Push Specification

```
spec:
  output:
    to:
      kind: "DockerImage"
      name: "my-registry.mycompany.com:5000/myimages/myimage:tag"
```

8.4.2. Output Image Environment Variables

Docker and Source strategy builds set the following environment variables on output images:

Variable	Description
OPENSIFT_BUILD_NAME	Name of the build
OPENSIFT_BUILD_NAMESPACE	Namespace of the build
OPENSIFT_BUILD_SOURCE	The source URL of the build
OPENSIFT_BUILD_REFERENCE	The Git reference used in the build
OPENSIFT_BUILD_COMMIT	Source commit used in the build

Additionally, any user-defined environment variable, for example those configured via [Source](#) or [Docker](#) strategy options, will also be part of the output image environment variable list.

8.4.3. Output Image Labels

Docker and Source builds set the following labels on output images:

Label	Description
io.openshift.build.commit.author	Author of the source commit used in the build
io.openshift.build.commit.date	Date of the source commit used in the build
io.openshift.build.commit.id	Hash of the source commit used in the build
io.openshift.build.commit.message	Message of the source commit used in the build

Label	Description
<code>io.openshift.build.commit.ref</code>	Branch or reference specified in the source
<code>io.openshift.build.source-location</code>	Source URL for the build

You can also use the `BuildConfig.spec.output.imageLabels` field to specify a list of custom labels that will be applied to each image built from the `BuildConfig`.

Custom Labels to be Applied to Built Images

```
spec:
  output:
    to:
      kind: "ImageStreamTag"
      name: "my-image:latest"
  imageLabels:
    - name: "vendor"
      value: "MyCompany"
    - name: "authoritative-source-url"
      value: "registry.mycompany.com"
```

8.4.4. Output Image Digest

Built images can be uniquely identified by their [digest](#), which can later be used to [pull the image by digest](#) regardless of its current tag.

Docker and **Source** builds store the digest in `Build.status.output.to.imageDigest` after the image is pushed to a registry. The digest is computed by the registry. Therefore, it may not always be present, for example when the registry did not return a digest, or when the builder image did not understand its format.

Built Image Digest After a Successful Push to the Registry

```
status:
  output:
    to:
      imageDigest:
sha256:29f5d56d12684887bdfa50dcd29fc31eea4aaf4ad3bec43daf19026a7ce69912
```

8.4.5. Using Docker Credentials for Private Registries

To push an image to a private Docker registry, credentials can be supplied using a secret. See [Build Inputs](#) for instructions.

8.5. BUILD STRATEGY OPTIONS

8.5.1. Source-to-Image Strategy Options

The following options are specific to the [S2I build strategy](#).

8.5.1.1. Force Pull

By default, if the builder image specified in the build configuration is available locally on the node, that image will be used. However, to override the local image and refresh it from the registry to which the image stream points, create a **BuildConfig** with the **forcePull** flag set to **true**:

```
strategy:
  sourceStrategy:
    from:
      kind: "ImageStreamTag"
      name: "builder-image:latest" ❶
    forcePull: true ❷
```

- ❶ The builder image being used, where the local version on the node may not be up to date with the version in the registry to which the image stream points.
- ❷ This flag causes the local builder image to be ignored and a fresh version to be pulled from the registry to which the image stream points. Setting **forcePull** to **false** results in the default behavior of honoring the image stored locally.

8.5.1.2. Incremental Builds

S2I can perform incremental builds, which means it reuses artifacts from previously-built images. To create an incremental build, create a **BuildConfig** with the following modification to the strategy definition:

```
strategy:
  sourceStrategy:
    from:
      kind: "ImageStreamTag"
      name: "incremental-image:latest" ❶
    incremental: true ❷
```

- ❶ Specify an image that supports incremental builds. Consult the documentation of the builder image to determine if it supports this behavior.
- ❷ This flag controls whether an incremental build is attempted. If the builder image does not support incremental builds, the build will still succeed, but you will get a log message stating the incremental build was not successful because of a missing *save-artifacts* script.



NOTE

See the [S2I Requirements](#) topic for information on how to create a builder image supporting incremental builds.

8.5.1.3. Extended Builds



NOTE

This feature is in technology preview. This means the API may change without notice or the feature may be removed entirely. For a supported mechanism to produce application images with runtime-only content, consider using the [Image Source](#) feature and defining two builds, one which produces an image containing the runtime artifacts and a second build which consumes the runtime artifacts from that image and adds them to a runtime-only image.

For compiled languages (Go, C, C++, Java, etc.) the dependencies necessary for compilation might increase the size of the image or introduce vulnerabilities that can be exploited.

To avoid these problems, S2I (Source-to-Image) introduces a two-image build process that allows an application to be built via the normal flow in a builder image, but then injects the resulting application artifacts into a runtime-only image for execution.

To offer flexibility in this process, S2I executes an `assemble-runtime` script inside the runtime image that allows further customization of the resulting runtime image.

More information about this can be found in the official [S2I extended builds documents](#).

This feature is available only for the source strategy.

```
strategy:
  type: "Source"
  sourceStrategy:
    from:
      kind: "ImageStreamTag"
      name: "builder-image:latest"
    runtimeImage: ❶
      kind: "ImageStreamTag"
      name: "runtime-image:latest"
    runtimeArtifacts: ❷
      - sourcePath: "/path/to/source"
        destinationDir: "path/to/destination"
```

- ❶ The runtime image that the artifacts should be copied to. This is the final image that the application will run on. This image should contain the minimum application dependencies to run the injected content from the builder image.
- ❷ The runtime artifacts are a mapping of artifacts produced in the builder image that should be injected into the runtime image. `sourcePath` can be the full path to a file or directory inside the builder image. `destinationDir` must be a directory inside the runtime image where the artifacts will be copied. This directory is relative to the specified `WORKDIR` inside that image.



NOTE

In the current implementation, you cannot have incremental extended builds thus, the `incremental` option is not valid with `runtimeImage`.

If the runtime image needs authentication to be pulled across OpenShift projects or from another private registry, the details can be specified [within the image pull secret configuration](#).

8.5.1.3.1. Testing your Application

Extended builds offer two ways of running tests against your application.

The first option is to install all test dependencies and run the tests inside your *builder image* since that image, in the context of extended builds, will not be pushed to a registry. This can be done as a part of the `assemble` script for the builder image.

The second option is to specify a script via the [postcommit hook](#). This is executed in an ephemeral container based on the runtime image, thus it is not committed to the image.

8.5.1.4. Overriding Builder Image Scripts

You can override the `assemble`, `run`, and `save-artifacts` [S2I scripts](#) provided by the builder image in one of two ways. Either:

1. Provide an `assemble`, `run`, and/or `save-artifacts` script in the `.s2i/bin` directory of your application source repository, or
2. Provide a URL of a directory containing the scripts as part of the strategy definition. For example:

```
strategy:
  sourceStrategy:
    from:
      kind: "ImageStreamTag"
      name: "builder-image:latest"
      scripts: "http://somehost.com/scripts_directory" 1
```

- 1** This path will have `run`, `assemble`, and `save-artifacts` appended to it. If any or all scripts are found they will be used in place of the same named script(s) provided in the image.



NOTE

Files located at the `scripts` URL take precedence over files located in `.s2i/bin` of the source repository. See the [S2I Requirements](#) topic and the [S2I documentation](#) for information on how S2I scripts are used.

8.5.1.5. Environment Variables

There are two ways to make environment variables available to the [source build](#) process and resulting image. [Environment files](#) and [BuildConfig environment](#) values. Variables provided will be present during the build process and in the output image.

8.5.1.5.1. Environment Files

Source build enables you to set environment values (one per line) inside your application, by specifying them in a `.s2i/environment` file in the source repository. The environment variables specified in this file are present during the build process and in the output image. The complete list of supported environment variables is available in the [documentation](#) for each image.

If you provide a `.s2i/environment` file in your source repository, S2I reads this file during the build. This allows customization of the build behavior as the `assemble` script may use these variables.

For example, if you want to disable assets compilation for your Rails application, you can add **DISABLE_ASSET_COMPILATION=true** in the **.s2i/environment** file to cause assets compilation to be skipped during the build.

In addition to builds, the specified environment variables are also available in the running application itself. For example, you can add **RAILS_ENV=development** to the **.s2i/environment** file to cause the Rails application to start in **development** mode instead of **production**.

8.5.1.5.2. BuildConfig Environment

You can add environment variables to the **sourceStrategy** definition of the **BuildConfig**. The environment variables defined there are visible during the **assemble** script execution and will be defined in the output image, making them also available to the **run** script and application code.

For example disabling assets compilation for your Rails application:

```
sourceStrategy:
...
  env:
    - name: "DISABLE_ASSET_COMPILATION"
      value: "true"
```

You can also manage environment variables defined in the **BuildConfig** with the **oc set env** command.

8.5.1.6. Adding Secrets via Web Console

To add a secret to your build configuration so that it can access a private repository:

1. Create a new OpenShift Container Platform project.
2. [Create a secret](#) that contains credentials for accessing a private source code repository.
3. Create a [Source-to-Image \(S2I\) build configuration](#).
4. On the build configuration editor page or in the **create app from builder image** page of the [web console](#), set the **Source Secret**.
5. Click the **Save** button.

8.5.1.6.1. Enabling Pulling and Pushing

Enable pulling to a private registry by setting the **Pull Secret** in the build configuration and enable pushing by setting the **Push Secret**.

8.5.1.7. Ignoring Source Files

Source to image supports a **.s2iignore** file, which contains a list of file patterns that should be ignored. Files in the build working directory, as provided by the various [input sources](#), that match a pattern found in the **.s2iignore** file will not be made available to the **assemble** script.

For more details on the format of the **.s2iignore** file, see the [source-to-image documentation](#).

8.5.2. Docker Strategy Options

The following options are specific to the [Docker build strategy](#).

8.5.2.1. FROM Image

The **FROM** instruction of the *Dockerfile* will be replaced by the **from** of the **BuildConfig**:

```
strategy:
  dockerStrategy:
    from:
      kind: "ImageStreamTag"
      name: "debian:latest"
```

8.5.2.2. Dockerfile Path

By default, Docker builds use a Dockerfile (named *Dockerfile*) located at the root of the context specified in the **BuildConfig.spec.source.contextDir** field.

The **dockerfilePath** field allows the build to use a different path to locate your Dockerfile, relative to the **BuildConfig.spec.source.contextDir** field. It can be simply a different file name other than the default *Dockerfile* (for example, *MyDockerfile*), or a path to a Dockerfile in a subdirectory (for example, *dockerfiles/app1/Dockerfile*):

```
strategy:
  dockerStrategy:
    dockerfilePath: dockerfiles/app1/Dockerfile
```

8.5.2.3. No Cache

Docker builds normally reuse cached layers found on the host performing the build. Setting the **noCache** option to **true** forces the build to ignore cached layers and rerun all steps of the *Dockerfile*:

```
strategy:
  dockerStrategy:
    noCache: true
```

8.5.2.4. Force Pull

By default, if the builder image specified in the build configuration is available locally on the node, that image will be used. However, to override the local image and refresh it from the registry to which the image stream points, create a **BuildConfig** with the **forcePull** flag set to **true**:

```
strategy:
  dockerStrategy:
    forcePull: true 1
```

- 1** This flag causes the local builder image to be ignored, and a fresh version to be pulled from the registry to which the image stream points. Setting **forcePull** to **false** results in the default behavior of honoring the image stored locally.

8.5.2.5. Environment Variables

To make environment variables available to the [Docker build](#) process and resulting image, you can add environment variables to the `dockerStrategy` definition of the `BuildConfig`.

The environment variables defined there are inserted as a single `ENV` Dockerfile instruction right after the `FROM` instruction, so that it can be referenced later on within the Dockerfile.

The variables are defined during build and stay in the output image, therefore they will be present in any container that runs that image as well.

For example, defining a custom HTTP proxy to be used during build and runtime:

```
dockerStrategy:
  ...
  env:
    - name: "HTTP_PROXY"
      value: "http://myproxy.net:5187/"
```

Cluster administrators can also [configure global build settings using Ansible](#).

You can also manage environment variables defined in the `BuildConfig` with the `oc set env` command.

8.5.2.6. Adding Secrets via Web Console

To add a secret to your build configuration so that it can access a private repository"

1. Create a new OpenShift Container Platform project.
2. [Create a secret](#) that contains credentials for accessing a private source code repository.
3. Create a [docker build configuration](#).
4. On the build configuration editor page or in the `fromimage` page of the [web console](#), set the **Source Secret**.
5. Click the **Save** button.

8.5.2.6.1. Enabling Pulling and Pushing

Enable pulling to a private registry by setting the `Pull Secret` in the build configuration and enable pushing by setting the `Push Secret`.

8.5.3. Custom Strategy Options

The following options are specific to the [Custom build strategy](#).

8.5.3.1. FROM Image

Use the `customStrategy.from` section to indicate the image to use for the custom build:

```
strategy:
  customStrategy:
    from:
      kind: "DockerImage"
      name: "openshift/sti-image-builder"
```

8.5.3.2. Exposing the Docker Socket

In order to allow the running of Docker commands and the building of container images from inside the container, the build container must be bound to an accessible socket. To do so, set the `exposeDockerSocket` option to `true`:

```
strategy:
  customStrategy:
    exposeDockerSocket: true
```

8.5.3.3. Secrets

In addition to [secrets](#) for [source](#) and [images](#) that can be added to all build types, custom strategies allow adding an arbitrary list of secrets to the builder pod.

Each secret can be mounted at a specific location:

```
strategy:
  customStrategy:
    secrets:
      - secretSource: 1
        name: "secret1"
        mountPath: "/tmp/secret1" 2
      - secretSource:
        name: "secret2"
        mountPath: "/tmp/secret2"
```

1 `secretSource` is a reference to a secret in the same namespace as the build.

2 `mountPath` is the path inside the custom builder where the secret should be mounted.

8.5.3.3.1. Adding Secrets via Web Console

To add a secret to your build configuration so that it can access a private repository:

1. Create a new OpenShift Container Platform project.
2. [Create a secret](#) that contains credentials for accessing a private source code repository.
3. Create a [custom build configuration](#).
4. On the build configuration editor page or in the `fromimage` page of the [web console](#), set the **Source Secret**.
5. Click the **Save** button.

8.5.3.3.2. Enabling Pulling and Pushing

Enable pulling to a private registry by setting the **Pull Secret** in the build configuration and enable pushing by setting the **Push Secret**.

8.5.3.4. Force Pull

By default, when setting up the build pod, the build controller checks if the image specified in the build configuration is available locally on the node. If so, that image will be used. However, to override the local image and refresh it from the registry to which the image stream points, create a **BuildConfig** with the **forcePull** flag set to **true**:

```
strategy:
  customStrategy:
    forcePull: true 1
```

- 1 This flag causes the local builder image to be ignored, and a fresh version to be pulled from the registry to which the image stream points. Setting **forcePull** to **false** results in the default behavior of honoring the image stored locally.

8.5.3.5. Environment Variables

To make environment variables available to the **Custom build** process, you can add environment variables to the **customStrategy** definition of the **BuildConfig**.

The environment variables defined there are passed to the pod that runs the custom build.

For example, defining a custom HTTP proxy to be used during build:

```
customStrategy:
  ...
  env:
    - name: "HTTP_PROXY"
      value: "http://myproxy.net:5187/"
```

Cluster administrators can also [configure global build settings using Ansible](#).

You can also manage environment variables defined in the **BuildConfig** with the **oc set env** command.

8.5.4. Pipeline Strategy Options

The following options are specific to the **Pipeline build strategy**.

8.5.4.1. Providing the Jenkinsfile

You can provide the Jenkinsfile in one of two ways:

1. Embed the Jenkinsfile in the build configuration.
2. Include in the build configuration a reference to the Git repository that contains the Jenkinsfile.

Embedded Definition

```
kind: "BuildConfig"
apiVersion: "v1"
metadata:
```

```

    name: "sample-pipeline"
spec:
  strategy:
    jenkinsPipelineStrategy:
      jenkinsfile: "node('agent') {\nstage
'build'\nopenshiftBuild(buildConfig: 'ruby-sample-build', showBuildLogs:
'true')\nstage 'deploy'\nopenshiftDeploy(deploymentConfig: 'frontend')\n}"

```

Reference to Git Repository

```

kind: "BuildConfig"
apiVersion: "v1"
metadata:
  name: "sample-pipeline"
spec:
  source:
    git:
      uri: "https://github.com/openshift/ruby-hello-world"
  strategy:
    jenkinsPipelineStrategy:
      jenkinsfilePath: some/repo/dir/filename 1

```

- 1** The optional `jenkinsfilePath` field specifies the name of the file to use, relative to the source `contextDir`. If `contextDir` is omitted, it defaults to the root of the repository. If `jenkinsfilePath` is omitted, it defaults to *Jenkinsfile*.

8.6. TRIGGERING BUILDS

8.6.1. Build Triggers Overview

When defining a **BuildConfig**, you can define triggers to control the circumstances in which the **BuildConfig** should be run. The following build triggers are available:

- [Webhook](#)
- [Image change](#)
- [Configuration change](#)

8.6.2. Webhook Triggers

Webhook triggers allow you to trigger a new build by sending a request to the OpenShift Container Platform API endpoint. You can define these triggers using [GitHub webhooks](#) or Generic webhooks.

8.6.2.1. GitHub Webhooks

[GitHub webhooks](#) handle the call made by GitHub when a repository is updated. When defining the trigger, you must specify a **secret**, which will be part of the URL you supply to GitHub when configuring the webhook. The secret ensures the uniqueness of the URL, preventing others from triggering the build. The following example is a trigger definition YAML within the **BuildConfig**:

```
type: "GitHub"
github:
  secret: "secret101"
```



NOTE

The secret field in webhook trigger configuration is not the same as `secret` field you encounter when configuring webhook in GitHub UI. The former is to make the webhook URL unique and hard to predict, the latter is an optional string field used to create HMAC hex digest of the body, which is sent as an `X-Hub-Signature` header.

The payload URL is returned as the GitHub Webhook URL by the `describe` command (see [Displaying Webhook URLs](#)), and is structured as follows:

```
http://<openshift_api_host:port>/oapi/v1/namespaces/<namespace>/buildconfigs/<name>/webhooks/<secret>/github
```

To configure a GitHub Webhook:

1. Describe the build configuration to get the webhook URL:

```
$ oc describe bc <name>
```

2. Copy the webhook URL.
3. Follow the [GitHub setup instructions](#) to paste the webhook URL into your GitHub repository settings.



NOTE

[Gogs](#) supports the same webhook payload format as GitHub. Therefore, if you are using a Gogs server, you can define a GitHub webhook trigger on your `BuildConfig` and trigger it via your Gogs server also.

Given a file containing a valid JSON payload, you can manually trigger the webhook via `curl`:

```
$ curl -H "X-GitHub-Event: push" -H "Content-Type: application/json" -k -X POST --data-binary @github_payload_file.json https://<openshift_api_host:port>/oapi/v1/namespaces/<namespace>/buildconfigs/<name>/webhooks/<secret>/github
```

The `-k` argument is only necessary if your API server does not have a properly signed certificate.

8.6.2.2. Generic Webhooks

Generic webhooks are invoked from any system capable of making a web request. As with a GitHub webhook, you must specify a secret, which will be part of the URL that the caller must use to trigger the build. The secret ensures the uniqueness of the URL, preventing others from triggering the build. The following is an example trigger definition YAML within the `BuildConfig`:

```
type: "Generic"
generic:
```

```
secret: "secret101"
allowEnv: true 1
```

- 1** Set to **true** to allow a generic webhook to pass in environment variables.

To set up the caller, supply the calling system with the URL of the generic webhook endpoint for your build:

```
http://<openshift_api_host:port>/oapi/v1/namespaces/<namespace>/buildconfigs/<name>/webhooks/<secret>/generic
```

The caller must invoke the webhook as a **POST** operation.

To invoke the webhook manually you can use **curl**:

```
$ curl -X POST -k
https://<openshift_api_host:port>/oapi/v1/namespaces/<namespace>/buildconfigs/<name>/webhooks/<secret>/generic
```

The HTTP verb must be set to **POST**. The insecure **-k** flag is specified to ignore certificate validation. This second flag is not necessary if your cluster has properly signed certificates.

The endpoint can accept an optional payload with the following format:

```
git:
  uri: "<url to git repository>"
  ref: "<optional git reference>"
  commit: "<commit hash identifying a specific git commit>"
  author:
    name: "<author name>"
    email: "<author e-mail>"
  committer:
    name: "<committer name>"
    email: "<committer e-mail>"
  message: "<commit message>"
env: 1
  - name: "<variable name>"
    value: "<variable value>"
```


- 1** Similar to the **BuildConfig environment** variables, the environment variables defined here are made available to your build. If these variables collide with the **BuildConfig** environment variables, these variables take precedence. By default, environment variables passed via webhook are ignored. Set the **allowEnv** field to **true** on the webhook definition to enable this behavior.

To pass this payload using **curl**, define it in a file named **payload_file.yaml** and run:

```
$ curl -H "Content-Type: application/yaml" --data-binary
@payload_file.yaml -X POST -k
https://<openshift_api_host:port>/oapi/v1/namespaces/<namespace>/buildconfigs/<name>/webhooks/<secret>/generic
```

The arguments are the same as the previous example with the addition of a header and a payload. The

-H argument sets the **Content-Type** header to **application/yaml** or **application/json** depending on your payload format. The **--data-binary** argument is used to send a binary payload with newlines intact with the **POST** request.



NOTE

OpenShift Container Platform permits builds to be triggered via the generic webhook even if an invalid request payload is presented (for example, invalid content type, unparseable or invalid content, and so on). This behavior is maintained for backwards compatibility. If an invalid request payload is presented, OpenShift Container Platform returns a warning in JSON format as part of its **HTTP 200 OK** response.

8.6.2.3. Displaying Webhook URLs

Use the following command to display any webhook URLs associated with a build configuration:

```
$ oc describe bc <name>
```

If the above command does not display any webhook URLs, then no webhook trigger is defined for that build configuration.

8.6.3. Image Change Triggers

Image change triggers allow your build to be automatically invoked when a new version of an upstream image is available. For example, if a build is based on top of a RHEL image, then you can trigger that build to run any time the RHEL image changes. As a result, the application image is always running on the latest RHEL base image.

Configuring an image change trigger requires the following actions:

1. Define an **ImageStream** that points to the upstream image you want to trigger on:

```
kind: "ImageStream"
apiVersion: "v1"
metadata:
  name: "ruby-20-centos7"
```

This defines the image stream that is tied to a container image repository located at **<system-registry>/<namespace>/ruby-20-centos7**. The **<system-registry>** is defined as a service with the name **docker-registry** running in OpenShift Container Platform.

2. If an image stream is the base image for the build, set the **from** field in the build strategy to point to the image stream:

```
strategy:
  sourceStrategy:
    from:
      kind: "ImageStreamTag"
      name: "ruby-20-centos7:latest"
```

In this case, the **sourceStrategy** definition is consuming the **latest** tag of the image stream named **ruby-20-centos7** located within this namespace.

3. Define a build with one or more triggers that point to image streams:

```

type: "imageChange" ❶
imageChange: {}
type: "imagechange" ❷
imageChange:
  from:
    kind: "ImageStreamTag"
    name: "custom-image:latest"

```

- ❶ An image change trigger that monitors the **ImageStream** and **Tag** as defined by the build strategy's **from** field. The **imageChange** object here must be empty.
- ❷ An image change trigger that monitors an arbitrary image stream. The **imageChange** part in this case must include a **from** field that references the **ImageStreamTag** to monitor.

When using an image change trigger for the strategy image stream, the generated build is supplied with an immutable Docker tag that points to the latest image corresponding to that tag. This new image reference will be used by the strategy when it executes for the build.

For other image change triggers that do not reference the strategy image stream, a new build will be started, but the build strategy will not be updated with a unique image reference.

In the example above that has an image change trigger for the strategy, the resulting build will be:

```

strategy:
  sourceStrategy:
    from:
      kind: "DockerImage"
      name: "172.30.17.3:5001/mynamespace/ruby-20-centos7:<immutableid>"

```

This ensures that the triggered build uses the new image that was just pushed to the repository, and the build can be re-run any time with the same inputs.

In addition to setting the image field for all **Strategy** types, for custom builds, the **OPENSIFT_CUSTOM_BUILD_BASE_IMAGE** environment variable is checked. If it does not exist, then it is created with the immutable image reference. If it does exist then it is updated with the immutable image reference.

If a build is triggered due to a webhook trigger or manual request, the build that is created uses the **<immutableid>** resolved from the **ImageStream** referenced by the **Strategy**. This ensures that builds are performed using consistent image tags for ease of reproduction.



NOTE

Image streams that point to container images in [v1 Docker registries](#) only trigger a build once when the [image stream tag](#) becomes available and not on subsequent image updates. This is due to the lack of uniquely identifiable images in v1 Docker registries.

8.6.4. Configuration Change Triggers

A configuration change trigger allows a build to be automatically invoked as soon as a new **BuildConfig** is created. The following is an example trigger definition YAML within the **BuildConfig**:

type: "ConfigChange"



NOTE

Configuration change triggers currently only work when creating a new **BuildConfig**. In a future release, configuration change triggers will also be able to launch a build whenever a **BuildConfig** is updated.

8.7. BUILD HOOKS

8.7.1. Build Hooks Overview

Build hooks allow behavior to be injected into the build process.

The **postCommit** field of a **BuildConfig** object executes commands inside a temporary container that is running the build output image. The hook is executed immediately after the last layer of the image has been committed and before the image is pushed to a registry.

The current working directory is set to the image's **WORKDIR**, which is the default working directory of the container image. For most images, this is where the source code is located.

The hook fails if the script or command returns a non-zero exit code or if starting the temporary container fails. When the hook fails it marks the build as failed and the image is not pushed to a registry. The reason for failing can be inspected by looking at the build logs.

Build hooks can be used to run unit tests to verify the image before the build is marked complete and the image is made available in a registry. If all tests pass and the test runner returns with exit code 0, the build is marked successful. In case of any test failure, the build is marked as failed. In all cases, the build log will contain the output of the test runner, which can be used to identify failed tests.

The **postCommit** hook is not only limited to running tests, but can be used for other commands as well. Since it runs in a temporary container, changes made by the hook do not persist, meaning that the hook execution cannot affect the final image. This behavior allows for, among other uses, the installation and usage of test dependencies that are automatically discarded and will be not present in the final image.

8.7.2. Configuring Post Commit Build Hooks

There are different ways to configure the post build hook. All forms in the following examples are equivalent and execute **bundle exec rake test --verbose**:

- Shell script:

```
postCommit:
  script: "bundle exec rake test --verbose"
```

The **script** value is a shell script to be run with **/bin/sh -ic**. Use this when a shell script is appropriate to execute the build hook. For example, for running unit tests as above. To control the image entry point, or if the image does not have **/bin/sh**, use **command** and/or **args**.

**NOTE**

The additional `-i` flag was introduced to improve the experience working with CentOS and RHEL images, and may be removed in a future release.

- Command as the image entry point:

```
postCommit:
  command: ["/bin/bash", "-c", "bundle exec rake test --verbose"]
```

In this form, `command` is the command to run, which overrides the image entry point in the `exec` form, as documented in the [Dockerfile reference](#). This is needed if the image does not have `/bin/sh`, or if you do not want to use a shell. In all other cases, using `script` might be more convenient.

- Pass arguments to the default entry point:

```
postCommit:
  args: ["bundle", "exec", "rake", "test", "--verbose"]
```

In this form, `args` is a list of arguments that are provided to the default entry point of the image. The image entry point must be able to handle arguments.

- Shell script with arguments:

```
postCommit:
  script: "bundle exec rake test $1"
  args: ["--verbose"]
```

Use this form if you need to pass arguments that would otherwise be hard to quote properly in the shell script. In the `script`, `$0` will be `/bin/sh` and `$1`, `$2`, etc, are the positional arguments from `args`.

- Command with arguments:

```
postCommit:
  command: ["bundle", "exec", "rake", "test"]
  args: ["--verbose"]
```

This form is equivalent to appending the arguments to `command`.

**NOTE**

Providing both `script` and `command` simultaneously creates an invalid build hook.

8.7.2.1. Using the CLI

The `oc set build-hook` command can be used to set the build hook for a build configuration.

To set a command as the post-commit build hook:

```
$ oc set build-hook bc/mybc \
  --post-commit \
  --command \
```

```
-- bundle exec rake test --verbose
```

To set a script as the post-commit build hook:

```
$ oc set build-hook bc/mybc --post-commit --script="bundle exec rake test
--verbose"
```

8.8. BUILD RUN POLICY

8.8.1. Build Run Policy Overview

The build run policy describes the order in which the builds created from the build configuration should run. This can be done by changing the value of the `runPolicy` field in the `spec` section of the **Build** specification.

It is also possible to change the `runPolicy` value for existing build configurations.

- Changing **Parallel** to **Serial** or **SerialLatestOnly** and triggering a new build from this configuration will cause the new build to wait until all parallel builds complete as the serial build can only run alone.
- Changing **Serial** to **SerialLatestOnly** and triggering a new build will cause cancellation of all existing builds in queue, except the currently running build and the most recently created build. The newest build will execute next.

8.8.2. Serial Run Policy

Setting the `runPolicy` field to **Serial** will cause all new builds created from the **Build** configuration to be run sequentially. That means there will be only one build running at a time and every new build will wait until the previous build completes. Using this policy will result in consistent and predictable build output. This is the default `runPolicy`.

Triggering three builds from the **sample-build** configuration, using the **Serial** policy will result in:

NAME	TYPE	FROM	STATUS	STARTED	DURATION
sample-build-1	Source	Git@e79d887	Running	13 seconds ago	13s
sample-build-2	Source	Git	New		
sample-build-3	Source	Git	New		

When the **sample-build-1** build completes, the **sample-build-2** build will run:

NAME	TYPE	FROM	STATUS	STARTED	DURATION
sample-build-1	Source	Git@e79d887	Completed	43 seconds ago	34s
sample-build-2	Source	Git@1aa381b	Running	2 seconds ago	2s
sample-build-3	Source	Git	New		

8.8.3. SerialLatestOnly Run Policy

Setting the `runPolicy` field to **SerialLatestOnly** will cause all new builds created from the **Build** configuration to be run sequentially, same as using the **Serial** run policy. The difference is that when

a currently running build completes, the next build that will run is the latest build created. In other words, you do not wait for the queued builds to run, as they are skipped. Skipped builds are marked as **Cancelled**. This policy can be used for fast, iterative development.

Triggering three builds from the **sample-build** configuration, using the **SerialLatestOnly** policy will result in:

NAME	TYPE	FROM	STATUS	STARTED
DURATION				
sample-build-1	Source	Git@e79d887	Running	13 seconds ago 13s
sample-build-2	Source	Git	Cancelled	
sample-build-3	Source	Git	New	

The **sample-build-2** build will be canceled (skipped) and the next build run after **sample-build-1** completes will be the **sample-build-3** build:

NAME	TYPE	FROM	STATUS	STARTED
DURATION				
sample-build-1	Source	Git@e79d887	Completed	43 seconds ago 34s
sample-build-2	Source	Git	Cancelled	
sample-build-3	Source	Git@1aa381b	Running	2 seconds ago 2s

8.8.4. Parallel Run Policy

Setting the **runPolicy** field to **Parallel** causes all new builds created from the **Build** configuration to be run in parallel. This can produce unpredictable results, as the first created build can complete last, which will replace the pushed container image produced by the last build which completed earlier.

Use the parallel run policy in cases where you do not care about the order in which the builds will complete.

Triggering three builds from the **sample-build** configuration, using the **Parallel** policy will result in three simultaneous builds:

NAME	TYPE	FROM	STATUS	STARTED
DURATION				
sample-build-1	Source	Git@e79d887	Running	13 seconds ago 13s
sample-build-2	Source	Git@a76d881	Running	15 seconds ago 3s
sample-build-3	Source	Git@689d111	Running	17 seconds ago 3s

The completion order is not guaranteed:

NAME	TYPE	FROM	STATUS	STARTED
DURATION				
sample-build-1	Source	Git@e79d887	Running	13 seconds ago 13s
sample-build-2	Source	Git@a76d881	Running	15 seconds ago 3s
sample-build-3	Source	Git@689d111	Completed	17 seconds ago 5s

8.9. ADVANCED BUILD OPERATIONS

8.9.1. Setting Build Resources

By default, builds are completed by pods using unbound resources, such as memory and CPU. These resources can be limited by specifying resource limits in a project's default container limits.

You can also limit resource use by specifying resource limits as part of the build configuration. In the following example, each of the **resources**, **cpu**, and **memory** parameters are optional:

```
apiVersion: "v1"
kind: "BuildConfig"
metadata:
  name: "sample-build"
spec:
  resources:
    limits:
      cpu: "100m" ❶
      memory: "256Mi" ❷
```

❶ **cpu** is in CPU units: **100m** represents 0.1 CPU units ($100 * 1e-3$).

❷ **memory** is in bytes: **256Mi** represents 268435456 bytes ($256 * 2^{20}$).

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

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

```
resources:
  requests: ❶
    cpu: "100m"
    memory: "256Mi"
```

❶ The **requests** object contains the list of resources that correspond to the list of resources in the quota.

- A **limit range** defined in your project, where the defaults from the **LimitRange** object apply to pods created during the build process.

Otherwise, build pod creation will fail, citing a failure to satisfy quota.

8.9.2. Setting Maximum Duration

When defining a **BuildConfig**, you can define its maximum duration by setting the **completionDeadlineSeconds** field. It is specified in seconds and is not set by default. When not set, there is no maximum duration enforced.

The maximum duration is counted from the time when a build pod gets scheduled in the system, and defines how long it can be active, including the time needed to pull the builder image. After reaching the specified timeout, the build is terminated by OpenShift Container Platform.

The following example shows the part of a **BuildConfig** specifying **completionDeadlineSeconds** field for 30 minutes:

```
spec:
  completionDeadlineSeconds: 1800
```

8.9.3. Assigning Builds to Specific Nodes


Builds can be targeted to run on specific nodes by specifying labels in the `nodeSelector` field of a build configuration. The `nodeSelector` value is a set of key/value pairs that are matched to `node` labels when scheduling the build pod.

```
apiVersion: "v1"
kind: "BuildConfig"
metadata:
  name: "sample-build"
spec:
  nodeSelector: 1
    key1: value1
    key2: value2
```

- 1 Builds associated with this build configuration will run only on nodes with the `key1=value1` and `key2=value2` labels.

The `nodeSelector` value can also be controlled by cluster-wide default and override values. Defaults will only be applied if the build configuration does not define any key/value pairs for the `nodeSelector` and also does not define an explicitly empty map value of `nodeSelector: {}`. Override values will replace values in the build configuration on a key by key basis.

See [Configuring Global Build Defaults and Overrides](#) for more information.



NOTE

If the specified `NodeSelector` cannot be matched to a node with those labels, the build still stay in the **Pending** state indefinitely.

8.9.4. Chaining Builds


For compiled languages (Go, C, C++, Java, etc.), including the dependencies necessary for compilation in the application image might increase the size of the image or introduce vulnerabilities that can be exploited.

To avoid these problems, two builds can be chained together: one that produces the compiled artifact, and a second build that places that artifact in a separate image that runs the artifact. In the following example, a [Source-to-Image](#) build is combined with a [Docker](#) build to compile an artifact that is then placed in a separate runtime image.



NOTE

Although this example chains a Source-to-Image build and a Docker build, the first build can use any strategy that will produce an image containing the desired artifacts, and the second build can use any strategy that can consume input content from an image.



