



OpenShift Container Platform 3.3

Installation and Configuration

OpenShift Container Platform 3.3 Installation and Configuration

OpenShift Container Platform 3.3 Installation and Configuration

OpenShift Container Platform 3.3 Installation and Configuration

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

OpenShift Installation and Configuration topics cover the basics of installing and configuring OpenShift in your environment. Use these topics for the one-time tasks required to get OpenShift up and running.

Table of Contents

CHAPTER 1. OVERVIEW	16
CHAPTER 2. INSTALLING A CLUSTER	17
2.1. PLANNING	17
2.1.1. Initial Planning	17
2.1.2. Installation Methods	17
2.1.3. Sizing Considerations	17
2.1.4. Environment Scenarios	18
2.1.4.1. Single Master and Multiple Nodes	18
2.1.4.2. Single Master, Multiple etcd, and Multiple Nodes	19
2.1.4.3. Multiple Masters Using Native HA	19
2.1.4.4. Stand-alone Registry	20
2.1.5. RPM vs Containerized	20
2.2. PREREQUISITES	20
2.2.1. System Requirements	21
2.2.1.1. Red Hat Subscriptions	21
2.2.1.2. Minimum Hardware Requirements	21
2.2.1.3. Production Level Hardware Requirements	22
2.2.1.4. Configuring Core Usage	23
2.2.1.5. SELinux	23
2.2.1.6. NTP	23
2.2.1.7. Security Warning	23
2.2.2. Environment Requirements	24
2.2.2.1. DNS	24
2.2.2.1.1. Configuring Hosts to Use DNS	25
2.2.2.1.2. Disabling dnsmasq	26
2.2.2.1.3. Configuring a DNS Wildcard	26
2.2.2.2. Network Access	27
2.2.2.2.1. NetworkManager	27
2.2.2.2.2. Required Ports	27
2.2.2.3. Persistent Storage	30
2.2.2.4. Cloud Provider Considerations	30
2.2.2.4.1. Configuring a Security Group	30
2.2.2.4.2. Overriding Detected IP Addresses and Host Names	31
2.2.2.4.3. Post-Installation Configuration for Cloud Providers	33
2.3. HOST PREPARATION	33
2.3.1. Operating System Requirements	33
2.3.2. Host Registration	33
2.3.3. Installing Base Packages	34
2.3.4. Installing Docker	35
2.3.5. Configuring Docker Storage	36
2.3.5.1. Reconfiguring Docker Storage	39
2.3.5.2. Managing Container Logs	39
2.3.5.3. Viewing Available Container Logs	40
2.3.6. Ensuring Host Access	40
2.3.7. Setting Global Proxy Values	41
2.3.8. What's Next?	41
2.4. CONTAINERIZED COMPONENTS	41
2.4.1. Overview	42
2.4.2. Required Images	42
2.4.3. Starting and Stopping Containers	43

2.4.4. File Paths	43
2.4.5. Storage Requirements	43
2.4.6. Open vSwitch SDN Initialization	43
2.5. QUICK INSTALLATION	43
2.5.1. Overview	43
2.5.2. Before You Begin	44
2.5.3. Running an Interactive Installation	45
2.5.4. Defining an Installation Configuration File	45
2.5.5. Running an Unattended Installation	47
2.5.6. Verifying the Installation	47
2.5.7. Uninstalling OpenShift Container Platform	48
2.5.8. What's Next?	48
2.6. ADVANCED INSTALLATION	48
2.6.1. Overview	48
2.6.2. Before You Begin	48
2.6.3. Configuring Ansible	49
2.6.3.1. Configuring Host Variables	49
Image Version Policy	50
2.6.3.2. Configuring Cluster Variables	51
2.6.3.3. Configuring a Registry Location	54
2.6.3.4. Configuring Global Proxy Options	55
2.6.3.5. Configuring Schedulability on Masters	56
2.6.3.6. Configuring Node Host Labels	57
2.6.3.6.1. Configuring Dedicated Infrastructure Nodes	57
2.6.3.7. Configuring Session Options	58
2.6.3.8. Configuring Custom Certificates	58
2.6.3.9. Configuring Deployment Type	59
2.6.4. Configuring Cluster Metrics	59
2.6.4.1. Metrics Storage	60
2.6.5. Single Master Examples	60
2.6.6. Multiple Masters Examples	63
2.6.7. Running the Advanced Installation	68
2.6.8. Verifying the Installation	69
2.6.9. Optionally Securing Builds	70
2.6.10. Uninstalling OpenShift Container Platform	70
2.6.10.1. Uninstalling Nodes	71
2.6.11. Known Issues	72
2.6.12. What's Next?	72
2.7. DISCONNECTED INSTALLATION	72
2.7.1. Overview	72
2.7.2. Prerequisites	73
2.7.3. Required Software and Components	73
2.7.3.1. Syncing Repositories	73
2.7.3.2. Syncing Images	74
2.7.3.3. Preparing Images for Export	76
2.7.4. Repository Server	77
2.7.4.1. Placing the Software	77
2.7.5. OpenShift Container Platform Systems	78
2.7.5.1. Building Your Hosts	78
2.7.5.2. Connecting the Repositories	78
2.7.5.3. Host Preparation	78
2.7.6. Installing OpenShift Container Platform	78
2.7.6.1. Importing OpenShift Container Platform Containerized Components	78

2.7.6.2. Running the OpenShift Container Platform Installer	79
2.7.6.3. Creating the Internal Docker Registry	79
2.7.7. Post-Installation Changes	79
2.7.7.1. Re-tagging S2I Builder Images	79
2.7.7.2. Configuring a Registry Location	80
2.7.7.3. Creating an Administrative User	81
2.7.7.4. Modifying the Security Policies	81
2.7.7.5. Editing the Image Stream Definitions	82
2.7.7.6. Loading the Container Images	83
2.7.8. Installing a Router	83
2.8. INSTALLING A STAND-ALONE DEPLOYMENT OF OPENSIFT CONTAINER REGISTRY	83
2.8.1. About OpenShift Container Registry	83
2.8.2. Minimum Hardware Requirements	84
2.8.3. Supported System Topologies	85
2.8.4. Host Preparation	85
2.8.5. Installation Methods	85
2.8.5.1. Quick Installation for Stand-alone OpenShift Container Registry	85
2.8.5.2. Advanced Installation for Stand-alone OpenShift Container Registry	85
CHAPTER 3. SETTING UP THE REGISTRY	89
3.1. REGISTRY OVERVIEW	89
3.1.1. About the Registry	89
3.1.2. Integrated or Stand-alone Registries	89
3.2. DEPLOYING A REGISTRY ON EXISTING CLUSTERS	89
3.2.1. Overview	89
3.2.2. Deploying the Registry	89
3.2.3. Deploying the Registry as a DaemonSet	90
3.2.4. Registry Compute Resources	90
3.2.5. Storage for the Registry	90
3.2.5.1. Production Use	91
3.2.5.1.1. Use Amazon S3 as a Storage Back-end	91
3.2.5.2. Non-Production Use	92
3.2.6. Enabling the Registry Console	92
3.2.6.1. Deploying the Registry Console	93
3.2.6.2. Securing the Registry Console	93
3.2.6.3. Troubleshooting the Registry Console	95
3.2.6.3.1. Debug Mode	95
3.2.6.3.2. Display SSL Certificate Path	95
3.3. ACCESSING THE REGISTRY	95
3.3.1. Viewing Logs	96
3.3.2. File Storage	96
3.3.3. Accessing the Registry Directly	98
3.3.3.1. User Prerequisites	98
3.3.3.2. Logging in to the Registry	99
3.3.3.3. Pushing and Pulling Images	99
3.4. SECURING AND EXPOSING THE REGISTRY	100
3.4.1. Securing the Registry	100
3.4.2. Exposing the Registry	103
3.4.2.1. Exposing a Secure Registry	105
3.4.2.2. Exposing a Non-Secure Registry	105
3.5. EXTENDED REGISTRY CONFIGURATION	106
3.5.1. Maintaining the Registry IP Address	107
3.5.2. Whitelisting Docker Registries	107

3.5.3. Overriding the Registry Configuration	108
3.5.4. Registry Configuration Reference	110
3.5.4.1. Log	110
3.5.4.2. Hooks	110
3.5.4.3. Storage	110
3.5.4.4. Auth	111
3.5.4.5. Middleware	112
3.5.4.5.1. CloudFront Middleware	112
3.5.4.5.2. Overriding Middleware Configuration Options	113
3.5.4.5.3. Image Pullthrough	114
3.5.4.5.4. Manifest V2 Schema 2 Support	114
3.5.4.6. Reporting	115
3.5.4.7. HTTP	115
3.5.4.8. Notifications	115
3.5.4.9. Redis	115
3.5.4.10. Health	116
3.5.4.11. Proxy	116
3.6. KNOWN ISSUES	116
3.6.1. Overview	116
3.6.2. Image Push Errors with Scaled Registry Using Shared NFS Volume	116
3.6.3. Pull of Internally Managed Image Fails with "not found" Error	117
3.6.4. Image Push Fails with "500 Internal Server Error" on S3 Storage	117
3.6.5. Image Pruning Fails	117
CHAPTER 4. SETTING UP A ROUTER	119
4.1. ROUTER OVERVIEW	119
4.1.1. About Routers	119
4.1.2. Router Service Account	119
4.2. USING THE DEFAULT HAProxy ROUTER	119
4.2.1. Overview	119
4.2.2. Creating a Router	120
4.2.3. Other Basic Router Commands	120
4.2.4. Filtering Routes to Specific Routers	121
4.2.5. Highly-Available Routers	121
4.2.6. Customizing the Router Service Ports	122
4.2.7. Working With Multiple Routers	122
4.2.8. Adding a Node Selector to a Deployment Configuration	122
4.2.9. Using Router Shards	123
4.2.9.1. Creating Router Shards	125
4.2.9.2. Modifying Router Shards	126
4.2.9.3. Using Namespace Router Shards	128
4.2.10. Customizing the Default Routing Subdomain	129
4.2.11. Forcing Route Host Names to a Custom Routing Subdomain	129
4.2.12. Using Wildcard Certificates	129
4.2.13. Manually Redeploy Certificates	130
4.2.14. Using Secured Routes	131
4.2.15. Using the Container Network Stack	132
4.2.16. Exposing Router Metrics	133
4.2.17. Preventing Connection Failures During Restarts	134
4.2.18. Protecting Against DDoS Attacks	135
4.3. DEPLOYING A CUSTOMIZED HAProxy ROUTER	136
4.3.1. Overview	136
4.3.2. Obtaining the Router Configuration Template	136