OPENSIFT_466208_0218

The first build takes the application source and produces an image containing a WAR file. The image is pushed to the `artifact-image` image stream. The path of the output artifact will depend on the **`assemble`** script of the Source-to-Image builder used. In this case, it will be output to **`/wildfly/standalone/deployments/ROOT.war`**.

```

apiVersion: v1
kind: BuildConfig
metadata:
  name: artifact-build
spec:
  output:
    to:
      kind: ImageStreamTag
      name: artifact-image:latest
  source:
    git:
      uri: https://github.com/openshift/openshift-jee-sample.git
  strategy:
    sourceStrategy:
      from:
        kind: ImageStreamTag
        name: wildfly:10.1
        namespace: openshift

```

The second build uses **Image Source** with a path to the WAR file inside the output image from the first build. An inline **`Dockerfile`** copies that WAR file into a runtime image.

```

apiVersion: v1
kind: BuildConfig
metadata:
  name: image-build
spec:
  output:

```

```

to:
  kind: ImageStreamTag
  name: image-build:latest
source:
  dockerfile: |-
    FROM jee-runtime:latest
    COPY ROOT.war /deployments/ROOT.war
  images:
    - from: ❶
      kind: ImageStreamTag
      name: artifact-image:latest
      paths: ❷
      - sourcePath: /wildfly/standalone/deployments/ROOT.war
        destinationDir: "."
  strategy:
    dockerStrategy:
      from: ❸
      kind: ImageStreamTag
      name: jee-runtime:latest
  triggers:
    - imageChange: {}
      type: ImageChange

```

- ❶ **from** specifies that the Docker build should include the output of the image from the **artifact-image** image stream, which was the target of the previous build.
- ❷ **paths** specifies which paths from the target image to include in the current Docker build.
- ❸ The runtime image is used as the source image for the Docker build.

The result of this setup is that the output image of the second build does not need to contain any of the build tools that are needed to create the WAR file. Also, because the second build contains an [image change trigger](#), whenever the first build is run and produces a new image with the binary artifact, the second build is automatically triggered to produce a runtime image that contains that artifact. Therefore, both builds behave as a single build with two stages.

8.10. BUILD TROUBLESHOOTING

8.10.1. Requested Access to Resources Denied

Issue

A build fails with:

```
requested access to the resource is denied
```

Resolution

You have exceeded one of the [image quotas](#) set on your project. Check your current quota and verify the limits applied and storage in use:

```
$ oc describe quota
```



CHAPTER 9. DEPLOYMENTS

9.1. HOW DEPLOYMENTS WORK

9.1.1. What Is a Deployment?

OpenShift Container Platform deployments provide fine-grained management over common user applications. They are described using three separate API objects:

- A deployment configuration, which describes the desired state of a particular component of the application as a pod template.
- One or more replication controllers, which contain a point-in-time record of the state of a deployment configuration as a pod template.
- One or more pods, which represent an instance of a particular version of an application.



IMPORTANT

Users do not need to manipulate replication controllers or pods owned by deployment configurations. The deployment system ensures changes to deployment configurations are propagated appropriately. If the existing deployment strategies are not suited for your use case and you have the need to run manual steps during the lifecycle of your deployment, then you should consider creating a [custom strategy](#).

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

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

Features provided by the deployment system:

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

9.1.2. Creating a Deployment Configuration

Deployment configurations are **deploymentConfig** OpenShift Container Platform API resources which can be managed with the **oc** command like any other resource. The following is an example of a **deploymentConfig** resource:

```
kind: "DeploymentConfig"
apiVersion: "v1"
metadata:
  name: "frontend"
spec:
  template: ❶
    metadata:
      labels:
        name: "frontend"
    spec:
      containers:
        - name: "helloworld"
          image: "openshift/origin-ruby-sample"
          ports:
            - containerPort: 8080
              protocol: "TCP"
      replicas: 5 ❷
      triggers:
        - type: "ConfigChange" ❸
        - type: "ImageChange" ❹
          imageChangeParams:
            automatic: true
            containerNames:
              - "helloworld"
            from:
              kind: "ImageStreamTag"
              name: "origin-ruby-sample:latest"
      strategy: ❺
        type: "Rolling"
      paused: false ❻
      revisionHistoryLimit: 2 ❼
      minReadySeconds: 0 ❽
```

- ❶ The pod template of the **frontend** deployment configuration describes a simple Ruby application.
- ❷ There will be 5 replicas of **frontend**.
- ❸ A [configuration change trigger](#) causes a new replication controller to be created any time the pod template changes.
- ❹ An [image change trigger](#) trigger causes a new replication controller to be created each time a new version of the **origin-ruby-sample:latest** image stream tag is available.
- ❺ The [Rolling strategy](#) is the default way of deploying your pods. May be omitted.
- ❻ Pause a deployment configuration. This disables the functionality of all triggers and allows for multiple changes on the pod template before actually rolling it out.
- ❼ Revision history limit is the limit of old replication controllers you want to keep around for rolling back. May be omitted. If omitted, old replication controllers will not be cleaned up.

- 8** Minimum seconds to wait (after the readiness checks succeed) for a pod to be considered available. The default value is 0.

9.2. BASIC DEPLOYMENT OPERATIONS

9.2.1. Starting a Deployment

You can start a new deployment process manually using the web console, or from the CLI:

```
$ oc deploy --latest dc/<name>
```



NOTE

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

9.2.2. Viewing a Deployment

To get basic information about all the available revisions of your application:

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


This will show details about all recently created replication controllers for the provided deployment configuration, including any currently running deployment process.

You can view details specific to a revision by using the `--revision` flag:

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

For more detailed information about a deployment configuration and its latest revision:

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



NOTE

The [web console](#) shows deployments in the **Browse** tab.

9.2.3. Rolling Back a Deployment

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

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

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

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

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

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



NOTE

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

9.2.4. Executing Commands Inside a Container

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

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

```
...
spec:
  containers:
  -
    name: <container_name>
    image: 'image'
    command:
    - '<command>'
    args:
    - '<argument_1>'
    - '<argument_2>'
    - '<argument_3>'
  ...
```

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

```
...
spec:
  containers:
  -
    name: example-spring-boot
    image: 'image'
    command:
    - java
    args:
    - '-jar'
    - /opt/app-root/springboots2idemo.jar
  ...
```

9.2.5. Viewing Deployment Logs

To stream the logs of the latest revision for a given deployment configuration:

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

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

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

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

For more options on retrieving logs see:

```
$ oc logs --help
```

9.2.6. Setting Deployment Triggers

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



WARNING

If no triggers are defined on a deployment configuration, a **ConfigChange** trigger is added by default. If triggers are defined as an empty field, deployments must be [started manually](#).

9.2.6.1. Configuration Change Trigger

The **ConfigChange** trigger results in a new replication controller whenever changes are detected in the pod template of the deployment configuration.



NOTE

If a **ConfigChange** trigger is defined on a deployment configuration, the first replication controller will be automatically created soon after the deployment configuration itself is created and it is not paused.

Example 9.1. A ConfigChange Trigger

```
triggers:
  - type: "ConfigChange"
```

9.2.6.2. ImageChange Trigger


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

Example 9.2. An ImageChange Trigger

```
triggers:
- type: "ImageChange"
  imageChangeParams:
    automatic: true 1
    from:
      kind: "ImageStreamTag"
      name: "origin-ruby-sample:latest"
      namespace: "myproject"
    containerNames:
      - "helloworld"
```

1 If the `imageChangeParams.automatic` field is set to `false`, the trigger is disabled.

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



NOTE

If an **ImageChange** trigger is defined on a deployment configuration (with a **ConfigChange** trigger and `automatic=false`, or with `automatic=true`) and the **ImageStreamTag** pointed by the **ImageChange** trigger does not exist yet, then the initial deployment process will automatically start as soon as an image is imported or pushed by a build to the **ImageStreamTag**.

9.2.6.2.1. Using the Command Line

The `oc set triggers` command can be used to set a deployment trigger for a deployment configuration. For the example above, you can set the **ImageChangeTrigger** by using the following command:

```
$ oc set triggers dc/frontend --from-image=myproject/origin-ruby-sample:latest -c helloworld
```

For more information, see:

```
$ oc set triggers --help
```

9.2.7. Setting Deployment Resources

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

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

In the following example, each of **resources**, **cpu**, and **memory** is optional:

```
type: "Recreate"
resources:
  limits:
    cpu: "100m" ❶
    memory: "256Mi" ❷
```

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

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

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

```
type: "Recreate"
resources:
  requests: ❶
    cpu: "100m"
    memory: "256Mi"
```

- ❶ The **requests** object contains the list of resources that correspond to the list of resources in the quota.

See [Quotas and Limit Ranges](#) to learn more about compute resources and the differences between requests and limits.

- A [limit range](#) defined in your project, where the defaults from the `LimitRange` object apply to pods created during the deployment process.

Otherwise, deploy pod creation will fail, citing a failure to satisfy quota.

9.2.8. Manual Scaling

In addition to rollbacks, you can exercise fine-grained control over the number of replicas from the web console, or by using the `oc scale` command. For example, the following command sets the replicas in the deployment configuration **frontend** to 3.

```
$ oc scale dc frontend --replicas=3
```

The number of replicas eventually propagates to the desired and current state of the deployment configured by the deployment configuration **frontend**.



NOTE

Pods can also be autoscaled using the `oc autoscale` command. See [Pod Autoscaling](#) for more details.

9.2.9. Assigning Pods to Specific Nodes

You can use node selectors in conjunction with labeled nodes to control pod placement.



NOTE

OpenShift Container Platform administrators can assign labels [during an advanced installation](#), or [added to a node after installation](#).

Cluster administrators [can set the default node selector](#) for your project in order to restrict pod placement to specific nodes. As an OpenShift Container Platform developer, you can set a node selector on a pod configuration to restrict nodes even further.

To add a node selector when creating a pod, edit the pod configuration, and add the `nodeSelector` value. This can be added to a single pod configuration, or in a pod template:

```
apiVersion: v1
kind: Pod
spec:
  nodeSelector:
    disktype: ssd
  ...
```

Pods created when the node selector is in place are assigned to nodes with the specified labels.

The labels specified here are used in conjunction with the labels [added by a cluster administrator](#).

For example, if a project has the `type=user-node` and `region=east` labels added to a project by the cluster administrator, and you add the above `disktype: ssd` label to a pod, the pod will only ever be scheduled on nodes that have all three labels.



NOTE

Labels can only be set to one value, so setting a node selector of `region=west` in a pod configuration that has `region=east` as the administrator-set default, results in a pod that will never be scheduled.

9.2.10. Running a Pod with a Different Service Account

You can run a pod with a service account other than the default:

1. Edit the deployment configuration:

```
$ oc edit dc/<deployment_config>
```

2. Add the `serviceAccount` and `serviceAccountName` parameters to the `spec` field, and specify the service account you want to use:

```
spec:
  securityContext: {}
  serviceAccount: <service_account>
  serviceAccountName: <service_account>
```


9.2.11. Adding Secrets to Deployment Configurations from the Web Console

Add a secret to your deployment configuration so that it can access a private repository.

1. Create a new OpenShift Container Platform project.
2. [Create a secret](#) that contains credentials for accessing a private image repository.
3. Create a deployment configuration.
4. On the deployment configuration editor page or in the `fromimage` page of the [web console](#), set the **Pull Secret**.
5. Click the **Save** button.

9.3. DEPLOYMENT STRATEGIES

9.3.1. What Are Deployment Strategies?

A deployment strategy determines the deployment process, and is defined by the deployment configuration. Each application has different requirements for availability (and other considerations) during deployments. OpenShift Container Platform provides strategies to support a variety of deployment scenarios.

A deployment strategy uses [readiness checks](#) to determine if a new pod is ready for use. If a readiness check fails, the deployment configuration will retry to run the pod until it times out. The default timeout is `10m`, a value set in `TimeoutSeconds` in `dc.spec.strategy.*params`.

The [Rolling strategy](#) is the default strategy used if no strategy is specified on a deployment configuration.

9.3.2. Rolling Strategy

A rolling deployment slowly replaces instances of the previous version of an application with instances of the new version of the application. A rolling deployment typically waits for new pods to become **ready** via a **readiness check** before scaling down the old components. If a significant issue occurs, the rolling deployment can be aborted.

9.3.2.1. Canary Deployments

All rolling deployments in OpenShift Container Platform are *canary* deployments; a new version (the canary) is tested before all of the old instances are replaced. If the readiness check never succeeds, the canary instance is removed and the deployment configuration will be automatically rolled back. The readiness check is part of the application code, and may be as sophisticated as necessary to ensure the new instance is ready to be used. If you need to implement more complex checks of the application (such as sending real user workloads to the new instance), consider implementing a custom deployment or using a blue-green deployment strategy.

9.3.2.2. When to Use a Rolling Deployment

- When you want to take no downtime during an application update.
- When your application supports having old code and new code running at the same time.

A rolling deployment means you to have both old and new versions of your code running at the same

time. This typically requires that your application handle [N-1 compatibility](#), that data stored by the new version can be read and handled (or gracefully ignored) by the old version of the code. This can take many forms – data stored on disk, in a database, in a temporary cache, or that is part of a user’s browser session. While most web applications can support rolling deployments, it is important to test and design your application to handle it.


The following is an example of the Rolling strategy:

```
strategy:
  type: Rolling
  rollingParams:
    updatePeriodSeconds: 1 1
    intervalSeconds: 1 2
    timeoutSeconds: 120 3
    maxSurge: "20%" 4
    maxUnavailable: "10%" 5
    pre: {} 6
    post: {}
```

- 1 The time to wait between individual pod updates. If unspecified, this value defaults to 1.
- 2 The time to wait between polling the deployment status after update. If unspecified, this value defaults to 1.
- 3 The time to wait for a scaling event before giving up. Optional; the default is 600. Here, *giving up* means automatically rolling back to the previous complete deployment.
- 4 **maxSurge** is optional and defaults to 25% if not specified. See the information below the following procedure.
- 5 **maxUnavailable** is optional and defaults to 25% if not specified. See the information below the following procedure.
- 6 **pre** and **post** are both [lifecycle hooks](#).

The Rolling strategy will:

1. Execute any **pre** lifecycle hook.
2. Scale up the new replication controller based on the surge count.
3. Scale down the old replication controller based on the max unavailable count.
4. Repeat this scaling until the new replication controller has reached the desired replica count and the old replication controller has been scaled to zero.
5. Execute any **post** lifecycle hook.



IMPORTANT

When scaling down, the Rolling strategy waits for pods to become ready so it can decide whether further scaling would affect availability. If scaled up pods never become ready, the deployment process will eventually time out and result in a deployment failure.

The `maxUnavailable` parameter is the maximum number of pods that can be unavailable during the update. The `maxSurge` parameter is the maximum number of pods that can be scheduled above the original number of pods. Both parameters can be set to either a percentage (e.g., 10%) or an absolute value (e.g., 2). The default value for both is 25%.

These parameters allow the deployment to be tuned for availability and speed. For example:

- `maxUnavailable=0` and `maxSurge=20%` ensures full capacity is maintained during the update and rapid scale up.
- `maxUnavailable=10%` and `maxSurge=0` performs an update using no extra capacity (an in-place update).
- `maxUnavailable=10%` and `maxSurge=10%` scales up and down quickly with some potential for capacity loss.

Generally, if you want fast rollouts, use `maxSurge`. If you need to take into account resource quota and can accept partial unavailability, use `maxUnavailable`.

9.3.2.3. Rolling Example

Rolling deployments are the default in OpenShift Container Platform. To see a rolling update, follow these steps:

1. Create an application based on the example deployment images found in [DockerHub](#):

```
$ oc new-app openshift/deployment-example
```

If you have the router installed, make the application available via a route (or use the service IP directly)

```
$ oc expose svc/deployment-example
```

Browse to the application at `deployment-example.<project>.<router_domain>` to verify you see the v1 image.

2. Scale the deployment configuration up to three replicas:

```
$ oc scale dc/deployment-example --replicas=3
```

3. Trigger a new deployment automatically by tagging a new version of the example as the `latest` tag:

```
$ oc tag deployment-example:v2 deployment-example:latest
```

4. In your browser, refresh the page until you see the v2 image.
5. If you are using the CLI, the following command will show you how many pods are on version 1 and how many are on version 2. In the web console, you should see the pods slowly being added to v2 and removed from v1.

```
$ oc describe dc deployment-example
```

During the deployment process, the new replication controller is incrementally scaled up. Once the

new pods are marked as **ready** (by passing their readiness check), the deployment process will continue. If the pods do not become ready, the process will abort, and the deployment configuration will be rolled back to its previous version.

9.3.3. Recreate Strategy

The Recreate strategy has basic rollout behavior and supports [lifecycle hooks](#) for injecting code into the deployment process.

The following is an example of the Recreate strategy:

```
strategy:
  type: Recreate
  recreateParams: ❶
    pre: {} ❷
    mid: {}
    post: {}
```

- ❶ `recreateParams` are optional.
- ❷ `pre`, `mid`, and `post` are [lifecycle hooks](#).

The Recreate strategy will:

1. Execute any `pre` lifecycle hook.
2. Scale down the previous deployment to zero.
3. Execute any `mid` lifecycle hook.
4. Scale up the new deployment.
5. Execute any `post` lifecycle hook.



IMPORTANT

During scale up, if the replica count of the deployment is greater than one, the first replica of the deployment will be validated for readiness before fully scaling up the deployment. If the validation of the first replica fails, the deployment will be considered a failure.

9.3.3.1. When to Use a Recreate Deployment

- When you must run migrations or other data transformations before your new code starts.
- When you do not support having new and old versions of your application code running at the same time.
- When you want to use a RWO volume, which is not supported being shared between multiple replicas.

A recreate deployment incurs downtime because, for a brief period, no instances of your application are running. However, your old code and new code do not run at the same time.

9.3.4. Custom Strategy

The Custom strategy allows you to provide your own deployment behavior.

The following is an example of the Custom strategy:

```
strategy:
  type: Custom
  customParams:
    image: organization/strategy
    command: [ "command", "arg1" ]
    environment:
      - name: ENV_1
        value: VALUE_1
```

In the above example, the **organization/strategy** container image provides the deployment behavior. The optional **command** array overrides any **CMD** directive specified in the image's *Dockerfile*. The optional environment variables provided are added to the execution environment of the strategy process.

Additionally, OpenShift Container Platform provides the following environment variables to the deployment process:

Environment Variable	Description
OPENSIFT_DEPLOYMENT_NAME	The name of the new deployment (a replication controller).
OPENSIFT_DEPLOYMENT_NAMESPACE	The name space of the new deployment.

The replica count of the new deployment will initially be zero. The responsibility of the strategy is to make the new deployment active using the logic that best serves the needs of the user.

Learn more about [advanced deployment strategies](#).

Alternatively, use **customParams** to inject the custom deployment logic into the existing deployment strategies. Provide a custom shell script logic and call the **openshift-deploy** binary. Users do not have to supply their custom deployer container image, but the default OpenShift Container Platform deployer image will be used instead:

```
strategy:
  type: Rolling
  customParams:
    command:
      - /bin/sh
      - -c
      - |
        set -e
        openshift-deploy --until=50%
        echo Halfway there
        openshift-deploy
        echo Complete
```

This will result in following deployment:

```
Started deployment #2
--> Scaling up custom-deployment-2 from 0 to 2, scaling down custom-
deployment-1 from 2 to 0 (keep 2 pods available, don't exceed 3 pods)
    Scaling custom-deployment-2 up to 1
--> Reached 50% (currently 50%)
Halfway there
--> Scaling up custom-deployment-2 from 1 to 2, scaling down custom-
deployment-1 from 2 to 0 (keep 2 pods available, don't exceed 3 pods)
    Scaling custom-deployment-1 down to 1
    Scaling custom-deployment-2 up to 2
    Scaling custom-deployment-1 down to 0
--> Success
Complete
```

If the custom deployment strategy process requires access to the OpenShift Container Platform API or the Kubernetes API the container that executes the strategy can use the service account token available inside the container for authentication.

9.3.5. Lifecycle Hooks

The [Recreate](#) and [Rolling](#) strategies support lifecycle hooks, which allow behavior to be injected into the deployment process at predefined points within the strategy:

The following is an example of a **pre** lifecycle hook:

```
pre:
  failurePolicy: Abort
  execNewPod: {} 1
```

1 `execNewPod` is a [pod-based lifecycle hook](#).

Every hook has a **failurePolicy**, which defines the action the strategy should take when a hook failure is encountered:

Abort	The deployment process will be considered a failure if the hook fails.
Retry	The hook execution should be retried until it succeeds.
Ignore	Any hook failure should be ignored and the deployment should proceed.

Hooks have a type-specific field that describes how to execute the hook. Currently, [pod-based hooks](#) are the only supported hook type, specified by the `execNewPod` field.

9.3.5.1. Pod-based Lifecycle Hook

Pod-based lifecycle hooks execute hook code in a new pod derived from the template in a deployment configuration.

The following simplified example deployment configuration uses the [Rolling strategy](#). Triggers and some other minor details are omitted for brevity:

```
kind: DeploymentConfig
apiVersion: v1
metadata:
  name: frontend
spec:
  template:
    metadata:
      labels:
        name: frontend
    spec:
      containers:
        - name: helloworld 1
          image: openshift/origin-ruby-sample
      replicas: 5
      selector:
        name: frontend
      strategy:
        type: Rolling
        rollingParams:
          pre:
            failurePolicy: Abort
            execNewPod:
              containerName: helloworld 1
              command: [ "/usr/bin/command", "arg1", "arg2" ] 2
              env: 3
                - name: CUSTOM_VAR1
                  value: custom_value1
              volumes:
                - data 4
```

- 1** The `helloworld` name refers to `spec.template.spec.containers[0].name`.
- 2** This `command` overrides any `ENTRYPOINT` defined by the `openshift/origin-ruby-sample` image.
- 3** `env` is an optional set of environment variables for the hook container.
- 4** `volumes` is an optional set of volume references for the hook container.

In this example, the `pre` hook will be executed in a new pod using the `openshift/origin-ruby-sample` image from the `helloworld` container. The hook pod will have the following properties:

- The hook command will be `/usr/bin/command arg1 arg2`.
- The hook container will have the `CUSTOM_VAR1=custom_value1` environment variable.
- The hook failure policy is `Abort`, meaning the deployment process will fail if the hook fails.
- The hook pod will inherit the `data` volume from the deployment configuration pod.

9.3.5.2. Using the Command Line

The `oc set deployment-hook` command can be used to set the deployment hook for a deployment configuration. For the example above, you can set the pre-deployment hook with the following command:

```
$ oc set deployment-hook dc/frontend --pre -c helloworld -e
CUSTOM_VAR1=custom_value1 \
-v data --failure-policy=abort -- /usr/bin/command arg1 arg2
```

9.4. ADVANCED DEPLOYMENT STRATEGIES

9.4.1. Blue-Green Deployment

Blue-green deployments involve running two versions of an application at the same time and moving production traffic from the old version to the new version. There are several ways to implement a blue-green deployment in OpenShift Container Platform.

9.4.1.1. When to Use a Blue-Green Deployment

Use a blue-green deployment when you want to test a new version of your application in a production environment before moving traffic to it.

Blue-green deployments make switching between two different versions of your application easy. However, since many applications depend on persistent data, you will need to have an application that supports [N-1 compatibility](#) if you share a database, or implement a live data migration between your database, store, or disk if you choose to create two copies of your data layer.

9.4.1.2. Blue-Green Deployment Example

In order to maintain control over two distinct groups of instances (old and new versions of the code), the blue-green deployment is best represented with multiple deployment configurations.

9.4.1.2.1. Using a Route and Two Services

A route points to a service, and can be changed to point to a different service at any time. As a developer, test the new version of your code by connecting to the new service before your production traffic is routed to it. Routes are intended for web (HTTP and HTTPS) traffic, so this technique is best suited for web applications.

1. Create two copies of the example application:

```
$ oc new-app openshift/deployment-example:v1 --name=example-green
$ oc new-app openshift/deployment-example:v2 --name=example-blue
```

This will create two independent application components: one running the **v1** image under the **example-green** service, and one using the **v2** image under the **example-blue** service.

2. Create a route that points to the old service:

```
$ oc expose svc/example-green --name=bluegreen-example
```

3. Browse to the application at **bluegreen-example.<project>.<router_domain>** to verify you see the **v1** image.

**NOTE**

On versions of OpenShift Container Platform older than v3.0.1, this command will generate a route at `example-green.<project>.<router_domain>`, not the above location.

4. Edit the route and change the service name to **example-blue**:

```
$ oc patch route/bluegreen-example -p '{"spec":{"to":{"name":"example-blue"}}}'
```

5. In your browser, refresh the page until you see the v2 image.

9.4.2. A/B Deployment

A/B deployments generally imply running two (or more) versions of the application code or application configuration at the same time for testing or experimentation purposes.

The simplest form of an A/B deployment is to divide production traffic between two or more distinct **shards** – a single group of instances with homogeneous configuration and code.

More complicated A/B deployments may involve a specialized proxy or load balancer that assigns traffic to specific shards based on information about the user or application (all "test" users get sent to the B shard, but regular users get sent to the A shard).

A/B deployments can be considered similar to A/B testing, although an A/B deployment implies multiple versions of code and configuration, whereas A/B testing often uses one code base with application specific checks.

9.4.2.1. When to Use an A/B Deployment

- When you want to test multiple versions of code or configuration, but are not planning to roll one out in preference to the other.
- When you want to have different configuration in different regions.

An A/B deployment groups different configuration and code – multiple shards – together under a single logical endpoint. Generally, these deployments, if they access persistent data, should properly deal with N-1 compatibility (the more shards you have, the more possible versions you have running). Use this pattern when you need separate internal configuration and code, but end users should not be aware of the changes.

9.4.2.2. A/B Deployment Example

All A/B deployments are composite deployment types consisting of multiple deployment configurations.

9.4.2.2.1. One Service, Multiple Deployment Configurations

OpenShift Container Platform, through labels and deployment configurations, supports multiple simultaneous shards being exposed through the same service. To the consuming user, the shards are invisible. An example of the simplest possible sharding is described below:

1. Create the first shard of the application based on the example deployment images:

■

```
$ oc new-app openshift/deployment-example --name=ab-example-a --  
labels=ab-example=true SUBTITLE="shard A"
```

2. Edit the newly created shard to set a label **ab-example=true** that will be common to all shards:

```
$ oc edit dc/ab-example-a
```

In the editor, add the line **ab-example: "true"** underneath **spec.selector** and **spec.template.metadata.labels** alongside the existing **deploymentconfig=ab-example-a** label. Save and exit the editor.

3. Trigger a re-deployment of the first shard to pick up the new labels:

```
$ oc deploy ab-example-a --latest
```

4. Create a service that uses the common label:

```
$ oc expose dc/ab-example-a --name=ab-example --selector=ab-  
example=true
```

If you have the router installed, make the application available via a route (or use the service IP directly):

```
$ oc expose svc/ab-example
```

Browse to the application at **ab-example.<project>.<router_domain>** to verify you see the **v1** image.

5. Create a second shard based on the same source image as the first shard but different tagged version, and set a unique value:

```
$ oc new-app openshift/deployment-example:v2 --name=ab-example-b --  
labels=ab-example=true SUBTITLE="shard B" COLOR="red"
```

6. Edit the newly created shard to set a label **ab-example=true** that will be common to all shards:

```
$ oc edit dc/ab-example-b
```

In the editor, add the line **ab-example: "true"** underneath **spec.selector** and **spec.template.metadata.labels** alongside the existing **deploymentconfig=ab-example-b** label. Save and exit the editor.

7. Trigger a re-deployment of the second shard to pick up the new labels:

```
$ oc deploy ab-example-b --latest
```

8. At this point, both sets of pods are being served under the route. However, since both browsers (by leaving a connection open) and the router (by default, through a cookie) will attempt to preserve your connection to a back-end server, you may not see both shards being returned to you. To force your browser to one or the other shard, use the **scale** command:

```
$ oc scale dc/ab-example-a --replicas=0
```

Refreshing your browser should show **v2** and **shard B** (in red).

```
$ oc scale dc/ab-example-a --replicas=1; oc scale dc/ab-example-b --replicas=0
```

Refreshing your browser should show **v1** and **shard A** (in blue).

If you trigger a deployment on either shard, only the pods in that shard will be affected. You can easily trigger a deployment by changing the **SUBTITLE** environment variable in either deployment config `oc edit dc/ab-example-a` or `oc edit dc/ab-example-b`. You can add additional shards by repeating steps 5-7.



NOTE

These steps will be simplified in future versions of OpenShift Container Platform.

9.4.3. Proxy Shard / Traffic Splitter

In production environments, you can precisely control the distribution of traffic that lands on a particular shard. When dealing with large numbers of instances, you can use the relative scale of individual shards to implement percentage based traffic. That combines well with a **proxy shard**, which forwards or splits the traffic it receives to a separate service or application running elsewhere.

In the simplest configuration, the proxy would forward requests unchanged. In more complex setups, you can duplicate the incoming requests and send to both a separate cluster as well as to a local instance of the application, and compare the result. Other patterns include keeping the caches of a DR installation warm, or sampling incoming traffic for analysis purposes.

While an implementation is beyond the scope of this example, any TCP (or UDP) proxy could be run under the desired shard. Use the `oc scale` command to alter the relative number of instances serving requests under the proxy shard. For more complex traffic management, consider customizing the OpenShift Container Platform router with proportional balancing capabilities.

9.4.4. N-1 Compatibility

Applications that have new code and old code running at the same time must be careful to ensure that data written by the new code can be read by the old code. This is sometimes called *schema evolution* and is a complex problem.

For some applications, the period of time that old code and new code is running side by side is short, so bugs or some failed user transactions are acceptable. For others, the failure pattern may result in the entire application becoming non-functional.

One way to validate N-1 compatibility is to use an A/B deployment. Run the old code and new code at the same time in a controlled way in a test environment, and verify that traffic that flows to the new deployment does not cause failures in the old deployment.

9.4.5. Graceful Termination

OpenShift Container Platform and Kubernetes give application instances time to shut down before removing them from load balancing rotations. However, applications must ensure they cleanly terminate user connections as well before they exit.

On shutdown, OpenShift Container Platform will send a **TERM** signal to the processes in the container. Application code, on receiving **SIGTERM**, should stop accepting new connections. This will ensure that load balancers route traffic to other active instances. The application code should then wait until all open connections are closed (or gracefully terminate individual connections at the next opportunity) before exiting.

After the graceful termination period expires, a process that has not exited will be sent the **KILL** signal, which immediately ends the process. The `terminationGracePeriodSeconds` attribute of a pod or pod template controls the graceful termination period (default 30 seconds) and may be customized per application as necessary.

9.5. KUBERNETES DEPLOYMENTS SUPPORT

9.5.1. New Object Type: Deployments

In the upstream Kubernetes project, a new first-class object type called *deployments* was added in version 1.2. This object type (referred to here as *Kubernetes deployments* for distinction) serves as a descendant of the deployment configuration object type.

Support for Kubernetes deployments is available as a [Technology Preview](#) feature.

Like deployment configurations, Kubernetes deployments describe the desired state of a particular component of an application as a pod template. Kubernetes deployments create *replica sets* (an iteration of [replication controllers](#)), which orchestrate pod lifecycles.

For example, this definition of a Kubernetes deployment creates a replica set to bring up one **hello-openshift** pod:

Example Kubernetes Deployment Definition *hello-openshift-deployment.yaml*

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: hello-openshift
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: hello-openshift
    spec:
      containers:
      - name: hello-openshift
        image: openshift/hello-openshift:latest
        ports:
        - containerPort: 80
```

After saving the definition to a local file, you could then use it to create a Kubernetes deployment:

```
$ oc create -f hello-openshift-deployment.yaml
```

You can use the CLI to inspect and operate on Kubernetes deployments and replica sets like other object types, as described in [Common Operations](#), like `get` and `describe`. For the object type, use `deployments` or `deploy` for Kubernetes deployments and `replicasets` or `rs` for replica sets.

See the Kubernetes documentation for more details about [Deployments](#) and [Replica Sets](#), substituting `oc` for `kubectl` in CLI usage examples.

9.5.2. Kubernetes Deployments vs Deployment Configurations

Because deployment configurations existed in OpenShift Container Platform prior to deployments being added in Kubernetes 1.2, the latter object type naturally diverges slightly from the former. The long-term goal in OpenShift Container Platform is to reach full feature parity in Kubernetes deployments and switch to using them as a single object type that provides fine-grained management over applications.

Kubernetes deployments are supported to ensure upstream projects and examples that use the new object type can run smoothly on OpenShift Container Platform. Given the current feature set of Kubernetes deployments, you may want to use them instead of deployment configurations in OpenShift Container Platform if you do not plan to use any of the following in particular:

- [image streams](#)
- [lifecycle hooks](#)
- [Custom deployment strategies](#)

The following sections go into more details on the differences between the two object types to further help you decide when you might want to use Kubernetes deployments over deployment configurations.

9.5.2.1. Deployment Configuration-Specific Features

9.5.2.1.1. Automatic Rollbacks

Kubernetes deployments do not support automatically rolling back to the last successfully deployed replica set in case of a failure. This feature should be added soon.

9.5.2.1.2. Triggers

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

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

At the moment, Kubernetes deployments do not support `ImageChange` triggers. A generic triggering mechanism has been proposed upstream, but it is unknown if and when it may be accepted. Eventually, a OpenShift Container Platform-specific mechanism could be implemented to layer on top of Kubernetes deployments, but it would be more desirable for it to exist as part of the Kubernetes core.

9.5.2.1.3. Lifecycle Hooks

Kubernetes deployments do not support any lifecycle hooks.

9.5.2.1.4. Custom Strategies

Kubernetes deployments do not yet support user-specified Custom deployment strategies yet.

9.5.2.1.5. Canary Deployments

Kubernetes deployments do not yet run canaries as part of a new rollout.

9.5.2.1.6. Test Deployments

Kubernetes deployments do not support running test tracks.

9.5.2.2. Kubernetes Deployment-Specific Features

9.5.2.2.1. Rollover

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

Deployment configurations can have at most one deployer pod running, otherwise multiple deployers end up fighting with each other trying to scale up what they think should be the newest replication controller. Because of this, only two replication controllers can be active at any point in time. Ultimately, this translates to faster rapid rollouts for Kubernetes deployments.

9.5.2.2.2. Proportional Scaling

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

Deployment configurations cannot be scaled when a rollout is ongoing because the deployment configuration controller will end up fighting with the deployer process about the size of the new replication controller.

9.5.2.2.3. Pausing Mid-rollout

Kubernetes deployments can be paused at any point in time, meaning you can also pause ongoing rollouts. On the other hand, you cannot pause deployer pods currently, so if you try to pause a deployment configuration in the middle of a rollout, the deployer process will not be affected and will continue until it finishes.

CHAPTER 10. TEMPLATES

10.1. OVERVIEW

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

You can create a list of objects from a template [using the CLI](#) or, if a [template has been uploaded](#) to your project or the global template library, [using the web console](#). For a curated set of templates, see the [OpenShift Image Streams and Templates library](#).

10.2. UPLOADING A TEMPLATE

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

To upload a template to your current project's template library, pass the JSON or YAML file with the following command:

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

You can upload a template to a different project using the `-n` option with the name of the project:


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

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

10.3. CREATING FROM TEMPLATES USING THE WEB CONSOLE

To create the objects from an [uploaded template](#) using the web console:

1. While in the desired project, click **Add to Project**:




2. Select a template from the list of templates in your project, or provided by the global template library:


[Browse Catalog](#) [Deploy Image](#) [Import YAML / JSON](#)


Choose from web frameworks, databases, and other components to add content to your project.


Filter by name or description


Languages



Java



JavaScript


.NET


Perl


PHP


Python


Ruby


Technologies

Continuous Integration & Deployment
 Automate the build, test, and deployment of your application with each new code revision.

Data Stores
 Store and manage collections of data.




Single Sign-On
 A centralized authentication server for users to log in, log out, register, and manage user accounts for applications and RESTful web services.

3. Modify template parameters in the template creation screen:



rails-postgresql-example 1
 Namespace: openshift
 An example Rails application with a PostgreSQL database

Images 2

-  \${NAMESPACE}/ruby:2.3
-  \${NAME}:latest
-  postgresql:9.5

Parameters 3

*** Name**

 The name assigned to all of the frontend objects defined in this template.

*** Namespace**

 The OpenShift Namespace where the ImageStream resides.

*** Memory Limit**

 Maximum amount of memory the Rails container can use.

*** Memory Limit (PostgreSQL)**

 Maximum amount of memory the PostgreSQL container can use.

*** Git Repository URL**

 The URL of the repository with your application source code.

Git Reference

 Set this to a branch name, tag or other ref of your repository if you are not using the default branch.

Context Directory

 Set this to the relative path to your project if it is not in the root of your repository.

Application Hostname

 The exposed hostname that will route to the Rails service, if left blank a value will be defaulted.

GitHub Webhook Secret

GITHUB WEBHOOK SECRET

A secret string used to configure the GitHub webhook.

Secret Key

Your secret key for verifying the integrity of signed cookies.

*** Application Username**

The application user that is used within the sample application to authorize access on pages.

*** Application Password**

The application password that is used within the sample application to authorize access on pages.

*** Rails Environment**

Environment under which the sample application will run. Could be set to production, development or test.

*** Database Service Name**

Database Username

Database Password

*** Database Name**

Maximum Database Connections

Shared Buffer Amount

Custom RubyGems Mirror URL

The custom RubyGems mirror URL

Labels [About labels](#)

The following labels are being added automatically. If you want to override them, you can do so below.

template	rails-postgresql-example
app	rails-postgresql-example

Each label is applied to each created resource.

Name	Value
<input type="text"/>	<input type="text"/>

[Add label](#)

Template name and description.

Container images included in the template.

Parameters defined by the template. You can edit values for parameters defined in the template [here](#).

[Labels](#) to assign to all items included in the template. You can add and edit labels for objects.

10.4. CREATING FROM TEMPLATES USING THE CLI

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

10.4.1. Labels

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

There is also the ability to add labels in the template from the command line.

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

10.4.2. Parameters

The list of parameters that you can override are listed in the [parameters section of the template](#). You can list them with the CLI by using the following command and specifying the file to be used:

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

Alternatively, if the template is already uploaded:

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

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

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

```

expression          user[A-Z0-9]{3}
POSTGRESQL_PASSWORD database password
expression          [a-zA-Z0-9]{8}
POSTGRESQL_DATABASE database name
root
POSTGRESQL_MAX_CONNECTIONS database max connections
10
POSTGRESQL_SHARED_BUFFERS database shared buffers
12MB

```

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

10.4.3. Generating a List of Objects

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

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

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

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

You can create objects from a template by processing the template and piping the output to **oc create**:

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

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

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

You can override any [parameter](#) values defined in the file by adding the **-p** option for each **<name>=<value>** pair you want to override. A parameter reference may appear in any text field inside the template items.

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

Example 10.1. Creating a List of Objects from a Template

```

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

```

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

```

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

```

```
| oc create -f -
```

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

```
$ cat postgres.env
POSTGRESQL_USER=bob
POSTGRESQL_DATABASE=mydatabase
$ oc process -f my-rails-postgresql --param-file=postgres.env
```

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

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

10.5. MODIFYING AN UPLOADED TEMPLATE

You can edit a template that has already been uploaded to your project by using the following command:

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

10.6. USING THE INSTANT APP AND QUICKSTART TEMPLATES

OpenShift Container Platform provides a number of default Instant App and Quickstart templates to make it easy to quickly get started creating a new application for different languages. Templates are provided for Rails (Ruby), Django (Python), Node.js, CakePHP (PHP), and Dancer (Perl). Your cluster administrator should have created these templates in the default, global **openshift** project so you have access to them. You can list the available default Instant App and Quickstart templates with:

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


If they are not available, direct your cluster administrator to the [Loading the Default Image Streams and Templates](#) topic.

By default, the templates build using a public source repository on [GitHub](#) that contains the necessary application code. In order to be able to modify the source and build your own version of the application, you must:

1. Fork the repository referenced by the template's default **SOURCE_REPOSITORY_URL** parameter.
2. Override the value of the **SOURCE_REPOSITORY_URL** parameter when creating from the template, specifying your fork instead of the default value.

By doing this, the build configuration created by the template will now point to your fork of the application code, and you can modify the code and rebuild the application at will.

A walkthrough of this process using the web console is provided in [Getting Started for Developers: Web Console](#).

**NOTE**

Some of the Instant App and Quickstart templates define a database [deployment configuration](#). The configuration they define uses ephemeral storage for the database content. These templates should be used for demonstration purposes only as all database data will be lost if the database pod restarts for any reason.

10.7. WRITING TEMPLATES

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

Example 10.2. A Simple Template Object Definition (YAML)

```

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

```

10.7.1. Description

The template description covers information that informs users what your template does and helps them find it when searching in the web console. In addition to general descriptive information, it includes a set of tags. Useful tags include the name of the language your template is related to (e.g., `java`, `php`, `ruby`, etc.).

Example 10.3. Template Description Metadata

```

kind: "Template"
apiVersion: "v1"
metadata:
  name: "cakephp-mysql-example" ❶
  annotations:
    openshift.io/display-name: "CakePHP MySQL Example" ❷
    description: "An example CakePHP application with a MySQL
database.\n\nFor more information see
https://github.com/openshift/cakephp-ex" ❸
    tags: "instant-app,php,cakephp,mysql" ❹
    iconClass: "icon-php" ❺
  message: "Your admin credentials are
${ADMIN_USERNAME}:${ADMIN_PASSWORD}" ❻

```

- ❶ The unique name of the template.
- ❷ A brief, user-friendly name, which can be employed by user interfaces.
- ❸ A description of the template. Include enough detail that the user will understand what is being deployed and any caveats they need to know before deploying. It should also provide links to additional information, such as a **README** file. Newline characters `\n` can be included to create paragraphs.
- ❹ Tags to be associated with the template for searching and grouping. Add tags that will include it into one of the provided catalog categories. Refer to the `id` and `categoryAliases` in `CATALOG_CATEGORIES` in the console's [constants file](#). The categories can also be [customized](#) for the whole cluster.
- ❺ An icon to be displayed with your template in the web console. Choose from our existing [logo icons](#) when possible. You can also use icons from [FontAwesome](#). Alternatively, provide icons through [CSS customizations](#) that can be added to an OpenShift Container Platform cluster that uses your template. You must specify an icon class that exists, or it will prevent falling back to the generic icon.
- ❻ An instructional message that is displayed when this template is instantiated. This field should inform the user how to use the newly created resources. Parameter substitution is performed on the message before being displayed so that generated credentials and other parameters can be included in the output. Include links to any next-steps documentation that users should follow.

10.7.2. Labels

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

Example 10.4. Template Object Labels

```

kind: "Template"
apiVersion: "v1"

```

```
...
labels:
  template: "cakephp-mysql-example" 1
```

- 1** A label that will be applied to all objects created from this template.

10.7.3. Parameters

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

Parameters can be referenced in two ways:

- As a string value by placing values in the form `${PARAMETER_NAME}` in any string field in the template.
- As a json/yaml value by placing values in the form `${{PARAMETER_NAME}}` in place of any field in the template.

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

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

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

A default value can be provided, which is used if the user does not supply a different value:

Example 10.5. Setting an Explicit Value as the Default Value

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

Parameter values can also be generated based on rules specified in the parameter definition:

Example 10.6. Generating a Parameter Value

```
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}`.

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

Example 10.7. A full template with parameter definitions and references

```
kind: Template
apiVersion: v1
objects:
- kind: BuildConfig
  apiVersion: v1
  metadata:
    name: cakephp-mysql-example
    annotations:
      description: Defines how to build the application
  spec:
    source:
      type: Git
      git:
        uri: "${SOURCE_REPOSITORY_URL}" ❶
        ref: "${SOURCE_REPOSITORY_REF}"
        contextDir: "${CONTEXT_DIR}"
- kind: DeploymentConfig
  apiVersion: v1
  metadata:
    name: frontend
  spec:
    replicas: "${REPLICA_COUNT}" ❷
parameters:
- name: SOURCE_REPOSITORY_URL ❸
  displayName: Source Repository URL ❹
  description: The URL of the repository with your application source
  code ❺
  value: https://github.com/openshift/cakephp-ex.git ❻
  required: true ❼
- name: GITHUB_WEBHOOK_SECRET
  description: A secret string used to configure the GitHub webhook
  generate: expression ❽
  from: "[a-zA-Z0-9]{40}" ❾
- name: REPLICA_COUNT
  description: Number of replicas to run
  value: "2"
```



```

    required: true
    message: "... The GitHub webhook secret is ${GITHUB_WEBHOOK_SECRET} ..."

```

10

- 1 This value will be replaced with the value of the **SOURCE_REPOSITORY_URL** parameter when the template is instantiated.
- 2 This value will be replaced with the unquoted value of the **REPLICA_COUNT** parameter when the template is instantiated.
- 3 The name of the parameter. This value is used to reference the parameter within the template.
- 4 The user-friendly name for the parameter. This will be displayed to users.
- 5 A description of the parameter. Provide more detailed information for the purpose of the parameter, including any constraints on the expected value. Descriptions should use complete sentences to follow the console's [text standards](#). Don't make this a duplicate of the display name.
- 6 A default value for the parameter which will be used if the user does not override the value when instantiating the template. Avoid using default values for things like passwords, instead use generated parameters in combination with Secrets.
- 7 Indicates this parameter is required, meaning the user cannot override it with an empty value. If the parameter does not provide a default or generated value, the user must supply a value.
- 8 A parameter which has its value generated.
- 9 The input to the generator. In this case, the generator will produce a 40 character alphanumeric value including upper and lowercase characters.
- 10 Parameters can be included in the template message. This informs the user about generated values.

10.7.4. Object List

The main portion of the template is the list of objects which will be created when the template is instantiated. This can be any [valid API object](#), such as a **BuildConfig**, **DeploymentConfig**, **Service**, etc. The object will be created exactly as defined here, with any parameter values substituted in prior to creation. The definition of these objects can reference parameters defined earlier.

```

kind: "Template"
apiVersion: "v1"
objects:
  - kind: "Service" 1
    apiVersion: "v1"
    metadata:
      name: "cakephp-mysql-example"
      annotations:
        description: "Exposes and load balances the application pods"
    spec:
      ports:

```

```

- name: "web"
  port: 8080
  targetPort: 8080
selector:
  name: "cakephp-mysql-example"

```

- 1 The definition of a **Service** which will be created by this template.



NOTE

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

10.7.5. Other Recommendations

- Group related services together in the management console by adding the `service.alpha.openshift.io/dependencies` annotation to the Service object in your template.

Example 10.8. Group the Frontend and Database Services Together on the Management Console Overview

```

kind: "Template"
apiVersion: "v1"
objects:
- kind: "Service"
  apiVersion: "v1"
  metadata:
    name: "frontend"
    annotations:
      "service.alpha.openshift.io/dependencies": "[{\"name\": \"database\", \"kind\": \"Service\"}]"
  ...
- kind: "Service"
  apiVersion: "v1"
  metadata:
    name: "database"

```

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

10.7.6. Creating a Template from Existing Objects

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

```
$ oc export all --as-template=<template_name> > <template_filename>
```

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

The object types included in **oc export all** are:

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

CHAPTER 11. OPENING A REMOTE SHELL TO CONTAINERS

11.1. OVERVIEW

The `oc rsh` command allows you to locally access and manage tools that are on the system. The secure shell (SSH) is the underlying technology and industry standard that provides a secure connection to the application. Access to applications with the shell environment is protected and restricted with Security-Enhanced Linux (SELinux) policies.

11.2. START A SECURE SHELL SESSION

Open a remote shell session to a container:

```
$ oc rsh <pod>
```

While in the remote shell, you can issue commands as if you are inside the container and perform local operations like monitoring, debugging, and using CLI commands specific to what is running in the container.

For example, in a MySQL container, you can count the number of records in the database by invoking the `mysql` command, then using the prompt to type in the `SELECT` command. You can also use commands like `ps(1)` and `ls(1)` for validation.

BuildConfigs and **DeployConfigs** map out how you want things to look and pods (with containers inside) are created and dismantled as needed. Your changes are not persistent. If you make changes directly within the container and that container is destroyed and rebuilt, your changes will no longer exist.



NOTE

`oc exec` can be used to execute a command remotely. However, the `oc rsh` command provides an easier way to keep a remote shell open persistently.

11.3. SECURE SHELL SESSION HELP

For help with usage, options, and to see examples:

```
$ oc rsh -h
```

CHAPTER 12. SERVICE ACCOUNTS

12.1. OVERVIEW

When a person uses the OpenShift Container Platform CLI or web console, their API token authenticates them to the OpenShift API. However, when a regular user's credentials are not available, it is common for components to make API calls independently. For example:

- Replication controllers make API calls to create or delete pods.
- Applications inside containers could make API calls for discovery purposes.
- External applications could make API calls for monitoring or integration purposes.

Service accounts provide a flexible way to control API access without sharing a regular user's credentials.

12.2. USER NAMES AND GROUPS

Every service account has an associated user name that can be granted roles, just like a regular user. The user name is derived from its project and name:

```
system:serviceaccount:<project>:<name>
```

For example, to add the **view** role to the **robot** service account in the **top-secret** project:

```
$ oc policy add-role-to-user view system:serviceaccount:top-secret:robot
```

IMPORTANT

If you want to grant access to a specific service account in a project, you can use the **-z** flag. From the project to which the service account belongs, use the **-z** flag and specify the **<serviceaccount_name>**. This is highly recommended, as it helps prevent typos and ensures that access is granted only to the specified service account. For example:

```
$ oc policy add-role-to-user <role_name> -z
<serviceaccount_name>
```

If not in the project, use the **-n** option to indicate the project namespace it applies to, as shown in the examples below.

Every service account is also a member of two groups:

system:serviceaccount

Includes all service accounts in the system.

system:serviceaccount:<project>

Includes all service accounts in the specified project.

For example, to allow all service accounts in all projects to view resources in the **top-secret** project:

```
$ oc policy add-role-to-group view system:serviceaccount -n top-secret
```

To allow all service accounts in the **managers** project to edit resources in the **top-secret** project:

```
$ oc policy add-role-to-group edit system:serviceaccount:managers -n top-secret
```

12.3. DEFAULT SERVICE ACCOUNTS AND ROLES

Three service accounts are automatically created in every project:

Service Account	Usage
builder	Used by build pods. It is given the system:image-builder role, which allows pushing images to any image stream in the project using the internal Docker registry.
deployer	Used by deployment pods and is given the system:deployer role, which allows viewing and modifying replication controllers and pods in the project.
default	Used to run all other pods unless they specify a different service account.

All service accounts in a project are given the **system:image-puller** role, which allows pulling images from any image stream in the project using the internal Docker registry.

12.4. MANAGING SERVICE ACCOUNTS

Service accounts are API objects that exist within each project. To manage service accounts, you can use the **oc** command with the **sa** or **serviceaccount** object type or use the web console.

To get a list of existing service accounts in the current project:

```
$ oc get sa
NAME          SECRETS  AGE
builder       2        2d
default       2        2d
deployer      2        2d
```

To create a new service account:

```
$ oc create sa robot
serviceaccount "robot" created
```

As soon as a service account is created, two secrets are automatically added to it:

- an API token
- credentials for the OpenShift Container Registry

These can be seen by describing the service account:

```
$ oc describe sa robot
Name: robot
```

```

Namespace: project1
Labels:    <none>
Annotations: <none>

Image pull secrets: robot-dockercfg-qzbhb

Mountable secrets:  robot-token-f4khf
                    robot-dockercfg-qzbhb

Tokens:             robot-token-f4khf
                    robot-token-z8h44

```

The system ensures that service accounts always have an API token and registry credentials.

The generated API token and registry credentials do not expire, but they can be revoked by deleting the secret. When the secret is deleted, a new one is automatically generated to take its place.

12.5. ENABLING SERVICE ACCOUNT AUTHENTICATION

Service accounts authenticate to the API using tokens signed by a private RSA key. The authentication layer verifies the signature using a matching public RSA key.

To enable service account token generation, update the `serviceAccountConfig` stanza in the `/etc/origin/master/master-config.yml` file on the master to specify a `privateKeyFile` (for signing), and a matching public key file in the `publicKeyFiles` list:

```

serviceAccountConfig:
  ...
  masterCA: ca.crt ❶
  privateKeyFile: serviceaccount.private.key ❷
  publicKeyFiles:
  - serviceaccount.public.key ❸
  - ...

```

- ❶ CA file used to validate the API server's serving certificate.
- ❷ Private RSA key file (for token signing).
- ❸ Public RSA key files (for token verification). If private key files are provided, then the public key component is used. Multiple public key files can be specified, and a token will be accepted if it can be validated by one of the public keys. This allows rotation of the signing key, while still accepting tokens generated by the previous signer.

12.6. MANAGED SERVICE ACCOUNTS

Service accounts are required in each project to run builds, deployments, and other pods. The `managedNames` setting in the `/etc/origin/master/master-config.yml` file on the master controls which service accounts are automatically created in every project:

```

serviceAccountConfig:
  ...
  managedNames: ❶
  - builder ❷

```

- **deployer** ³
- **default** ⁴
- ...

- ¹ List of service accounts to automatically create in every project.
- ² A **builder** service account in each project is required by build pods, and is given the **system:image-builder** role, which allows pushing images to any image stream in the project using the internal container registry.
- ³ A **deployer** service account in each project is required by deployment pods, and is given the **system:deployer** role, which allows viewing and modifying replication controllers and pods in the project.
- ⁴ A **default** service account is used by all other pods unless they specify a different service account.

All service accounts in a project are given the **system:image-puller** role, which allows pulling images from any image stream in the project using the internal container registry.

12.7. INFRASTRUCTURE SERVICE ACCOUNTS

Several infrastructure controllers run using service account credentials. The following service accounts are created in the OpenShift Container Platform infrastructure project (**openshift-infra**) at server start, and given the following roles cluster-wide:

Service Account	Description
replication-controller	Assigned the system:replication-controller role
deployment-controller	Assigned the system:deployment-controller role
build-controller	Assigned the system:build-controller role. Additionally, the build-controller service account is included in the privileged security context constraint in order to create privileged build pods.

To configure the project where those service accounts are created, set the **openshiftInfrastructureNamespace** field in the `/etc/origin/master/master-config.yml` file on the master:

```
policyConfig:
  ...
  openshiftInfrastructureNamespace: openshift-infra
```

12.8. SERVICE ACCOUNTS AND SECRETS

Set the **limitSecretReferences** field in the `/etc/origin/master/master-config.yml` file on the master to **true** to require pod secret references to be whitelisted by their service accounts. Set its value to **false** to allow pods to reference any secret in the project.


```
serviceAccountConfig:
  ...
  limitSecretReferences: false
```

12.9. MANAGING ALLOWED SECRETS

In addition to providing API credentials, a pod's service account determines which secrets the pod is allowed to use.

Pods use secrets in two ways:

- image pull secrets, providing credentials used to pull images for the pod's containers
- mountable secrets, injecting the contents of secrets into containers as files

To allow a secret to be used as an image pull secret by a service account's pods, run:

```
$ oc secrets link --for=pull <serviceaccount-name> <secret-name>
```

To allow a secret to be mounted by a service account's pods, run:

```
$ oc secrets link --for=mount <serviceaccount-name> <secret-name>
```



NOTE

Limiting secrets to only the service accounts that reference them is disabled by default. This means that if `serviceAccountConfig.limitSecretReferences` is set to **false** (the default setting) in the master configuration file, mounting secrets to a service account's pods with the `--for=mount` option is not required. However, using the `--for=pull` option to enable using an image pull secret is required, regardless of the `serviceAccountConfig.limitSecretReferences` value.

This example creates and adds secrets to a service account:

```
$ oc secrets new secret-plans plan1.txt plan2.txt
secret/secret-plans

$ oc secrets new-dockercfg my-pull-secret \
  --docker-username=mastermind \
  --docker-password=12345 \
  --docker-email=mastermind@example.com
secret/my-pull-secret

$ oc secrets link robot secret-plans --for=mount

$ oc secrets link robot my-pull-secret --for=pull

$ oc describe serviceaccount robot
Name:          robot
Labels:        <none>
Image pull secrets: robot-dockercfg-624cx
                  my-pull-secret
```

```
Mountable secrets:  robot-token-uzkbh
                   robot-dockercfg-624cx
                   secret-plans

Tokens:             robot-token-8bhpp
                   robot-token-uzkbh
```

12.10. USING A SERVICE ACCOUNT'S CREDENTIALS INSIDE A CONTAINER

When a pod is created, it specifies a service account (or uses the default service account), and is allowed to use that service account's API credentials and referenced secrets.

A file containing an API token for a pod's service account is automatically mounted at */var/run/secrets/kubernetes.io/serviceaccount/token*.

That token can be used to make API calls as the pod's service account. This example calls the *users/~* API to get information about the user identified by the token:

```
$ TOKEN="$(cat /var/run/secrets/kubernetes.io/serviceaccount/token)"

$ curl --cacert /var/run/secrets/kubernetes.io/serviceaccount/ca.crt \
    "https://openshift.default.svc.cluster.local/oapi/v1/users/~" \
    -H "Authorization: Bearer $TOKEN"

kind: "User"
apiVersion: "v1"
metadata:
  name: "system:serviceaccount:top-secret:robot"
  selflink: "/oapi/v1/users/system:serviceaccount:top-secret:robot"
  creationTimestamp: null
identities: null
groups:
  - "system:serviceaccount"
  - "system:serviceaccount:top-secret"
```

12.11. USING A SERVICE ACCOUNT'S CREDENTIALS EXTERNALLY

The same token can be distributed to external applications that need to authenticate to the API.

Use the following syntax to view a service account's API token:

```
$ oc describe secret <secret-name>
```

For example:

```
$ oc describe secret robot-token-uzkbh -n top-secret
Name:  robot-token-uzkbh
Labels:  <none>
Annotations:  kubernetes.io/service-account.name=robot, kubernetes.io/service-account.uid=49f19e2e-16c6-11e5-afdc-3c970e4b7ffe
```

```
Type: kubernetes.io/service-account-token
```

```
Data
```

```
token: eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9...
```

```
$ oc login --token=eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9...
```

```
Logged into "https://server:8443" as "system:serviceaccount:top-secret:robot" using the token provided.
```

```
You don't have any projects. You can try to create a new project, by running
```

```
$ oc new-project <projectname>
```

```
$ oc whoami
```

```
system:serviceaccount:top-secret:robot
```

CHAPTER 13. MANAGING IMAGES

13.1. OVERVIEW

An [image stream](#) comprises any number of [container images](#) identified by tags. It presents a single virtual view of related images, similar to a Docker image repository.

By watching an image stream, builds and deployments can receive notifications when new images are added or modified and react by performing a build or deployment, respectively.

There are many ways you can interact with images and set up image streams, depending on where the images' registries are located, any authentication requirements around those registries, and how you want your builds and deployments to behave. The following sections cover a range of these topics.

13.2. TAGGING IMAGES

Before working with OpenShift Container Platform image streams and their tags, it helps to first understand image tags in the context of container images generally.


Container images can have names added to them that make it more intuitive to determine what they contain, called a *tag*. Using a tag to specify the version of what is contained in the image is a common use case. If you have an image named `ruby`, you could have a tag named `2.0` for 2.0 version of Ruby, and another named `latest` to indicate literally the latest built image in that repository overall.

When interacting directly with images using the `docker` CLI, the `docker tag` command can add tags, which essentially adds an alias to an image that can consist of several parts. Those parts can include:

```
<registry_server>/<user_name>/<image_name>:<tag>
```

The `<user_name>` part in the above could also refer to a [project](#) or [namespace](#) if the image is being stored in an OpenShift Container Platform environment with an internal registry (the OpenShift Container Registry).

OpenShift Container Platform provides the `oc tag` command, which is similar to the `docker tag` command, but operates on image streams instead of directly on images.



NOTE

See Red Hat Enterprise Linux 7's [Getting Started with Containers](#) documentation for more about tagging images directly using the `docker` CLI.

13.2.1. Adding Tags to Image Streams

Keeping in mind that an image stream in OpenShift Container Platform comprises zero or more container images identified by tags, you can add tags to an image stream using the `oc tag` command:

```
$ oc tag <source> <destination>
```

For example, to configure the `ruby` imagestream's `static-2.0` tag to always refer to the current image for the `ruby` imagestream's `2.0` tag:

```
$ oc tag ruby:2.0 ruby:static-2.0
```

This creates a new imagestream tag named **static-2.0** in the **ruby** imagestream. The new tag directly reference the image id that the **ruby:2.0** imagestream tag pointed to at the time **oc tag** was run, and the image it points to never changes.

There are different types of tags available. The default behavior uses a *permanent* tag, which points to a specific image in time; even when the source changes, the new (destination) tag does not change.

A *tracking* tag means the destination tag's metadata is updated during the import of the source tag. To ensure the destination tag is updated whenever the source tag changes, use the **--alias=true** flag:

```
$ oc tag --alias=true <source> <destination>
```

You can also add the **--scheduled=true** flag to have the destination tag be refreshed (i.e., re-imported) periodically. The period is [configured globally](#) at the system level.

See [Importing Tag and Image Metadata](#) for more details.

If you want to instruct Docker to always fetch the tagged image from the integrated registry, use **--reference-policy=local**. The registry uses the [pull-through feature](#) pull-through feature to serve the image to the client. By default, the image blobs are mirrored locally by the registry. As a result, they can be pulled more quickly the next time they are needed. The flag also allows for pulling from insecure registries without a need to supply **--insecure-registry** to the Docker daemon as long as the image stream has an [insecure annotation](#) or the tag has an [insecure import policy](#).

13.2.2. Recommended Tagging Conventions

Images evolve over time and their tags reflect this. An image tag always points to the latest image built.

If there is too much information embedded in a tag name (for example, **v2.0.1-may-2016**), the tag points to just one revision of an image and is never updated. Using default image pruning options, such an image is never removed. In very large clusters, the schema of creating new tags for every revised image could eventually fill up the etcd datastore with excess tag metadata for images that are long outdated.

Instead, if the tag is named **v2.0**, more image revisions are more likely. This results in longer [tag history](#) and, therefore, the image pruner is more likely to remove old and unused images. Refer to [Pruning Images](#) for more information.

Although tag naming convention is up to you, here are a few examples in the format **<image_name>: <image_tag>**:

Table 13.1. Image Tag Naming Conventions

Description	Example
Revision	myimage:v2.0.1
Architecture	myimage:v2.0-x86_64
Base image	myimage:v1.2-centos7
Latest (potentially unstable)	myimage:latest

Description	Example
Latest stable	<code>myimage:stable</code>

If you require dates in tag names, periodically inspect old and unsupported images and `istags` and remove them. Otherwise, you might experience increasing resource usage caused by old images.

13.2.3. Removing Tags from Image Streams

To remove a tag completely from an image stream run:

```
$ oc delete istag/ruby:latest
```

or:

```
$ oc tag -d ruby:latest
```

13.2.4. Referencing Images in Image Streams

Images can be referenced in image streams using the following reference types:

- An **ImageStreamTag** is used to reference or retrieve an image for a given image stream and tag. It uses the following convention for its name:

```
<image_stream_name>:<tag>
```


- An **ImageStreamImage** is used to reference or retrieve an image for a given image stream and image name. It uses the following convention for its name:

```
<image_stream_name>@<id>
```

The `<id>` is an immutable identifier for a specific image, also called a digest.

- A **DockerImage** is used to reference or retrieve an image for a given external registry. It uses standard Docker *pull specification* for its name, e.g.:

```
openshift/ruby-20-centos7:2.0
```



NOTE

When no tag is specified, it is assumed the **latest** tag is used.

You can also reference a third-party registry:

```
registry.access.redhat.com/rhel7:latest
```

Or an image with a digest:

```
centos/ruby-22-
centos7@sha256:3a335d7d8a452970c5b4054ad7118ff134b3a6b50a2bb6d0c07c7
46e8986b28e
```

When viewing example image stream definitions, such as the [example CentOS image streams](#), you may notice they contain definitions of `ImageStreamTag` and references to `DockerImage`, but nothing related to `ImageStreamImage`.

This is because the `ImageStreamImage` objects are automatically created in OpenShift Container Platform whenever you import or tag an image into the image stream. You should never have to explicitly define an `ImageStreamImage` object in any image stream definition that you use to create image streams.

You can view an image's object definition by retrieving an `ImageStreamImage` definition using the image stream name and ID:

```
$ oc export isimage <image_stream_name>@<id>
```



NOTE

You can find valid `<id>` values for a given image stream by running:

```
$ oc describe is <image_stream_name>
```

For example, from the `ruby` image stream asking for the `ImageStreamImage` with the name and ID of `ruby@3a335d7`:

Example 13.1. Definition of an Image Object Retrieved via `ImageStreamImage`

```
$ oc export isimage ruby@3a335d7

apiVersion: v1
image:
  dockerImageLayers:
    - name:
sha256:a3ed95caeb02ffe68cdd9fd84406680ae93d633cb16422d00e8a7c22955b46d4
  size: 0
    - name:
sha256:ee1dd2cb6df21971f4af6de0f1d7782b81fb63156801cfde2bb47b4247c23c29
  size: 196634330
    - name:
sha256:a3ed95caeb02ffe68cdd9fd84406680ae93d633cb16422d00e8a7c22955b46d4
  size: 0
    - name:
sha256:a3ed95caeb02ffe68cdd9fd84406680ae93d633cb16422d00e8a7c22955b46d4
  size: 0
    - name:
sha256:ca062656bff07f18bff46be00f40cfbb069687ec124ac0aa038fd676cfaea092
  size: 177723024
    - name:
sha256:63d529c59c92843c395befd065de516ee9ed4995549f8218eac6ff088bfa6b6e
  size: 55679776
  dockerImageMetadata:
```

```

Architecture: amd64
Author: SoftwareCollections.org <sclorg@redhat.com>
Config:
  Cmd:
    - /bin/sh
    - -c
    - $STI_SCRIPTS_PATH/usage
  Entrypoint:
    - container-entrypoint
  Env:
    - 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: {}
  Image:
d9c3abc5456a9461954ff0de8ae25e0e016aad35700594714d42b687564b1f51
  Labels:
    build-date: 2015-12-23
    io.k8s.description: Platform for building and running Ruby 2.2
applications
    io.k8s.display-name: Ruby 2.2
    io.openshift.builder-base-version: 8d95148
    io.openshift.builder-version:
8847438ba06307f86ac877465eadc835201241df
    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
ContainerConfig: {}
Created: 2016-01-26T21:07:27Z
DockerVersion: 1.8.2-e17
Id: 57b08d979c86f4500dc8cad639c9518744c8dd39447c055a3517dc9c18d6fccd
Parent:
d9c3abc5456a9461954ff0de8ae25e0e016aad35700594714d42b687564b1f51
  Size: 430037130
  apiVersion: "1.0"
  kind: DockerImage
  dockerImageMetadataVersion: "1.0"
  dockerImageReference: centos/ruby-22-
centos7@sha256:3a335d7d8a452970c5b4054ad7118ff134b3a6b50a2bb6d0c07c746e8
986b28e
  metadata:
    creationTimestamp: 2016-01-29T13:17:45Z
    name:
sha256:3a335d7d8a452970c5b4054ad7118ff134b3a6b50a2bb6d0c07c746e8986b28e

```



```

    resourceVersion: "352"
    uid: af2e7a0c-c68a-11e5-8a99-525400f25e34
  kind: ImageStreamImage
  metadata:
    creationTimestamp: null
    name: ruby@3a335d7
    namespace: openshift
    selflink: /oapi/v1/namespaces/openshift/imagestreamimages/ruby@3a335d7

```

13.3. IMAGE PULL POLICY

Each container in a pod has a container image. Once you have created an image and pushed it to a registry, you can then refer to it in the pod.

When OpenShift Container Platform creates containers, it uses the container's **imagePullPolicy** to determine if the image should be pulled prior to starting the container. There are three possible values for **imagePullPolicy**:

- **Always** - always pull the image.
- **IfNotPresent** - only pull the image if it does not already exist on the node.
- **Never** - never pull the image.

If a container's **imagePullPolicy** parameter is not specified, OpenShift Container Platform sets it based on the image's tag:

1. If the tag is **latest**, OpenShift Container Platform defaults **imagePullPolicy** to **Always**.
2. Otherwise, OpenShift Container Platform defaults **imagePullPolicy** to **IfNotPresent**.

13.4. ACCESSING THE INTERNAL REGISTRY

You can access OpenShift Container Platform's internal registry directly to push or pull images. For example, this could be helpful if you wanted to [create an image stream by manually pushing an image](#), or just to **docker pull** an image directly.

The internal registry authenticates using the same **tokens** as the OpenShift Container Platform API. To perform a **docker login** against the internal registry, you can choose any user name and email, but the password must be a valid OpenShift Container Platform token.

To log into the internal registry:

1. Log in to OpenShift Container Platform:

```
$ oc login
```

2. Get your access token:

```
$ oc whoami -t
```

3. Log in to the internal registry using the token. You must have **docker** installed on your system:

```
■
```

```
$ docker login -u <user_name> -e <email_address> \
  -p <token_value> <registry_server>:<port>
```

**NOTE**

Contact your cluster administrator if you do not know the registry IP or host name and port to use.

In order to pull an image, the authenticated user must have **get** rights on the requested **imagestreams/layers**. In order to push an image, the authenticated user must have **update** rights on the requested **imagestreams/layers**.

By default, all service accounts in a project have rights to pull any image in the same project, and the **builder** service account has rights to push any image in the same project.

13.5. USING IMAGE PULL SECRETS

[Docker registries](#) can be secured to prevent unauthorized parties from accessing certain images. If you are [using OpenShift Container Platform's internal registry](#) and are pulling from image streams located in the same project, then your pod's service account should already have the correct permissions and no additional action should be required.

However, for other scenarios, such as referencing images across OpenShift Container Platform projects or from secured registries, then additional configuration steps are required. The following sections detail these scenarios and their required steps.

13.5.1. Allowing Pods to Reference Images Across Projects

When using the internal registry, to allow pods in **project-a** to reference images in **project-b**, a service account in **project-a** must be bound to the **system:image-puller** role in **project-b**:

```
$ oc policy add-role-to-user \
  system:image-puller system:serviceaccount:project-a:default \
  --namespace=project-b
```

After adding that role, the pods in **project-a** that reference the default service account is able to pull images from **project-b**.

To allow access for any service account in **project-a**, use the group:

```
$ oc policy add-role-to-group \
  system:image-puller system:serviceaccounts:project-a \
  --namespace=project-b
```

13.5.2. Allowing Pods to Reference Images from Other Secured Registries

The **.dockercfg** file (or **\$HOME/.docker/config.json** for newer Docker clients) is a Docker credentials file that stores your information if you have previously logged into a secured or insecure registry.

To pull a secured container image that is not from OpenShift Container Platform's internal registry, you must create a *pull secret* from your Docker credentials and add it to your service account.

If you already have a **.dockercfg** file for the secured registry, you can create a secret from that file by running:

```
$ oc secrets new <pull_secret_name> .dockercfg=<path/to/.dockercfg>
```

Or if you have a **\$HOME/.docker/config.json** file:

```
$ oc secrets new <pull_secret_name> .dockerconfigjson=
<path/to/.docker/config.json>
```

If you do not already have a Docker credentials file for the secured registry, you can create a secret by running:

```
$ oc secrets new-dockercfg <pull_secret_name> \
  --docker-server=<registry_server> --docker-username=<user_name> \
  --docker-password=<password> --docker-email=<email>
```

To use a secret for pulling images for pods, you must add the secret to your service account. The name of the service account in this example should match the name of the service account the pod uses; **default** is the default service account:

```
$ oc secrets link default <pull_secret_name> --for=pull
```

To use a secret for pushing and pulling build images, the secret must be mountable inside of a pod. You can do this by running:

```
$ oc secrets link builder <pull_secret_name>
```

13.6. IMPORTING TAG AND IMAGE METADATA

An image stream can be configured to import tag and image metadata from an image repository in an external Docker image registry. You can do this using a few different methods.