4.3.3. Modifying the Router Configuration Template	137
4.3.3.1. Background	137
4.3.3.2. Go Template Actions	137
4.3.3.3. Router Provided Information	138
4.3.3.4. Annotations	142
4.3.3.5. Environment Variables	143
4.3.3.6. Example Usage	143
4.3.4. Using a ConfigMap to Replace the Router Configuration Template	145
4.3.5. Using Stick Tables	146
4.3.6. Rebuilding Your Router	148
4.4. CONFIGURING THE HAProxy ROUTER TO USE THE PROXY PROTOCOL	148
4.4.1. Overview	148
4.4.2. Why Use the PROXY Protocol?	149
4.4.3. Using the PROXY Protocol	149
4.5. USING THE F5 ROUTER PLUG-IN	153
4.5.1. Overview	153
4.5.2. Deploying the F5 Router	154
CHAPTER 5. UPGRADING A CLUSTER	155
5.1. OVERVIEW	155
5.1.1. In-place or Blue-Green Upgrades	155
5.2. PERFORMING AUTOMATED IN-PLACE CLUSTER UPGRADES	155
5.2.1. Overview	155
5.2.2. Preparing for an Automated Upgrade	156
5.2.3. Using the Installer to Upgrade	157
5.2.4. Running the Upgrade Playbook Directly	157
5.2.4.1. Upgrading to OpenShift Container Platform 3.3	158
5.2.4.2. Upgrading to OpenShift Container Platform 3.3 Asynchronous Releases	158
5.2.5. Upgrading the EFK Logging Stack	159
5.2.6. Upgrading Cluster Metrics	159
5.2.7. Verifying the Upgrade	159
5.3. PERFORMING MANUAL IN-PLACE CLUSTER UPGRADES	160
5.3.1. Overview	160
5.3.2. Preparing for a Manual Upgrade	160
5.3.3. Upgrading Master Components	161
5.3.4. Updating Policy Definitions	163
5.3.5. Upgrading Nodes	164
5.3.6. Upgrading the Router	166
5.3.7. Upgrading the Registry	167
5.3.7.1. Updating Custom Registry Configuration Files	167
5.3.7.2. Enforcing Quota in the Registry	168
5.3.8. Updating the Default Image Streams and Templates	168
5.3.9. Importing the Latest Images	170
5.3.10. Upgrading the EFK Logging Stack	171
5.3.11. Upgrading Cluster Metrics	172
5.3.12. Additional Manual Steps Per Release	173
5.3.13. Verifying the Upgrade	173
5.4. BLUE-GREEN DEPLOYMENTS	174
5.4.1. Overview	174
5.4.2. Preparing for Upgrade	174
5.4.3. Warming the New Nodes	175
5.4.4. Node Evacuation	176
5.5. OPERATING SYSTEM UPDATES AND UPGRADES	176

5.5.1. Updating and Upgrading the Operating System	176
CHAPTER 6. DOWNGRADING OPENSIFT	178
6.1. OVERVIEW	178
6.2. VERIFYING BACKUPS	178
6.3. SHUTTING DOWN THE CLUSTER	178
6.4. REMOVING RPMS	179
6.5. DOWNGRADING DOCKER	179
6.6. REINSTALLING RPMS	180
6.7. BRINGING OPENSIFT CONTAINER PLATFORM SERVICES BACK ONLINE	181
6.8. VERIFYING THE DOWNGRADE	181
CHAPTER 7. MASTER AND NODE CONFIGURATION	182
7.1. OVERVIEW	182
7.2. MASTER CONFIGURATION FILES	182
7.2.1. Admission Control Configuration	182
7.2.2. Asset Configuration	183
7.2.3. Authentication and Authorization Configuration	184
7.2.4. Controller Configuration	184
7.2.5. etcd Configuration	184
7.2.6. Grant Configuration	185
7.2.7. Image Configuration	186
7.2.8. Kubernetes Master Configuration	186
7.2.9. Network Configuration	187
7.2.10. OAuth Authentication Configuration	188
7.2.11. Project Configuration	189
7.2.12. Scheduler Configuration	189
7.2.13. Security Allocator Configuration	189
7.2.14. Service Account Configuration	190
7.2.15. Serving Information Configuration	191
7.2.16. Volume Configuration	192
7.2.17. Audit Configuration	192
7.3. NODE CONFIGURATION FILES	193
7.3.1. Pod and Node Configuration	195
7.3.2. Docker Configuration	195
7.3.3. Parallel Image Pulls with Docker 1.9+	195
7.4. PASSWORDS AND OTHER SENSITIVE DATA	196
7.5. CREATING NEW CONFIGURATION FILES	197
7.6. LAUNCHING SERVERS USING CONFIGURATION FILES	197
7.7. CONFIGURING LOGGING LEVELS	198
CHAPTER 8. ADDING HOSTS TO AN EXISTING CLUSTER	201
8.1. OVERVIEW	201
8.2. ADDING HOSTS USING THE QUICK INSTALLER TOOL	201
8.3. ADDING HOSTS USING THE ADVANCED INSTALL	201
CHAPTER 9. LOADING THE DEFAULT IMAGE STREAMS AND TEMPLATES	205
9.1. OVERVIEW	205
9.2. OFFERINGS BY SUBSCRIPTION TYPE	205
9.2.1. OpenShift Container Platform Subscription	205
9.2.2. xPaaS Middleware Add-on Subscriptions	206
9.3. BEFORE YOU BEGIN	206
9.4. PREREQUISITES	206
9.5. CREATING IMAGE STREAMS FOR OPENSIFT CONTAINER PLATFORM IMAGES	207

9.6. CREATING IMAGE STREAMS FOR XPAAS MIDDLEWARE IMAGES	207
9.7. CREATING DATABASE SERVICE TEMPLATES	207
9.8. CREATING INSTANT APP AND QUICKSTART TEMPLATES	208
9.9. WHAT'S NEXT?	209
CHAPTER 10. CONFIGURING CUSTOM CERTIFICATES	210
10.1. OVERVIEW	210
10.2. CONFIGURING CUSTOM CERTIFICATES WITH ANSIBLE	210
10.3. CONFIGURING CUSTOM CERTIFICATES	210
CHAPTER 11. REDEPLOYING CERTIFICATES	212
11.1. OVERVIEW	212
11.2. CHECKING CERTIFICATE EXPIRATIONS	212
11.2.1. Role Variables	212
11.2.2. Running Certificate Expiration Playbooks	213
Other Example Playbooks	214
11.2.3. Output Formats	214
HTML Report	214
JSON Report	215
11.3. REDEPLOYING CERTIFICATES	215
11.3.1. Redeploying All Certificates Using the Current OpenShift Container Platform and etcd CA	216
11.3.2. Redeploying a New or Custom OpenShift Container Platform CA	216
11.3.3. Redeploying a New etcd CA	217
11.3.4. Redeploying Master Certificates Only	217
11.3.5. Redeploying etcd Certificates Only	218
11.3.6. Redeploying Node Certificates Only	218
11.3.7. Redeploying Registry or Router Certificates Only	218
11.3.7.1. Redeploying Registry Certificates Only	218
11.3.7.2. Redeploying Router Certificates Only	219
11.3.8. Redeploying Custom Registry or Router Certificates	219
11.3.8.1. Redeploying Registry Certificates Manually	219
11.3.8.2. Redeploying Router Certificates Manually	220
CHAPTER 12. CONFIGURING AUTHENTICATION AND USER AGENT	223
12.1. OVERVIEW	223
12.2. IDENTITY PROVIDER PARAMETERS	223
12.3. CONFIGURING IDENTITY PROVIDERS	224
12.3.1. Configuring identity providers with Ansible	225
12.3.2. Configuring identity providers in the master configuration file	225
12.3.3. Allow all	226
12.3.4. Deny all	227
12.3.5. HTPasswd	227
12.3.6. Keystone	229
12.3.7. LDAP authentication	230
12.3.8. Basic authentication (remote)	233
12.3.9. Request header	234
12.3.10. GitHub	241
12.3.11. GitLab	242
12.3.12. Google	243
12.3.13. OpenID connect	244
12.4. TOKEN OPTIONS	247
12.5. GRANT OPTIONS	248
12.6. SESSION OPTIONS	248
12.7. PREVENTING CLI VERSION MISMATCH WITH USER AGENT	249