- You can manually import tag and image information with the **oc import-image** command using the **--from** option:

```
$ oc import-image <image_stream_name>[:<tag>] --from=
<docker_image_repo> --confirm
```

For example:

```
$ oc import-image my-ruby --from=docker.io/openshift/ruby-20-centos7
--confirm
The import completed successfully.

Name:    my-ruby
Created:  Less than a second ago
Labels:  <none>
Annotations:  openshift.io/image.dockerRepositoryCheck=2016-05-
06T20:59:30Z
Docker Pull Spec: 172.30.94.234:5000/demo-project/my-ruby
```

Tag Spec	Created	PullSpec	Image
latest	docker.io/openshift/ruby-20-centos7	Less than a second ago	docker.io/openshift/ruby-20-centos7@sha256:772c5bf9b2d1e8... <same>

You can also add the `--all` flag to import all tags for the image instead of just `latest`.

- Like most objects in OpenShift Container Platform, you can also write and save a JSON or YAML definition to a file then create the object using the CLI. Set the `spec.dockerImageRepository` field to the Docker pull spec for the image:

```
apiVersion: "v1"
kind: "ImageStream"
metadata:
  name: "my-ruby"
spec:
  dockerImageRepository: "docker.io/openshift/ruby-20-centos7"
```

Then create the object:

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

When you create an image stream that references an image in an external Docker registry, OpenShift Container Platform communicates with the external registry within a short amount of time to get up to date information about the image.

After the tag and image metadata is synchronized, the image stream object would look similar to the following:

```
apiVersion: v1
kind: ImageStream
metadata:
  name: my-ruby
  namespace: demo-project
  selflink: /oapi/v1/namespaces/demo-project/imagestreams/my-ruby
  uid: 5b9bd745-13d2-11e6-9a86-0ada84b8265d
  resourceVersion: '4699413'
  generation: 2
  creationTimestamp: '2016-05-06T21:34:48Z'
  annotations:
    openshift.io/image.dockerRepositoryCheck: '2016-05-06T21:34:48Z'
spec:
  dockerImageRepository: docker.io/openshift/ruby-20-centos7
  tags:
  -
    name: latest
    annotations: null
    from:
      kind: DockerImage
      name: 'docker.io/openshift/ruby-20-centos7:latest'
      generation: 2
      importPolicy: { }
status:
  dockerImageRepository: '172.30.94.234:5000/demo-project/my-ruby'
  tags:
  -
```

```

    tag: latest
    items:
      -
        created: '2016-05-06T21:34:48Z'
        dockerImageReference: 'docker.io/openshift/ruby-20-
centos7@sha256:772c5bf9b2d1e8e80742ed75aab05820419dc4532fa6d7ad8a1efddda54
93dc3'
        image:
          'sha256:772c5bf9b2d1e8e80742ed75aab05820419dc4532fa6d7ad8a1efddda5493dc3'
          generation: 2

```

You can set a tag to query external registries at a scheduled interval to synchronize tag and image metadata by setting the `--scheduled=true` flag with the `oc tag` command as mentioned in [Adding Tags to Image Streams](#).

Alternatively, you can set `importPolicy.scheduled` to `true` in the tag's definition:

```

apiVersion: v1
kind: ImageStream
metadata:
  name: ruby
spec:
  tags:
    - from:
        kind: DockerImage
        name: openshift/ruby-20-centos7
      name: latest
      importPolicy:
        scheduled: true

```

13.6.1. Importing Images from Insecure Registries

An image stream can be configured to import tag and image metadata from insecure image registries, such as those signed with a self-signed certificate or using plain HTTP instead of HTTPS.

To configure this, add the `openshift.io/image.insecureRepository` annotation and set it to `true`. This setting bypasses certificate validation when connecting to the registry:

```

kind: ImageStream
apiVersion: v1
metadata:
  name: ruby
  annotations:
    openshift.io/image.insecureRepository: "true" ❶
spec:
  dockerImageRepository: my.repo.com:5000/myimage


```

❶ Set the `openshift.io/image.insecureRepository` annotation to `true`



IMPORTANT

This option instructs integrated registry to fall back to an insecure transport for any external image tagged in the image stream when serving it, which is dangerous. If possible, avoid this risk by [marking just an istag as insecure](#).



IMPORTANT

The above definition only affects importing tag and image metadata. For this image to be used in the cluster (e.g., to be able to do a `docker pull`), one of the following must be true:

1. Each node has Docker configured with the `--insecure-registry` flag matching the registry part of the `dockerImageRepository`. See [Host Preparation](#) for more information.
2. Each `istag` specification must have `referencePolicy.type` set to `Local`. See [Reference Policy](#) for more information.

13.6.1.1. ImageStream Tag Policies

13.6.1.1.1. Insecure Tag Import Policy

The above annotation applies to all images and tags of a particular `ImageStream`. For a finer-grained control, policies may be set on [istags](#). Set `importPolicy.insecure` in the tag's definition to `true` to allow a fall-back to insecure transport just for images under this tag.



NOTE

The fall-back to insecure transport for an image under particular `istag` is enabled either when the image stream is annotated as insecure or the `istag` has insecure import policy. The `importPolicy.insecure` set to false can not override the image stream annotation.`

13.6.1.1.2. Reference Policy

The Reference Policy allows you to specify where the image consumers pulls from. It is only applicable to remote images (those imported from external registries). There are two options to choose from, `Local` and `Source`.

The `Source` policy instructs clients to pull directly from the source registry of the image. The integrated registry is not involved unless the image is managed by the cluster. (It is not an external image.) This is the default policy.

The `Local` policy instructs clients to always pull from the integrated registry. This is useful if you want to pull from external insecure registries without modifying Docker daemon settings.

The [pull-through feature](#) of the registry serves the remote image to the client. Additionally, all the blobs are mirrored for faster access later.

You can set the policy in a specification of image stream tag as `referencePolicy.type`.

Example 13.2. Example of Insecure Tag with a Local Reference Policy

```

kind: ImageStream
apiVersion: v1
metadata:
  name: ruby
  tags:
  - from:
      kind: DockerImage
      name: my.repo.com:5000/myimage
      name: mytag
      importPolicy:
        insecure: true 1
      referencePolicy:
        type: Local 2

```

- 1 Set tag `mytag` to use an insecure connection to that registry.
- 2 Set tag `mytag` to use integrated registry for pulling external images. If the reference policy type is set to `Source`, clients fetch the image directly from `my.repo.com:5000/myimage`.

13.6.2. Importing Images from Private Registries

An image stream can be configured to import tag and image metadata from private image registries, requiring authentication.

To configure this, you need to create a [secret](#), which is used to store your credentials. See [Allowing Pods to Reference Images from Other Secured Registries](#) for instructions on creating a secret using `oc create secret` command.

After the secret is configured, proceed with creating the new image stream or using the `oc import-image` command. During the import process, OpenShift Container Platform picks up the secrets and provide them to the remote party.



NOTE

When importing from an insecure registry, the registry URL defined in the secret must include the `:80` port suffix or the secret is not used when attempting to import from the registry.

13.6.3. Adding Trusted Certificates for External Registries

If the registry you are importing from is using a certificate that is not signed by a standard certificate authority, you need to explicitly configure the system to trust the registry's certificate or signing authority. This can be done by adding the CA certificate or registry certificate to the host system running the registry import controller (typically the master node).

The certificate or CA certificate must be added to `/etc/pki/tls/certs` or `/etc/pki/ca-trust`, respectively, on the host system. The `update-ca-trust` command also needs to be run on Red Hat distributions followed by a restart of the master service to pick up the certificate changes.

13.6.4. Importing Images Across Projects

An image stream can be configured to import tag and image metadata from the internal registry, but from a different project. The recommended method for this is to use the `oc tag` command as shown in [Adding Tags to Image Streams](#):

```
$ oc tag <source_project>/<image_stream>:<tag> <new_image_stream>:
<new_tag>
```

Another method is to import the image from the other project manually using the pull spec:



WARNING

The following method is strongly discouraged and should be used only if the former using `oc tag` is insufficient.

1. First, add the necessary [policy](#) to access the other project:

```
$ oc policy add-role-to-group \
  system:image-puller \
  system:serviceaccounts:<destination_project> \
  -n <source_project>
```

This allows `<destination_project>` to pull images from `<source_project>`.

2. With the policy in place, you can import the image manually:

```
$ oc import-image <new_image_stream> --confirm \
  --from=<docker_registry>/<source_project>/<image_stream>
```

13.6.5. Creating an Image Stream by Manually Pushing an Image

An image stream can also be automatically created by manually pushing an image to the internal registry. This is only possible when using an OpenShift Container Platform internal registry.

Before performing this procedure, the following must be satisfied:

- The destination project you push to must already exist.
- The user must be authorized to `{get, update} "imagestream/layers"` in that project. In addition, since the image stream does not already exist, the user must be authorized to `{create} "imagestream"` in that project. If you are a project administrator, then you would have these permissions.



NOTE

The `system:image-pusher` role does not grant permission to create new image streams, only to push images to existing image streams, so it cannot be used to push images to image streams that do not yet exist unless additional permissions are also granted to the user.

To create an image stream by manually pushing an image:

1. First, [log in to the internal registry](#).
2. Then, tag your image using the appropriate internal registry location. For example, if you had already pulled the `docker.io/centos:centos7` image locally:

```
$ docker tag docker.io/centos:centos7 172.30.48.125:5000/test/my-image
```

3. Finally, push the image to your internal registry. For example:

```
$ docker push 172.30.48.125:5000/test/my-image
The push refers to a repository [172.30.48.125:5000/test/my-image]
(len: 1)
c8a648134623: Pushed
2bf4902415e3: Pushed
latest: digest:
sha256:be8bc4068b2f60cf274fc216e4caba6aa845fff5fa29139e6e7497bb57e48d67 size: 6273
```

4. Verify that the image stream was created:

```
$ oc get is
NAME          DOCKER REPO          TAGS          UPDATED
my-image      172.30.48.125:5000/test/my-image  latest        3 seconds ago
```

13.7. WRITING IMAGE STREAM DEFINITIONS

You can define image streams by writing the image stream definition for the entire image stream. This allows you to distribute the definition to different clusters without running `oc` commands.

An image stream definition specifies information about the image stream and the specific tags to be imported.

Definition of an Image Stream Object

```
apiVersion: v1
kind: ImageStream
metadata:
  name: ruby
  annotations:
    openshift.io/display-name: Ruby 1
spec:
  tags:
    - name: '2.0' 2
      annotations:
        openshift.io/display-name: Ruby 2.0 3
      description: >- 4
        Build and run Ruby 2.0 applications on CentOS 7. For more
        information
        about using this builder image, including OpenShift
        considerations,
```

```

    see
    https://github.com/sclorg/s2i-ruby-
container/tree/master/2.0/README.md.
    iconClass: icon-ruby 5
    sampleRepo: 'https://github.com/openshift/ruby-ex.git' 6
    tags: 'builder,ruby' 7
    version: '2.0' 8
  from:
    kind: DockerImage 9
    name: 'docker.io/openshift/ruby-20-centos7:latest' 10

```

- 1 A brief, user-friendly name for the whole image stream.
- 2 The tag is referred to as the version. Tags appear in a drop-down menu.
- 3 A user-friendly name for this tag within the image stream. This should be brief and include version information when appropriate.
- 4 A description of the tag, which includes enough detail for users to understand what the image is providing. It can include links to additional instructions. Limit the description to a few sentences.
- 5 The icon to show for this tag. Pick from our existing [logo icons](#) when possible. Icons from [FontAwesome](#) and [Patternfly](#) can also be used. Alternatively, provide icons through [CSS customizations](#) that can be added to an OpenShift Container Platform cluster that uses your image stream. You must specify an icon class that exists, or it prevents falling back to the generic icon.
- 6 A URL to a source repository that works with this builder image tag and results in a sample running application.
- 7 Categories that the image stream tag is associated with. The builder tag is required for it to show up in the catalog. Add tags that associates it with one of the provided catalog categories. Refer to the `id` and `categoryAliases` in `CATALOG_CATEGORIES` in the console's [constants file](#). The categories can also be [customized](#) for the whole cluster.
- 8 Languages this image supports. This value is used during `oc new-app` invocations to try to match potential builder images to the provided source repository.

Version information for this tag.

- 9 The type of object this image stream tag is referencing. Valid values are: `DockerImage`, `ImageStreamTag`, and `ImageStreamImage`.
- 10 The object this image stream tag imports.

CHAPTER 14. IMAGE SIGNATURES


14.1. OVERVIEW

Container image signing on Red Hat Enterprise Linux (RHEL) systems provides a means of:

- Validating where a container image came from,
- Checking that the image has not been tampered with, and
- Setting policies to determine which validated images can be pulled to a host.

For a more complete understanding of the architecture of container image signing on RHEL systems, see the [Container Image Signing Integration Guide](#).

The OpenShift Container Registry allows the ability to store signatures via REST API. The `oc` CLI can be used to verify image signatures, with their validated displayed in the web console or CLI.



NOTE

Initial support for storing image signatures was added in OpenShift Container Platform 3.3.

14.2. SIGNING IMAGES USING ATOMIC CLI

OpenShift Container Platform does not automate image signing. Signing requires a developer's private GPG key, typically stored securely on a workstation. This document describes that workflow.

The `atomic` command line interface (CLI), version 1.12.5 or greater, provides commands for signing container images, which can be pushed to an OpenShift Container Registry. The `atomic` CLI is available on Red Hat-based distributions: RHEL, Centos, and Fedora.

You can sign an image at push time:

```
$ atomic push [--sign-by ...] <image>
```

Or you can sign an image that has already been pushed to the registry:

```
$ atomic sign [--sign-by ...] <image>
```

For full details on how to set up and perform image signing using the `atomic` CLI, see the [RHEL Atomic Host Managing Containers: Signing Container Images](#) documentation.

A specific example workflow of working with the `atomic` CLI and an OpenShift Container Registry is documented in the [Container Image Signing Integration Guide](#).

14.3. ACCESSING IMAGE SIGNATURES USING REGISTRY API

The OpenShift Container Registry provides an `extensions` endpoint that allows you to write and read image signatures. The image signatures are stored in the OpenShift Container Platform key-value store via the Docker Registry API.

**NOTE**

This endpoint is experimental and not supported by the upstream Docker Registry project. See the [upstream API documentation](#) for general information about the Docker Registry API.

14.3.1. Writing Image Signatures

In order to add a new signature to the image, you can use the HTTP **PUT** method to send a JSON payload to the **extensions** endpoint:

```
PUT /extensions/v2/<namespace>/<name>/signatures/<digest>
```

```
$ curl -X PUT --data @signature.json http://<user>:
<token>@<registry_endpoint>:5000/extensions/v2/<namespace>/<name>/signatur
es/sha256:<digest>
```

The JSON payload with the signature content should have the following structure:

```
{
  "version": 2,
  "type": "atomic",
  "name":
    "sha256:4028782c08eae4a8c9a28bf661c0a8d1c2fc8e19dbaae2b018b21011197e1484@c
    ddeb7006d914716e2728000746a0b23",
  "content": "<cryptographic_signature>"
}
```

The **name** field contains the name of the image signature, which must be unique and in the format **<digest>@<name>**. The **<digest>** represents an image name and the **<name>** is the name of the signature. The signature name must be 32 characters long. The **<cryptographic_signature>** must follow the specification documented in the [containers/image](#) library.

**NOTE**

In order to attach the signature to the image, the user must have the **system:image-signer** cluster role. Cluster administrators can add this using:

```
$ oc adm policy add-cluster-role-to-user system:image-signer
<user_name>
```

14.3.2. Reading Image Signatures

Assuming a signed image has already been pushed into the OpenShift Container Registry, you can read the signatures using the following command:

```
GET /extensions/v2/<namespace>/<name>/signatures/<digest>
```

```
$ curl http://<user>:
<token>@<registry_endpoint>:5000/extensions/v2/<namespace>/<name>/signatur
es/sha256:<digest>
```

The `<namespace>` represents the OpenShift Container Platform project name or registry repository name and the `<name>` refers to the name of the image repository. The `digest` represents the SHA-256 checksum of the image.

If the given image contains the signature data, the output of the command above should produce following JSON response:

```
{
  "signatures": [
    {
      "version": 2,
      "type": "atomic",
      "name":
"sha256:4028782c08eae4a8c9a28bf661c0a8d1c2fc8e19dbaae2b018b21011197e1484@c
ddeb7006d914716e2728000746a0b23",
      "content": "<cryptographic_signature>"
    }
  ]
}
```

The `name` field contains the name of the image signature, which must be unique and in the format `<digest>@<name>`. The `<digest>` represents an image name and the `<name>` is the name of the signature. The signature name must be 32 characters long. The `<cryptographic_signature>` must follow the specification documented in the [containers/image](#) library.

14.4. VERIFYING IMAGE SIGNATURES USING OPENSIFT CLI

You can verify the signatures of an image imported to an OpenShift Container Registry using the `oc adm verify-image-signature` command. This command verifies if the image identity contained in the image signature can be trusted by using the public GPG key to verify the signature itself then match the provided expected identity with the identity (the pull spec) of the given image.

By default, this command uses the public GPG keyring located in `$GNUPGHOME/pubring.gpg`, typically in path `~/.gnupg`. By default, this command does not save the result of the verification back to the image object. To do so, you must specify the `--save` flag.



NOTE

To modify the image signature verification status, your account must have to have permissions to edit image objects, for example using the `image-auditor` role.

Using the `--save` flag on already verified image together with invalid GPG key or invalid expected identity causes the saved verification status to be removed, and the image will become unverified.

To verify an image signature:

```
$ oc adm verify-image-signature <image> --expected-identity=<pull_spec> [-
-save] [options]
```

CHAPTER 15. QUOTAS AND LIMIT RANGES

15.1. OVERVIEW


Using [quotas](#) and [limit ranges](#), cluster administrators can set constraints to limit the number of objects or amount of compute resources that are used in your project. This helps cluster administrators better manage and allocate resources across all projects, and ensure that no projects are using more than is appropriate for the cluster size.

As a developer, you can also set [requests and limits on compute resources](#) at the pod and container level.

The following sections help you understand how to check on your quota and limit range settings, what sorts of things they can constrain, and how you can request or limit compute resources in your own pods and containers.

15.2. QUOTAS

A resource quota, defined by a `ResourceQuota` object, provides constraints that limit aggregate resource consumption per project. It can limit the quantity of objects that can be created in a project by type, as well as the total amount of compute resources and storage that may be consumed by resources in that project.



NOTE

Quotas are set by cluster administrators and are scoped to a given project.

15.2.1. Viewing Quotas

You can view usage statistics related to any hard limits defined in a project's quota by navigating in the web console to the project's **Quota** page.

You can also use the CLI to view quota details:

1. First, get the list of quotas defined in the project. For example, for a project called **demoproject**:

```
$ oc get quota -n demoproject
NAME                AGE
besteffort          11m
compute-resources   2m
core-object-counts  29m
```

2. Then, describe the quota you are interested in, for example the **core-object-counts** quota:

```
$ oc describe quota core-object-counts -n demoproject
Name:      core-object-counts
Namespace: demoproject
Resource   Used Hard
-----
configmaps 3 10
persistentvolumeclaims 0 4
replicationcontrollers 3 20
secrets    9 10
```

```
services 2 10
```

Full quota definitions can be viewed by running `oc export` on the object. The following show some sample quota definitions:

object-counts.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: core-object-counts
spec:
  hard:
    configmaps: "10" ❶
    persistentvolumeclaims: "4" ❷
    replicationcontrollers: "20" ❸
    secrets: "10" ❹
    services: "10" ❺
```

- ❶ The total number of **ConfigMap** objects that can exist in the project.
- ❷ The total number of persistent volume claims (PVCs) that can exist in the project.
- ❸ The total number of replication controllers that can exist in the project.
- ❹ The total number of secrets that can exist in the project.
- ❺ The total number of services that can exist in the project.

openshift-object-counts.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: openshift-object-counts
spec:
  hard:
    openshift.io/imagestreams: "10" ❶
```

- ❶ The total number of image streams that can exist in the project.

compute-resources.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources
spec:
  hard:
    pods: "4" ❶
    requests.cpu: "1" ❷
```

```

requests.memory: 1Gi 3
limits.cpu: "2" 4
limits.memory: 2Gi 5

```

- 1 The total number of pods in a non-terminal state that can exist in the project.
- 2 Across all pods in a non-terminal state, the sum of CPU requests cannot exceed 1 core.
- 3 Across all pods in a non-terminal state, the sum of memory requests cannot exceed 1Gi.
- 4 Across all pods in a non-terminal state, the sum of CPU limits cannot exceed 2 cores.
- 5 Across all pods in a non-terminal state, the sum of memory limits cannot exceed 2Gi.

besteffort.yaml

```

apiVersion: v1
kind: ResourceQuota
metadata:
  name: besteffort
spec:
  hard:
    pods: "1" 1
  scopes:
    - BestEffort 2

```

- 1 The total number of pods in a non-terminal state with **BestEffort** quality of service that can exist in the project.
- 2 Restricts the quota to only matching pods that have **BestEffort** quality of service for either memory or CPU.

compute-resources-long-running.yaml

```

apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources-long-running
spec:
  hard:
    pods: "4" 1
    limits.cpu: "4" 2
    limits.memory: "2Gi" 3
  scopes:
    - NotTerminating 4

```

- 1 The total number of pods in a non-terminal state.
- 2 Across all pods in a non-terminal state, the sum of CPU limits cannot exceed this value.
- 3 Across all pods in a non-terminal state, the sum of memory limits cannot exceed this value.

- 4 Restricts the quota to only matching pods where `spec.activeDeadlineSeconds` is set to `nil`. Build pods will fall under `NotTerminating` unless the `RestartNever` policy is applied.

compute-resources-time-bound.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: compute-resources-time-bound
spec:
  hard:
    pods: "2" 1
    limits.cpu: "1" 2
    limits.memory: "1Gi" 3
  scopes:
    - Terminating 4
```

- 1 The total number of pods in a non-terminal state.
- 2 Across all pods in a non-terminal state, the sum of CPU limits cannot exceed this value.
- 3 Across all pods in a non-terminal state, the sum of memory limits cannot exceed this value.
- 4 Restricts the quota to only matching pods where `spec.activeDeadlineSeconds >=0`. For example, this quota would charge for build or deployer pods, but not long running pods like a web server or database.

storage-consumption.yaml

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: storage-consumption
spec:
  hard:
    persistentvolumeclaims: "10" 1
    requests.storage: "50Gi" 2
```

- 1 The total number of persistent volume claims in a project
- 2 Across all persistent volume claims in a project, the sum of storage requested cannot exceed this value.

15.2.2. Resources Managed by Quota

The following describes the set of compute resources and object types that may be managed by a quota.



NOTE

A pod is in a terminal state if `status.phase` in (`Failed`, `Succeeded`) is true.

Table 15.1. Compute Resources Managed by Quota

Resource Name	Description
cpu	The sum of CPU requests across all pods in a non-terminal state cannot exceed this value. cpu and requests.cpu are the same value and can be used interchangeably.
memory	The sum of memory requests across all pods in a non-terminal state cannot exceed this value. memory and requests.memory are the same value and can be used interchangeably.
requests.cpu	The sum of CPU requests across all pods in a non-terminal state cannot exceed this value. cpu and requests.cpu are the same value and can be used interchangeably.
requests.memory	The sum of memory requests across all pods in a non-terminal state cannot exceed this value. memory and requests.memory are the same value and can be used interchangeably.
requests.storage	The sum of storage requests across all persistent volume claims cannot exceed this value. storage and requests.storage are the same value and can be used interchangeably.
limits.cpu	The sum of CPU limits across all pods in a non-terminal state cannot exceed this value.
limits.memory	The sum of memory limits across all pods in a non-terminal state cannot exceed this value.

Table 15.2. Storage Resources Managed by Quota

Resource Name	Description
requests.storage	The sum of storage requests across all persistent volume claims in any state cannot exceed this value.
persistentvolumeclaims	The total number of persistent volume claims that can exist in the project.
<storage-class-name>.storageclass.storage.k8s.io/requests.storage	The sum of storage requests across all persistent volume claims in any state that have a matching storage class, cannot exceed this value.
<storage-class-name>.storageclass.storage.k8s.io/persistentvolumeclaims	The total number of persistent volume claims with a matching storage class that can exist in the project.

Table 15.3. Object Counts Managed by Quota

Resource Name	Description
pods	The total number of pods in a non-terminal state that can exist in the project.
replicationcontrollers	The total number of replication controllers that can exist in the project.
resourcequotas	The total number of resource quotas that can exist in the project.
services	The total number of services that can exist in the project.
secrets	The total number of secrets that can exist in the project.
configmaps	The total number of ConfigMap objects that can exist in the project.
persistentvolumeclaims	The total number of persistent volume claims that can exist in the project.
openshift.io/imagestreams	The total number of image streams that can exist in the project.

15.2.3. Quota Scopes

Each quota can have an associated set of *scopes*. A quota will only measure usage for a resource if it matches the intersection of enumerated scopes.

Adding a scope to a quota restricts the set of resources to which that quota can apply. Specifying a resource outside of the allowed set results in a validation error.

Scope	Description
Terminating	Match pods where <code>spec.activeDeadlineSeconds >= 0</code> .
NotTerminating	Match pods where <code>spec.activeDeadlineSeconds</code> is <code>nil</code> .
BestEffort	Match pods that have best effort quality of service for either cpu or memory . See the Quality of Service Classes for more on committing compute resources.
NotBestEffort	Match pods that do not have best effort quality of service for cpu and memory .

A **BestEffort** scope restricts a quota to limiting the following resources:

- **pods**

A **Terminating**, **NotTerminating**, and **NotBestEffort** scope restricts a quota to tracking the following resources:

- `pods`
- `memory`
- `requests.memory`
- `limits.memory`
- `cpu`
- `requests.cpu`
- `limits.cpu`

15.2.4. Quota Enforcement

After a resource quota for a project is first created, the project restricts the ability to create any new resources that may violate a quota constraint until it has calculated updated usage statistics.

After a quota is created and usage statistics are updated, the project accepts the creation of new content. When you create or modify resources, your quota usage is incremented immediately upon the request to create or modify the resource.

When you delete a resource, your quota use is decremented during the next full recalculation of quota statistics for the project. If project modifications exceed a quota usage limit, the server denies the action. An appropriate error message is returned explaining the quota constraint violated, and what your currently observed usage stats are in the system.

15.2.5. Requests vs Limits

When allocating [compute resources](#), each container may specify a request and a limit value each for CPU and memory. Quotas can restrict any of these values.

If the quota has a value specified for `requests.cpu` or `requests.memory`, then it requires that every incoming container make an explicit request for those resources. If the quota has a value specified for `limits.cpu` or `limits.memory`, then it requires that every incoming container specify an explicit limit for those resources.

See [Compute Resources](#) for more on setting requests and limits in pods and containers.

15.3. LIMIT RANGES

A limit range, defined by a `LimitRange` object, enumerates [compute resource constraints](#) in a [project](#) at the pod, container, image, image stream, and persistent volume claim level, and specifies the amount of resources that a pod, container, image, image stream, or persistent volume claim can consume.

All resource create and modification requests are evaluated against each `LimitRange` object in the project. If the resource violates any of the enumerated constraints, then the resource is rejected. If the resource does not set an explicit value, and if the constraint supports a default value, then the default value is applied to the resource.



NOTE

Limit ranges are set by cluster administrators and are scoped to a given project.

15.3.1. Viewing Limit Ranges

You can view any limit ranges defined in a project by navigating in the web console to the project's **Quota** page.

You can also use the CLI to view limit range details:

1. First, get the list of limit ranges defined in the project. For example, for a project called **demoproject**:

```
$ oc get limits -n demoproject
NAME                AGE
resource-limits     6d
```

2. Then, describe the limit range you are interested in, for example the **resource-limits** limit range:

\$ oc describe limits resource-limits -n demoproject							
Name:				resource-limits			
Namespace:				demoproject			
Type				Resource	Min		
Max	Default	Request	Default	Limit	Max	Limit/Request	Ratio
----				-----			---
Pod				cpu			200m 2
-		-		-			
Pod				memory			6Mi
1Gi	-		-		-		
Container				cpu			100m 2
200m		300m		10			
Container				memory			4Mi
1Gi	100Mi		200Mi		-		
openshift.io/Image				storage			-
1Gi	-		-		-		
openshift.io/ImageStream				openshift.io/image			- 12
-		-		-			
openshift.io/ImageStream				openshift.io/image-tags			- 10
-		-		-			

Full limit range definitions can be viewed by running `oc export` on the object. The following shows an example limit range definition:

Example 15.1. Limit Range Object Definition

```
apiVersion: "v1"
kind: "LimitRange"
metadata:
  name: "core-resource-limits" ❶
spec:
  limits:
    - type: "Pod"
      max:
        cpu: "2" ❷
        memory: "1Gi" ❸
      min:
```

```

    cpu: "200m" 4
    memory: "6Mi" 5
- type: "Container"
  max:
    cpu: "2" 6
    memory: "1Gi" 7
  min:
    cpu: "100m" 8
    memory: "4Mi" 9
  default:
    cpu: "300m" 10
    memory: "200Mi" 11
  defaultRequest:
    cpu: "200m" 12
    memory: "100Mi" 13
  maxLimitRequestRatio:
    cpu: "10" 14

```

- 1 The name of the limit range object.
- 2 The maximum amount of CPU that a pod can request on a node across all containers.
- 3 The maximum amount of memory that a pod can request on a node across all containers.
- 4 The minimum amount of CPU that a pod can request on a node across all containers.
- 5 The minimum amount of memory that a pod can request on a node across all containers.
- 6 The maximum amount of CPU that a single container in a pod can request.
- 7 The maximum amount of memory that a single container in a pod can request.
- 8 The minimum amount of CPU that a single container in a pod can request.
- 9 The minimum amount of memory that a single container in a pod can request.
- 10 The default amount of CPU that a container will be limited to use if not specified.
- 11 The default amount of memory that a container will be limited to use if not specified.
- 12 The default amount of CPU that a container will request to use if not specified.
- 13 The default amount of memory that a container will request to use if not specified.
- 14 The maximum amount of CPU burst that a container can make as a ratio of its limit over request.

For more information on how CPU and memory are measured, see [Compute Resources](#).

15.3.2. Container Limits

Supported Resources:

- CPU
- Memory

Supported Constraints:

Per container, the following must hold true if specified:

Table 15.4. Container

Constraint	Behavior
Min	<p>Min[resource] less than or equal to container.resources.requests[resource] (required) less than or equal to container/resources.limits[resource] (optional)</p> <p>If the configuration defines a min CPU, then the request value must be greater than the CPU value. A limit value does not need to be specified.</p>
Max	<p>container.resources.limits[resource] (required) less than or equal to Max[resource]</p> <p>If the configuration defines a max CPU, then you do not need to define a request value, but a limit value does need to be set that satisfies the maximum CPU constraint.</p>
MaxLimitRequestRatio	<p>MaxLimitRequestRatio[resource] less than or equal to (container.resources.limits[resource] / container.resources.requests[resource])</p> <p>If a configuration defines a maxLimitRequestRatio value, then any new containers must have both a request and limit value. Additionally, OpenShift Container Platform calculates a limit to request ratio by dividing the limit by the request. This value should be a non-negative integer greater than 1.</p> <p>For example, if a container has cpu: 500 in the limit value, and cpu: 100 in the request value, then its limit to request ratio for cpu is 5. This ratio must be less than or equal to the maxLimitRequestRatio.</p>

Supported Defaults:

Default[resource]

Defaults **container.resources.limit[resource]** to specified value if none.

Default Requests[resource]

Defaults **container.resources.requests[resource]** to specified value if none.

15.3.3. Pod Limits

Supported Resources:

- CPU
- Memory

Supported Constraints:

Across all containers in a pod, the following must hold true:

Table 15.5. Pod

Constraint	Enforced Behavior
Min	<code>Min[resource]</code> less than or equal to <code>container.resources.requests[resource]</code> (required) less than or equal to <code>container.resources.limits[resource]</code> (optional)
Max	<code>container.resources.limits[resource]</code> (required) less than or equal to <code>Max[resource]</code>
MaxLimitRequestRatio	<code>MaxLimitRequestRatio[resource]</code> less than or equal to (<code>container.resources.limits[resource] / container.resources.requests[resource]</code>)

15.4. COMPUTE RESOURCES

Each container running on a node consumes compute resources, which are measurable quantities that can be requested, allocated, and consumed.

When authoring a pod configuration file, you can optionally specify how much CPU and memory (RAM) each container needs in order to better schedule pods in the cluster and ensure satisfactory performance.

CPU is measured in units called millicores. Each node in a cluster inspects the operating system to determine the amount of CPU cores on the node, then multiplies that value by 1000 to express its total capacity. For example, if a node has 2 cores, the node's CPU capacity would be represented as 2000m. If you wanted to use 1/10 of a single core, it would be represented as 100m.

Memory is measured in bytes. In addition, it may be used with SI suffices (E, P, T, G, M, K) or their power-of-two-equivalents (Ei, Pi, Ti, Gi, Mi, Ki).

```

apiVersion: v1
kind: Pod
spec:
  containers:
  - image: openshift/hello-openshift
    name: hello-openshift
    resources:
      requests:
        cpu: 100m ①
        memory: 200Mi ②
      limits:
        cpu: 200m ③
        memory: 400Mi ④

```

① The container requests 100m cpu.

② The container requests 200Mi memory.

- 3 The container limits 200m cpu.
- 4 The container limits 400Mi memory.

15.4.1. CPU Requests

Each container in a pod can specify the amount of CPU it requests on a node. The scheduler uses CPU requests to find a node with an appropriate fit for a container.

The CPU request represents a minimum amount of CPU that your container may consume, but if there is no contention for CPU, it can use all available CPU on the node. If there is CPU contention on the node, CPU requests provide a relative weight across all containers on the system for how much CPU time the container may use.

On the node, CPU requests map to Kernel CFS shares to enforce this behavior.

15.4.2. Viewing Compute Resources

To view compute resources for a pod:

```
$ oc describe pod ruby-hello-world-tfjxt
Name:          ruby-hello-world-tfjxt
Namespace:     default
Image(s):      ruby-hello-world
Node:          /
Labels:        run=ruby-hello-world
Status:        Pending
Reason:
Message:
IP:
Replication Controllers:  ruby-hello-world (1/1 replicas created)
Containers:
  ruby-hello-world:
    Container ID:
    Image ID:
    Image:        ruby-hello-world
    QoS Tier:
      cpu:  Burstable
      memory: Burstable
    Limits:
      cpu:    200m
      memory: 400Mi
    Requests:
      cpu:        100m
      memory:     200Mi
    State:        Waiting
    Ready:        False
    Restart Count: 0
    Environment Variables:
```

15.4.3. CPU Limits

Each container in a pod can specify the amount of CPU it is limited to use on a node. CPU limits control the maximum amount of CPU that your container may use independent of contention on the node. If a

container attempts to exceed the specified limit, the system will throttle the container. This allows the container to have a consistent level of service independent of the number of pods scheduled to the node.

15.4.4. Memory Requests

By default, a container is able to consume as much memory on the node as possible. In order to improve placement of pods in the cluster, specify the amount of memory required for a container to run. The scheduler will then take available node memory capacity into account prior to binding your pod to a node. A container is still able to consume as much memory on the node as possible even when specifying a request.

15.4.5. Memory Limits

If you specify a memory limit, you can constrain the amount of memory the container can use. For example, if you specify a limit of 200Mi, a container will be limited to using that amount of memory on the node. If the container exceeds the specified memory limit, it will be terminated and potentially restarted dependent upon the container restart policy.

15.4.6. Quality of Service Tiers

A compute resource is classified with a *quality of service* (QoS) based on the specified request and limit value.

Quality of Service	Description
BestEffort	Provided when a request and limit are not specified.
Burstable	Provided when a request is specified that is less than an optionally specified limit.
Guaranteed	Provided when a limit is specified that is equal to an optionally specified request.

A container may have a different quality of service for each compute resource. For example, a container can have **Burstable** CPU and **Guaranteed** memory qualities of service.

The quality of service has different impacts on different resources, depending on whether the resource is compressible or not. CPU is a compressible resource, whereas memory is an incompressible resource.

With CPU Resources:

- A **BestEffort** CPU container is able to consume as much CPU as is available on a node but runs with the lowest priority.
- A **Burstable** CPU container is guaranteed to get the minimum amount of CPU requested, but it may or may not get additional CPU time. Excess CPU resources are distributed based on the amount requested across all containers on the node.
- A **Guaranteed** CPU container is guaranteed to get the amount requested and no more, even if there are additional CPU cycles available. This provides a consistent level of performance independent of other activity on the node.

With Memory Resources:

- A **BestEffort** memory container is able to consume as much memory as is available on the node, but there are no guarantees that the scheduler will place that container on a node with enough memory to meet its needs. In addition, a **BestEffort** container has the greatest chance of being killed if there is an out of memory event on the node.
- A **Burstable** memory container is scheduled on the node to get the amount of memory requested, but it may consume more. If there is an out of memory event on the node, **Burstable** containers are killed after **BestEffort** containers when attempting to recover memory.
- A **Guaranteed** memory container gets the amount of memory requested, but no more. In the event of an out of memory event, it will only be killed if there are no more **BestEffort** or **Burstable** containers on the system.


15.4.7. Specifying Compute Resources via CLI

To specify compute resources via the CLI:

```
$ oc run ruby-hello-world --image=ruby-hello-world --
limits=cpu=200m,memory=400Mi --requests=cpu=100m,memory=200Mi
```

15.4.8. Opaque Integer Resources

Opaque integer resources allow cluster operators to provide new node-level resources that would be otherwise unknown to the system. Users can consume these resources in pod specifications, similar to CPU and memory. The scheduler performs resource accounting so that no more than the available amount is simultaneously allocated to pods.



NOTE

Opaque integer resources are Alpha currently, and only resource accounting is implemented. There is no resource quota or limit range support for these resources, and they have no impact on QoS.

Opaque integer resources are called *opaque* because OpenShift Container Platform does not know what the resource is, but will schedule a pod on a node only if enough of that resource is available. They are called *integer resources* because they must be available, or *advertised*, in integer amounts. The API server restricts quantities of these resources to whole numbers. Examples of *valid* quantities are **3**, **3000m**, and **3Ki**.

The cluster administrator is usually responsible for creating the resources and making them available. For more information on creating opaque integer resources, see [Opaque Integer Resources](#) in the Administrator Guide.

To consume an opaque integer resource in a pod, edit the pod to include the name of the opaque resource as a key in the `spec.containers[].resources.requests` field.

For example: The following pod requests two CPUs and one `foo` (an opaque resource).

```
apiVersion: v1
kind: Pod
```

```
metadata:
  name: my-pod
spec:
  containers:
  - name: my-container
    image: myimage
    resources:
      requests:
        cpu: 2
        pod.alpha.kubernetes.io/opaque-int-resource-foo: 1
```

The pod will be scheduled only if all of the resource requests are satisfied (including CPU, memory, and any opaque resources). The pod will remain in the **PENDING** state while the resource request cannot be met by any node.

```
Conditions:
  Type      Status
  PodSchede False
...
Events:
  FirstSeen    LastSeen  Count  From              SubObjectPath Type      Reason
Message
-----
-----
14s          0s        6      default-scheduler Warning    FailedScheduling  No nodes
are available that match all of the following predicates:: Insufficient
pod.alpha.kubernetes.io/opaque-int-resource-foo (1).
```

15.5. PROJECT RESOURCE LIMITS

[Resource limits can be set per-project](#) by cluster administrators. Developers do not have the ability to create, edit, or delete these limits, but can [view them](#) for projects they have access to.

CHAPTER 16. GETTING TRAFFIC INTO A CLUSTER

16.1. GETTING TRAFFIC INTO A CLUSTER

OpenShift Container Platform provides multiple methods for communicating from outside the cluster with services running in the cluster.



NOTE

The procedures in this section require prerequisites performed by the cluster administrator.

Administrators can expose a service endpoint that external traffic can reach, by assigning a unique external IP address to that service from a range of external IP addresses. Administrators can designate a range of addresses using a CIDR notation, which allows an application user to make a request against the cluster for an external IP address.

Each IP address should be assigned to only one service to ensure that each service has a unique endpoint. Potential port clashes are handled on a first-come, first-served basis.

The recommendation, in order of preference, is:

- If you have HTTP/HTTPS, use a [router](#).
- If you have a TLS-encrypted protocol other than HTTPS (for example, TLS with the SNI header), use a [router](#).
- Otherwise, use a [Load Balancer](#), an [External IP](#), or a [NodePort](#).

Method	Purpose
Use a router	Allows access to HTTP/HTTPS traffic and TLS-encrypted protocols other than HTTPS (for example, TLS with the SNI header).
Automatically Assign a Public IP Using a Load Balancer Service	Allows traffic to non-standard ports through an IP address assigned from a pool.
Manually assign an external IP to a service	Allows traffic to non-standard ports through a specific IP address.
Configure a NodePort	Expose a service on all nodes in the cluster.

16.2. USING A ROUTER TO GET TRAFFIC INTO THE CLUSTER

16.2.1. Overview

Using a router is the most common way [to allow external access to an OpenShift Container Platform cluster](#).

A [router](#) is configured to accept external requests and proxy them based on the configured [routes](#). This is limited to HTTP/HTTPS(SNI)/TLS(SNI), which covers web applications.

16.2.2. Administrator Prerequisites

Before starting this procedure, the administrator must:

- Set up the external port to the cluster networking environment so that requests can reach the cluster. For example, names can be configured into DNS to point to specific nodes or other IP addresses in the cluster. The [DNS wildcard](#) feature can be used to configure a subset of names to an IP address in the cluster. This allows the users to set up routes within the cluster without further administrator attention.
- Make sure that the local firewall on each node permits the request to reach the IP address.
- Configure the OpenShift Container Platform cluster to [use an identity provider](#) that allows appropriate user access.
- Make sure there is at least one user with cluster admin role. To add this role to a user, run the following command:

```
oc adm policy add-cluster-role-to-user cluster-admin username
```

- Have an OpenShift Container Platform cluster with at least one master and at least one node and a system outside the cluster that has network access to the cluster. This procedure assumes that the external system is on the same subnet as the cluster. The additional networking required for external systems on a different subnet is out-of-scope for this topic.

16.2.2.1. Defining the Public IP Range

The first step in allowing access to a service is to define an external IP address range in the master configuration file:

1. Log into OpenShift Container Platform as a user with the cluster admin role.

```
$ oc login
Authentication required (openshift)
Username: admin
Password:
Login successful.
```

```
You have access to the following projects and can switch between
them with 'oc project <projectname>':
* default
Using project "default".
```

2. Configure the `externalIPNetworkCIDRs` parameter in the `/etc/origin/master/master-config.yaml` file as shown:

```
networkConfig:
  externalIPNetworkCIDRs:
    - <ip_address>/<cidr>
```

For example:

■

```
networkConfig:
  externalIPNetworkCIDRs:
    - 192.168.120.0/24
```

3. Restart the OpenShift Container Platform master service to apply the changes.

```
# systemctl restart atomic-openshift-master
```

CAUTION

The IP address pool must terminate at one or more nodes in the cluster.

16.2.3. Create a Project and Service

If the project and service that you want to expose do not exist, first create the project, then the service.

If the project and service already exist, go to the next step: **Expose the Service to Create a Route**.

1. Log into OpenShift Container Platform.
2. Create a new project for your service:

```
$ oc new-project <project_name>
```

For example:

```
$ oc new-project external-ip
```

3. Use the `oc new-app` command to [create a service](#):
For example:

```
$ oc new-app \
  -e MYSQL_USER=admin \
  -e MYSQL_PASSWORD=redhat \
  -e MYSQL_DATABASE=mysqlldb \
  registry.access.redhat.com/openshift3/mysql-55-rhel7
```

4. Run the following command to see that the new service is created:

```
oc get svc
NAME                                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
mysql-55-rhel7                     172.30.131.89   <none>           3306/TCP         13m
```

By default, the new service does not have an external IP address.

16.2.4. Expose the Service to Create a Route

You must [expose the service as a route](#) using the `oc expose` command.

To expose the service:

1. Log into OpenShift Container Platform.

2. Log into the project where the service you want to expose is located.

```
$ oc project project1
```

3. Run the following command to expose the route:

```
oc expose service <service-name>
```

For example:

```
oc expose service mysql-55-rhel7  
route "mysql-55-rhel7" exposed
```

4. On the master, use a tool, such as cURL, to make sure you can reach the service using the cluster IP address for the service:

```
curl <pod-ip>:<port>
```

For example:

```
curl 172.30.131.89:3306
```

The examples in this section use a MySQL service, which requires a client application. If you get a string of characters with the **Got packets out of order** message, you are connected to the service.

If you have a MySQL client, log in with the standard CLI command:

```
$ mysql -h 172.30.131.89 -u admin -p  
Enter password:  
Welcome to the MariaDB monitor.  Commands end with ; or \g.  
  
MySQL [(none)]>
```

16.2.5. Configure the Router

Work with your administrator to configure a router to accept external requests and proxy them based on the configured routes.

The administrator can create a [wildcard DNS](#) entry and then set up a router. Then, you can self-service the edge router without having to contact the administrators.

The router has controls to allow the administrator to specify whether the users can self-provision host names or the host names require a specific pattern.

When a set of routes is created in various projects, the overall set of routes is available to the set of routers. Each router admits (or selects) routes from the set of routes. By default, all routers admit all routes.

Routers that have permission to view all of the [labels](#) in all projects can select routes to admit based on the labels. This is called [router sharding](#). This is useful when balancing incoming traffic load among a set of routers and when isolating traffic to a specific router. For example, company A goes to one router and company B to another.

Since a router runs on a specific node, when it or the node fails traffic ingress stops. The impact of this can be reduced by creating redundant routers on different nodes and using [high availability](#) to switch the router IP address when a node fails.

16.2.6. Configure IP Failover using VIPs

Optionally, an administrator can configure IP failover.

IP failover manages a pool of Virtual IP (VIP) addresses on a set of nodes. Every VIP in the set is serviced by a node selected from the set. As long as a single node is available, the VIPs will be served. There is no way to explicitly distribute the VIPs over the nodes. As such, there may be nodes with no VIPs and other nodes with multiple VIPs. If there is only one node, all VIPs will be on it.

The VIPs must be routable from outside the cluster.

To configure IP failover:

1. On the master, make sure the `ipfailover` service account has sufficient security privileges:

```
oc adm policy add-scc-to-user privileged -z ipfailover
```

2. Run the following command to create the IP failover:

```
oc adm ipfailover --virtual-ips=<exposed-ip-address> --watch-port=
<exposed-port> --replicas=<number-of-pods> --create
```

For example:

```
oc adm ipfailover --virtual-ips="172.30.233.169" --watch-port=32315
--replicas=4 --create
--> Creating IP failover ipfailover ...
    serviceaccount "ipfailover" created
    deploymentconfig "ipfailover" created
--> Success
```

16.3. USING A LOAD BALANCER TO GET TRAFFIC INTO THE CLUSTER

16.3.1. Overview

If you do not need a specific external IP address, you can configure a load balancer service to allow external access to an OpenShift Container Platform cluster.

A load balancer service allocates a unique IP from a configured pool. The load balancer has a single edge router IP (which can be a [virtual IP \(VIP\)](#), but is still a single machine for initial load balancing).

This process involves the following:

- [The administrator performs the prerequisites](#) ;
- [The developer creates a project and service](#) , if the service to be exposed does not exist;
- [The developer exposes the service to create a route](#) .
- [The developer creates the Load Balancer Service](#) .

- [The network administrator configures networking to the service](#) .

16.3.2. Administrator Prerequisites

Before starting this procedure, the administrator must:

- Set up the external port to the cluster networking environment so that requests can reach the cluster. For example, names can be configured into DNS to point to specific nodes or other IP addresses in the cluster. The [DNS wildcard](#) feature can be used to configure a subset of names to an IP address in the cluster. This allows the users to set up routes within the cluster without further administrator attention.
- Make sure that the local firewall on each node permits the request to reach the IP address.
- Configure the OpenShift Container Platform cluster to [use an identity provider](#) that allows appropriate user access.
- Make sure there is at least one user with cluster admin role. To add this role to a user, run the following command:

```
oc adm policy add-cluster-role-to-user cluster-admin username
```

- Have an OpenShift Container Platform cluster with at least one master and at least one node and a system outside the cluster that has network access to the cluster. This procedure assumes that the external system is on the same subnet as the cluster. The additional networking required for external systems on a different subnet is out-of-scope for this topic.

16.3.2.1. Defining the Public IP Range

The first step in allowing access to a service is to define an external IP address range in the master configuration file:

1. Log into OpenShift Container Platform as a user with the cluster admin role.

```
$ oc login
Authentication required (openshift)
Username: admin
Password:
Login successful.

You have access to the following projects and can switch between
them with 'oc project <projectname>':
* default
Using project "default".
```

2. Configure the **externalIPNetworkCIDRs** parameter in the **/etc/origin/master/master-config.yaml** file as shown:

```
networkConfig:
  externalIPNetworkCIDRs:
  - <ip_address>/<cidr>
```

For example:

```
networkConfig:
```

```
externalIPNetworkCIDRs:
- 192.168.120.0/24
```

3. Restart the OpenShift Container Platform master service to apply the changes.

```
# systemctl restart atomic-openshift-master
```

CAUTION

The IP address pool must terminate at one or more nodes in the cluster.

16.3.3. Create a Project and Service

If the project and service that you want to expose do not exist, first create the project, then the service.

If the project and service already exist, go to the next step: **Expose the Service to Create a Route** .

1. Log into OpenShift Container Platform.
2. Create a new project for your service:

```
$ oc new-project <project_name>
```

For example:

```
$ oc new-project external-ip
```

3. Use the `oc new-app` command to [create a service](#):

For example:

```
$ oc new-app \
  -e MYSQL_USER=admin \
  -e MYSQL_PASSWORD=redhat \
  -e MYSQL_DATABASE=mysqlldb \
  registry.access.redhat.com/openshift3/mysql-55-rhel7
```

4. Run the following command to see that the new service is created:

```
oc get svc
NAME                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
mysql-55-rhel7      172.30.131.89   <none>           3306/TCP         13m
```

By default, the new service does not have an external IP address.

16.3.4. Expose the Service to Create a Route

You must [expose the service as a route](#) using the `oc expose` command.

To expose the service:

1. Log into OpenShift Container Platform.

2. Log into the project where the service you want to expose is located.

```
$ oc project project1
```

3. Run the following command to expose the route:

```
oc expose service <service-name>
```

For example:

```
oc expose service mysql-55-rhel7
route "mysql-55-rhel7" exposed
```

4. On the master, use a tool, such as cURL, to make sure you can reach the service using the cluster IP address for the service:

```
curl <pod-ip>:<port>
```

For example:

```
curl 172.30.131.89:3306
```

The examples in this section use a MySQL service, which requires a client application. If you get a string of characters with the **Got packets out of order** message, you are connected to the service.

If you have a MySQL client, log in with the standard CLI command:

```
$ mysql -h 172.30.131.89 -u admin -p
Enter password:
Welcome to the MariaDB monitor.  Commands end with ; or \g.

MySQL [(none)]>
```

Then, perform the following tasks:

- [Create the Load Balancer Service](#)
- [Configure networking](#)
- [Configure IP Failover](#)

16.3.5. Create the Load Balancer Service

To create a load balancer service:

1. Log into OpenShift Container Platform.
2. Load the project where the service you want to expose is located. If the project or service does not exist, see [Create a Project and Service](#).

```
$ oc project project1
```

3. Open a text file on the master node and paste the following text, editing the file as needed:

Example 16.1. Sample load balancer configuration file

```
apiVersion: v1
kind: Service
metadata:
  name: egress-2 ❶
spec:
  ports:
    - name: db
      port: 3306 ❷
  loadBalancerIP:
  type: LoadBalancer ❸
  selector:
    name: mysql ❹
```

- ❶ Enter a descriptive name for the load balancer service.
- ❷ Enter the same port that the service you want to expose is listening on.
- ❸ Enter **loadbalancer** as the type.
- ❹ Enter the name of the service.

4. Save and exit the file.
5. Run the following command to create the service:

```
oc create -f <file-name>
```

For example:

```
oc create -f mysql-lb.yaml
```

6. Execute the following command to view the new service:

```
oc get svc
NAME                                CLUSTER-IP          EXTERNAL-IP
PORT(S)                            AGE
egress-2                           172.30.236.167      172.29.121.74,172.29.121.74
3306/TCP                            6s
```

Note that the service has an external IP address automatically assigned.

7. On the master, use a tool, such as cURL, to make sure you can reach the service using the public IP address:

```
$ curl <public-ip>:<port>
```

++ For example:

```
$ curl 172.29.121.74:3306
```

The examples in this section use a MySQL service, which requires a client application. If you get a string of characters with the **Got packets out of order** message, you are connecting with the service:

If you have a MySQL client, log in with the standard CLI command:

```
$ mysql -h 172.30.131.89 -u admin -p
Enter password:
Welcome to the MariaDB monitor.  Commands end with ; or \g.

MySQL [(none)]>
```

16.3.6. Configuring Networking

The following steps are general guidelines for configuring the networking required to access the exposed service from other nodes. As network environments vary, consult your network administrator for specific configurations that need to be made within your environment.

These steps assume that all of the systems are on the same subnet.

On the Node:

1. Restart the network to make sure the network is up.

```
$ service network restart
Restarting network (via systemctl): [ OK ]
```

If the network is not up, you will receive error messages such as **Network is unreachable** when executing the following commands.

2. Add a route between the IP address of the exposed service on the master and the IP address of the master host:

```
$ route add -net 172.29.121.74 netmask 255.255.0.0 gw 10.16.41.22
dev eth0
```

3. Use a tool, such as cURL, to make sure you can reach the service using the public IP address:

```
$ curl <public-ip>:<port>
```

For example:

```
curl 172.29.121.74:3306
```

If you get a string of characters with the **Got packets out of order** message, your service is accessible from the node.

On the system that is not in the cluster:

1. Restart the network to make sure the network is up.

```
$ service network restart
```

```
Restarting network (via systemctl): [ OK ]
```

If the network is not up, you will receive error messages such as **Network is unreachable** when executing the following commands.

2. Add a route between the IP address of the exposed service on master and the IP address of the master host:

```
$ route add -net 172.29.121.74 netmask 255.255.0.0 gw 10.16.41.22
dev eth0
```

3. Make sure you can reach the service using the public IP address:

```
$ curl <public-ip>:<port>
```

For example:

```
curl 172.29.121.74:3306
```

If you get a string of characters with the **Got packets out of order** message, your service is accessible outside the cluster.

16.3.7. Configure IP Failover using VIPs

Optionally, an administrator can configure IP failover.

IP failover manages a pool of Virtual IP (VIP) addresses on a set of nodes. Every VIP in the set is serviced by a node selected from the set. As long as a single node is available, the VIPs will be served. There is no way to explicitly distribute the VIPs over the nodes. As such, there may be nodes with no VIPs and other nodes with multiple VIPs. If there is only one node, all VIPs will be on it.

The VIPs must be routable from outside the cluster.

To configure IP failover:

1. On the master, make sure the **ipfailover** service account has sufficient security privileges:

```
oc adm policy add-scc-to-user privileged -z ipfailover
```

2. Run the following command to create the IP failover:

```
oc adm ipfailover --virtual-ips=<exposed-ip-address> --watch-port=
<exposed-port> --replicas=<number-of-pods> --create
```

For example:

```
oc adm ipfailover --virtual-ips="172.30.233.169" --watch-port=32315
--replicas=4 --create
--> Creating IP failover ipfailover ...
    serviceaccount "ipfailover" created
    deploymentconfig "ipfailover" created
--> Success
```

16.4. USING A SERVICE EXTERNAL IP TO GET TRAFFIC INTO THE CLUSTER


16.4.1. Overview

One method to expose a service is to assign an external IP access directly to the service you want to make accessible from outside the cluster.

Make sure you have created a range of IP addresses to use, as shown in [Defining the Public IP Address Range](#).

By setting an external IP on the service, OpenShift Container Platform sets up IP table rules to allow traffic arriving at any cluster node that is targeting that IP address to be sent to one of the internal pods. This is similar to the internal service IP addresses, but the external IP tells OpenShift Container Platform that this service should also be exposed externally at the given IP. The administrator must assign the IP address to a host (node) interface on one of the nodes in the cluster. Alternatively, the address can be used as a [virtual IP \(VIP\)](#).

These IPs are not managed by OpenShift Container Platform and administrators are responsible for ensuring that traffic arrives at a node with this IP.



NOTE

The following is a non-HA solution and does not configure [IP failover](#). IP failover is required to make the service highly-available.

This process involves the following:

- [The administrator performs the prerequisites](#) ;
- [The developer creates a project and service](#) , if the service to be exposed does not exist;
- [The developer exposes the service to create a route](#) .
- [The developer assigns the IP address to the service](#) .
- [The network administrator configures networking to the service](#) .

16.4.2. Administrator Prerequisites

Before starting this procedure, the administrator must:

- Set up the external port to the cluster networking environment so that requests can reach the cluster. For example, names can be configured into DNS to point to specific nodes or other IP addresses in the cluster. The [DNS wildcard](#) feature can be used to configure a subset of names to an IP address in the cluster. This allows the users to set up routes within the cluster without further administrator attention.
- Make sure that the local firewall on each node permits the request to reach the IP address.
- Configure the OpenShift Container Platform cluster to [use an identity provider](#) that allows appropriate user access.
- Make sure there is at least one user with cluster admin role. To add this role to a user, run the following command:


```
oc adm policy add-cluster-role-to-user cluster-admin username
```

- Have an OpenShift Container Platform cluster with at least one master and at least one node and a system outside the cluster that has network access to the cluster. This procedure assumes that the external system is on the same subnet as the cluster. The additional networking required for external systems on a different subnet is out-of-scope for this topic.

16.4.2.1. Defining the Public IP Range

The first step in allowing access to a service is to define an external IP address range in the master configuration file:

1. Log into OpenShift Container Platform as a user with the cluster admin role.

```
$ oc login
Authentication required (openshift)
Username: admin
Password:
Login successful.

You have access to the following projects and can switch between
them with 'oc project <projectname>':
* default
Using project "default".
```

2. Configure the `externalIPNetworkCIDRs` parameter in the `/etc/origin/master/master-config.yaml` file as shown:

```
networkConfig:
  externalIPNetworkCIDRs:
    - <ip_address>/<cidr>
```

For example:

```
networkConfig:
  externalIPNetworkCIDRs:
    - 192.168.120.0/24
```

3. Restart the OpenShift Container Platform master service to apply the changes.

```
# systemctl restart atomic-openshift-master
```

CAUTION

The IP address pool must terminate at one or more nodes in the cluster.

16.4.3. Create a Project and Service

If the project and service that you want to expose do not exist, first create the project, then the service.

If the project and service already exist, go to the next step: **Expose the Service to Create a Route**.

1. Log into OpenShift Container Platform.

2. Create a new project for your service:

```
$ oc new-project <project_name>
```

For example:

```
$ oc new-project external-ip
```

3. Use the `oc new-app` command to [create a service](#):

For example:

```
$ oc new-app \  
  -e MYSQL_USER=admin \  
  -e MYSQL_PASSWORD=redhat \  
  -e MYSQL_DATABASE=mysqlldb \  
  registry.access.redhat.com/openshift3/mysql-55-rhel7
```

4. Run the following command to see that the new service is created:

```
oc get svc  
NAME                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE  
mysql-55-rhel7      172.30.131.89   <none>           3306/TCP         13m
```

By default, the new service does not have an external IP address.

16.4.4. Expose the Service to Create a Route

You must [expose the service as a route](#) using the `oc expose` command.

To expose the service:

1. Log into OpenShift Container Platform.

2. Log into the project where the service you want to expose is located.

```
$ oc project project1
```

3. Run the following command to expose the route:

```
oc expose service <service-name>
```

For example:

```
oc expose service mysql-55-rhel7  
route "mysql-55-rhel7" exposed
```

4. On the master, use a tool, such as `cURL`, to make sure you can reach the service using the cluster IP address for the service:

```
curl <pod-ip>:<port>
```

For example:

```
curl 172.30.131.89:3306
```

The examples in this section use a MySQL service, which requires a client application. If you get a string of characters with the **Got packets out of order** message, you are connected to the service.

If you have a MySQL client, log in with the standard CLI command:

```
$ mysql -h 172.30.131.89 -u admin -p
Enter password:
Welcome to the MariaDB monitor.  Commands end with ; or \g.

MySQL [(none)]>
```

Then, perform the following tasks:

- [Assign an IP Address to the Service](#)
- [Configure networking](#)
- [Configure IP Failover](#)

16.4.5. Assigning an IP Address to the Service

To assign an external IP address to a service:

1. Log into OpenShift Container Platform.
2. Load the project where the service you want to expose is located. If the project or service does not exist, see [Create a Project and Service](#) in the Prerequisites.
3. Run the following command to assign an external IP address to the service you want to access. Use an IP address from the [external IP address range](#):

```
oc patch svc <name> -p '{"spec":{"externalIPs":["<ip_address>"]}}'
```

The **<name>** is the name of the service and **-p** indicates a patch to be applied to the service JSON file. The expression in the brackets will assign the specified IP address to the specified service.

For example:

```
oc patch svc mysql-55-rhel7 -p '{"spec":{"externalIPs":
["192.174.120.10"]}}'

"mysql-55-rhel7" patched
```

4. Run the following command to see that the service has a public IP:

```
oc get svc
NAME                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
mysql-55-rhel7      172.30.131.89   192.174.120.10   3306/TCP         13m
```

5. On the master, use a tool, such as cURL, to make sure you can reach the service using the public IP address:

```
$ curl <public-ip>:<port>
```

For example:

```
curl 192.168.120.10:3306
```

If you get a string of characters with the **Got packets out of order** message, you are connected to the service.

If you have a MySQL client, log in with the standard CLI command:


```
$ mysql -h 192.168.120.10 -u admin -p
Enter password:
Welcome to the MariaDB monitor. Commands end with ; or \g.

MySQL [(none)]>
```

16.4.6. Configuring Networking

After the external IP address is assigned, you need to create routes to that IP.

The following steps are general guidelines for configuring the networking required to access the exposed service from other nodes. As network environments vary, consult your network administrator for specific configurations that need to be made within your environment.



NOTE

These steps assume that all of the systems are on the same subnet.

On the master:

1. Restart the network to make sure the network is up.

```
$ service network restart
Restarting network (via systemctl): [ OK ]
```

If the network is not up, you will receive error messages such as **Network is unreachable** when running the following commands.

2. Run the following command with the external IP address of the service you want to expose and device name associated with the host IP from the **ifconfig** command output:

```
$ ip address add <external-ip> dev <device>
```

For example:

```
$ ip address add 192.168.120.10 dev eth0
```

If you need to, run the following command to obtain the IP address of the host server where the master resides:

```
$ ifconfig
```

Look for the device that is listed similar to: **UP, BROADCAST, RUNNING, MULTICAST**.

```
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
      inet 10.16.41.22 netmask 255.255.248.0 broadcast
10.16.47.255
      ...
```

3. Add a route between the IP address of the host where the master resides and the gateway IP address of the master host:

```
$ route add -net <host_ip_address> netmask <netmask> gw
<gateway_ip_address> dev <device>
```

For example:

```
$ route add -host 10.16.41.22 gw 10.16.41.254 dev eth0
```

The **netstat -nr** command provides the gateway IP address:

```
$ netstat -nr
Kernel IP routing table
Destination      Gateway          Genmask         Flags   MSS Window
  irtt Iface
0.0.0.0          10.16.41.254    0.0.0.0         UG        0 0
0 eth0
```

4. Add a route between the IP address of the exposed service and the IP address of the master host:

```
$ route add -net 192.174.120.0/24 gw 10.16.41.22 eth0
```

On the Node:

1. Restart the network to make sure the network is up.

```
$ service network restart
Restarting network (via systemctl): [ OK ]
```

If the network is not up, you will receive error messages such as **Network is unreachable** when executing the following commands.

2. Add a route between IP address of the host where the node is located and the gateway IP of the node host:

```
$ route add -net 10.16.40.0 netmask 255.255.248.0 gw 10.16.47.254
eth0
```

The **ifconfig** command displays the host IP:

```
ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
```

```

    inet 10.16.41.71 netmask 255.255.255.0 broadcast
10.19.41.255

```

The **netstat -nr** command displays the gateway IP:

```

netstat -nr
Kernel IP routing table
Destination      Gateway          Genmask          Flags      MSS Window
 irtt Iface
0.0.0.0          10.16.41.254    0.0.0.0          UG          0 0
0 eth0

```

3. Add a route between the IP address of the exposed service and the IP address of the host system where the master node resides:

```

$ route add -net 192.174.120.0 netmask 255.255.255.0 gw 10.16.41.22
dev eth0

```

4. Use a tool, such as cURL, to make sure you can reach the service using the public IP address:

```

$ curl <public-ip>:<port>

```

For example:

```

curl 192.168.120.10:3306

```

If you get a string of characters with the **Got packets out of order** message, your service is accessible from the node.

On the system that is not in the cluster:

1. Restart the network to make sure the network is up.

```

$ service network restart
Restarting network (via systemctl): [ OK ]

```

If the network is not up, you will receive error messages such as **Network is unreachable** when executing the following commands.

2. Add a route between the IP address of the remote host and the gateway IP of the remote host:

```

$ route add -net 10.16.64.0 netmask 255.255.248.0 gw 10.16.71.254
eno1

```

3. Add a route between the IP address of the exposed service on master and the IP address of the master host:

```

$ route add -net 192.174.120.0 netmask 255.255.255.0 gw 10.16.41.22

```

4. Use a tool, such as cURL, to make sure you can reach the service using the public IP address:

```

$ curl <public-ip>:<port>

```

For example:

```
curl 192.168.120.10:3306
```

If you get a string of characters with the **Got packets out of order** message, your service is accessible outside the cluster.

16.4.7. Configure IP Failover using VIPs

Optionally, an administrator can configure IP failover.

IP failover manages a pool of Virtual IP (VIP) addresses on a set of nodes. Every VIP in the set is serviced by a node selected from the set. As long as a single node is available, the VIPs will be served. There is no way to explicitly distribute the VIPs over the nodes. As such, there may be nodes with no VIPs and other nodes with multiple VIPs. If there is only one node, all VIPs will be on it.

The VIPs must be routable from outside the cluster.

To configure IP failover:

1. On the master, make sure the **ipfailover** service account has sufficient security privileges:

```
oc adm policy add-scc-to-user privileged -z ipfailover
```

2. Run the following command to create the IP failover:

```
oc adm ipfailover --virtual-ips=<exposed-ip-address> --watch-port=
<exposed-port> --replicas=<number-of-pods> --create
```

For example:

```
oc adm ipfailover --virtual-ips="172.30.233.169" --watch-port=32315
--replicas=4 --create
--> Creating IP failover ipfailover ...
    serviceaccount "ipfailover" created
    deploymentconfig "ipfailover" created
--> Success
```

16.5. USING A NODEPORT TO GET TRAFFIC INTO THE CLUSTER

16.5.1. Overview

Use NodePorts [to expose the service](#) nodePort on all nodes in the cluster.

Using NodePorts requires additional port resources.

A node port exposes the service on a static port on the node IP address.

NodePorts are in the 30000-32767 range by default, which means a NodePort is unlikely to match a service's intended port (for example, 8080 may be exposed as 31020).

The administrator must ensure the external IPs are routed to the nodes and local firewall rules on all nodes allow access to the open port.

NodePorts and external IPs are independent and both can be used concurrently.

16.5.2. Administrator Prerequisites

Before starting this procedure, the administrator must:

- Set up the external port to the cluster networking environment so that requests can reach the cluster. For example, names can be configured into DNS to point to specific nodes or other IP addresses in the cluster. The [DNS wildcard](#) feature can be used to configure a subset of names to an IP address in the cluster. This allows the users to set up routes within the cluster without further administrator attention.
- Make sure that the local firewall on each node permits the request to reach the IP address.
- Configure the OpenShift Container Platform cluster to [use an identity provider](#) that allows appropriate user access.
- Make sure there is at least one user with cluster admin role. To add this role to a user, run the following command:

```
oc adm policy add-cluster-role-to-user cluster-admin username
```

- Have an OpenShift Container Platform cluster with at least one master and at least one node and a system outside the cluster that has network access to the cluster. This procedure assumes that the external system is on the same subnet as the cluster. The additional networking required for external systems on a different subnet is out-of-scope for this topic.

16.5.3. Configuring the Service

You specify a port number for the nodePort when you create or modify a service. If you didn't manually specify a port, system will allocate one for you.

1. Log into the master node.
2. If the project you want to use does not exist, create a new project for your service:

```
$ oc new-project <project_name>
```

For example:

```
$ oc new-project external-ip
```

3. Edit the service definition to specify `spec.type: NodePort` and optionally specify a port in the in the 30000-32767 range.

```
apiVersion: v1
kind: Service
metadata:
  name: mysql
  labels:
    name: mysql
spec:
  type: NodePort
  ports:
    - port: 3036
```



```

        nodePort: 30036
        name: http
    selector:
        name: mysql

```

4. Execute the following command to [create the service](#):

```
$ oc new-app <file-name>
```

For example:

```
oc new-app mysql.yaml
```

5. Execute the following command to see that the new service is created:

```
oc get svc
```

NAME	CLUSTER_IP	EXTERNAL_IP	PORT(S)
AGE			
mysql	172.30.89.219	<nodes>	3036:30036/TCP
2m			

Note that the external IP is listed as **<nodes>** and the node ports are listed.

You should be able to access the service using the **<NodeIP>:<NodePort>** address.

CHAPTER 17. ROUTES

17.1. OVERVIEW

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

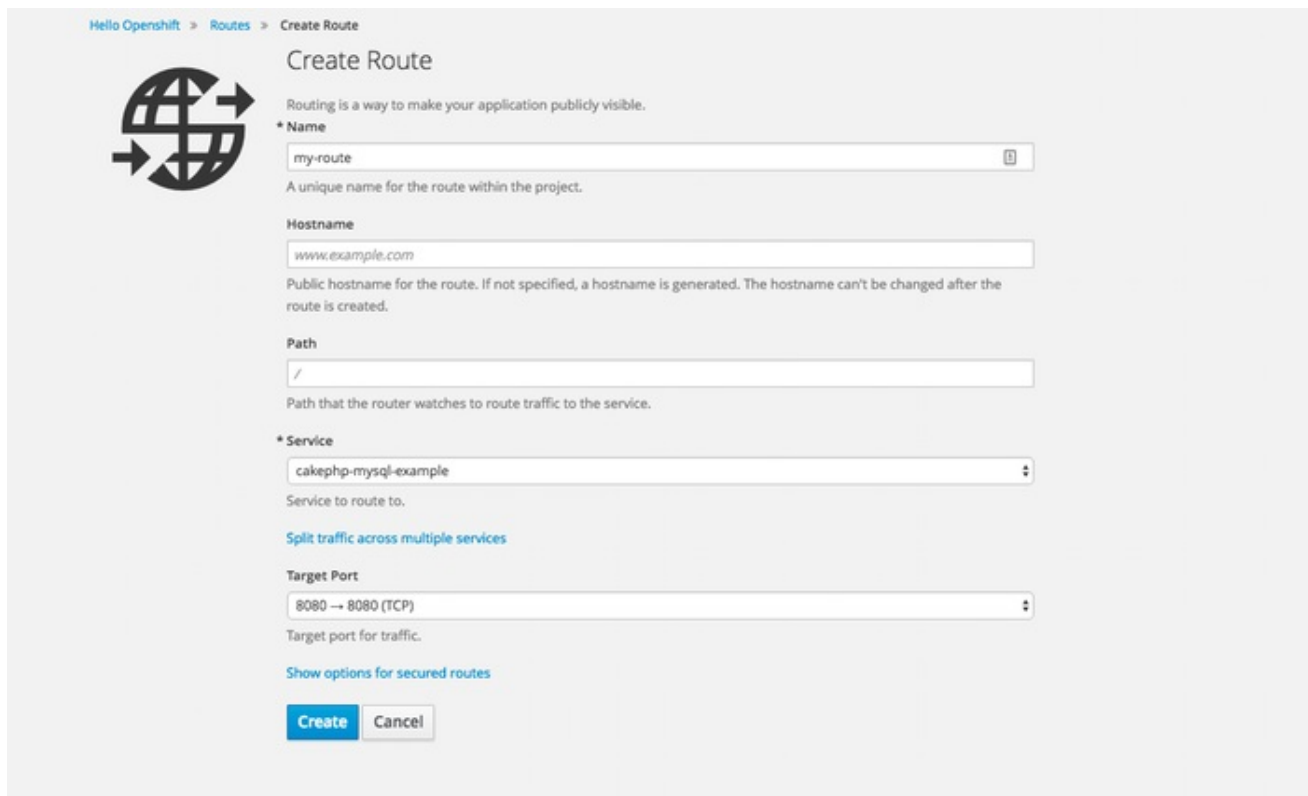
DNS resolution for a host name is handled separately from routing; your administrator may have configured a cloud domain that will always correctly resolve to the OpenShift Container Platform router, or if using an unrelated host name you may need to modify its DNS records independently to resolve to the router.

17.2. CREATING ROUTES

You can create unsecured and secured routes using the web console or the CLI.

Using the web console, you can navigate to the **Browse → Routes** page, then click **Create Route** to define and create a route in your project:

Figure 17.1. Creating a Route Using the Web Console



Hello OpenShift > Routes > Create Route

Create Route

Routing is a way to make your application publicly visible.

* Name
my-route
A unique name for the route within the project.

Hostname
www.example.com
Public hostname for the route. If not specified, a hostname is generated. The hostname can't be changed after the route is created.

Path
/
Path that the router watches to route traffic to the service.

* Service
cakephp-mysql-example
Service to route to.
[Split traffic across multiple services](#)

Target Port
8080 → 8080 (TCP)
Target port for traffic.
[Show options for secured routes](#)

Create **Cancel**

Using the CLI, the following example creates an unsecured route:

```
$ oc expose svc/frontend --hostname=www.example.com
```

The new route inherits the name from the service unless you specify one using the `--name` option.

YAML Definition of the Unsecured Route Created Above

```
apiVersion: v1
kind: Route
```

```

metadata:
  name: frontend
spec:
  host: www.example.com
  path: "/test" ❶
  to:
    kind: Service
    name: frontend

```

- ❶ For [path-based routing](#), specify a path component that can be compared against a URL.

For information on configuring routes using the CLI, see [Route Types](#).

Unsecured routes are the default configuration, and are therefore the simplest to set up. However, [secured routes](#) offer security for connections to remain private. To create a secured HTTPS route encrypted with a key and certificate (PEM-format files which you must generate and sign separately), you can use the `create route` command and optionally provide certificates and a key.



NOTE

[TLS](#) is the replacement of SSL for HTTPS and other encrypted protocols.