CHAPTER 13. SYNCING GROUPS WITH LDAP	251
13.1. OVERVIEW	251
13.2. CONFIGURING LDAP SYNC	251
13.2.1. LDAP Client Configuration	251
13.2.2. LDAP Query Definition	252
13.2.3. User-Defined Name Mapping	253
13.3. RUNNING LDAP SYNC	253
13.4. RUNNING A GROUP PRUNING JOB	254
13.5. SYNC EXAMPLES	254
13.5.1. RFC 2307	255
13.5.1.1. RFC2307 with User-Defined Name Mappings	257
13.5.2. RFC 2307 with User-Defined Error Tolerances	259
13.5.3. Active Directory	261
13.5.4. Augmented Active Directory	263
13.6. NESTED MEMBERSHIP SYNC EXAMPLE	266
13.7. LDAP SYNC CONFIGURATION SPECIFICATION	269
13.7.1. v1.LDAPSyncConfig	269
13.7.2. v1.StringSource	271
13.7.3. v1.LDAPQuery	271
13.7.4. v1.RFC2307Config	272
13.7.5. v1.ActiveDirectoryConfig	274
13.7.6. v1.AugmentedActiveDirectoryConfig	275
CHAPTER 14. CONFIGURING LDAP FAILOVER	276
14.1. PREREQUISITES FOR CONFIGURING BASIC REMOTE AUTHENTICATION	276
14.2. GENERATING AND SHARING CERTIFICATES WITH THE REMOTE BASIC AUTHENTICATION SERVER	276
14.3. CONFIGURING SSSD FOR LDAP FAILOVER	277
14.4. CONFIGURING APACHE TO USE SSSD	279
14.5. CONFIGURING OPENSIFT CONTAINER PLATFORM TO USE SSSD AS THE BASIC REMOTE AUTHENTICATION SERVER	282
CHAPTER 15. CONFIGURING THE SDN	284
15.1. OVERVIEW	284
15.2. CONFIGURING THE POD NETWORK WITH ANSIBLE	284
15.3. CONFIGURING THE POD NETWORK ON MASTERS	285
15.4. CONFIGURING THE POD NETWORK ON NODES	285
15.5. MIGRATING BETWEEN SDN PLUG-INS	285
15.6. EXTERNAL ACCESS TO THE CLUSTER NETWORK	286
15.7. USING FLANNEL	286
CHAPTER 16. CONFIGURING NUAGE SDN	288
16.1. NUAGE SDN AND OPENSIFT CONTAINER PLATFORM	288
16.2. DEVELOPER WORKFLOW	288
16.3. OPERATIONS WORKFLOW	288
16.4. INSTALLATION	288
16.4.1. Installation for a Single Master	288
16.4.2. Installation for Multiple Masters (HA)	289
CHAPTER 17. CONFIGURING FOR AWS	291
17.1. OVERVIEW	291
17.2. CONFIGURING AWS VARIABLES	291
17.3. CONFIGURING OPENSIFT CONTAINER PLATFORM MASTERS FOR AWS	291
17.3.1. Configuring OpenShift Container Platform for AWS with Ansible	291

17.3.2. Manually Configuring OpenShift Container Platform Masters for AWS	292
17.3.3. Manually Configuring OpenShift Container Platform Nodes for AWS	292
17.4. SETTING KEY VALUE ACCESS PAIRS	293
17.5. APPLYING CONFIGURATION CHANGES	293
CHAPTER 18. CONFIGURING FOR OPENSTACK	294
18.1. OVERVIEW	294
18.2. CONFIGURING OPENSTACK VARIABLES	294
18.3. CONFIGURING OPENSIFT CONTAINER PLATFORM MASTERS FOR OPENSTACK	294
18.3.1. Configuring OpenShift Container Platform for OpenStack with Ansible	294
18.3.2. Manually Configuring OpenShift Container Platform Masters for OpenStack	295
18.3.3. Manually Configuring OpenShift Container Platform Nodes for OpenStack	296
CHAPTER 19. CONFIGURING FOR GCE	297
19.1. OVERVIEW	297
19.2. CONFIGURING MASTERS	297
19.2.1. Configuring OpenShift Container Platform Masters for GCE with Ansible	297
19.2.2. Manually Configuring OpenShift Container Platform Masters for GCE	297
19.3. CONFIGURING NODES	298
19.4. CONFIGURING MULTIZONE SUPPORT IN A GCE DEPLOYMENT	298
CHAPTER 20. CONFIGURING PERSISTENT STORAGE	300
20.1. OVERVIEW	300
20.2. PERSISTENT STORAGE USING NFS	300
20.2.1. Overview	300
20.2.2. Provisioning	300
20.2.3. Enforcing Disk Quotas	302
20.2.4. NFS Volume Security	302
20.2.4.1. Group IDs	303
20.2.4.2. User IDs	304
20.2.4.3. SELinux	305
20.2.4.4. Export Settings	305
20.2.5. Reclaiming Resources	306
20.2.6. Automation	307
20.2.7. Additional Configuration and Troubleshooting	307
20.3. PERSISTENT STORAGE USING GLUSTERFS	307
20.3.1. Overview	307
20.3.1.1. Containerized Red Hat Gluster Storage	307
20.3.1.2. Dedicated Storage Cluster	308
20.3.2. Support Requirements	309
20.3.2.1. Supported Operating Systems	309
20.3.2.2. Environment Requirements	309
20.3.3. Provisioning	310
20.3.3.1. Creating Gluster Endpoints	311
20.3.3.2. Creating the Persistent Volume	312
20.3.3.3. Creating the Persistent Volume Claim	313
20.3.4. Gluster Volume Security	314
20.3.4.1. Group IDs	314
20.3.4.2. User IDs	315
20.3.4.3. SELinux	316
20.4. PERSISTENT STORAGE USING OPENSTACK CINDER	316
20.4.1. Overview	316
20.4.2. Provisioning	317
20.4.2.1. Creating the Persistent Volume	317

20.4.2.2. Volume Format	318
20.5. PERSISTENT STORAGE USING CEPH RADOS BLOCK DEVICE (RBD)	318
20.5.1. Overview	318
20.5.2. Provisioning	319
20.5.2.1. Creating the Ceph Secret	319
20.5.2.2. Creating the Persistent Volume	320
20.5.3. Ceph Volume Security	321
20.6. PERSISTENT STORAGE USING AWS ELASTIC BLOCK STORE	322
20.6.1. Overview	322
20.6.2. Provisioning	323
20.6.2.1. Creating the Persistent Volume	323
20.6.2.2. Volume Format	324
20.7. PERSISTENT STORAGE USING GCE PERSISTENT DISK	324
20.7.1. Overview	324
20.7.2. Provisioning	324
20.7.2.1. Creating the Persistent Volume	325
20.7.2.2. Volume Format	326
20.7.2.3. Multi-zone Configuration	326
20.8. PERSISTENT STORAGE USING ISCSI	327
20.8.1. Overview	327
20.8.2. Provisioning	327
20.8.2.1. Enforcing Disk Quotas	328
20.8.2.2. iSCSI Volume Security	328
20.9. PERSISTENT STORAGE USING FIBRE CHANNEL	328
20.9.1. Overview	328
20.9.2. Provisioning	328
20.9.2.1. Enforcing Disk Quotas	329
20.9.2.2. Fibre Channel Volume Security	329
20.10. DYNAMICALLY PROVISIONING PERSISTENT VOLUMES	329
20.10.1. Overview	329
20.10.2. Enabling Provisioner Plug-ins	330
20.10.3. Requesting Dynamically Provisioned Storage	331
20.10.3.1. Volume Owner Information	331
20.10.4. Volume Recycling	333
20.11. VOLUME SECURITY	333
20.11.1. Overview	333
20.11.2. SCCs, Defaults, and Allowed Ranges	333
20.11.3. Supplemental Groups	337
20.11.4. fsGroup	340
20.11.5. User IDs	342
20.11.6. SELinux Options	344
20.12. SELECTOR-LABEL VOLUME BINDING	345
20.12.1. Overview	345
20.12.2. Motivation	345
20.12.3. Deployment	346
20.12.3.1. Prerequisites	346
20.12.3.2. Define the Persistent Volume and Claim	346
20.12.3.3. Deploy the Persistent Volume and Claim	347
CHAPTER 21. PERSISTENT STORAGE EXAMPLES	348
21.1. OVERVIEW	348
21.2. SHARING AN NFS MOUNT ACROSS TWO PERSISTENT VOLUME CLAIMS	348
21.2.1. Overview	348

21.2.2. Creating the Persistent Volume	348
21.2.3. Creating the Persistent Volume Claim	349
21.2.4. Ensuring NFS Volume Access	350
21.2.5. Creating the Pod	351
21.2.6. Creating an Additional Pod to Reference the Same PVC	354
21.3. COMPLETE EXAMPLE USING CEPH RBD	357
21.3.1. Overview	357
21.3.2. Installing the ceph-common Package	357
21.3.3. Creating the Ceph Secret	357
21.3.4. Creating the Persistent Volume	358
21.3.5. Creating the Persistent Volume Claim	359
21.3.6. Creating the Pod	360
21.3.7. Defining Group and Owner IDs (Optional)	361
21.4. COMPLETE EXAMPLE USING GLUSTERFS	361
21.4.1. Overview	361
21.4.2. Installing the glusterfs-fuse Package	361
21.4.3. Creating the Gluster Endpoints and Gluster Service for Persistence	362
21.4.4. Creating the Persistent Volume	363
21.4.5. Creating the Persistent Volume Claim	364
21.4.6. Defining GlusterFS Volume Access	365
21.4.7. Creating the Pod using NGINX Web Server image	366
21.5. BACKING DOCKER REGISTRY WITH GLUSTERFS STORAGE	370
21.5.1. Overview	370
21.5.2. Prerequisites	370
21.5.3. Create the Gluster Persistent Volume	371
21.5.4. Attach the PVC to the Docker Registry	371
21.5.5. Known Issues	372
21.5.5.1. Pod Cannot Resolve the Volume Host	372
21.6. BINDING PERSISTENT VOLUMES BY LABELS	373
21.6.1. Overview	373
21.6.1.1. Assumptions	373
21.6.2. Defining Specifications	373
21.6.2.1. Persistent Volume with Labels	373
21.6.2.2. Persistent Volume Claim with Selectors	374
21.6.2.3. Volume Endpoints	375
21.6.2.4. Deploy the PV, PVC, and Endpoints	375
CHAPTER 22. WORKING WITH HTTP PROXIES	376
22.1. OVERVIEW	376
22.2. CONFIGURING NO_PROXY	376
22.3. CONFIGURING HOSTS FOR PROXIES	377
22.4. CONFIGURING HOSTS FOR PROXIES USING ANSIBLE	377
22.5. PROXYING DOCKER PULL	378
22.6. USING MAVEN BEHIND A PROXY	379
22.7. CONFIGURING S2I BUILDS FOR PROXIES	379
22.8. CONFIGURING DEFAULT TEMPLATES FOR PROXIES	379
22.9. SETTING PROXY ENVIRONMENT VARIABLES IN PODS	379
22.10. GIT REPOSITORY ACCESS	380
CHAPTER 23. CONFIGURING GLOBAL BUILD DEFAULTS AND OVERRIDES	381
23.1. OVERVIEW	381
23.2. SETTING GLOBAL BUILD DEFAULTS	381
23.2.1. Configuring Global Build Defaults with Ansible	381