```

$ oc create route edge --service=frontend \
  --cert=${MASTER_CONFIG_DIR}/ca.crt \
  --key=${MASTER_CONFIG_DIR}/ca.key \
  --ca-cert=${MASTER_CONFIG_DIR}/ca.crt \
  --hostname=www.example.com

```

YAML Definition of the Secured Route Created Above

```

apiVersion: v1
kind: Route
metadata:
  name: frontend
spec:
  host: www.example.com
  to:
    kind: Service
    name: frontend
  tls:
    termination: edge
    key: |-
      -----BEGIN PRIVATE KEY-----
      [...]
      -----END PRIVATE KEY-----
    certificate: |-
      -----BEGIN CERTIFICATE-----
      [...]
      -----END CERTIFICATE-----
    caCertificate: |-
      -----BEGIN CERTIFICATE-----
      [...]
      -----END CERTIFICATE-----

```

Currently, password protected key files are not supported. HAProxy prompts for a password upon starting and does not have a way to automate this process. To remove a passphrase from a keyfile, you can run:

```
# openssl rsa -in <passwordProtectedKey.key> -out <new.key>
```

You can create a secured route without specifying a key and certificate, in which case the [router's default certificate](#) will be used for TLS termination.



NOTE

TLS termination in OpenShift Container Platform 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 host name, resulting in validation errors.

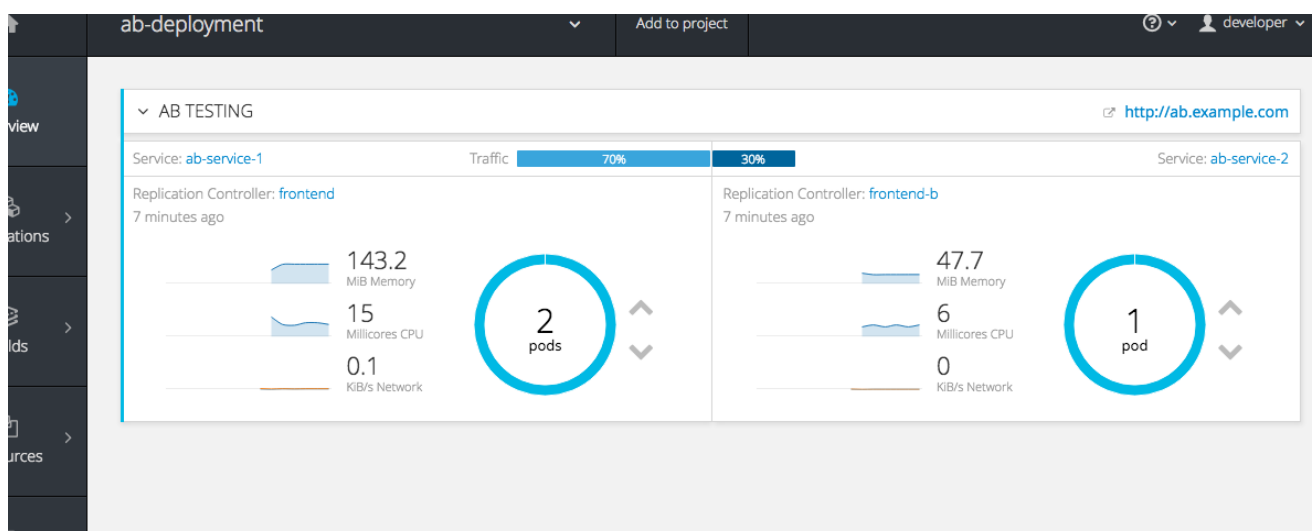
Further information on all types of [TLS termination](#) as well as [path-based routing](#) are available in the [Architecture section](#).

17.3. LOAD BALANCING FOR A/B TESTING

You can run two versions of an application, and, entirely within OpenShift Container Platform, control the percentage of traffic to and from each application for A/B testing. A/B testing is a method of comparing two versions of an application against each other to determine which one performs better.

Previously, A/B testing only worked by adding or removing more pods of every kind (A or B). However, this was not a scalable solution because for lower B percentages, you would create a large number of pods. Starting in 3.3, the HAProxy router now supports splitting the traffic coming to a route across multiple back end services via weighting.

The web console allows users to set the weighting and show balance between them:



If you have two deployments, A and B, or more, then create respective services for the pods in those deployments and use labels.

The **Route** resource now has an **alternateBackends** field, which you can use to specify **Service**. Use the **alternateBackends** and **To** fields to supply the route with all of the back end deployments grouped as services. Use the **weight** sub-field to specify a relative weight in integers ranging from 0

to 256. This value defaults to 100. The combined value of all the weights sets the relative proportions of traffic.

When you deploy the route, the router will balance the traffic according to the weights specified for the services.

To edit the route, run:

```
$ oc edit route <route-name>
```

Then, update the percentage/weight of the services in the **to** and **alternateBackends** fields.

CHAPTER 18. INTEGRATING EXTERNAL SERVICES

18.1. OVERVIEW

Many OpenShift Container Platform applications use external resources, such as external databases, or an external SaaS endpoint. These external resources can be modeled as native OpenShift Container Platform services, so that applications can work with them as they would any other internal service.

[Egress traffic](#) can be controlled by firewall rules or an Egress router. This permits having a static IP address for their application service.

18.2. DEFINING A SERVICE FOR AN EXTERNAL DATABASE

One of the most common types of external services is an external database. To support an external database, an application needs:

1. An endpoint to communicate with.
2. A set of credentials and coordinates, including:
 - A user name
 - A passphrase
 - A database name

The solution for integrating with an external database includes:

- A **Service** object to represent the SaaS provider as an OpenShift Container Platform service.
- One or more **Endpoints** for the service.
- Environment variables in the appropriate pods containing the credentials.

The following steps outline a scenario for integrating with an external MySQL database:

18.2.1. Step 1: Define a Service

You can define a service either by providing an IP address and endpoints, or by providing a Fully qualified domain name (FQDN).

18.2.1.1. Using an IP address

1. Create an [OpenShift Container Platform service](#) to represent your external database. This is similar to creating an internal service; the difference is in the service's **Selector** field. Internal OpenShift Container Platform services use the **Selector** field to associate pods with services using [labels](#). The **EndpointsController** system component synchronizes the endpoints for services that specify selectors with the pods that match the selector. The [service proxy](#) and OpenShift Container Platform [router](#) load-balance requests to the service amongst the service's endpoints.

Services that represent an external resource do not require associated pods. Instead, leave the **Selector** field unset. This represents the external service, making the **EndpointsController** ignore the service and allows you to specify endpoints manually:

```

kind: "Service"
apiVersion: "v1"
metadata:
  name: "external-mysql-service"
spec:
  ports:
    -
      name: "mysql"
      protocol: "TCP"
      port: 3306
      targetPort: 3306
      nodePort: 0
  selector: {} ❶

```

❶ The **selector** field to leave blank.

2. Next, create the required endpoints for the service. This gives the service proxy and router the location to send traffic directed to the service:

```

kind: "Endpoints"
apiVersion: "v1"
metadata:
  name: "external-mysql-service" ❶
subsets: ❷
  -
    addresses:
      -
        ip: "10.0.0.0" ❸
    ports:
      -
        port: 3306 ❹
        name: "mysql"

```

- ❶ The name of the **Service** instance, as defined in the previous step.
- ❷ Traffic to the service will be load-balanced between the supplied **Endpoints** if more than one is supplied.
- ❸ Endpoints IPs **cannot be** loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast (224.0.0.0/24).
- ❹ The **port** and **name** definition must match the **port** and **name** value in the service defined in the previous step.

18.2.1.2. Using an External Domain Name

Using external domain names make it easier to manage an external service linkage, because you do not have to worry about the external service's IP addresses changing.

ExternalName services do not have selectors, or any defined ports or endpoints, therefore, you can use an **ExternalName** service to direct traffic to an external service.

```
kind: "Service"
apiVersion: "v1"
metadata:
  name: "external-mysql-service"
spec:
  type: ExternalName
  externalName: example.domain.name
selector: {} ❶
```

- ❶ The `selector` field to leave blank.

Using an external domain name service tells the system that the DNS name in the `externalName` field (`example.domain.name` in the previous example) is the location of the resource that backs the service. When a DNS request is made against the Kubernetes DNS server, it returns the `externalName` in a CNAME record telling the client to look up the returned name to get the IP address.

18.2.2. Step 2: Consume a Service

Now that the service and endpoints are defined, give the appropriate pods access to the credentials to use the service by setting environment variables in the appropriate containers:

```
kind: "DeploymentConfig"
apiVersion: "v1"
metadata:
  name: "my-app-deployment"
spec: ❶
  strategy:
    type: "Rolling"
    rollingParams:
      updatePeriodSeconds: 1 ❷
      intervalSeconds: 1 ❸
      timeoutSeconds: 120
  replicas: 2
  selector:
    name: "frontend"
  template:
    metadata:
      labels:
        name: "frontend"
    spec:
      containers:
        -
          name: "helloworld"
          image: "origin-ruby-sample"
          ports:
            -
              containerPort: 3306
              protocol: "TCP"
          env:
            -
              name: "MYSQL_USER"
              value: "${MYSQL_USER}" ❹
```



```

-
  name: "MYSQL_PASSWORD"
  value: "${MYSQL_PASSWORD}" 5
-
  name: "MYSQL_DATABASE"
  value: "${MYSQL_DATABASE}" 6

```

- 1 Other fields on the `DeploymentConfig` are omitted
- 2 The time to wait between individual pod updates.
- 3 The time to wait between polling the deployment status after update.
- 4 The user name to use with the service.
- 5 The passphrase to use with the service.
- 6 The database name.

External Database Environment Variables

Using an external service in your application is similar to using an internal service. Your application will be assigned environment variables for the service and the additional environment variables with the credentials described in the previous step. For example, a MySQL container receives the following environment variables:

- `EXTERNAL_MYSQL_SERVICE_SERVICE_HOST=<ip_address>`
- `EXTERNAL_MYSQL_SERVICE_SERVICE_PORT=<port_number>`
- `MYSQL_USERNAME=<mysql_username>`
- `MYSQL_PASSWORD=<mysql_password>`
- `MYSQL_DATABASE_NAME=<mysql_database>`

The application is responsible for reading the coordinates and credentials for the service from the environment and establishing a connection with the database via the service.

18.3. EXTERNAL SAAS PROVIDER

A common type of external service is an external SaaS endpoint. To support an external SaaS provider, an application needs:

1. An endpoint to communicate with
2. A set of credentials, such as:
 - a. An API key
 - b. A user name
 - c. A passphrase

The following steps outline a scenario for integrating with an external SaaS provider:

18.3.1. Using an IP address and Endpoints

1. Create an [OpenShift Container Platform service](#) to represent the external service. This is similar to creating an internal service; however the difference is in the service's **Selector** field.

Internal OpenShift Container Platform services use the **Selector** field to associate pods with services using [labels](#). A system component called **EndpointsController** synchronizes the endpoints for services that specify selectors with the pods that match the selector. The [service proxy](#) and OpenShift Container Platform [router](#) load-balance requests to the service amongst the service's endpoints.

Services that represents an external resource do not require that pods be associated with it. Instead, leave the **Selector** field unset. This makes the **EndpointsController** ignore the service and allows you to specify endpoints manually:

```
kind: "Service"
apiVersion: "v1"
metadata:
  name: "example-external-service"
spec:
  ports:
  -
    name: "mysql"
    protocol: "TCP"
    port: 3306
    targetPort: 3306
    nodePort: 0
  selector: {} ❶
```

- ❶ The **selector** field to leave blank.

2. Next, create endpoints for the service containing the information about where to send traffic directed to the service proxy and the router:

```
kind: "Endpoints"
apiVersion: "v1"
metadata:
  name: "example-external-service" ❶
subsets: ❷
- addresses:
  - ip: "10.10.1.1"
  ports:
  - name: "mysql"
    port: 3306
```

- ❶ The name of the **Service** instance.

- ❷ Traffic to the service is load-balanced between the **subsets** supplied here.

3. Now that the service and endpoints are defined, give pods the credentials to use the service by setting environment variables in the appropriate containers:

```
kind: "DeploymentConfig"
```

```

apiVersion: "v1"
metadata:
  name: "my-app-deployment"
spec: ❶
  strategy:
    type: "Rolling"
    rollingParams:
      timeoutSeconds: 120
  replicas: 1
  selector:
    name: "frontend"
  template:
    metadata:
      labels:
        name: "frontend"
    spec:
      containers:
        -
          name: "helloworld"
          image: "openshift/openshift/origin-ruby-sample"
          ports:
            -
              containerPort: 3306
              protocol: "TCP"
          env:
            -
              name: "SAAS_API_KEY" ❷
              value: "<SaaS service API key>"
            -
              name: "SAAS_USERNAME" ❸
              value: "<SaaS service user>"
            -
              name: "SAAS_PASSPHRASE" ❹
              value: "<SaaS service passphrase>"

```

- ❶ Other fields on the **DeploymentConfig** are omitted.
- ❷ **SAAS_API_KEY**: The API key to use with the service.
- ❸ **SAAS_USERNAME**: The user name to use with the service.
- ❹ **SAAS_PASSPHRASE**: The passphrase to use with the service.

These variables get added to the containers as environment variables. Using environment variables allows service-to-service communication and it may or may not require additional parameters such as API keys, user name and password authentication, or certificates.

External SaaS Provider Environment Variables

Similarly, when using an internal service, your application is assigned environment variables for the service and the additional environment variables with the credentials described in the previous steps. In the previous example, the container receives the following environment variables:

- **EXAMPLE_EXTERNAL_SERVICE_SERVICE_HOST=<ip_address>**

- `EXAMPLE_EXTERNAL_SERVICE_SERVICE_PORT=<port_number>`
- `SAAS_API_KEY=<saas_api_key>`
- `SAAS_USERNAME=<saas_username>`
- `SAAS_PASSPHRASE=<saas_passphrase>`

The application reads the coordinates and credentials for the service from the environment and establishes a connection with the service.

18.3.2. Using an External Domain Name

`ExternalName` services do not have selectors, or any defined ports or endpoints. You can use an `ExternalName` service to assign traffic to an external service outside the cluster.

```
kind: "Service"
apiVersion: "v1"
metadata:
  name: "external-mysql-service"
spec:
  type: ExternalName
  externalName: example.domain.name
  selector: {} ❶
```

- ❶ The `selector` field to leave blank.

Using an `ExternalName` service maps the service to the value of the `externalName` field (`example.domain.name` in the previous example), by automatically injecting a CNAME record, mapping the service name directly to an outside DNS address, and bypassing the need for endpoint records.

CHAPTER 19. SECRETS

19.1. USING SECRETS

This topic discusses important properties of secrets and provides an overview on how developers can use them.

The `Secret` object type provides a mechanism to hold sensitive information such as passwords, OpenShift Container Platform client configuration files, `dockercfg` files, private source repository credentials, and so on. Secrets decouple sensitive content from the pods. You can mount secrets into containers using a volume plug-in or the system can use secrets to perform actions on behalf of a pod.

YAML Secret Object Definition

```
apiVersion: v1
kind: Secret
metadata:
  name: test-secret
  namespace: my-namespace
type: Opaque ❶
data: ❷
  username: dmFsdWUtMQ0K ❸
  password: dmFsdWUtMg0KDQo=
stringData: ❹
  hostname: myapp.mydomain.com ❺
```

- ❶ Indicates the [structure of the secret's key names and values](#) .
- ❷ The allowable format for the keys in the `data` field must meet the guidelines in the `DNS_SUBDOMAIN` value in [the Kubernetes identifiers glossary](#).
- ❸ The value associated with keys in the the `data` map must be base64 encoded.
- ❹❺ Entries in the `stringData` map are converted to base64 and the entry will then be moved to the `data` map automatically. This field is write-only; the value will only be returned via the `data` field.

The value associated with keys in the the `stringData` map is made up of plain text strings.

1. Create the secret from your local `.docker/config.json` file:

```
$ oc create secret generic dockerhub \
  --from-file=.dockerconfigjson=<path/to/.docker/config.json> \
  --type=kubernetes.io/dockerconfigjson
```

This command generates a JSON specification of the secret named `dockerhub` and creates the object.

YAML Opaque Secret Object Definition

```
apiVersion: v1
kind: Secret
```


19.1.3. Types of Secrets

The value in the `type` field indicates the structure of the secret's key names and values. The type can be used to enforce the presence of user names and keys in the secret object. If you do not want validation, use the `opaque` type, which is the default.

Specify one of the following types to trigger minimal server-side validation to ensure the presence of specific key names in the secret data:

- `kubernetes.io/service-account-token`. Uses a [service account token](#).
- `kubernetes.io/dockercfg`. Uses the [.dockercfg file](#) for required Docker credentials.
- `kubernetes.io/dockerconfigjson`. Uses the [.docker/config.json file](#) for required Docker credentials.
- `kubernetes.io/basic-auth`. Use with [Basic Authentication](#).
- `kubernetes.io/ssh-auth`. Use with [SSH Key Authentication](#).
- `kubernetes.io/tls`. Use with [TLS certificate authorities](#)

Specify `type= Opaque` if you do not want validation, which means the secret does not claim to conform to any convention for key names or values. An *opaque* secret, allows for unstructured `key:value` pairs that can contain arbitrary values.



NOTE

You can specify other arbitrary types, such as `example.com/my-secret-type`. These types are not enforced server-side, but indicate that the creator of the secret intended to conform to the key/value requirements of that type.

For examples of different secret types, see the [code samples](#) in *Using Secrets*.

19.1.4. Updating Secrets

When you modify the value of a secret, the value (used by an already running pod) will not dynamically change. To change a secret, you must delete the original pod and create a new pod (perhaps with an identical PodSpec).

Updating a secret follows the same workflow as deploying a new container image. You can use the `kubectl rolling-update` command.

The `resourceVersion` value in a secret is not specified when it is referenced. Therefore, if a secret is updated at the same time as pods are starting, then the version of the secret will be used for the pod will not be defined.



NOTE

Currently, it is not possible to check the resource version of a secret object that was used when a pod was created. It is planned that pods will report this information, so that a controller could restart ones using a old `resourceVersion`. In the interim, do not update the data of existing secrets, but create new ones with distinct names.

19.2. SECRETS IN VOLUMES AND ENVIRONMENT VARIABLES

See [examples](#) of YAML files with secret data.

After you [create a secret](#), you can:

1. Create the pod to reference your secret:

```
$ oc create -f <your_yaml_file>.yaml
```

2. Get the logs:

```
$ oc logs secret-example-pod
```

3. Delete the pod:

```
$ oc delete pod secret-example-pod
```

19.3. IMAGE PULL SECRETS

See [Using Image Pull Secrets](#) for more information.

19.4. SOURCE CLONE SECRETS

See [Build Inputs](#) for more information about using source clone secrets during a build.

19.5. SERVICE SERVING CERTIFICATE SECRETS

Service serving certificate secrets are intended to support complex middleware applications that need out-of-the-box certificates. It has the same settings as the server certificates generated by the administrator tooling for nodes and masters.

To secure communication to your service, have the cluster generate a signed serving certificate/key pair into a secret in your namespace. To do this, set the `service.alpha.openshift.io/serving-cert-secret-name` annotation on your service with the value set to the name you want to use for your secret. Then, your `PodSpec` can mount that secret. When it is available, your pod will run. The certificate will be good for the internal service DNS name, `<service.name>.<service.namespace>.svc`.

The certificate and key are in PEM format, stored in `tls.crt` and `tls.key` respectively. The certificate/key pair is automatically replaced when it gets close to expiration. View the expiration date in the `service.alpha.openshift.io/expiry` annotation on the secret, which is in RFC3339 format.

Other pods can trust cluster-created certificates (which are only signed for internal DNS names), by using the CA bundle in the `/var/run/secrets/kubernetes.io/serviceaccount/service-ca.crt` file that is automatically mounted in their pod.

The signature algorithm for this feature is `x509.SHA256WithRSA`. To manually rotate, delete the generated secret. A new certificate is created.

19.6. RESTRICTIONS

To use a secret, a pod needs to reference the secret. A secret can be used with a pod in three ways:

- to populate environment variables for containers.
- as files in a volume mounted on one or more of its containers.
- by kubelet when pulling images for the pod.

Volume type secrets write data into the container as a file using the volume mechanism.

imagePullSecrets use service accounts for the automatic injection of the secret into all pods in a namespace.

When a template contains a secret definition, the only way for the template to use the provided secret is to ensure that the secret volume sources are validated and that the specified object reference actually points to an object of type **Secret**. Therefore, a secret needs to be created before any pods that depend on it. The most effective way to ensure this is to have it get injected automatically through the use of a service account.

Secret API objects reside in a namespace. They can only be referenced by pods in that same namespace.

Individual secrets are limited to 1MB in size. This is to discourage the creation of large secrets that would exhaust apiserver and kubelet memory. However, creation of a number of smaller secrets could also exhaust memory.

19.6.1. Secret Data Keys

Secret keys must be in a DNS subdomain.

19.7. EXAMPLES

Example 19.1. YAML Secret That Will Create Four Files

```
apiVersion: v1
kind: Secret
metadata:
  name: test-secret
data:
  username: dmFsdWUtMQ0K ①
  password: dmFsdWUtMQ0KDQo= ②
stringData:
  hostname: myapp.mydomain.com ③
secret.properties: |- ④
  property1=valueA
  property2=valueB
```

- ① File contains decoded values.
- ② File contains decoded values.
- ③ File contains the provided string.
- ④ File contains the provided data.

Example 19.2. YAML of a Pod Populating Files in a Volume with Secret Data

```

apiVersion: v1
kind: Pod
metadata:
  name: secret-example-pod
spec:
  containers:
    - name: secret-test-container
      image: busybox
      command: [ "/bin/sh", "-c", "cat /etc/secret-volume/*" ]
      volumeMounts:
        # name must match the volume name below
        - name: secret-volume
          mountPath: /etc/secret-volume
          readOnly: true
  volumes:
    - name: secret-volume
      secret:
        secretName: test-secret
  restartPolicy: Never

```

Example 19.3. YAML of a Pod Populating Environment Variables with Secret Data

```

apiVersion: v1
kind: Pod
metadata:
  name: secret-example-pod
spec:
  containers:
    - name: secret-test-container
      image: busybox
      command: [ "/bin/sh", "-c", "export" ]
      env:
        - name: TEST_SECRET_USERNAME_ENV_VAR
          valueFrom:
            secretKeyRef:
              name: test-secret
              key: username
  restartPolicy: Never

```

19.8. TROUBLESHOOTING

If a [service certificate generations](#) fails with (service's `service.alpha.openshift.io/serving-cert-generation-error` annotation contains):

```
secret/ssl-key references serviceUID 62ad25ca-d703-11e6-9d6f-0e9c0057b608,
which does not match 77b6dd80-d716-11e6-9d6f-0e9c0057b60
```

The service that generated the certificate no longer exists, or has a different `serviceUID`. You must force certificates regeneration by removing the old secret, and clearing the following annotations on

the service `service.alpha.openshift.io/serving-cert-generation-error`,
`service.alpha.openshift.io/serving-cert-generation-error-num`:

```
$ oc delete secret <secret_name>
$ oc annotate service <service_name> service.alpha.openshift.io/serving-
cert-generation-error-
$ oc annotate service <service_name> service.alpha.openshift.io/serving-
cert-generation-error-num-
```



NOTE

The command removing annotation has a - after the annotation name to be removed.

CHAPTER 20. CONFIGMAPS

20.1. OVERVIEW

Many applications require configuration using some combination of configuration files, command line arguments, and environment variables. These configuration artifacts should be decoupled from image content in order to keep containerized applications portable.

The **ConfigMap** object provides mechanisms to inject containers with configuration data while keeping containers agnostic of OpenShift Container Platform. A **ConfigMap** can be used to store fine-grained information like individual properties or coarse-grained information like entire configuration files or JSON blobs.

The **ConfigMap** API object holds key-value pairs of configuration data that can be consumed in pods or used to store configuration data for system components such as controllers. **ConfigMap** is similar to [secrets](#), but designed to more conveniently support working with strings that do not contain sensitive information.

For example:

Example 20.1. ConfigMap Object Definition

```
kind: ConfigMap
apiVersion: v1
metadata:
  creationTimestamp: 2016-02-18T19:14:38Z
  name: example-config
  namespace: default
data: ①
  example.property.1: hello
  example.property.2: world
  example.property.file: |-
    property.1=value-1
    property.2=value-2
    property.3=value-3
```

① Contains the configuration data.

Configuration data can be consumed in pods in a variety of ways. A **ConfigMap** can be used to:

1. Populate the value of environment variables.
2. Set command-line arguments in a container.
3. Populate configuration files in a volume.

Both users and system components may store configuration data in a **ConfigMap**.

20.2. CREATING CONFIGMAPS

You can use the following command to create a **ConfigMap** easily from directories, specific files, or literal values:

```
$ oc create configmap <configmap_name> [options]
```

The following sections cover the different ways you can create a **ConfigMap**.

20.2.1. Creating from Directories

Consider a directory with some files that already contain the data with which you want to populate a **ConfigMap**:

```
$ ls example-files
game.properties
ui.properties

$ cat example-files/game.properties
enemies=aliens
lives=3
enemies.cheat=true
enemies.cheat.level=noGoodRotten
secret.code.passphrase=UUDDLRLRBABAS
secret.code.allowed=true
secret.code.lives=30

$ cat example-files/ui.properties
color.good=purple
color.bad=yellow
allow.textmode=true
how.nice.to.look=fairlyNice
```

You can use the following command to create a **ConfigMap** holding the content of each file in this directory:

```
$ oc create configmap game-config \
  --from-file=example-files/
```

When the **--from-file** option points to a directory, each file directly in that directory is used to populate a key in the **ConfigMap**, where the name of the key is the file name, and the value of the key is the content of the file.

For example, the above command creates the following **ConfigMap**:

```
$ oc describe configmaps game-config
Name:          game-config
Namespace:     default
Labels:        <none>
Annotations:   <none>

Data

game.properties: 121 bytes
ui.properties:   83 bytes
```

You can see the two keys in the map are created from the file names in the directory specified in the command. Because the content of those keys may be large, the output of `oc describe` only shows the names of the keys and their sizes.

If you want to see the values of the keys, you can `oc get` the object with the `-o` option:

```
$ oc get configmaps game-config -o yaml

apiVersion: v1
data:
  game.properties: |-
    enemies=aliens
    lives=3
    enemies.cheat=true
    enemies.cheat.level=noGoodRotten
    secret.code.passphrase=UUDDLRLRBABAS
    secret.code.allowed=true
    secret.code.lives=30
  ui.properties: |
    color.good=purple
    color.bad=yellow
    allow.textmode=true
    how.nice.to.look=fairlyNice
kind: ConfigMap
metadata:
  creationTimestamp: 2016-02-18T18:34:05Z
  name: game-config
  namespace: default
  resourceVersion: "407"-
  selflink: /api/v1/namespaces/default/configmaps/game-config
  uid: 30944725-d66e-11e5-8cd0-68f728db1985
```

20.2.2. Creating from Files

You can also pass the `--from-file` option with a specific file, and pass it multiple times to the CLI. The following yields equivalent results to the [Creating from Directories](#) example:

1. Create the **ConfigMap** specifying a specific file:

```
$ oc create configmap game-config-2 \
  --from-file=example-files/game.properties \
  --from-file=example-files/ui.properties
```

2. Verify the results:

```
$ oc get configmaps game-config-2 -o yaml

apiVersion: v1
data:
  game.properties: |-
    enemies=aliens
    lives=3
    enemies.cheat=true
    enemies.cheat.level=noGoodRotten
    secret.code.passphrase=UUDDLRLRBABAS
```

```

    secret.code.allowed=true
    secret.code.lives=30
  ui.properties: |
    color.good=purple
    color.bad=yellow
    allow.textmode=true
    how.nice.to.look=fairlyNice
kind: ConfigMap
metadata:
  creationTimestamp: 2016-02-18T18:52:05Z
  name: game-config-2
  namespace: default
  resourceVersion: "516"
  selflink: /api/v1/namespaces/default/configmaps/game-config-2
  uid: b4952dc3-d670-11e5-8cd0-68f728db1985

```

You can also set the key to use for an individual file with the `--from-file` option by passing an expression of `key=value`. For example:

1. Create the **ConfigMap** specifying a key-value pair:

```

$ oc create configmap game-config-3 \
  --from-file=game-special-key=example-files/game.properties

```

2. Verify the results:

```

$ oc get configmaps game-config-3 -o yaml

apiVersion: v1
data:
  game-special-key: |-
    enemies=aliens
    lives=3
    enemies.cheat=true
    enemies.cheat.level=noGoodRotten
    secret.code.passphrase=UDDDLRLRBABAS
    secret.code.allowed=true
    secret.code.lives=30
kind: ConfigMap
metadata:
  creationTimestamp: 2016-02-18T18:54:22Z
  name: game-config-3
  namespace: default
  resourceVersion: "530"
  selflink: /api/v1/namespaces/default/configmaps/game-config-3
  uid: 05f8da22-d671-11e5-8cd0-68f728db1985

```

20.2.3. Creating from Literal Values

You can also supply literal values for a **ConfigMap**. The `--from-literal` option takes a `key=value` syntax that allows literal values to be supplied directly on the command line:

1. Create the **ConfigMap** specifying a literal value:

```

$ oc create configmap special-config \

```

```
--from-literal=special.how=very \
--from-literal=special.type=charm
```

2. Verify the results:

```
$ oc get configmaps special-config -o yaml

apiVersion: v1
data:
  special.how: very
  special.type: charm
kind: ConfigMap
metadata:
  creationTimestamp: 2016-02-18T19:14:38Z
  name: special-config
  namespace: default
  resourceVersion: "651"
  selflink: /api/v1/namespaces/default/configmaps/special-config
  uid: dadce046-d673-11e5-8cd0-68f728db1985
```

20.3. USE CASES: CONSUMING CONFIGMAPS IN PODS

The following sections describe some uses cases when consuming **ConfigMap** objects in pods.

20.3.1. Consuming in Environment Variables

A **ConfigMap** can be used to populate the value of command line arguments. For example, consider the following **ConfigMap**:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config
  namespace: default
data:
  special.how: very
  special.type: charm
```

You can consume the keys of this **ConfigMap** in a pod using **configMapKeyRef** sections:

```
apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  containers:
  - name: test-container
    image: gcr.io/google_containers/busybox
    command: [ "/bin/sh", "-c", "env" ]
    env:
    - name: SPECIAL_LEVEL_KEY
      valueFrom:
        configMapKeyRef:
```



```

        name: special-config
        key: special.how
    - name: SPECIAL_TYPE_KEY
      valueFrom:
        configMapKeyRef:
          name: special-config
          key: special.type
  restartPolicy: Never

```

When this pod is run, its output will include the following lines:

```

SPECIAL_LEVEL_KEY=very
SPECIAL_TYPE_KEY=charm

```

20.3.2. Setting Command-line Arguments

A **ConfigMap** can also be used to set the value of the command or arguments in a container. This is accomplished using the Kubernetes substitution syntax `$(VAR_NAME)`. Consider the following **ConfigMap**:

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config
  namespace: default
data:
  special.how: very
  special.type: charm

```

To inject values into the command line, you must consume the keys you want to use as environment variables, as in the [Consuming in Environment Variables](#) use case. Then you can refer to them in a container's command using the `$(VAR_NAME)` syntax.

```

apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  containers:
    - name: test-container
      image: gcr.io/google_containers/busybox
      command: [ "/bin/sh", "-c", "echo $(SPECIAL_LEVEL_KEY)
$(SPECIAL_TYPE_KEY)" ]
      env:
        - name: SPECIAL_LEVEL_KEY
          valueFrom:
            configMapKeyRef:
              name: special-config
              key: special.how
        - name: SPECIAL_TYPE_KEY
          valueFrom:
            configMapKeyRef:

```

```
        name: special-config
        key: special.type
    restartPolicy: Never
```

When this pod is run, the output from the **test-container** container will be:

```
very charm
```

20.3.3. Consuming in Volumes

A **ConfigMap** can also be consumed in volumes. Returning again to the following example **ConfigMap**:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config
  namespace: default
data:
  special.how: very
  special.type: charm
```

You have a couple different options for consuming this **ConfigMap** in a volume. The most basic way is to populate the volume with files where the key is the file name and the content of the file is the value of the key:

```
apiVersion: v1
kind: Pod
metadata:
  name: dapi-test-pod
spec:
  containers:
    - name: test-container
      image: gcr.io/google_containers/busybox
      command: [ "/bin/sh", "cat", "/etc/config/special.how" ]
      volumeMounts:
        - name: config-volume
          mountPath: /etc/config
  volumes:
    - name: config-volume
      configMap:
        name: special-config
      restartPolicy: Never
```

When this pod is run, the output will be:

```
very
```

You can also control the paths within the volume where **ConfigMap** keys are projected:

```
apiVersion: v1
kind: Pod
metadata:
```

```

  name: dapi-test-pod
spec:
  containers:
  - name: test-container
    image: gcr.io/google_containers/busybox
    command: [ "/bin/sh", "cat", "/etc/config/path/to/special-key" ]
    volumeMounts:
    - name: config-volume
      mountPath: /etc/config
  volumes:
  - name: config-volume
    configMap:
      name: special-config
      items:
      - key: special.how
        path: path/to/special-key
  restartPolicy: Never

```

When this pod is run, the output will be:

```
very
```

20.4. EXAMPLE: CONFIGURING REDIS

For a real-world example, you can configure Redis using a **ConfigMap**. To inject Redis with the recommended configuration for using Redis as a cache, the Redis configuration file should contain the following:

```

maxmemory 2mb
maxmemory-policy allkeys-lru

```

If your configuration file is located at **example-files/redis/redis-config**, create a **ConfigMap** with it:

1. Create the **ConfigMap** specifying the configuration file:

```

$ oc create configmap example-redis-config \
  --from-file=example-files/redis/redis-config

```

2. Verify the results:

```

$ oc get configmap example-redis-config -o yaml

apiVersion: v1
data:
  redis-config: |
    maxmemory 2mb
    maxmemory-policy allkeys-lru
kind: ConfigMap
metadata:
  creationTimestamp: 2016-04-06T05:53:07Z
  name: example-redis-config
  namespace: default
  resourceVersion: "2985"
  selflink: /api/v1/namespaces/default/configmaps/example-redis-

```

```
config
  uid: d65739c1-fbbb-11e5-8a72-68f728db1985
```

Now, create a pod that uses this **ConfigMap**:

1. Create a pod definition like the following and save it to a file, for example *redis-pod.yaml*:

```
apiVersion: v1
kind: Pod
metadata:
  name: redis
spec:
  containers:
  - name: redis
    image: kubernetes/redis:v1
    env:
    - name: MASTER
      value: "true"
    ports:
    - containerPort: 6379
    resources:
      limits:
        cpu: "0.1"
    volumeMounts:
    - mountPath: /redis-master-data
      name: data
    - mountPath: /redis-master
      name: config
  volumes:
  - name: data
    emptyDir: {}
  - name: config
    configMap:
      name: example-redis-config
      items:
      - key: redis-config
        path: redis.conf
```

2. Create the pod:

```
$ oc create -f redis-pod.yaml
```

The newly-created pod has a **ConfigMap** volume that places the **redis-config** key of the **example-redis-config ConfigMap** into a file called *redis.conf*. This volume is mounted into the */redis-master* directory in the Redis container, placing our configuration file at */redis-master/redis.conf*, which is where the image looks for the Redis configuration file for the master.

If you **oc exec** into this pod and run the **redis-cli** tool, you can check that the configuration was applied correctly:

```
$ oc exec -it redis redis-cli
127.0.0.1:6379> CONFIG GET maxmemory
1) "maxmemory"
2) "2097152"
```

```
127.0.0.1:6379> CONFIG GET maxmemory-policy
1) "maxmemory-policy"
2) "allkeys-lru"
```

20.5. RESTRICTIONS

A **ConfigMap** must be created before they are consumed in pods. Controllers can be written to tolerate missing configuration data; consult individual components configured via **ConfigMap** on a case-by-case basis.

ConfigMap objects reside in a project. They can only be referenced by pods in the same project.

The Kubelet only supports use of a **ConfigMap** for pods it gets from the API server. This includes any pods created using the CLI, or indirectly from a replication controller. It does not include pods created using the OpenShift Container Platform node's `--manifest-url` flag, its `--config` flag, or its REST API (these are not common ways to create pods).

CHAPTER 21. USING DAEMONSETS

21.1. OVERVIEW

A daemonset can be used to run replicas of a pod on specific or all nodes in an OpenShift Container Platform cluster.

Use daemonsets to create shared storage, run a logging pod on every node in your cluster, or deploy a monitoring agent on every node.

For more information on daemonsets, see the [Kubernetes documentation](#).

IMPORTANT

Daemonset scheduling is incompatible with project's default node selector. If you fail to disable it, the daemonset gets restricted by merging with the default node selector. This results in frequent pod recreates on the nodes that got unselected by the merged node selector, which in turn puts unwanted load on the cluster.

Therefore,

- Before you start using daemonsets, disable the default project-wide [node selector](#) in your namespace, by setting the namespace annotation `openshift.io/node-selector` to an empty string:

```
# oc patch namespace myproject -p \
  '{"metadata": {"annotations": {"openshift.io/node-selector": ""}}}'
```

- If you are creating a new project, overwrite the default node selector using `oc adm new-project --node-selector=""`.

21.2. CREATING DAEMONSETS

IMPORTANT

Before creating daemonsets, ensure you have been [given the required role by your OpenShift Container Platform administrator](#).

When creating daemonsets, the `nodeSelector` field is used to indicate the nodes on which the daemonset should deploy replicas.

1. Define the daemonset yaml file:

```
apiVersion: extensions/v1beta1
kind: DaemonSet
metadata:
  name: hello-daemonset
spec:
  selector:
    matchLabels:
      name: hello-daemonset 1
  template:
```

```

metadata:
  labels:
    name: hello-daemonset ❷
spec:
  nodeSelector: ❸
    type: infra
  containers:
  - image: openshift/hello-openshift
    imagePullPolicy: Always
    name: registry
    ports:
    - containerPort: 80
      protocol: TCP
    resources: {}
    terminationMessagePath: /dev/termination-log
  serviceAccount: default
  terminationGracePeriodSeconds: 10

```

- ❶ The label selector that determines which pods belong to the daemonset.
- ❷ The pod template's label selector. Must match the label selector above.
- ❸ The node selector that determines on which nodes pod replicas should be deployed.

2. Create the daemonset object:

```
oc create -f daemonset.yaml
```

3. To verify that the pods were created, and that each node has a pod replica:

a. Find the daemonset pods:

```

$ oc get pods
hello-daemonset-cx6md    1/1      Running    0          2m
hello-daemonset-e3md9    1/1      Running    0          2m

```

b. View the pods to verify the pod has been placed onto the node:

```

$ oc describe pod/hello-daemonset-cx6md | grep Node
Node:          openshift-node01.hostname.com/10.14.20.134
$ oc describe pod/hello-daemonset-e3md9 | grep Node
Node:          openshift-node02.hostname.com/10.14.20.137

```



IMPORTANT

- If you update a DaemonSet's pod template, the existing pod replicas are not affected.
- If you delete a DaemonSet and then create a new DaemonSet with a different template but the same label selector, it recognizes any existing pod replicas as having matching labels and thus does not update them or create new replicas despite a mismatch in the pod template.
- If you change node labels, the DaemonSet adds pods to nodes that match the new labels and deletes pods from nodes that do not match the new labels.

To update a DaemonSet, force new pod replicas to be created by deleting the old replicas or nodes.

CHAPTER 22. POD AUTOSCALING

22.1. OVERVIEW

A horizontal pod autoscaler, defined by a `HorizontalPodAutoscaler` object, specifies how the system should automatically increase or decrease the scale of a replication controller or deployment configuration, based on metrics collected from the pods that belong to that replication controller or deployment configuration.



NOTE

Horizontal pod autoscaling is supported starting in OpenShift Enterprise 3.1.1.

22.2. REQUIREMENTS FOR USING HORIZONTAL POD AUTOSCALERS

In order to use horizontal pod autoscalers, your cluster administrator must have [properly configured cluster metrics](#).

22.3. SUPPORTED METRICS

The following metrics are supported by horizontal pod autoscalers:

Table 22.1. Metrics

Metric	Description
CPU Utilization	Percentage of the requested CPU

22.4. AUTOSCALING

You can create a horizontal pod autoscaler with the `oc autoscale` command and specify the minimum and maximum number of pods you want to run, as well as the CPU utilization your pods should target.

After a horizontal pod autoscaler is created, it begins attempting to query Heapster for metrics on the pods. It may take one to two minutes before Heapster obtains the initial metrics.

After metrics are available in Heapster, the horizontal pod autoscaler computes the ratio of the current metric utilization with the desired metric utilization, and scales up or down accordingly. The scaling will occur at a regular interval, but it may take one to two minutes before metrics make their way into Heapster.

For replication controllers, this scaling corresponds directly to the replicas of the replication controller. For deployment configurations, scaling corresponds directly to the replica count of the deployment configuration. Note that autoscaling applies only to the latest deployment in the **Complete** phase.

OpenShift Container Platform automatically accounts for resources and prevents unnecessary autoscaling during resource spikes, such as during start up. Pods in the **unready** state have **0 CPU** usage when scaling up and the autoscaler ignores the pods when scaling down. Pods without known

metrics have **0% CPU** usage when scaling up and **100% CPU** when scaling down. This allows for more stability during the HPA decision. To use this feature, you must configure [readiness checks](#) to determine if a new pod is ready for use.

22.5. CREATING A HORIZONTAL POD AUTOSCALER

Use the `oc autoscale` command and specify at least the maximum number of pods you want to run at any given time. You can optionally specify the minimum number of pods and the average CPU utilization your pods should target, otherwise those are given default values from the OpenShift Container Platform server.

For example:

```
$ oc autoscale dc/frontend --min 1 --max 10 --cpu-percent=80
deploymentconfig "frontend" autoscaled
```

The above example creates a horizontal pod autoscaler with the following definition:

Example 22.1. Horizontal Pod Autoscaler Object Definition

```
apiVersion: extensions/v1beta1
kind: HorizontalPodAutoscaler
metadata:
  name: frontend ❶
spec:
  scaleRef:
    kind: DeploymentConfig ❷
    name: frontend ❸
    apiVersion: v1 ❹
    subresource: scale
  minReplicas: 1 ❺
  maxReplicas: 10 ❻
  cpuUtilization:
    targetPercentage: 80 ❼
```

- ❶ The name of this horizontal pod autoscaler object
- ❷ The kind of object to scale
- ❸ The name of the object to scale
- ❹ The API version of the object to scale
- ❺ The minimum number of replicas to which to scale down
- ❻ The maximum number of replicas to which to scale up
- ❼ The percentage of the requested CPU that each pod should ideally be using

22.6. VIEWING A HORIZONTAL POD AUTOSCALER

To view the status of a horizontal pod autoscaler:

```
$ oc get hpa/frontend
```

NAME	REFERENCE	TARGET
CURRENT	MINPODS	MAXPODS
AGE		
frontend	DeploymentConfig/default/frontend/scale	80%
1	10	8d

```
$ oc describe hpa/frontend
```

```
Name: frontend
Namespace: default
Labels: <none>
CreationTimestamp: Mon, 26 Oct 2015 21:13:47 -0400
Reference: DeploymentConfig/default/frontend/scale
Target CPU utilization: 80%
Current CPU utilization: 79%
Min pods: 1
Max pods: 10
```

CHAPTER 23. MANAGING VOLUMES

23.1. OVERVIEW

Containers are not persistent by default; on restart, their contents are cleared. Volumes are mounted file systems available to [pods](#) and their containers which may be backed by a number of host-local or network attached storage endpoints.

To ensure that the file system on the volume contains no errors and, if errors are present, to repair them when possible, OpenShift Container Platform invokes the `fsck` utility prior to the `mount` utility. This occurs when either adding a volume or updating an existing volume.

The simplest volume type is `emptyDir`, which is a temporary directory on a single machine. Administrators may also allow you to request a [persistent volume](#) that is automatically attached to your pods.



NOTE

`emptyDir` volume storage may be restricted by a quota based on the pod's `FSGroup`, if the `FSGroup` parameter is enabled by your cluster administrator.

You can use the CLI command `oc volume` to [add](#), [update](#), or [remove](#) volumes and volume mounts for any object that has a pod template like [replication controllers](#) or [deployment configurations](#). You can also [list](#) volumes in pods or any object that has a pod template.

23.2. GENERAL CLI USAGE

The `oc volume` command uses the following general syntax:

```
$ oc volume <object_selection> <operation> <mandatory_parameters>
<optional_parameters>
```

This topic uses the form `<object_type>/<name>` for `<object_selection>` in later examples. However, you can choose one of the following options:

Table 23.1. Object Selection

Syntax	Description	Example
<code><object_type> <name></code>	Selects <code><name></code> of type <code><object_type></code> .	<code>deploymentConfig registry</code>
<code><object_type>/<name></code>	Selects <code><name></code> of type <code><object_type></code> .	<code>deploymentConfig/registry</code>
<code><object_type> --selector=<object_label_selector></code>	Selects resources of type <code><object_type></code> that matched the given label selector.	<code>deploymentConfig --selector="name=registry"</code>
<code><object_type> --all</code>	Selects all resources of type <code><object_type></code> .	<code>deploymentConfig --all</code>

Syntax	Description	Example
-f or --filename=<file_name>	File name, directory, or URL to file to use to edit the resource.	-f registry-deployment-config.json

The *<operation>* can be one of **--add**, **--remove**, or **--list**.

Any *<mandatory_parameters>* or *<optional_parameters>* are specific to the selected operation and are discussed in later sections.

23.3. ADDING VOLUMES

To add a volume, a volume mount, or both to pod templates:

```
$ oc volume <object_type>/<name> --add [options]
```

Table 23.2. Supported Options for Adding Volumes

Option	Description	Default
--name	Name of the volume.	Automatically generated, if not specified.
-t, --type	Name of the volume source. Supported values: emptyDir , hostPath , secret , configmap , or persistentVolumeClaim .	emptyDir
-c, --containers	Select containers by name. It can also take wildcard '*' that matches any character.	'*'
-m, --mount-path	Mount path inside the selected containers.	
--path	Host path. Mandatory parameter for --type=hostPath .	
--secret-name	Name of the secret. Mandatory parameter for --type=secret .	
--configmap-name	Name of the configmap. Mandatory parameter for --type=configmap .	

Option	Description	Default
--claim-name	Name of the persistent volume claim. Mandatory parameter for --type=persistentVolumeClaim .	
--source	Details of volume source as a JSON string. Recommended if the desired volume source is not supported by --type . See available volume sources	
-o, --output	Display the modified objects instead of updating them on the server. Supported values: json , yaml .	
--output-version	Output the modified objects with the given version.	api-version

Examples

Add a new volume source **emptyDir** to deployment configuration **registry**:

```
$ oc volume dc/registry --add
```

Add volume **v1** with secret **\$secret** for replication controller **r1** and mount inside the containers at **/data**:

```
$ oc volume rc/r1 --add --name=v1 --type=secret --secret-name='$secret' --mount-path=/data
```

Add existing persistent volume **v1** with claim name **pvc1** to deployment configuration **dc.json** on disk, mount the volume on container **c1** at **/data**, and update the deployment configuration on the server:

```
$ oc volume -f dc.json --add --name=v1 --type=persistentVolumeClaim \
  --claim-name=pvc1 --mount-path=/data --containers=c1
```

Add volume **v1** based on Git repository **https://github.com/namespace1/project1** with revision **5125c45f9f563** for all replication controllers:

```
$ oc volume rc --all --add --name=v1 \
  --source='{ "gitRepo": {
    "repository": "https://github.com/namespace1/project1",
    "revision": "5125c45f9f563"
  } }'
```

23.4. UPDATING VOLUMES

Updating existing volumes or volume mounts is the same as [adding volumes](#), but with the `--overwrite` option:

```
$ oc volume <object_type>/<name> --add --overwrite [options]
```

Examples

Replace existing volume **v1** for replication controller **r1** with existing persistent volume claim **pvc1**:

```
$ oc volume rc/r1 --add --overwrite --name=v1 --type=persistentVolumeClaim
--claim-name=pvc1
```

Change deployment configuration **d1** mount point to **/opt** for volume **v1**:

```
$ oc volume dc/d1 --add --overwrite --name=v1 --mount-path=/opt
```

23.5. REMOVING VOLUMES

To remove a volume or volume mount from pod templates:

```
$ oc volume <object_type>/<name> --remove [options]
```

Table 23.3. Supported Options for Removing Volumes

Option	Description	Default
<code>--name</code>	Name of the volume.	
<code>-c, --containers</code>	Select containers by name. It can also take wildcard <code>'*'</code> that matches any character.	<code>'*'</code>
<code>--confirm</code>	Indicate that you want to remove multiple volumes at once.	
<code>-o, --output</code>	Display the modified objects instead of updating them on the server. Supported values: json , yaml .	
<code>--output-version</code>	Output the modified objects with the given version.	api-version

Examples

Remove a volume **v1** from deployment configuration **d1**:

```
$ oc volume dc/d1 --remove --name=v1
```

Unmount volume **v1** from container **c1** for deployment configuration **d1** and remove the volume **v1** if it is not referenced by any containers on **d1**:

```
$ oc volume dc/d1 --remove --name=v1 --containers=c1
```

Remove all volumes for replication controller **r1**:

```
$ oc volume rc/r1 --remove --confirm
```

23.6. LISTING VOLUMES

To list volumes or volume mounts for pods or pod templates:

```
$ oc volume <object_type>/<name> --list [options]
```

List volume supported options:

Option	Description	Default
--name	Name of the volume.	
-c, --containers	Select containers by name. It can also take wildcard '*' that matches any character.	'*'

Examples

List all volumes for pod **p1**:

```
$ oc volume pod/p1 --list
```

List volume **v1** defined on all deployment configurations:

```
$ oc volume dc --all --name=v1
```

23.7. SPECIFYING A SUB-PATH

Use the `volumeMounts.subPath` property to specify a `subPath` inside a volume instead of the volume's root. `subPath` allows you to share one volume for multiple uses in a single pod.

To view the list of files in the volume, run the `oc rsh` command:

```
$ oc rsh <pod>
sh-4.2$ ls /path/to/volume/subpath/mount
example_file1 example_file2 example_file3
```

Specify the `subPath`:

Example subPath Usage

```
apiVersion: v1
kind: Pod
metadata:
```



```
name: my-site
spec:
  containers:
    - name: mysql
      image: mysql
      volumeMounts:
        - mountPath: /var/lib/mysql
          name: site-data
          subPath: mysql 1
    - name: php
      image: php
      volumeMounts:
        - mountPath: /var/www/html
          name: site-data
          subPath: html 2
  volumes:
    - name: site-data
      persistentVolumeClaim:
        claimName: my-site-data
```


1 Databases are stored in the **mysql** folder.

2 HTML content is stored in the **html** folder.

CHAPTER 24. USING PERSISTENT VOLUMES

24.1. OVERVIEW

A **PersistentVolume** object is a storage resource in an OpenShift Container Platform cluster. Storage is provisioned by your cluster administrator by creating **PersistentVolume** objects from sources such as GCE Persistent Disk, AWS Elastic Block Store (EBS), and NFS mounts.



NOTE

The [Installation and Configuration Guide](#) provides instructions for cluster administrators on provisioning an OpenShift Container Platform cluster with persistent storage using [NFS](#), [GlusterFS](#), [Ceph RBD](#), [OpenStack Cinder](#), [AWS EBS](#), [GCE Persistent Disk](#), [iSCSI](#), and [Fibre Channel](#).

Storage can be made available to you by laying claims to the resource. You can make a request for storage resources using a **PersistentVolumeClaim** object; the claim is paired with a volume that generally matches your request.

24.2. REQUESTING STORAGE

You can request storage by creating **PersistentVolumeClaim** objects in your projects:

Persistent Volume Claim Object Definition

```
apiVersion: "v1"
kind: "PersistentVolumeClaim"
metadata:
  name: "claim1"
spec:
  accessModes:
    - "ReadWriteOnce"
  resources:
    requests:
      storage: "1Gi"
  volumeName: "pv0001"
```

24.3. VOLUME AND CLAIM BINDING

A **PersistentVolume** is a specific resource. A **PersistentVolumeClaim** is a request for a resource with specific attributes, such as storage size. In between the two is a process that matches a claim to an available volume and binds them together. This allows the claim to be used as a volume in a pod. OpenShift Container Platform finds the volume backing the claim and mounts it into the pod.

You can tell whether a claim or volume is bound by querying using the CLI:

```
$ oc get pvc
NAME          LABELS          STATUS    VOLUME
claim1        map[]          Bound     pv0001

$ oc get pv
NAME          LABELS          CAPACITY    ACCESSMODES
```

STATUS	CLAIM		
pv0001	map[]	5368709120	RWO
Bound	yournamespace / claim1		

24.4. CLAIMS AS VOLUMES IN PODS

A `PersistentVolumeClaim` is used by a pod as a volume. OpenShift Container Platform finds the claim with the given name in the same namespace as the pod, then uses the claim to find the corresponding volume to mount.

Pod Definition with a Claim

```
apiVersion: "v1"
kind: "Pod"
metadata:
  name: "mypod"
  labels:
    name: "frontendhttp"
spec:
  containers:
  -
    name: "myfrontend"
    image: openshift/hello-openshift
    ports:
    -
      containerPort: 80
      name: "http-server"
    volumeMounts:
    -
      mountPath: "/var/www/html"
      name: "pvol"
  volumes:
  -
    name: "pvol"
    persistentVolumeClaim:
      claimName: "claim1"
```

24.5. VOLUME AND CLAIM PRE-BINDING

If you know exactly what `PersistentVolume` you want your `PersistentVolumeClaim` to bind to, you can specify the PV in your PVC using the `volumeName` field. This method skips the normal matching and binding process. The PVC will only be able to bind to a PV that has the same name specified in `volumeName`. If such a PV with that name exists and is **Available**, the PV and PVC will be bound regardless of whether the PV satisfies the PVC's label selector, access modes, and resource requests.

Example 24.1. Persistent Volume Claim Object Definition with `volumeName`

```
apiVersion: "v1"
kind: "PersistentVolumeClaim"
metadata:
  name: "claim1"
spec:
  accessModes:
```

```

- "ReadWriteOnce"
resources:
  requests:
    storage: "1Gi"
  volumeName: "pv0001"

```



IMPORTANT

The ability to set `claimRefs` is a temporary workaround for the described use cases. A long-term solution for limiting who can claim a volume is in development.



NOTE

The cluster administrator should first consider configuring [selector-label volume binding](#) before resorting to setting `claimRefs` on behalf of users.

You may also want your cluster administrator to "reserve" the volume for only your claim so that nobody else's claim can bind to it before yours does. In this case, the administrator can specify the PVC in the PV using the `claimRef` field. The PV will only be able to bind to a PVC that has the same name and namespace specified in `claimRef`. The PVC's access modes and resource requests must still be satisfied in order for the PV and PVC to be bound, though the label selector is ignored.

Persistent Volume Object Definition with `claimRef`

```

apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv0001
spec:
  capacity:
    storage: 1Gi
  accessModes:
    - ReadWriteOnce
  nfs:
    path: /tmp
    server: 172.17.0.2
  persistentVolumeReclaimPolicy: Recycle
  claimRef:
    name: claim1
    namespace: default

```

Specifying a `volumeName` in your PVC does not prevent a different PVC from binding to the specified PV before yours does. Your claim will remain **Pending** until the PV is **Available**.

Specifying a `claimRef` in a PV does not prevent the specified PVC from being bound to a different PV. The PVC is free to choose another PV to bind to according to the normal binding process. Therefore, to avoid these scenarios and ensure your claim gets bound to the volume you want, you must ensure that both `volumeName` and `claimRef` are specified.

You can tell that your setting of `volumeName` and/or `claimRef` influenced the matching and binding process by inspecting a **Bound** PV and PVC pair for the `pv.kubernetes.io/bound-by-controller` annotation. The PVs and PVCs where you set the `volumeName` and/or `claimRef`


yourself will have no such annotation, but ordinary PVs and PVCs will have it set to "yes".

When a PV has its `claimRef` set to some PVC name and namespace, and is reclaimed according to a **Retain** or **Recycle** reclaim policy, its `claimRef` will remain set to the same PVC name and namespace even if the PVC or the whole namespace no longer exists.

CHAPTER 25. EXECUTING REMOTE COMMANDS

25.1. OVERVIEW

You can use the CLI to execute remote commands in a container. This allows you to run general Linux commands for routine operations in the container.



IMPORTANT

For security purposes, the `oc exec` command does not work when accessing privileged containers except when the command is executed by a `cluster-admin` user. See the [CLI operations topic](#) for more information.

25.2. BASIC USAGE

Support for remote container command execution is built into [the CLI](#):

```
$ oc exec <pod> [-c <container>] <command> [<arg_1> ... <arg_n>]
```

For example:

```
$ oc exec mypod date
Thu Apr 9 02:21:53 UTC 2015
```

25.3. PROTOCOL

Clients initiate the execution of a remote command in a container by issuing a request to the Kubernetes API server:

```
/proxy/minions/<node_name>/exec/<namespace>/<pod>/<container>?command=
<command>
```

In the above URL:

- `<node_name>` is the FQDN of the node.
- `<namespace>` is the namespace of the target pod.
- `<pod>` is the name of the target pod.
- `<container>` is the name of the target container.
- `<command>` is the desired command to be executed.

For example:

```
/proxy/minions/node123.openshift.com/exec/myns/mypod/mycontainer?
command=date
```

Additionally, the client can add parameters to the request to indicate if:

- the client should send input to the remote container's command (stdin).

- the client's terminal is a TTY.
- the remote container's command should send output from `stdout` to the client.
- the remote container's command should send output from `stderr` to the client.

After sending an `exec` request to the API server, the client upgrades the connection to one that supports multiplexed streams; the current implementation uses **SPDY**.

The client creates one stream each for `stdin`, `stdout`, and `stderr`. To distinguish among the streams, the client sets the `streamType` header on the stream to one of `stdin`, `stdout`, or `stderr`.

The client closes all streams, the upgraded connection, and the underlying connection when it is finished with the remote command execution request.



NOTE

Administrators can see the [Architecture](#) guide for more information.

CHAPTER 26. COPYING FILES TO OR FROM A CONTAINER

26.1. OVERVIEW

You can use the CLI to copy local files to or from a remote directory in a container. This is a useful tool for copying database archives to and from your pods for backup and restore purposes. It can also be used to copy source code changes into a running pod for development debugging, when the running pod supports hot reload of source files.

26.2. BASIC USAGE

Support for copying local files to or from a container is built into [the CLI](#):

```
$ oc rsync <source> <destination> [-c <container>]
```

For example, to copy a local directory to a pod directory:

```
$ oc rsync /home/user/source devpod1234:/src
```

Or to copy a pod directory to a local directory:

```
$ oc rsync devpod1234:/src /home/user/source
```

26.3. BACKING UP AND RESTORING DATABASES

Use `oc rsync` to copy database archives from an existing database container to a new database container's persistent volume directory.



NOTE

MySQL is used in the example below. Replace `mysql` | `MYSQL` with `pgsql` | `PGSQL` or `mongodb` | `MONGODB` and refer to [the migration guide](#) to find the exact commands for each of our supported database images. The example assumes an existing database container.

1. Back up the existing database from a running database pod:

```
$ oc rsh <existing db container>
# mkdir /var/lib/mysql/data/db_archive_dir
# mysqldump --skip-lock-tables -h ${MYSQL_SERVICE_HOST} -P
${MYSQL_SERVICE_PORT:-3306} \
-u ${MYSQL_USER} --password="${MYSQL_PASSWORD}" --all-databases >
/var/lib/mysql/data/db_archive_dir/all.sql
# exit
```

2. Remote sync the archive file to your local machine:

```
$ oc rsync <existing db container with db
archive>:/var/lib/mysql/data/db_archive_dir /tmp/.
```


3. Start a second MySQL pod into which to load the database archive file created above. The MySQL pod must have a unique **DATABASE_SERVICE_NAME**.

```
$ oc new-app mysql-persistent \
  -p MYSQL_USER=<archived mysql username> \
  -p MYSQL_PASSWORD=<archived mysql password> \
  -p MYSQL_DATABASE=<archived database name> \
  -p DATABASE_SERVICE_NAME='mysql2' ❶
$ oc rsync /tmp/db_archive_dir new_dbpod1234:/var/lib/mysql/data
$ oc rsh new_dbpod1234
```

❶ **mysql** is the default. In this example, **mysql2** is created.

4. Use the appropriate commands to restore the database in the new database container from the copied database archive directory:

MySQL

```
$ cd /var/lib/mysql/data/db_archive_dir
$ mysql -u root
$ source all.sql
$ GRANT ALL PRIVILEGES ON <dbname>.* TO '<your
username>'@'localhost'; FLUSH PRIVILEGES;
$ cd ../; rm -rf /var/lib/mysql/data/db_backup_dir
```

You now have two MySQL database pods running in your project with the archived database.

26.4. REQUIREMENTS

The **oc rsync** command uses the local **rsync** command if present on the client's machine. This requires that the remote container also have the **rsync** command.

If **rsync** is not found locally or in the remote container, then a tar archive will be created locally and sent to the container where **tar** will be used to extract the files. If **tar** is not available in the remote container, then the copy will fail.

The **tar** copy method does not provide the same functionality as **rsync**. For example, **rsync** creates the destination directory if it does not exist and will only send files that are different between the source and the destination.



NOTE

In Windows, the **cwRsync** client should be installed and added to the PATH for use with the **oc rsync** command.

26.5. SPECIFYING THE COPY SOURCE

The source argument of the **oc rsync** command must point to either a local directory or a pod directory. Individual files are not currently supported.

When specifying a pod directory the directory name must be prefixed with the pod name:

```
<pod name>:<dir>
```

Just as with standard `rsync`, if the directory name ends in a path separator (`/`), only the contents of the directory are copied to the destination. Otherwise, the directory itself is copied to the destination with all its contents.

26.6. SPECIFYING THE COPY DESTINATION

The destination argument of the `oc rsync` command must point to a directory. If the directory does not exist, but `rsync` is used for copy, the directory is created for you.

26.7. DELETING FILES AT THE DESTINATION

The `--delete` flag may be used to delete any files in the remote directory that are not in the local directory.

26.8. CONTINUOUS SYNCING ON FILE CHANGE

Using the `--watch` option causes the command to monitor the source path for any file system changes, and synchronizes changes when they occur. With this argument, the command runs forever.

Synchronization occurs after short quiet periods to ensure a rapidly changing file system does not result in continuous synchronization calls.

When using the `--watch` option, the behavior is effectively the same as manually invoking `oc rsync` repeatedly, including any arguments normally passed to `oc rsync`. Therefore, you can control the behavior via the same flags used with manual invocations of `oc rsync`, such as `--delete`.

26.9. ADVANCED RSYNC FEATURES

The `oc rsync` command exposes fewer command line options than standard `rsync`. In the case that you wish to use a standard `rsync` command line option which is not available in `oc rsync` (for example the `--exclude-from=FILE` option), it may be possible to use standard `rsync`'s `--rsh (-e)` option or `RSYNC_RSH` environment variable as a workaround, as follows:

```
$ rsync --rsh='oc rsh' --exclude-from=FILE SRC POD:DEST
```

or:

```
$ export RSYNC_RSH='oc rsh'
$ rsync --exclude-from=FILE SRC POD:DEST
```

Both of the above examples configure standard `rsync` to use `oc rsh` as its remote shell program to enable it to connect to the remote pod, and are an alternative to running `oc rsync`.

CHAPTER 27. PORT FORWARDING

27.1. OVERVIEW

OpenShift Container Platform takes advantage of a feature built-in to [Kubernetes](#) to support port forwarding to pods. See [Architecture](#) for more information.

You can use the CLI to forward one or more local ports to a pod. This allows you to listen on a given or random port locally, and have data forwarded to and from given ports in the pod.

27.2. BASIC USAGE

Support for port forwarding is built into [the CLI](#):

```
$ oc port-forward <pod> [<local_port>:]<remote_port> [... [<local_port_n>:]<remote_port_n>]
```

The CLI listens on each local port specified by the user, forwarding via the [protocol](#) described below.

Ports may be specified using the following formats:

5000	The client listens on port 5000 locally and forwards to 5000 in the pod.
6000:5000	The client listens on port 6000 locally and forwards to 5000 in the pod.
:5000 or 0:5000	The client selects a free local port and forwards to 5000 in the pod.

For example, to listen on ports **5000** and **6000** locally and forward data to and from ports **5000** and **6000** in the pod, run:

```
$ oc port-forward <pod> 5000 6000
```

To listen on port **8888** locally and forward to **5000** in the pod, run:

```
$ oc port-forward <pod> 8888:5000
```

To listen on a free port locally and forward to **5000** in the pod, run:

```
$ oc port-forward <pod> :5000
```

Or, alternatively:

```
$ oc port-forward <pod> 0:5000
```

27.3. PROTOCOL

Clients initiate port forwarding to a pod by issuing a request to the Kubernetes API server:

■

```
/proxy/minions/<node_name>/portForward/<namespace>/<pod>
```

In the above URL:

- **<node_name>** is the FQDN of the node.
- **<namespace>** is the namespace of the target pod.
- **<pod>** is the name of the target pod.


For example:

```
/proxy/minions/node123.openshift.com/portForward/myns/mypod
```

After sending a port forward request to the API server, the client upgrades the connection to one that supports multiplexed streams; the current implementation uses [SPDY](#).

The client creates a stream with the **port** header containing the target port in the pod. All data written to the stream is delivered via the Kubelet to the target pod and port. Similarly, all data sent from the pod for that forwarded connection is delivered back to the same stream in the client.

The client closes all streams, the upgraded connection, and the underlying connection when it is finished with the port forwarding request.



NOTE

Administrators can see the [Architecture](#) guide for more information.

CHAPTER 28. SHARED MEMORY

28.1. OVERVIEW

There are two types of shared memory objects in Linux: System V and POSIX. The containers in a pod share the IPC namespace of the pod infrastructure container and so are able to share the System V shared memory objects. This document describes how they can also share POSIX shared memory objects.

28.2. POSIX SHARED MEMORY

POSIX shared memory requires that a tmpfs be mounted at `/dev/shm`. The containers in a pod do not share their mount namespaces so we use volumes to provide the same `/dev/shm` into each container in a pod. The following example shows how to set up POSIX shared memory between two containers.

shared-memory.yaml

```
---
apiVersion: v1
id: hello-openshift
kind: Pod
metadata:
  name: hello-openshift
  labels:
    name: hello-openshift
spec:
  volumes:
    - name: dshm
      emptyDir:
        medium: Memory
  containers:
    - image: kubernetes/pause
      name: hello-container1
      ports:
        - containerPort: 8080
          hostPort: 6061
      volumeMounts:
        - mountPath: /dev/shm
          name: dshm
    - image: kubernetes/pause
      name: hello-container2
      ports:
        - containerPort: 8081
          hostPort: 6062
      volumeMounts:
        - mountPath: /dev/shm
          name: dshm
```

❶ specifies the tmpfs volume **dshm**.

❷ enables POSIX shared memory for **hello-container1** via **dshm**.

❸ enables POSIX shared memory for **hello-container2** via **dshm**.

Create the pod using the *shared-memory.yaml* file:

```
$ oc create -f shared-memory.yaml
```

CHAPTER 29. APPLICATION HEALTH

29.1. OVERVIEW

In software systems, components can become unhealthy due to transient issues (such as temporary connectivity loss), configuration errors, or problems with external dependencies. OpenShift Container Platform applications have a number of options to detect and handle unhealthy containers.

29.2. CONTAINER HEALTH CHECKS USING PROBES

A probe is a Kubernetes action that periodically performs diagnostics on a running container. Currently, two types of probes exist, each serving a different purpose:

Liveness Probe	A liveness probe checks if the container in which it is configured is still running. If the liveness probe fails, the kubelet kills the container, which will be subjected to its restart policy. Set a liveness check by configuring the <code>template.spec.containers.livenessprobe</code> stanza of a pod configuration.
Readiness Probe	A readiness probe determines if a container is ready to service requests. If the readiness probe fails a container, the endpoints controller ensures the container has its IP address removed from the endpoints of all services. A readiness probe can be used to signal to the endpoints controller that even though a container is running, it should not receive any traffic from a proxy. Set a readiness check by configuring the <code>template.spec.containers.readinessprobe</code> stanza of a pod configuration.

The exact timing of a probe is controlled by two fields, both expressed in units of seconds:

Field	Description
<code>initialDelaySeconds</code>	How long to wait after the container starts to begin the probe.
<code>timeoutSeconds</code>	How long to wait for the probe to finish (default: 1). If this time is exceeded, OpenShift Container Platform considers the probe to have failed.