23.2.2. Manually Setting Global Build Defaults	382
23.3. SETTING GLOBAL BUILD OVERRIDES	383
CHAPTER 24. CONFIGURING PIPELINE EXECUTION	385
24.1. OVERVIEW	385
CHAPTER 25. CONFIGURING ROUTING	387
25.1. OVERVIEW	387
25.2. CONFIGURING ROUTE TIMEOUTS	387
25.3. CONFIGURING NATIVE CONTAINER ROUTING	387
25.3.1. Network Overview	387
CHAPTER 26. ROUTING FROM EDGE LOAD BALANCERS	390
26.1. OVERVIEW	390
26.2. INCLUDING THE LOAD BALANCER IN THE SDN	390
26.3. ESTABLISHING A TUNNEL USING A RAMP NODE	390
26.3.1. Configuring a Highly-Available Ramp Node	393
CHAPTER 27. AGGREGATING CONTAINER LOGS	394
27.1. OVERVIEW	394
27.2. PRE-DEPLOYMENT CONFIGURATION	394
27.3. SPECIFYING DEPLOYER PARAMETERS	395
27.4. DEPLOYING THE EFK STACK	398
27.5. UNDERSTANDING AND ADJUSTING THE DEPLOYMENT	399
27.5.1. Ops Cluster	399
27.5.2. Elasticsearch	399
27.5.3. Fluentd	404
27.5.4. Kibana	409
27.5.5. Curator	409
27.5.5.1. Creating the Curator Configuration	411
27.6. CLEANUP	411
27.7. UPGRADING	411
27.8. TROUBLESHOOTING KIBANA	412
27.9. SENDING LOGS TO AN EXTERNAL ELASTICSEARCH INSTANCE	413
27.10. PERFORMING ADMINISTRATIVE ELASTICSEARCH OPERATIONS	414
CHAPTER 28. AGGREGATE LOGGING SIZING GUIDELINES	416
28.1. OVERVIEW	416
28.2. INSTALLATION	416
28.3. SYSTEMD-JOURNALD AND RSYSLOG	418
28.4. SCALING UP EFK LOGGING	419
28.5. STORAGE CONSIDERATIONS	420
CHAPTER 29. ENABLING CLUSTER METRICS	422
29.1. OVERVIEW	422
29.2. BEFORE YOU BEGIN	422
29.3. SERVICE ACCOUNTS	422
29.3.1. Metrics Deployer Service Account	423
29.3.2. Heapster Service Account	423
29.4. METRICS DATA STORAGE	423
29.4.1. Persistent Storage	423
29.4.2. Capacity Planning for Cluster Metrics	424
29.4.3. Non-Persistent Storage	425
29.5. METRICS DEPLOYER	426
29.5.1. Using Secrets	426

29.5.1.1. Providing Your Own Certificates	426
29.5.1.2. Using the Router's Default Certificate	426
29.5.1.3. Deployer Secret Options	427
29.5.2. Modifying the Deployer Template	427
29.5.2.1. Deployer Template Parameters	427
29.6. DEPLOYING THE METRIC COMPONENTS	429
29.6.1. Metrics Deployer Validations	430
29.7. SETTING THE METRICS PUBLIC URL	431
29.8. ACCESSING HAWKULAR METRICS DIRECTLY	432
29.8.1. OpenShift Container Platform Projects and Hawkular Tenants	432
29.8.2. Authorization	432
29.9. SCALING OPENSIFT CONTAINER PLATFORM METRICS PODS	432
29.9.1. Prerequisites	433
29.9.2. Scaling the Cassandra Components	433
29.10. CLEANUP	434
CHAPTER 30. CUSTOMIZING THE WEB CONSOLE	435
30.1. OVERVIEW	435
30.2. LOADING EXTENSION SCRIPTS AND STYLESHEETS	435
30.2.1. Setting Extension Properties	436
30.2.2. Customizing the Logo	436
30.2.3. Changing Links to Documentation	437
30.2.4. Adding or Changing Links to Download the CLI	437
30.2.5. Customizing the About Page	437
30.2.6. Configuring Navigation Menus	438
30.2.6.1. Top Navigation Dropdown Menus	438
30.2.6.2. Project Left Navigation	439
30.2.7. Enabling Features in Technology Preview	441
30.3. SERVING STATIC FILES	442
30.3.1. Enabling HTML5 Mode	442
30.4. CUSTOMIZING THE LOGIN PAGE	442
30.4.1. Example Usage	443
30.5. CUSTOMIZING THE OAUTH ERROR PAGE	443
30.6. CHANGING THE LOGOUT URL	443
30.7. CONFIGURING WEB CONSOLE CUSTOMIZATIONS WITH ANSIBLE	444
CHAPTER 31. REVISION HISTORY: INSTALLATION AND CONFIGURATION	446
31.1. WED MAR 07 2018	446
31.2. FRI AUG 25 2017	446
31.3. TUE AUG 08 2017	446
31.4. FRI JUL 28 2017	446
31.5. THU JUL 27 2017	446
31.6. WED JUL 12 2017	447
31.7. TUE JUN 13 2017	447
31.8. WED MAY 31 2017	447
31.9. THU MAY 25 2017	447
31.10. MON MAY 15 2017	448
31.11. TUE MAY 09 2017	448
31.12. TUE APR 25 2017	448
31.13. WED APR 12 2017	448
31.14. MON APR 03 2017	449
31.15. MON MAR 27 2017	449
31.16. MON MAR 20 2017	450

31.17. TUE MAR 14 2017	450
31.18. TUE MAR 07 2017	451
31.19. THU FEB 16 2017	451
31.20. MON FEB 06 2017	451
31.21. TUE JAN 31 2017	452
31.22. MON JAN 30 2017	453
31.23. WED JAN 25 2017	453
31.24. WED JAN 18 2017	453
31.25. MON JAN 16 2017	454
31.26. MON JAN 09 2017	454
31.27. TUE DEC 20 2016	454
31.28. TUE DEC 13 2016	454
31.29. MON DEC 05 2016	455
31.30. MON NOV 21 2016	455
31.31. MON NOV 14 2016	455
31.32. MON NOV 07 2016	456
31.33. TUE NOV 01 2016	456
31.34. THU OCT 27 2016	456
31.35. MON OCT 17 2016	456
31.36. TUE OCT 11 2016	457
31.37. TUE OCT 04 2016	457
31.38. TUE SEP 27 2016	458

CHAPTER 1. OVERVIEW

OpenShift Container Platform Installation and Configuration topics cover the basics of installing and configuring OpenShift Container Platform in your environment. Configuration, management, and logging are also covered. Use these topics for the one-time tasks required to quickly set up your OpenShift Container Platform environment and configure it based on your organizational needs.

For day to day cluster administration tasks, see [Cluster Administration](#).

CHAPTER 2. INSTALLING A CLUSTER

2.1. PLANNING

2.1.1. Initial Planning

For production environments, several factors influence installation. Consider the following questions as you read through the documentation:

- *Which installation method do you want to use?* The [Installation Methods](#) section provides some information about the quick and advanced installation methods.
- *How many hosts do you require in the cluster?* The [Environment Scenarios](#) section provides multiple examples of Single Master and Multiple Master configurations.
- *How many pods are required in your cluster?* The [Sizing Considerations](#) section provides limits for nodes and pods so you can calculate how large your environment needs to be.
- *Is [high availability](#) required?* High availability is recommended for fault tolerance. In this situation, you might aim to use the [Multiple Masters Using Native HA](#) example as a basis for your environment.
- *Which installation type do you want to use: [RPM](#) or [containerized](#)?* Both installations provide a working OpenShift Container Platform environment, but you might have a preference for a particular method of installing, managing, and updating your services.
- *Is my installation supported if integrating with other technologies?* See the [OpenShift Container Platform Tested Integrations](#) for a list of tested integrations.

2.1.2. Installation Methods

Both the quick and advanced installation methods are supported for development and production environments. If you want to quickly get OpenShift Container Platform up and running to try out for the first time, use the quick installer and let the interactive CLI guide you through the configuration options relevant to your environment.

For the most control over your cluster's configuration, you can use the advanced installation method. This method is particularly suited if you are already familiar with Ansible. However, following along with the OpenShift Container Platform documentation should equip you with enough information to reliably deploy your cluster and continue to manage its configuration post-deployment using the provided Ansible playbooks directly.

If you install initially using the quick installer, you can always further tweak your cluster's configuration and adjust the number of hosts in the cluster using the same installer tool. If you wanted to later switch to using the advanced method, you can create an inventory file for your configuration and carry on that way.

2.1.3. Sizing Considerations

Determine how many nodes and pods you require for your OpenShift Container Platform cluster. Cluster scalability correlates to the number of pods in a cluster environment. That number influences the other numbers in your setup.

The following table provides the maximum sizing limits for nodes and pods:

Type	Maximum
Maximum nodes per cluster	1000
Maximum pods per cluster	120,000
Maximum pods per node	250
Maximum pods per core	10



IMPORTANT

Oversubscribing the physical resources on a node affects resource guarantees the Kubernetes scheduler makes during pod placement. Learn what measures you can take to [avoid memory swapping](#).

Determine how many pods are expected to fit per node:

$$\text{Maximum Pods per Cluster} / \text{Expected Pods per Node} = \text{Total Number of Nodes}$$

Example Scenario

If you want to scope your cluster for 2200 pods per cluster, you would need at least 9 nodes, assuming that there are 250 maximum pods per node:

$$2200 / 250 = 8.8$$

If you increase the number of nodes to 20, then the pod distribution changes to 110 pods per node:

$$2200 / 20 = 110$$

2.1.4. Environment Scenarios

This section outlines different examples of scenarios for your OpenShift Container Platform environment. Use these scenarios as a basis for planning your own OpenShift Container Platform cluster.



NOTE

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

2.1.4.1. Single Master and Multiple Nodes

The following table describes an example environment for a single [master](#) (with embedded **etcd**) and two [nodes](#):

Host Name	Infrastructure Component to Install
master.example.com	Master and node

Host Name	Infrastructure Component to Install
node1.example.com	Node
node2.example.com	

2.1.4.2. Single Master, Multiple etcd, and Multiple Nodes

The following table describes an example environment for a single [master](#), three [etcd](#) hosts, and two [nodes](#):

Host Name	Infrastructure Component to Install
master.example.com	Master and node
etcd1.example.com	etcd
etcd2.example.com	
etcd3.example.com	
node1.example.com	Node
node2.example.com	



NOTE

When specifying multiple **etcd** hosts, external **etcd** is installed and configured. Clustering of OpenShift Container Platform's embedded **etcd** is not supported.