Both probes can be configured in three ways:

HTTP Checks

The kubelet uses a web hook to determine the healthiness of the container. The check is deemed successful if the HTTP response code is between 200 and 399. The following is an example of a readiness check using the HTTP checks method:

Example 29.1. Readiness HTTP check

```
...
readinessProbe:
  httpGet:
```

```
    path: /healthz
    port: 8080
    initialDelaySeconds: 15
    timeoutSeconds: 1
    ...
```

A HTTP check is ideal for applications that return HTTP status codes when completely initialized.

Container Execution Checks

The kubelet executes a command inside the container. Exiting the check with status 0 is considered a success. The following is an example of a liveness check using the container execution method:

Example 29.2. Liveness Container Execution Check

```
    ...
    livenessProbe:
      exec:
        command:
        - cat
        - /tmp/health
      initialDelaySeconds: 15
      timeoutSeconds: 1
    ...
```

TCP Socket Checks

The kubelet attempts to open a socket to the container. The container is only considered healthy if the check can establish a connection. The following is an example of a liveness check using the TCP socket check method:

Example 29.3. Liveness TCP Socket Check

```
    ...
    livenessProbe:
      tcpSocket:
        port: 8080
      initialDelaySeconds: 15
      timeoutSeconds: 1
    ...
```

A TCP socket check is ideal for applications that do not start listening until initialization is complete.

For more information on health checks, see the [Kubernetes documentation](#).

CHAPTER 30. EVENTS

30.1. OVERVIEW

Events in OpenShift Container Platform are modeled based on events that happen to API objects in an OpenShift Container Platform cluster. Events allow OpenShift Container Platform to record information about real-world events in a resource-agnostic manner. They also allow developers and administrators to consume information about system components in a unified way.

30.2. VIEWING EVENTS WITH THE CLI

You can get a list of events in a given project using the following command:

```
$ oc get events [-n <project>]
```

30.3. VIEWING EVENTS IN THE CONSOLE

You can see events in your project from the web console from the **Browse** → **Events** page. Many other objects, such as pods and deployments, have their own **Events** tab as well, which shows events related to that object.

30.4. COMPREHENSIVE LIST OF EVENTS

This section describes the events of OpenShift Container Platform.

Table 30.1. Configuration Events

Name	Description
FailedValidation	Failed pod configuration validation.

Table 30.2. Container Events

Name	Description
BackOff	Back-off restarting failed the container.
Created	Container created.
Failed	Pull/Create/Start failed.
Killing	Killing the container.
Started	Container started.
Preempting	Preempting other pods.

Name	Description
ExceededGracePeriod	Container runtime did not stop the pod within specified grace period.

Table 30.3. Health Events

Name	Description
Unhealthy	Container is unhealthy.

Table 30.4. Image Events

Name	Description
BackOff	Back off Ctr Start, image pull.
ErrImageNeverPull	The image's NeverPull Policy is violated.
Failed	Failed to pull the image.
InspectFailed	Failed to inspect the image.
Pulled	Successfully pulled the image or the container image is already present on the machine.
Pulling	Pulling the image.

Table 30.5. Image Manager Events

Name	Description
FreeDiskSpaceFailed	Free disk space failed.
InvalidDiskCapacity	Invalid disk capacity.

Table 30.6. Node Events

Name	Description
FailedMount	Volume mount failed.

Name	Description
HostNetworkNotSupported	Host network not supported.
HostPortConflict	Host/port conflict.
InsufficientFreeCPU	Insufficient free CPU.
InsufficientFreeMemory	Insufficient free memory.
KubeletSetupFailed	Kubelet setup failed.
NilShaper	Undefined shaper.
NodeNotReady	Node is not ready.
NodeNotSchedulable	Node is not schedulable.
NodeReady	Node is ready.
NodeSchedulable	Node is schedulable.
NodeSelectorMismatching	Node selector mismatch.
OutOfDisk	Out of disk.
Rebooted	Node rebooted.
Starting	Starting kubelet.
FailedAttachVolume	Failed to attach volume.
FailedDetachVolume	Failed to detach volume.
VolumeResizeFailed	Failed to expand/reduce volume.

Name	Description
VolumeResizeSuccessful	Successfully expanded/reduced volume.
FileSystemResizeFailed	Failed to expand/reduce file system.
FileSystemResizeSuccessful	Successfully expanded/reduced file system.
FailedUnmount	Failed to unmount volume.
FailedMapVolume	Failed to map a volume.
FailedUnmapDevice	Failed unmaped device.
AlreadyMountedVolume	Volume is already mounted.
SuccessfulDetachVolume	Volume is successfully detached.
SuccessfulMountVolume	Volume is successfully mounted.
SuccessfulUnmountVolume	Volume is successfully unmounted.
ContainerGCFailed	Container garbage collection failed.
ImageGCFailed	Image garbage collection failed.
FailedNodeAllocatableEnforcement	Failed to enforce System Reserved Cgroup limit.
NodeAllocatableEnforced	Enforced System Reserved Cgroup limit.
UnsupportedMountOption	Unsupported mount option.

Name	Description
SandboxChanged	Pod sandbox changed.
FailedCreatePodSandbox	Failed to create pod sandbox.
FailedPodSandboxStatus	Failed pod sandbox status.

Table 30.7. Pod Worker Events

Name	Description
FailedSync	Pod sync failed.

Table 30.8. System Events

Name	Description
SystemOOM	There is an OOM (out of memory) situation on the cluster.

Table 30.9. Pod Events

Name	Description
FailedKillPod	Failed to stop a pod.
FailedCreatePodContainer	Failed to create a pod container.
Failed	Failed to make pod data directories.
NetworkNotReady	Network is not ready.
FailedCreate	Error creating: <error-msg> .
SuccessfulCreate	Created pod: <pod-name> .
FailedDelete	Error deleting: <error-msg> .
SuccessfulDelete	Deleted pod: <pod-id> .

Table 30.10. Horizontal Pod AutoScaler Events

Name	Description
SelectorRequired	Selector is required.
InvalidSelector	Could not convert selector into a corresponding internal selector object.
FailedGetObjectMetric	HPA was unable to compute the replica count.
InvalidMetricSourceType	Unknown metric source type.
ValidMetricFound	HPA was able to successfully calculate a replica count.
FailedConvertHPA	Failed to convert the given HPA.
FailedGetScale	HPA controller was unable to get the target's current scale.
SucceededGetScale	HPA controller was able to get the target's current scale.
FailedComputeMetricsReplicas	Failed to compute desired number of replicas based on listed metrics.
FailedRescale	New size: <size>; reason: <msg>; error: <error - msg>.
SuccessfulRescale	New size: <size>; reason: <msg>.
FailedUpdateStatus	Failed to update status.

Table 30.11. Network Events (openshift-sdn)

Name	Description
Starting	Starting OpenShift-SDN.
NetworkFailed	The pod's network interface has been lost and the pod will be stopped.

Table 30.12. Network Events (kube-proxy)

Name	Description
NeedPods	The service-port <code><serviceName>:<port></code> needs pods.

Table 30.13. Volume Events

Name	Description
FailedBinding	There are no persistent volumes available and no storage class is set.
VolumeMismatch	Volume size or class is different from what is requested in claim.
VolumeFailedRecycle	Error creating recycler pod.
VolumeRecycled	Occurs when volume is recycled.
RecyclerPod	Occurs when pod is recycled.
VolumeDelete	Occurs when volume is deleted.
VolumeFailedDelete	Error when deleting the volume.
ExternalProvisioning	Occurs when volume for the claim is provisioned either manually or via external software.
ProvisioningFailed	Failed to provision volume.
ProvisioningCleanupFailed	Error cleaning provisioned volume.
ProvisioningSucceeded	Occurs when the volume is provisioned successfully.
WaitForFirstConsumer	Delay binding until pod scheduling.

Table 30.14. Lifecycle hooks

Name	Description
FailedPostStartHook	Handler failed for pod start.
FailedPreStopHook	Handler failed for pre-stop.
UnfinishedPreStopHook	Pre-stop hook unfinished.

Table 30.15. Deployments

Name	Description
DeploymentCancellationFailed	Failed to cancel deployment.
DeploymentCancelled	Cancelled deployment.
DeploymentCreated	Created new replication controller.
IngressIPRangeFull	No available ingress IP to allocate to service.

Table 30.16. Scheduler Events

Name	Description
FailedScheduling	Failed to schedule pod: <pod - namespace>/<pod - name> . This event is raised for multiple reasons, for example: AssumePodVolumes failed, Binding rejected etc.
Preempted	By <preemptor - namespace>/<preemptor - name> on node <node - name> .
Scheduled	Successfully assigned <pod - name> to <node - name> .

Table 30.17. DaemonSet Events

Name	Description
SelectingAll	This daemon set is selecting all pods. A non-empty selector is required.
FailedPlacement	Failed to place pod on <node - name> .

Name	Description
FailedDaemonPod	Found failed daemon pod <pod-name> on node <node-name> , will try to kill it.

Table 30.18. LoadBalancer Service Events

Name	Description
CreatingLoadBalancerFailed	Error creating load balancer.
DeletingLoadBalancer	Deleting load balancer.
EnsuringLoadBalancer	Ensuring load balancer.
EnsuredLoadBalancer	Ensured load balancer.
UnavailableLoadBalancer	There are no available nodes for LoadBalancer service.
LoadBalancerSourceRanges	Lists the new LoadBalancerSourceRanges . For example, <old-source-range> → <new-source-range> .
LoadbalancerIP	Lists the new IP address. For example, <old-ip> → <new-ip> .
ExternalIP	Lists external IP address. For example, Added: <external-ip> .
UID	Lists the new UID. For example, <old-service-uid> → <new-service-uid> .
ExternalTrafficPolicy	Lists the new ExternalTrafficPolicy . For example, <old-policy> → <new-policy> .
HealthCheckNodePort	Lists the new HealthCheckNodePort . For example, <old-node-port> → <new-node-port> .
UpdatedLoadBalancer	Updated load balancer with new hosts.
LoadBalancerUpdateFailed	Error updating load balancer with new hosts.

Name	Description
DeletingLoad Balancer	Deleting load balancer.
DeletingLoad BalancerFailed	Error deleting load balancer.
DeletedLoadBalancer	Deleted load balancer.

CHAPTER 31. DOWNWARD API

31.1. OVERVIEW

The downward API is a mechanism that allows containers to consume information about API objects without coupling to OpenShift Container Platform. Such information includes the pod's name, namespace, and resource values. Containers can consume information from the downward API using environment variables or a volume plug-in.

31.2. SELECTING FIELDS

Fields within the pod are selected using the **FieldRef** API type. **FieldRef** has two fields:

Field	Description
fieldPath	The path of the field to select, relative to the pod.
apiVersion	The API version to interpret the fieldPath selector within.

Currently, the valid selectors in the v1 API include:

Selector	Description
metadata.name	The pod's name. This is supported in both environment variables and volumes.
metadata.namespace	The pod's namespace. This is supported in both environment variables and volumes.
metadata.labels	The pod's labels. This is only supported in volumes and not in environment variables.
metadata.annotations	The pod's annotations. This is only supported in volumes and not in environment variables.
status.podIP	The pod's IP. This is only supported in environment variables and not volumes.

The **apiVersion** field, if not specified, defaults to the API version of the enclosing pod template.

31.3. CONSUMING THE CONTAINER VALUES USING THE DOWNWARD API

31.3.1. Using Environment Variables

One mechanism for consuming the downward API is using a container's environment variables. The

EnvVar type's **valueFrom** field (of type **EnvVarSource**) is used to specify that the variable's value should come from a **FieldRef** source instead of the literal value specified by the **value** field. In the future, additional sources may be supported; currently the source's **fieldRef** field is used to select a field from the downward API.

Only constant attributes of the pod can be consumed this way, as environment variables cannot be updated once a process is started in a way that allows the process to be notified that the value of a variable has changed. The fields supported using environment variables are:

- Pod name
- Pod namespace

1. Create a *pod.json* file:

```
apiVersion: v1
kind: Pod
metadata:
  name: dapi-env-test-pod
spec:
  containers:
    - name: env-test-container
      image: gcr.io/google_containers/busybox
      command: [ "/bin/sh", "-c", "env" ]
      env:
        - name: MY_POD_NAME
          valueFrom:
            fieldRef:
              fieldPath: metadata.name
        - name: MY_POD_NAMESPACE
          valueFrom:
            fieldRef:
              fieldPath: metadata.namespace
      restartPolicy: Never
```

2. Create the pod from the *pod.json* file:

```
$ oc create -f pod.json
```

3. Check the container's logs for the **MY_POD_NAME** and **MY_POD_NAMESPACE** values:

```
$ oc logs -p dapi-env-test-pod
```

31.3.2. Using the Volume Plug-in

Another mechanism for consuming the downward API is using a volume plug-in. The downward API volume plug-in creates a volume with configured fields projected into files. The **metadata** field of the **VolumeSource** API object is used to configure this volume. The plug-in supports the following fields:

- Pod name
- Pod namespace
- Pod annotations

- Pod labels

Example 31.1. Downward API Volume Plug-in Configuration

```
spec:
  volumes:
    - name: podinfo
      metadata: ❶
        items: ❷
          - name: "labels" ❸
            fieldRef:
              fieldPath: metadata.labels ❹
```

- ❶ The `metadata` field of the volume source configures the downward API volume.
- ❷ The `items` field holds a list of fields to project into the volume.
- ❸ The name of the file to project the field into.
- ❹ The selector of the field to project.

For example:

1. Create a ***volume-pod.json*** file:

```
kind: Pod
apiVersion: v1
metadata:
  labels:
    zone: us-east-coast
    cluster: downward-api-test-cluster1
    rack: rack-123
  name: dapi-volume-test-pod
  annotations:
    annotation1: "345"
    annotation2: "456"
spec:
  containers:
    - name: volume-test-container
      image: gcr.io/google_containers/busybox
      command: ["sh", "-c", "cat /etc/labels /etc/annotations"]
      volumeMounts:
        - name: podinfo
          mountPath: /etc
          readOnly: false
  volumes:
    - name: podinfo
      metadata:
        items:
          - name: "labels"
            fieldRef:
              fieldPath: metadata.labels
          - name: "annotations"
```

```
        fieldRef:
          fieldPath: metadata.annotations
restartPolicy: Never
```

2. Create the pod from the ***volume-pod.json*** file:

```
$ oc create -f volume-pod.json
```

3. Check the container's logs and verify the presence of the configured fields:

```
$ oc logs -p dapi-volume-test-pod
cluster=downward-api-test-cluster1
rack=rack-123
zone=us-east-coast
annotation1=345
annotation2=456
kubernetes.io/config.source=api
```

31.4. CONSUMING CONTAINER RESOURCES USING THE DOWNWARD API

When creating pods, you can use the downward API to inject information about computing resource requests and limits so that image and application authors can correctly create an image for specific environments.

You can do this using both the [environment variable](#) and [volume plug-in](#) methods.

31.4.1. Using Environment Variables

1. When creating a pod configuration, specify environment variables that correspond to the contents of the **resources** field in the **spec.container** field:

```
...
spec:
  containers:
    - name: test-container
      image: gcr.io/google_containers/busybox:1.24
      command: [ "/bin/sh", "-c", "env" ]
      resources:
        requests:
          memory: "32Mi"
          cpu: "125m"
        limits:
          memory: "64Mi"
          cpu: "250m"
      env:
        - name: MY_CPU_REQUEST
          valueFrom:
            resourceFieldRef:
              resource: requests.cpu
        - name: MY_CPU_LIMIT
          valueFrom:
            resourceFieldRef:
              resource: limits.cpu
```

```

- name: MY_MEM_REQUEST
  valueFrom:
    resourceFieldRef:
      resource: requests.memory
- name: MY_MEM_LIMIT
  valueFrom:
    resourceFieldRef:
      resource: limits.memory
...

```

If the resource limits are not included in the container configuration, the downward API defaults to the node's CPU and memory allocatable values.

2. Create the pod from the *pod.json* file:

```
$ oc create -f pod.json
```

31.4.2. Using the Volume Plug-in

1. When creating a pod configuration, use the `spec.volumes.downwardAPI.items` field to describe the desired resources that correspond to the `spec.resources` field:

```

....
spec:
  containers:
    - name: client-container
      image: gcr.io/google_containers/busybox:1.24
      command: ["sh", "-c", "while true; do echo; if [[ -e
/etcd/cpu_limit ]]; then cat /etcd/cpu_limit; fi; if [[ -e
/etcd/cpu_request ]]; then cat /etcd/cpu_request; fi; if [[ -e
/etcd/mem_limit ]]; then cat /etcd/mem_limit; fi; if [[ -e
/etcd/mem_request ]]; then cat /etcd/mem_request; fi; sleep 5; done"]
      resources:
        requests:
          memory: "32Mi"
          cpu: "125m"
        limits:
          memory: "64Mi"
          cpu: "250m"
      volumeMounts:
        - name: podinfo
          mountPath: /etc
          readOnly: false
  volumes:
    - name: podinfo
      downwardAPI:
        items:
          - path: "cpu_limit"
            resourceFieldRef:
              containerName: client-container
              resource: limits.cpu
          - path: "cpu_request"
            resourceFieldRef:
              containerName: client-container
              resource: requests.cpu

```

```
- path: "mem_limit"
  resourceFieldRef:
    containerName: client-container
    resource: limits.memory
- path: "mem_request"
  resourceFieldRef:
    containerName: client-container
    resource: requests.memory
....
```

If the resource limits are not included in the container configuration, the downward API defaults to the node's CPU and memory allocatable values.

2. Create the pod from the ***volume-pod.json*** file:

```
$ oc create -f volume-pod.json
```


CHAPTER 32. MANAGING ENVIRONMENT VARIABLES

32.1. SETTING AND UNSETTING ENVIRONMENT VARIABLES

OpenShift Container Platform provides the `oc set env` command to set or unset environment variables for objects that have a pod [template](#), such as replication controllers or deployment configurations. It can also list environment variables in pods or any object that has a pod template. This command can also be used on `BuildConfig` objects.

32.2. LIST ENVIRONMENT VARIABLES

To list environment variables in pods or pod templates:

```
$ oc set env <object-selection> --list [<common-options>]
```

This example lists all environment variables for pod `p1`:

```
$ oc set env pod/p1 --list
```

32.3. SET ENVIRONMENT VARIABLES

To set environment variables in the pod templates:

```
$ oc set env <object-selection> KEY_1=VAL_1 ... KEY_N=VAL_N [<set-env-  
options>] [<common-options>]
```

Set environment options:

Option	Description
<code>-e, --env=<KEY>=<VAL></code>	Set given key value pairs of environment variables.
<code>--overwrite</code>	Confirm updating existing environment variables.

In the following example, both commands modify environment variable `STORAGE` in the deployment config `registry`. The first adds, with value `/data`. The second updates, with value `/opt`.

```
$ oc set env dc/registry STORAGE=/data
$ oc set env dc/registry --overwrite STORAGE=/opt
```

The following example finds environment variables in the current shell whose names begin with `RAILS_` and adds them to the replication controller `r1` on the server:

```
$ env | grep RAILS_ | oc set env rc/r1 -e -
```

The following example does not modify the replication controller defined in file `rc.json`. Instead, it writes a YAML object with updated environment `STORAGE=/local` to new file `rc.yaml`.

```
$ oc set env -f rc.json STORAGE=/opt -o yaml > rc.yaml
```

32.3.1. Automatically Added Environment Variables

Table 32.1. Automatically Added Environment Variables

Variable Name
<code><SVCNAME>_SERVICE_HOST</code>
<code><SVCNAME>_SERVICE_PORT</code>

Example Usage

The service **KUBERNETES** which exposes TCP port 53 and has been allocated cluster IP address 10.0.0.11 produces the following environment variables:

```
KUBERNETES_SERVICE_PORT=53
MYSQL_DATABASE=root
KUBERNETES_PORT_53_TCP=tcp://10.0.0.11:53
KUBERNETES_SERVICE_HOST=10.0.0.11
```



NOTE

Use the **oc rsh** command to SSH into your container and run **oc set env** to list all available variables.

32.4. UNSET ENVIRONMENT VARIABLES

To unset environment variables in the pod templates:

```
$ oc set env <object-selection> KEY_1- ... KEY_N- [<common-options>]
```



IMPORTANT

The trailing hyphen (-, U+2D) is required.

This example removes environment variables **ENV1** and **ENV2** from deployment config **d1**:

```
$ oc set env dc/d1 ENV1- ENV2-
```

This removes environment variable **ENV** from all replication controllers:

```
$ oc set env rc --all ENV-
```

This removes environment variable **ENV** from container **c1** for replication controller **r1**:

```
$ oc set env rc r1 --containers='c1' ENV-
```

CHAPTER 33. JOBS

33.1. OVERVIEW

A job, in contrast to a [replication controller](#), runs a pod with any number of replicas to completion. A job tracks the overall progress of a task and updates its status with information about active, succeeded, and failed pods. Deleting a job will clean up any pod replicas it created. Jobs are part of the Kubernetes API, which can be managed with `oc` commands like other [object types](#).

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

33.2. CREATING A JOB

A job configuration consists of the following key parts:

- A pod template, which describes the application the pod will create.
- An optional `parallelism` parameter, which specifies how many pod replicas running in parallel should execute a job. If not specified, this defaults to the value in the `completions` parameter.
- An optional `completions` parameter, specifying how many concurrently running pods should execute a job. If not specified, this value defaults to one.

The following is an example of a `job` resource:

```
apiVersion: batch/v1
kind: Job
metadata:
  name: pi
spec:
  parallelism: 1
  completions: 1
  template:
    metadata:
      name: pi
    spec:
      containers:
      - name: pi
        image: perl
        command: ["perl", "-Mbignum=bpi", "-wle", "print bpi(2000)"]
        restartPolicy: OnFailure
```

1. Optional value for how many pod replicas a job should run in parallel; defaults to `completions`.
2. Optional value for how many successful pod completions are needed to mark a job completed; defaults to one.
3. Template for the pod the controller creates.
4. The restart policy of the pod. This does not apply to the job controller. See [Section 33.2.1, “Known Limitations”](#) for details.

You can also create and launch a job from a single command using `oc run`. The following command creates and launches the same job as specified in the previous example:

```
$ oc run pi --image=perl --replicas=1 --restart=OnFailure \
  --command -- perl -Mbignum=bpi -wle 'print bpi(2000)'
```

33.2.1. Known Limitations

The job specification restart policy only applies to the *Pods*, and not the *job controller*. However, the job controller is hard-coded to keep retrying jobs to completion.

As such, `restartPolicy: Never` or `--restart=Never` results in the same behavior as `restartPolicy: OnFailure` or `--restart=OnFailure`. That is, when a job fails it is restarted automatically until it succeeds (or is manually discarded). The policy only sets which subsystem performs the restart.

With the `Never` policy, the *job controller* performs the restart. With each attempt, the job controller increments the number of failures in the job status and create new pods. This means that with each failed attempt, the number of pods increases.

With the `OnFailure` policy, *kubelet* performs the restart. Each attempt does not increment the number of failures in the job status. In addition, *kubelet* will retry failed jobs starting pods on the same nodes.

33.3. SCALING A JOB

A job can be scaled up or down by using the `oc scale` command with the `--replicas` option, which, in the case of jobs, modifies the `spec.parallelism` parameter. This will result in modifying the number of pod replicas running in parallel, executing a job.

The following command uses the example job above, and sets the `parallelism` parameter to three:

```
$ oc scale job pi --replicas=3
```



NOTE


Scaling replication controllers also uses the `oc scale` command with the `--replicas` option, but instead changes the `replicas` parameter of a replication controller configuration.

33.4. SETTING MAXIMUM DURATION

When defining a *Job*, you can define its maximum duration by setting the `activeDeadlineSeconds` field. It is specified in seconds and is not set by default. When not set, there is no maximum duration enforced.

The maximum duration is counted from the time when a first pod gets scheduled in the system, and defines how long a job can be active. It tracks overall time of an execution and is irrelevant to the number of completions (number of pod replicas needed to execute a task). After reaching the specified timeout, the job is terminated by OpenShift Container Platform.

The following example shows the part of a *Job* specifying `activeDeadlineSeconds` field for 30 minutes:



```
spec:  
  activeDeadlineSeconds: 1800
```

CHAPTER 34. CRON JOBS

34.1. OVERVIEW

A *cron job* builds on a regular [job](#) by allowing you to specifically schedule how the job should be run. Cron jobs are part of the [Kubernetes](#) API, which can be managed with `oc` commands like other [object types](#).



NOTE

As of OpenShift Container Platform 3.3.1, Cron Jobs is a feature in [Technology Preview](#).



WARNING

A cron job creates a job object approximately once per execution time of its schedule, but there are circumstances in which it will fail to create a job or two jobs might be created. As a result, jobs should be idempotent.

34.2. CREATING A CRON JOB

A cron job configuration consists of the following key parts:

- A schedule specified in [cron format](#).
- A job template used when creating the next job.
- An optional deadline (in seconds) for starting the job if it misses its scheduled time for any reason. Missed jobs executions will be counted as failed ones. If not specified, there is no deadline.
- **ConcurrencyPolicy:** An optional concurrency policy, specifying how to treat concurrent jobs within a cron job. Only one of the following concurrent policies may be specified. If not specified, this defaults to allowing concurrent executions.
 - **Allow** allows Cron Jobs to run concurrently.
 - **Forbid** forbids concurrent runs, skipping the next run if the previous has not finished yet.
 - **Replace** cancels the currently running job and replaces it with a new one.
- An optional flag allowing the suspension of a cron job. If set to `true`, all subsequent executions will be suspended.

The following is an example of a `CronJob` resource:

```
apiVersion: batch/v2alpha1
kind: CronJob
metadata:
  name: pi
```

```
spec:
  schedule: "*/1 * * * *" ❶
  jobTemplate: ❷
    spec:
      template:
        metadata:
          labels: ❸
            parent: "cronjobpi"
        spec:
          containers:
            - name: pi
              image: perl
              command: ["perl", "-Mbignum=bpi", "-wle", "print bpi(2000)"]
          restartPolicy: OnFailure ❹
```

1. Schedule for the job. In this example, the job will run every minute.
2. Job template. This is similar to the [job example](#).
3. Sets a label for jobs spawned by this cron job.
4. The restart policy of the pod. This does not apply to the job controller. See [Known Issues and Limitations](#) for details.

You can also create and launch a cron job from a single command using `oc run`. The following command creates and launches the same cron job as specified in the previous example:

```
$ oc run pi --image=perl --schedule='*/1 * * * *' \
  --restart=OnFailure --labels parent="cronjobpi" \
  --command -- perl -Mbignum=bpi -wle 'print bpi(2000)'
```

With `oc run`, the `--schedule` option accepts schedules in [cron format](#).



NOTE

When creating a cron job, `oc run` only supports the **Never** or **OnFailure** restart policies (`--restart`).

TIP

Delete cron jobs that you no longer need:

```
$ oc delete cronjob/<cron_job_name>
```

Doing this prevents them from generating unnecessary artifacts.

34.3. CLEANING UP AFTER A CRON JOB

Cron jobs can leave behind artifact resources such as jobs or pods. Check if any remain:

```
$ oc get jobs
$ oc get pods
```

All artifacts left over from a job execution use the job name as their prefix. For example, given the cron job example:

```
$ oc get jobs
NAME                                DESIRED    SUCCESSFUL    AGE
pi-1497848100                      1          1             1m
pi-1497848160                      1          1             49s

$ oc get pods
NAME                                READY      STATUS      RESTARTS    AGE
pi-1497848100-lxs4k                0/1       Completed   0           2m
pi-1497848160-6r0c8                0/1       Completed   0           59s
```

Delete each artifact if you no longer need them. To delete all jobs spawned by a cron job, specify the label set during cron job creation:

```
$ oc delete jobs -l <label>
```

For example, to delete only the jobs generated by the cron job example:

```
$ oc delete jobs -l parent=cronjobpi
job "pi-1497848100" deleted
job "pi-1497848160" deleted
```


CHAPTER 35. CREATE FROM URL

35.1. OVERVIEW

Create From URL is a function that allows you to construct a URL from an image stream, image tag, or template.

Create from URL only works with image streams or templates from namespaces that have been explicitly whitelisted. The whitelist contains the `openshift` namespace by default. To add namespaces to the whitelist, see [Configuring the Create From URL Namespace Whitelist](#).

You can define custom buttons.



These buttons leverage a defined URL pattern with an appropriate query string. The user is prompted to select the project. Then, the Create from URL workflow continues.

35.2. USING AN IMAGE STREAM AND IMAGE TAG

35.2.1. Query String Parameters

Name	Description	Required	Schema	Default
<code>imageStream</code>	The value <code>metadata.name</code> as defined in the image stream to be used.	true	string	
<code>imageTag</code>	The value <code>spec.tags.name</code> as defined in the image stream to be used.	true	string	
<code>namespace</code>	The name of the namespace containing the image stream and image tag to use.	false	string	<code>openshift</code>
<code>name</code>	Identifies the resources created for this application.	false	string	

Name	Description	Required	Schema	Default
sourceURI	The Git repository URL containing the application source code.	false	string	
sourceRef	The branch, tag, or commit for the application source code specified in sourceURI .	false	string	
contextDir	The subdirectory for the application source code specified in sourceURI , used as the context directory for the build.	false	string	

**NOTE**

[Reserved characters](#) in parameter values should be URL encoded.


35.2.1.1. Example

```
create?
imageStream=nodejs&imageTag=4&name=nodejs&sourceURI=https%3A%2F%2Fgithub.com%2Fopenshift%2Fnodejs-ex.git&sourceRef=master&contextDir=%2F
```

35.3. USING A TEMPLATE**35.3.1. Query String Parameters**

Name	Description	Required	Schema	Default
template	The value of metadata.name as defined in the template to be used.	true	string	

Name	Description	Required	Schema	Default
templateParamsMap	A JSON parameters map containing the template parameter name and corresponding value you wish to override.	false	JSON	
namespace	The name of the namespace containing the template to use.	false	string	openshift

**NOTE**

Reserved characters in parameter values should be URL encoded.

35.3.1.1. Example

```
create?template=nodejs-mongodb-example&templateParamsMap=
{"SOURCE_REPOSITORY_URL"%3A"https%3A%2F%2Fgithub.com%2Fopenshift%2Fnodejs-
ex.git"}
```

CHAPTER 36. REVISION HISTORY: DEVELOPER GUIDE

36.1. FRI FEB 23 2018

Affected Topic	Description of Change
Service Accounts	Noted the importance of <code>-z</code> flag usage when granting access to service accounts .

36.2. FRI FEB 16 2018

Affected Topic	Description of Change
Service Accounts	Changed <code>serviceaccounts</code> to <code>serviceaccount</code> in the User Names and Groups section.
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.
Build Inputs	Added an example of a Dockerfile referencing secret data in the Docker Strategy section.
Build Inputs	Updated <code>git clone</code> behavior in the Git Source section.

36.3. TUE FEB 06 2018

Affected Topic	Description of Change
Service Accounts	Changed <code>serviceaccounts</code> to <code>serviceaccount</code> in the User Names and Groups section.
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.
Build Inputs	Added an example of a Dockerfile referencing secret data in the Docker Strategy section.
Build Inputs	Updated <code>git clone</code> behavior in the Git Source section.

36.4. THU JAN 25 2018

Affected Topic	Description of Change
Service Accounts	Changed serviceaccounts to serviceaccount in the User Names and Groups section.
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.
Build Inputs	Added an example of a Dockerfile referencing secret data in the Docker Strategy section.
Build Inputs	Updated git clone behavior in the Git Source section.

36.5. MON JAN 08 2018

Affected Topic	Description of Change
Service Accounts	Changed serviceaccounts to serviceaccount in the User Names and Groups section.
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.

36.6. FRI DEC 22 2017

Affected Topic	Description of Change
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.

36.7. MON DEC 11 2017

Affected Topic	Description of Change
Service Accounts	Changed serviceaccounts to serviceaccount in the User Names and Groups section.

Affected Topic	Description of Change
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.

36.8. TUE NOV 21 2017

Affected Topic	Description of Change
Service Accounts	Changed serviceaccounts to serviceaccount in the User Names and Groups section.
Application Life Cycle Management → Creating New Applications	New section on Searching for Images, Templates, and Other Inputs when creating new applications.

36.9. FRI NOV 10 2017

Affected Topic	Description of Change
Service Accounts	Changed serviceaccounts to serviceaccount in the User Names and Groups section.

36.10. FRI NOV 03 2017

Affected Topic	Description of Change
Managing Volumes	Added - -configmap-name to the Supported Options for Adding Volumes table.

36.11. MON OCT 16 2017

Affected Topic	Description of Change
Managing Images	Clarified wording around Image Streams and added new section on Adding Trusted Certificates for External Registries .

36.12. MON SEP 18 2017

Affected Topic	Description of Change
Build Inputs	Clarified the Automatically Adding a Source Clone Secret to a Build Configuration section.
Integrating External Services	Added context and use cases for the <code>ExternalName</code> parameter.
Pod Autoscaling	Added information on how resource startup spikes are handled to avoid autoscaling .
Managing Images	Noted that <code>importPolicy.scheduled</code> is disabled by default.

36.13. WED SEP 06 2017

Affected Topic	Description of Change
Routes	Added a file path to the unsecured route CLI example in the Creating Routes section.

36.14. FRI AUG 25 2017

Affected Topic	Description of Change
Using Daemonsets	Removed Technology Preview notice.

36.15. TUE JUL 18 2017

Affected Topic	Description of Change
Image Signatures	New topic about signing and verifying container image signatures.

36.16. WED JUL 12 2017

Affected Topic	Description of Change
Developer Guide → Integrating External Services	Added <code>updatePeriodSeconds</code> and <code>intervalSeconds</code> to the Consume a Service example with definitions.
Developer Guide → Deployment Strategies	Added <code>updatePeriodSeconds</code> and <code>intervalSeconds</code> to the Rolling Deployment example with definitions.

36.17. WED JUL 05 2017

Affected Topic	Description of Change
Downward API	Updated annotation formatting in the Using the Volume Plug-in section.
Creating New Applications	Added Specifying Build Environment Variables section.

36.18. TUE JUN 27 2017

Affected Topic	Description of Change
Authorization	Added the Authorization topic.

36.19. MON JUN 19 2017

Affected Topic	Description of Change
Build Inputs	Added proxy warning for Jenkins in the Using a Proxy section.

36.20. TUE JUN 13 2017

Affected Topic	Description of Change
Managing Images	Renamed section to Recommended Tagging Conventions and updated Image Tag Naming Conventions table.

36.21. MON MAY 15 2017

Affected Topic	Description of Change
Environment Variables	Added Ignoring Source Files section.

36.22. WED APR 12 2017

OpenShift Container Platform 3.5 Initial Release

Affected Topic	Description of Change
Application Life Cycle Management → Creating New Applications	Added goolang to the Languages Detected by new-app table.
	Added information about <code>--param-file</code> and <code>--env-file</code> <code>oc</code> command arguments to the Template Parameters section.
	Added information about the pipeline build strategy.
	Noted that any BuildConfig objects created as part of <code>oc new-app</code> processing will not be updated with environment variables passed via the <code>--e -env</code> argument.
Builds → Build Inputs	Clarified when URLs can be used with binary builds.
	Added details about the <code>build.openshift.io/source-secret-match-uri-</code> annotation.
	Added information about using URLs with <code>oc start-build --from-file</code> and <code>oc start build --from-dir</code> .
Templates	Added link to OpenShift Image Streams and Templates library in the Overview section.
	Updated the NOTE box about updating template namespace behavior in the Object List section to reflect parameter support.
	Added information about <code>--param-file</code> and working with a large number of parameters to the Parameters section.
Managing Images	Added new sections about the Insecure Tag Import Policy and Reference Policy .
Quotas and Limit Ranges	Added a new section Opaque Integer Resources.
Secrets	Added information about expiration to the Service Serving Certificate Secrets section.
Cron Jobs	Changed instances of "Scheduled jobs" to "Cron jobs".
Create from URL	New topic on Create from URL, a function that allows you to construct a URL from an image stream, image tag, or template.