2.1.4.3. Multiple Masters Using Native HA

The following describes an example environment for three [masters](#), one HAProxy load balancer, three [etcd](#) hosts, and two [nodes](#) using the **native** HA method:

Host Name	Infrastructure Component to Install
master1.example.com	Master (clustered using native HA) and node
master2.example.com	
master3.example.com	
lb.example.com	HAProxy to load balance API master endpoints
etcd1.example.com	etcd

Host Name	Infrastructure Component to Install
etcd2.example.com	
etcd3.example.com	
node1.example.com	Node
node2.example.com	

**NOTE**

When specifying multiple **etcd** hosts, external **etcd** is installed and configured. Clustering of OpenShift Container Platform's embedded **etcd** is not supported.

2.1.4.4. Stand-alone Registry

You can also install OpenShift Container Platform to act as a stand-alone registry using the OpenShift Container Platform's integrated registry. See [Installing a Stand-alone Registry](#) for details on this scenario.

2.1.5. RPM vs Containerized

An RPM installation installs all services through package management and configures services to run within the same user space, while a containerized installation configures installs services using container images and runs separate services in individual containers.

The default method for installing OpenShift Container Platform on Red Hat Enterprise Linux (RHEL) uses RPMs. Alternatively, you can use the containerized method, which deploys containerized OpenShift Container Platform master and node components. When targeting a RHEL Atomic Host system, the containerized method is the only available option, and is automatically selected for you based on the detection of the */run/ostree-booted* file.

The following table outlines the differences between the RPM and Containerized methods:

Type	RPM	Containerized
Installation Method	Packages via yum	Container images via docker
Service Management	systemd	docker and systemd units
Operating System	Red Hat Enterprise Linux	Red Hat Enterprise Linux or Red Hat Atomic Host

The [Containerized Installation Preparation](#) section provides more details on configuring your installation to use containerized services.


2.2. PREREQUISITES

2.2.1. System Requirements

The following sections identify the hardware specifications and system-level requirements of all hosts within your OpenShift Container Platform environment.

2.2.1.1. Red Hat Subscriptions

You must have an active OpenShift Container Platform subscription on your Red Hat account to proceed. If you do not, contact your sales representative for more information.



IMPORTANT

OpenShift Container Platform 3.3 requires Docker 1.10.

2.2.1.2. Minimum Hardware Requirements

The system requirements vary per host type:

Masters	<ul style="list-style-type: none"> • Physical or virtual system, or an instance running on a public or private IaaS. • Base OS: RHEL 7.1 or later with "Minimal" installation option, or RHEL Atomic Host 7.2.6 or later. • 2 vCPU. • Minimum 16 GB RAM. • Minimum 40 GB hard disk space for the file system containing <code>/var/</code>.
Nodes	<ul style="list-style-type: none"> • Physical or virtual system, or an instance running on a public or private IaaS. • Base OS: RHEL 7.1 or later with "Minimal" installation option, or RHEL Atomic Host 7.2.6 or later. • NetworkManager 1.0 or later • 1 vCPU. • Minimum 8 GB RAM. • Minimum 15 GB hard disk space for the file system containing <code>/var/</code>. • An additional minimum 15 GB unallocated space to be used for Docker's storage back end; see Configuring Docker Storage.
Separate etcd Nodes	<ul style="list-style-type: none"> • Minimum 20 GB hard disk space for etcd data. • Consult Hardware Recommendations to properly size your etcd nodes. • Currently, OpenShift Container Platform stores image, build, and deployment metadata in etcd. You must periodically prune old resources. If you are planning to leverage a large number of images/builds/deployments, place etcd on machines with large amounts of memory and fast SSD drives.

**IMPORTANT**

OpenShift Container Platform only supports servers with the x86_64 architecture.

**NOTE**

Meeting the `/var/` file system sizing requirements in RHEL Atomic Host requires making changes to the default configuration. See [Managing Storage in Red Hat Enterprise Linux Atomic Host](#) for instructions on configuring this during or after installation.

2.2.1.3. Production Level Hardware Requirements

Test or sample environments function with the minimum requirements. For production environments, the following recommendations apply:

Master Hosts

In a highly available OpenShift Container Platform cluster with a separate etcd cluster, a master host should have, in addition to the minimum requirements in the table above, 1 CPU core and 1.5 GB of memory for each 1000 pods. Therefore, the recommended size of a master host in an OpenShift Container Platform cluster of 2000 pods would be the minimum requirements of 2 CPU cores and 16 GB of RAM, plus 2 CPU cores and 3 GB of RAM, totaling 4 CPU cores and 19 GB of RAM.

When planning an environment with multiple masters, a minimum of three etcd hosts and a load-balancer between the master hosts are required.

The OpenShift Container Platform master caches deserialized versions of resources aggressively to ease CPU load. However, in smaller clusters of less than 1000 pods, this cache can waste a lot of memory for negligible CPU load reduction. The default cache size is 50,000 entries, which, depending on the size of your resources, can grow to occupy 1 to 2 GB of memory. This cache size can be reduced using the following setting in `/etc/origin/master/master-config.yaml`:

```
kubernetesMasterConfig:
  apiServerArguments:
    deserialization-cache-size:
      - "1000"
```


Node Hosts

The size of a node host depends on the expected size of its workload. As an OpenShift Container Platform cluster administrator, you will need to calculate the expected workload, then add about 10 percent for overhead. For production environments, allocate enough resources so that a node host failure does not affect your maximum capacity.

Use the above with the following table to plan the maximum loads for nodes and pods:

Host	Sizing Recommendation
Maximum nodes per cluster	1000
Maximum pods per cluster	120000
Maximum pods per nodes	250

Host	Sizing Recommendation
Maximum pods per core	10



IMPORTANT

Oversubscribing the physical resources on a node affects resource guarantees the Kubernetes scheduler makes during pod placement. Learn what measures you can take to [avoid memory swapping](#).

2.2.1.4. Configuring Core Usage

By default, OpenShift Container Platform masters and nodes use all available cores in the system they run on. You can choose the number of cores you want OpenShift Container Platform to use by setting the [GOMAXPROCS environment variable](#).

For example, run the following before starting the server to make OpenShift Container Platform only run on one core:

```
# export GOMAXPROCS=1
```

2.2.1.5. SELinux

Security-Enhanced Linux (SELinux) must be enabled on all of the servers before installing OpenShift Container Platform or the installer will fail. Also, configure **SELINUXTYPE=targeted** in the */etc/selinux/config* file:

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#     enforcing - SELinux security policy is enforced.
#     permissive - SELinux prints warnings instead of enforcing.
#     disabled - No SELinux policy is loaded.
SELINUX=enforcing
# SELINUXTYPE= can take one of these three values:
#     targeted - Targeted processes are protected,
#     minimum - Modification of targeted policy. Only selected processes
are protected.
#     mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

2.2.1.6. NTP

You must enable Network Time Protocol (NTP) to prevent masters and nodes in the cluster from going out of sync. Set **openshift_clock_enabled** to **true** in the Ansible playbook to enable NTP on masters and nodes in the cluster during Ansible installation.

```
# openshift_clock_enabled=true
```

2.2.1.7. Security Warning

OpenShift Container Platform runs [containers](#) on your hosts, and in some cases, such as build

operations and the registry service, it does so using privileged containers. Furthermore, those containers access your host's Docker daemon and perform **docker build** and **docker push** operations. As such, you should be aware of the inherent security risks associated with performing **docker run** operations on arbitrary images as they effectively have root access.

For more information, see these articles:

- <http://opensource.com/business/14/7/docker-security-selinux>
- <https://docs.docker.com/engine/security/security/>


To address these risks, OpenShift Container Platform uses [security context constraints](#) that control the actions that pods can perform and what it has the ability to access.

2.2.2. Environment Requirements

The following section defines the requirements of the environment containing your OpenShift Container Platform configuration. This includes networking considerations and access to external services, such as Git repository access, storage, and cloud infrastructure providers.

2.2.2.1. DNS

OpenShift Container Platform requires a fully functional DNS server in the environment. This is ideally a separate host running DNS software and can provide name resolution to hosts and containers running on the platform.



IMPORTANT

Adding entries into the **/etc/hosts** file on each host is not enough. This file is not copied into containers running on the platform.

Key components of OpenShift Container Platform run themselves inside of containers and use the following process for name resolution:

1. By default, containers receive their DNS configuration file (**/etc/resolv.conf**) from their host.
2. OpenShift Container Platform then inserts one DNS value into the pods (above the node's nameserver values). That value is defined in the **/etc/origin/node/node-config.yaml** file by the **dnsIP** parameter, which by default is set to the address of the host node because the host is using **dnsmasq**.
3. If the **dnsIP** parameter is omitted from the **node-config.yaml** file, then the value defaults to the kubernetes service IP, which is the first nameserver in the pod's **/etc/resolv.conf** file.

As of OpenShift Container Platform 3.2, **dnsmasq** is automatically configured on all masters and nodes. The pods use the nodes as their DNS, and the nodes forward the requests. By default, **dnsmasq** is configured on the nodes to listen on port 53, therefore the nodes cannot run any other type of DNS application.

**NOTE**

Previously, in OpenShift Container Platform 3.1, a DNS server could not be installed on a master node, because it ran its own internal DNS server. Now, with master nodes using **dnsmasq**, SkyDNS is now configured to listen on port 8053 so that **dnsmasq** can run on the masters. Note that these DNS changes (**dnsmasq** configured on nodes and the SkyDNS port change) only apply to new installations of OpenShift Container Platform 3.2. Clusters upgraded to OpenShift Container Platform 3.2 from a previous version do not currently have these changes applied during the upgrade process.

**NOTE**

NetworkManager is required on the nodes in order to populate **dnsmasq** with the DNS IP addresses.

The following is an example set of DNS records for the [Single Master and Multiple Nodes](#) scenario:

```
master    A    10.64.33.100
node1     A    10.64.33.101
node2     A    10.64.33.102
```

If you do not have a properly functioning DNS environment, you could experience failure with:

- Product installation via the reference Ansible-based scripts
- Deployment of the infrastructure containers (registry, routers)
- Access to the OpenShift Container Platform web console, because it is not accessible via IP address alone

2.2.2.1.1. Configuring Hosts to Use DNS

Make sure each host in your environment is configured to resolve hostnames from your DNS server. The configuration for hosts' DNS resolution depend on whether DHCP is enabled. If DHCP is:

- Disabled, then configure your network interface to be static, and add DNS nameservers to NetworkManager.
- Enabled, then the NetworkManager dispatch script automatically configures DNS based on the DHCP configuration. Optionally, you can add a value to **dnsIP** in the **node-config.yaml** file to prepend the pod's **resolv.conf** file. The second nameserver is then defined by the host's first nameserver. By default, this will be the IP address of the node host.

**NOTE**

For most configurations, do not set the **openshift_dns_ip** option during the advanced installation of OpenShift Container Platform (using Ansible), because this option overrides the default IP address set by **dnsIP**.

Instead, allow the installer to configure each node to use **dnsmasq** and forward requests to SkyDNS or the external DNS provider. If you do set the **openshift_dns_ip** option, then it should be set either with a DNS IP that queries SkyDNS first, or to the SkyDNS service or endpoint IP (the Kubernetes service IP).

To verify that hosts can be resolved by your DNS server:

1. Check the contents of **/etc/resolv.conf**:

```
$ cat /etc/resolv.conf
# Generated by NetworkManager
search example.com
nameserver 10.64.33.1
# nameserver updated by /etc/NetworkManager/dispatcher.d/99-origin-
dns.sh
```

In this example, 10.64.33.1 is the address of our DNS server.

2. Test that the DNS servers listed in **/etc/resolv.conf** are able to resolve host names to the IP addresses of all masters and nodes in your OpenShift Container Platform environment:

```
$ dig <node_hostname> @<IP_address> +short
```

For example:

```
$ dig master.example.com @10.64.33.1 +short
10.64.33.100
$ dig node1.example.com @10.64.33.1 +short
10.64.33.101
```

2.2.2.1.2. Disabling dnsmasq

If you want to disable **dnsmasq** (for example, if your **/etc/resolv.conf** is managed by a configuration tool other than NetworkManager), then set **openshift_use_dnsmasq** to **false** in the Ansible playbook.

However, certain containers do not properly move to the next nameserver when the first issues **SERVFAIL**. Red Hat Enterprise Linux (RHEL)-based containers do not suffer from this, but certain versions of **uclibc** and **musl** do.

2.2.2.1.3. Configuring a DNS Wildcard

Optionally, configure a wildcard for the router to use, so that you do not need to update your DNS configuration when new routes are added.

A wildcard for a DNS zone must ultimately resolve to the IP address of the OpenShift Container Platform [router](#).

For example, create a wildcard DNS entry for **cloudapps** that has a low time-to-live value (TTL) and points to the public IP address of the host where the router will be deployed:

```
*.cloudapps.example.com. 300 IN A 192.168.133.2
```

In almost all cases, when referencing VMs you must use host names, and the host names that you use must match the output of the **hostname -f** command on each node.

**WARNING**

In your `/etc/resolv.conf` file on each node host, ensure that the DNS server that has the wildcard entry is not listed as a nameserver or that the wildcard domain is not listed in the search list. Otherwise, containers managed by OpenShift Container Platform may fail to resolve host names properly.

2.2.2.2. Network Access

A shared network must exist between the master and node hosts. If you plan to configure [multiple masters for high-availability](#) using the [advanced installation method](#), you must also select an IP to be configured as your [virtual IP](#) (VIP) during the installation process. The IP that you select must be routable between all of your nodes, and if you configure using a FQDN it should resolve on all nodes.

2.2.2.2.1. NetworkManager

NetworkManager, a program for providing detection and configuration for systems to automatically connect to the network, is required.

2.2.2.2.2. Required Ports

The OpenShift Container Platform installation automatically creates a set of internal firewall rules on each host using `iptables`. However, if your network configuration uses an external firewall, such as a hardware-based firewall, you must ensure infrastructure components can communicate with each other through specific ports that act as communication endpoints for certain processes or services.

Ensure the following ports required by OpenShift Container Platform are open on your network and configured to allow access between hosts. Some ports are optional depending on your configuration and usage.

Table 2.1. Node to Node

4789	UDP	Required for SDN communication between pods on separate hosts.
-------------	-----	--

Table 2.2. Nodes to Master

53 or 8053	TCP/ UDP	Required for DNS resolution of cluster services (SkyDNS). Installations prior to 3.2 or environments upgraded to 3.2 use port 53. New installations will use 8053 by default so that dnsmasq may be configured.
4789	UDP	Required for SDN communication between pods on separate hosts.
443 or 8443	TCP	Required for node hosts to communicate to the master API, for the node hosts to post back status, to receive tasks, and so on.

Table 2.3. Master to Node

4789	UDP	Required for SDN communication between pods on separate hosts.
10250	TCP	The master proxies to node hosts via the Kubelet for oc commands.

**NOTE**

In the following table, **(L)** indicates the marked port is also used in *loopback mode*, enabling the master to communicate with itself.

In a single-master cluster:

- Ports marked with **(L)** must be open.
- Ports not marked with **(L)** need not be open.

In a multiple-master cluster, all the listed ports must be open.

Table 2.4. Master to Master

53 (L) or 8053 (L)	TCP/ UDP	Required for DNS resolution of cluster services (SkyDNS). Installations prior to 3.2 or environments upgraded to 3.2 use port 53. New installations will use 8053 by default so that dnsmasq may be configured.
2049 (L)	TCP/ UDP	Required when provisioning an NFS host as part of the installer.
2379	TCP	Used for standalone etcd (clustered) to accept changes in state.
2380	TCP	etcd requires this port be open between masters for leader election and peering connections when using standalone etcd (clustered).
4001 (L)	TCP	Used for embedded etcd (non-clustered) to accept changes in state.
4789 (L)	UDP	Required for SDN communication between pods on separate hosts.

Table 2.5. External to Load Balancer

9000	TCP	If you choose the native HA method, optional to allow access to the HAProxy statistics page.
------	-----	---

Table 2.6. External to Master

443 or 8443	TCP	Required for node hosts to communicate to the master API, for node hosts to post back status, to receive tasks, and so on.
-------------	-----	--

Table 2.7. IaaS Deployments

22	TCP	Required for SSH by the installer or system administrator.
----	-----	--

53 or 8053	TCP/ UDP	Required for DNS resolution of cluster services (SkyDNS). Installations prior to 3.2 or environments upgraded to 3.2 use port 53. New installations will use 8053 by default so that dnsmasq may be configured. Only required to be internally open on master hosts.
80 or 443	TCP	For HTTP/HTTPS use for the router. Required to be externally open on node hosts, especially on nodes running the router.
1936	TCP	For router statistics use. Required to be open when running the template router to access statistics, and can be open externally or internally to connections depending on if you want the statistics to be expressed publicly.
4001	TCP	For embedded etcd (non-clustered) use. Only required to be internally open on the master host. 4001 is for server-client connections.
2379 and 2380	TCP	For standalone etcd use. Only required to be internally open on the master host. 2379 is for server-client connections. 2380 is for server-server connections, and is only required if you have clustered etcd.
4789	UDP	For VxLAN use (OpenShift Container Platform SDN). Required only internally on node hosts.
8443	TCP	For use by the OpenShift Container Platform web console, shared with the API server.
10250	TCP	For use by the Kubelet. Required to be externally open on nodes.

Notes

- In the above examples, port **4789** is used for User Datagram Protocol (UDP).
- When deployments are using the SDN, the pod network is accessed via a service proxy, unless it is accessing the registry from the same node the registry is deployed on.
- OpenShift Container Platform internal DNS cannot be received over SDN. Depending on the detected values of **openshift_facts**, or if the **openshift_ip** and **openshift_public_ip** values are overridden, it will be the computed value of **openshift_ip**. For non-cloud deployments, this will default to the IP address associated with the default route on the master host. For cloud deployments, it will default to the IP address associated with the first internal interface as defined by the cloud metadata.
- The master host uses port **10250** to reach the nodes and does not go over SDN. It depends on the target host of the deployment and uses the computed values of **openshift_hostname** and **openshift_public_hostname**.

Table 2.8. Aggregated Logging

9200	TCP	For Elasticsearch API use. Required to be internally open on any infrastructure nodes so Kibana is able to retrieve logs for display. It can be externally opened for direct access to Elasticsearch by means of a route. The route can be created using oc expose .
-------------	-----	---

9300	TCP	For Elasticsearch inter-cluster use. Required to be internally open on any infrastructure node so the members of the Elasticsearch cluster may communicate with each other.
-------------	-----	---

2.2.2.3. Persistent Storage

The Kubernetes [persistent volume](#) framework allows you to provision an OpenShift Container Platform cluster with persistent storage using networked storage available in your environment. This can be done after completing the initial OpenShift Container Platform installation depending on your application needs, giving users a way to request those resources without having any knowledge of the underlying infrastructure.

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 Elastic Block Store \(EBS\)](#), [GCE Persistent Disks](#), and [iSCSI](#).

2.2.2.4. Cloud Provider Considerations

There are certain aspects to take into consideration if installing OpenShift Container Platform on a cloud provider.

2.2.2.4.1. Configuring a Security Group

When installing on AWS or OpenStack, ensure that you set up the appropriate security groups. These are some ports that you should have in your security groups, without which the installation will fail. You may need more depending on the cluster configuration you want to install. For more information and to adjust your security groups accordingly, see [Required Ports](#) for more information.

All OpenShift Container Platform Hosts	<ul style="list-style-type: none"> • tcp/22 from host running the installer/Ansible
etcd Security Group	<ul style="list-style-type: none"> • tcp/2379 from masters • tcp/2380 from etcd hosts
Master Security Group	<ul style="list-style-type: none"> • tcp/8443 from 0.0.0.0/0 • tcp/53 from all OpenShift Container Platform hosts for environments installed prior to or upgraded to 3.2 • udp/53 from all OpenShift Container Platform hosts for environments installed prior to or upgraded to 3.2 • tcp/8053 from all OpenShift Container Platform hosts for new environments installed with 3.2 • udp/8053 from all OpenShift Container Platform hosts for new environments installed with 3.2

Node Security Group	<ul style="list-style-type: none"> • tcp/10250 from masters • udp/4789 from nodes
Infrastructure Nodes (ones that can host the OpenShift Container Platform router)	<ul style="list-style-type: none"> • tcp/443 from 0.0.0.0/0 • tcp/80 from 0.0.0.0/0

If configuring ELBs for load balancing the masters and/or routers, you also need to configure Ingress and Egress security groups for the ELBs appropriately.

2.2.2.4.2. Overriding Detected IP Addresses and Host Names

Some deployments require that the user override the detected host names and IP addresses for the hosts. To see the default values, run the **openshift_facts** playbook:

```
# ansible-playbook playbooks/byo/openshift_facts.yml
```

Now, verify the detected common settings. If they are not what you expect them to be, you can override them.

The [Advanced Installation](#) topic discusses the available Ansible variables in greater detail.

Variable	Usage
hostname	<ul style="list-style-type: none"> • Should resolve to the internal IP from the instances themselves. • openshift_hostname overrides.
ip	<ul style="list-style-type: none"> • Should be the internal IP of the instance. • openshift_ip will overrides.
public_hostname	<ul style="list-style-type: none"> • Should resolve to the external IP from hosts outside of the cloud. • Provider openshift_public_hostname overrides.
public_ip	<ul style="list-style-type: none"> • Should be the externally accessible IP associated with the instance. • openshift_public_ip overrides.

Variable	Usage
use_openshift_sdn	<ul style="list-style-type: none"> Should be true unless the cloud is GCE. openshift_use_openshift_sdn overrides.

**WARNING**

If **openshift_hostname** is set to a value other than the metadata-provided **private-dns-name** value, the native cloud integration for those providers will no longer work.

In AWS, situations that require overriding the variables include:

Variable	Usage
hostname	The user is installing in a VPC that is not configured for both DNS hostnames and DNS resolution .
ip	Possibly if they have multiple network interfaces configured and they want to use one other than the default. You must first set openshift_set_node_ip to True . Otherwise, the SDN would attempt to use the hostname setting or try to resolve the host name for the IP.
public_hostname	<ul style="list-style-type: none"> A master instance where the VPC subnet is not configured for Auto-assign Public IP. For external access to this master, you need to have an ELB or other load balancer configured that would provide the external access needed, or you need to connect over a VPN connection to the internal name of the host. A master instance where metadata is disabled. This value is not actually used by the nodes.
public_ip	<ul style="list-style-type: none"> A master instance where the VPC subnet is not configured for Auto-assign Public IP. A master instance where metadata is disabled. This value is not actually used by the nodes.

If setting **openshift_hostname** to something other than the metadata-provided **private-dns-name** value, the native cloud integration for those providers will no longer work.

For EC2 hosts in particular, they must be deployed in a VPC that has both **DNS host names** and **DNS resolution** enabled, and **openshift_hostname** should not be overridden.

2.2.2.4.3. Post-Installation Configuration for Cloud Providers

Following the installation process, you can configure OpenShift Container Platform for [AWS](#), [OpenStack](#), or [GCE](#).

2.3. HOST PREPARATION

2.3.1. Operating System Requirements

A base installation of RHEL 7.1 or later or RHEL Atomic Host 7.2.6 or later is required for master and node hosts. See the following documentation for the respective installation instructions, if required:

- [Red Hat Enterprise Linux 7 Installation Guide](#)
- [Red Hat Enterprise Linux Atomic Host 7 Installation and Configuration Guide](#)

2.3.2. Host Registration

Each host must be registered using Red Hat Subscription Manager (RHSM) and have an active OpenShift Container Platform subscription attached to access the required packages.

1. On each host, register with RHSM:

```
# subscription-manager register --username=<user_name> --password=
<password>
```

2. List the available subscriptions:

```
# subscription-manager list --available --matches '*OpenShift*'
```

3. In the output for the previous command, find the pool ID for an OpenShift Container Platform subscription and attach it:

```
# subscription-manager attach --pool=<pool_id>
```

4. Disable all yum repositories:

- a. Disable all the enabled RHSM repositories:

```
# subscription-manager repos --disable="*"
```

- b. List the remaining yum repositories and note their names under **repo id**, if any:

```
# yum repolist
```

- c. Use **yum-config-manager** to disable the remaining yum repositories:

■

```
# yum-config-manager --disable <repo_id>
```

Alternatively, disable all repositories:

```
yum-config-manager --disable \*
```

Note that this could take a few minutes if you have a large number of available repositories

5. Enable only the repositories required by OpenShift Container Platform 3.3:

```
# subscription-manager repos \
  --enable="rhel-7-server-rpms" \
  --enable="rhel-7-server-extras-rpms" \
  --enable="rhel-7-server-ose-3.3-rpms"
```

2.3.3. Installing Base Packages

For RHEL 7 systems:

1. Install the following base packages:

```
# yum install wget git net-tools bind-utils yum-utils iptables-
services bridge-utils bash-completion kexec-tools sos psacct
```

2. Update the system to the latest packages:

```
# yum update
```

3. Install the following package, which provides OpenShift Container Platform utilities and pulls in other tools required by the [quick](#) and [advanced installation](#) methods, such as Ansible and related configuration files:

```
# yum install atomic-openshift-utils
```

4. Downgrade the Ansible version 2.3.2. The 2.4 version installed by the **atomic-openshift-utils** package will cause the installer to fail. For example:

```
# yum downgrade ansible-2.3.2.0-2.el7.noarch
```

5. Install the following ***-excluder** packages on each RHEL 7 system, which helps ensure your systems stay on the correct versions of **atomic-openshift** and **docker** packages when you are not trying to upgrade, according to the OpenShift Container Platform version:

```
# yum install atomic-openshift-excluder atomic-openshift-docker-
excluder
```

6. The ***-excluder** packages add entries to the **exclude** directive in the host's **/etc/yum.conf** file when installed. Run the following command on each host to remove the **atomic-openshift** packages from the list for the duration of the installation.

```
# atomic-openshift-excluder unexclude
```

For RHEL Atomic Host 7 systems:

1. Ensure the host is up to date by upgrading to the latest Atomic tree if one is available:

```
# atomic host upgrade
```

2. After the upgrade is completed and prepared for the next boot, reboot the host:

```
# systemctl reboot
```

2.3.4. Installing Docker

At this point, you should install Docker on all master and node hosts. This allows you to configure your [Docker storage options](#) before installing OpenShift Container Platform.

1. For RHEL 7 systems, install Docker 1.10.



NOTE

On RHEL Atomic Host 7 systems, Docker should already be installed, configured, and running by default.

The **atomic-openshift-docker-excluder** package that was installed in [Installing Base Packages](#) should ensure that the correct version of Docker is installed in this step:

```
# yum install docker
```

After the package installation is complete, verify that version 1.10.3 was installed:

```
# docker version
```

2. Edit the `/etc/sysconfig/docker` file and add **--insecure-registry 172.30.0.0/16** to the **OPTIONS** parameter. For example:

```
OPTIONS='--selinux-enabled --insecure-registry 172.30.0.0/16'
```

If using the [Quick Installation](#) method, you can easily script a complete installation from a kickstart or cloud-init setup, change the default configuration file:


```
# sed -i '/OPTIONS=.*c\OPTIONS="--selinux-enabled --insecure-registry 172.30.0.0/16"' \
/etc/sysconfig/docker
```



NOTE

The [Advanced Installation](#) method automatically changes `/etc/sysconfig/docker`.

The **--insecure-registry** option instructs the Docker daemon to trust any Docker registry on the indicated subnet, rather than [requiring a certificate](#).



IMPORTANT

172.30.0.0/16 is the default value of the **servicesSubnet** variable in the **master-config.yaml** file. If this has changed, then the **--insecure-registry** value in the above step should be adjusted to match, as it is indicating the subnet for the registry to use. Note that the **openshift_portal_net** variable can be set in the Ansible inventory file and used during the [advanced installation](#) method to modify the **servicesSubnet** variable.



NOTE

After the initial OpenShift Container Platform installation is complete, you can choose to [secure the integrated Docker registry](#), which involves adjusting the **--insecure-registry** option accordingly.

2.3.5. Configuring Docker Storage

Containers and the images they are created from are stored in Docker's storage back end. This storage is ephemeral and separate from any [persistent storage](#) allocated to meet the needs of your applications.

For RHEL Atomic Host

The default storage back end for Docker on RHEL Atomic Host is a thin pool logical volume, which is supported for production environments. You must ensure that enough space is allocated for this volume per the Docker storage requirements mentioned in [System Requirements](#).

If you do not have enough allocated, see [Managing Storage with Docker Formatted Containers](#) for details on using **docker-storage-setup** and basic instructions on storage management in RHEL Atomic Host.

For RHEL

The default storage back end for Docker on RHEL 7 is a thin pool on loopback devices, which is not supported for production use and only appropriate for proof of concept environments. For production environments, you must create a thin pool logical volume and re-configure Docker to use that volume.

You can use the **docker-storage-setup** script included with Docker to create a thin pool device and configure Docker's storage driver. This can be done after installing Docker and should be done before creating images or containers. The script reads configuration options from the **/etc/sysconfig/docker-storage-setup** file and supports three options for creating the logical volume:

- **Option A)** Use an additional block device.
- **Option B)** Use an existing, specified volume group.
- **Option C)** Use the remaining free space from the volume group where your root file system is located.

Option A is the most robust option, however it requires adding an additional block device to your host before configuring Docker storage. Options B and C both require leaving free space available when provisioning your host.

1. Create the **docker-pool** volume using one of the following three options:

- **Option A) Use an additional block device.**

In `/etc/sysconfig/docker-storage-setup`, set **DEVS** to the path of the block device you wish to use. Set **VG** to the volume group name you wish to create; **docker-vg** is a reasonable choice. For example:

```
# cat <<EOF > /etc/sysconfig/docker-storage-setup
DEVS=/dev/vdc
VG=docker-vg
EOF
```

Then run **docker-storage-setup** and review the output to ensure the **docker-pool** volume was created:

```
# docker-storage-setup
[5/1868]
0
Checking that no-one is using this disk right now ...
OK

Disk /dev/vdc: 31207 cylinders, 16 heads, 63 sectors/track
sfdisk: /dev/vdc: unrecognized partition table type

Old situation:
sfdisk: No partitions found

New situation:
Units: sectors of 512 bytes, counting from 0

    Device Boot      Start         End      #sectors  Id System
/dev/vdc1           2048    31457279     31455232   8e  Linux LVM
/dev/vdc2              0          -           0    0  Empty
/dev/vdc3              0          -           0    0  Empty
/dev/vdc4              0          -           0    0  Empty
Warning: partition 1 does not start at a cylinder boundary
Warning: partition 1 does not end at a cylinder boundary
Warning: no primary partition is marked bootable (active)
This does not matter for LILO, but the DOS MBR will not boot this
disk.
Successfully wrote the new partition table

Re-reading the partition table ...

If you created or changed a DOS partition, /dev/foo7, say, then
use dd(1)
to zero the first 512 bytes: dd if=/dev/zero of=/dev/foo7 bs=512
count=1
(See fdisk(8).)
Physical volume "/dev/vdc1" successfully created
Volume group "docker-vg" successfully created
Rounding up size to full physical extent 16.00 MiB
Logical volume "docker-poolmeta" created.
Logical volume "docker-pool" created.
WARNING: Converting logical volume docker-vg/docker-pool and
docker-vg/docker-poolmeta to pool's data and metadata volumes.
```

```
THIS WILL DESTROY CONTENT OF LOGICAL VOLUME (filesystem etc.)
Converted docker-vg/docker-pool to thin pool.
Logical volume "docker-pool" changed.
```

- **Option B) Use an existing, specified volume group.**

In `/etc/sysconfig/docker-storage-setup`, set **VG** to the desired volume group. For example:

```
# cat <<EOF > /etc/sysconfig/docker-storage-setup
VG=docker-vg
EOF
```

Then run **docker-storage-setup** and review the output to ensure the **docker-pool** volume was created:

```
# docker-storage-setup
Rounding up size to full physical extent 16.00 MiB
Logical volume "docker-poolmeta" created.
Logical volume "docker-pool" created.
WARNING: Converting logical volume docker-vg/docker-pool and
docker-vg/docker-poolmeta to pool's data and metadata volumes.
THIS WILL DESTROY CONTENT OF LOGICAL VOLUME (filesystem etc.)
Converted docker-vg/docker-pool to thin pool.
Logical volume "docker-pool" changed.
```

- **Option C) Use the remaining free space from the volume group where your root file system is located.**

Verify that the volume group where your root file system resides has the desired free space, then run **docker-storage-setup** and review the output to ensure the **docker-pool** volume was created:

```
# docker-storage-setup
Rounding up size to full physical extent 32.00 MiB
Logical volume "docker-poolmeta" created.
Logical volume "docker-pool" created.
WARNING: Converting logical volume rhel/docker-pool and
rhel/docker-poolmeta to pool's data and metadata volumes.
THIS WILL DESTROY CONTENT OF LOGICAL VOLUME (filesystem etc.)
Converted rhel/docker-pool to thin pool.
Logical volume "docker-pool" changed.
```

2. Verify your configuration. You should have a **dm.thinpooldev** value in the `/etc/sysconfig/docker-storage` file and a **docker-pool** logical volume:

```
# cat /etc/sysconfig/docker-storage
DOCKER_STORAGE_OPTIONS=--storage-opt dm.fs=trfs --storage-opt
dm.thinpooldev=/dev/mapper/docker--vg-docker--pool

# lvs
  LV          VG   Attr          LSize   Pool Origin Data%  Meta%  Move
Log Cpy%Sync Convert
  docker-pool rhel twi-a-t--- 9.29g               0.00   0.12
```



IMPORTANT

Before using Docker or OpenShift Container Platform, verify that the **docker-pool** logical volume is large enough to meet your needs. The **docker-pool** volume should be 60% of the available volume group and will grow to fill the volume group via LVM monitoring.

3. Check if Docker is running:

```
# systemctl is-active docker
```

4. If Docker has not yet been started on the host, enable and start the service:

```
# systemctl enable docker
# systemctl start docker
```

If Docker is already running, re-initialize Docker:



WARNING

This will destroy any containers or images currently on the host.

```
# systemctl stop docker
# rm -rf /var/lib/docker/*
# systemctl restart docker
```

If there is any content in **/var/lib/docker/**, it must be deleted. Files will be present if Docker has been used prior to the installation of OpenShift Container Platform.

2.3.5.1. Reconfiguring Docker Storage

Should you need to reconfigure Docker storage after having created the **docker-pool**, you should first remove the **docker-pool** logical volume. If you are using a dedicated volume group, you should also remove the volume group and any associated physical volumes before reconfiguring **docker-storage-setup** according to the instructions above.

See [Logical Volume Manager Administration](#) for more detailed information on LVM management.

2.3.5.2. Managing Container Logs

Sometimes a container's log file (the **/var/lib/docker/containers/<hash>/<hash>-json.log** file on the node where the container is running) can increase to a problematic size. You can manage this by configuring Docker's **json-file** logging driver to restrict the size and number of log files.

Option	Purpose
--log-opt max-size	Sets the size at which a new log file is created.

Option	Purpose
<code>--log-opt max-file</code>	Sets the file on each host to configure the options.

For example, to set the maximum file size to 1MB and always keep the last three log files, edit the `/etc/sysconfig/docker` file to configure `max-size=1M` and `max-file=3`:

```
OPTIONS='--insecure-registry=172.30.0.0/16 --selinux-enabled --log-opt
max-size=1M --log-opt max-file=3'
```

Next, restart the Docker service:

```
# systemctl restart docker
```

2.3.5.3. Viewing Available Container Logs

Container logs are stored in the `/var/lib/docker/containers/<hash>/` directory on the node where the container is running. For example:

```
# ls -lh
/var/lib/docker/containers/f088349cceac173305d3e2c2e4790051799efe363842fda
b5732f51f5b001fd8/
total 2.6M
-rw-r--r--. 1 root root 5.6K Nov 24 00:12 config.json
-rw-r--r--. 1 root root 649K Nov 24 00:15
f088349cceac173305d3e2c2e4790051799efe363842fdab5732f51f5b001fd8-json.log
-rw-r--r--. 1 root root 977K Nov 24 00:15
f088349cceac173305d3e2c2e4790051799efe363842fdab5732f51f5b001fd8-
json.log.1
-rw-r--r--. 1 root root 977K Nov 24 00:15
f088349cceac173305d3e2c2e4790051799efe363842fdab5732f51f5b001fd8-
json.log.2
-rw-r--r--. 1 root root 1.3K Nov 24 00:12 hostconfig.json
drwx-----. 2 root root    6 Nov 24 00:12 secrets
```

See Docker's documentation for additional information on how to [Configure Logging Drivers](#).

2.3.6. Ensuring Host Access

The [quick](#) and [advanced installation](#) methods require a user that has access to all hosts. If you want to run the installer as a non-root user, passwordless **sudo** rights must be configured on each destination host.

For example, you can generate an SSH key on the host where you will invoke the installation process:

```
# ssh-keygen
```

Do **not** use a password.

An easy way to distribute your SSH keys is by using a **bash** loop:

```
# for host in master.example.com \
    node1.example.com \
    node2.example.com; \
do ssh-copy-id -i ~/.ssh/id_rsa.pub $host; \
done
```

Modify the host names in the above command according to your configuration.

2.3.7. Setting Global Proxy Values

The OpenShift Container Platform installer uses the proxy settings in the `_/etc/environment_` file.

Ensure the following domain suffixes and IP addresses are in the `/etc/environment` file in the `no_proxy` parameter:

- Master and node host names (domain suffix).
- Other internal host names (domain suffix).
- Etcd IP addresses (must be IP addresses and not host names, as **etcd** access is done by IP address).
- Docker registry IP address.
- Kubernetes IP address, by default 172.30.0.1. Must be the value set in the [openshift_portal_net](#) parameter in the Ansible inventory file, by default `/etc/ansible/hosts`.
- Kubernetes internal domain suffix: `cluster.local`.
- Kubernetes internal domain suffix: `.svc`.

The following example assumes `http_proxy` and `https_proxy` values are set:

```
no_proxy=.internal.example.com,10.0.0.1,10.0.0.2,10.0.0.3,.cluster.local,.svc,localhost,127.0.0.1,172.30.0.1
```



NOTE

Because **noproxy** does not support CIDR, you can use domain suffixes.

2.3.8. What's Next?

If you are interested in installing OpenShift Container Platform using the containerized method (optional for RHEL but required for RHEL Atomic Host), see [Containerized Components](#) to prepare your hosts.

When you are ready to proceed, you can install OpenShift Container Platform using the [quick installation](#) or [advanced installation](#) method.


If you are installing a stand-alone registry, continue with [Installing a Stand-alone Registry](#).

2.4. CONTAINERIZED COMPONENTS

2.4.1. Overview

This section explores some of the preparation required to install OpenShift Container Platform as a set of services within containers. This applies to hosts using either Red Hat Enterprise Linux or Red Hat Atomic Host.

- For the [quick installation method](#), you can choose between the RPM or containerized method on a per host basis during the interactive installation, or set the values manually in an [installation configuration file](#).
- For the [advanced installation method](#), you can set the Ansible variable **containerized=true** in an [inventory file](#) on a cluster-wide or per host basis.



NOTE

When installing an environment with multiple masters, the load balancer cannot be deployed by the installation process as a container. See [Advanced Installation](#) for load balancer requirements using the native HA method.

The following sections detail the preparation for a containerized OpenShift Container Platform installation.

2.4.2. Required Images

Containerized installations make use of the following images:

- **openshift3/ose**
- **openshift3/node**
- **openshift3/openswitch**
- **registry.access.redhat.com/rhel7/etcd**

By default, all of the above images are pulled from the Red Hat Registry at registry.access.redhat.com.

If you need to use a private registry to pull these images during the installation, you can specify the registry information ahead of time. For the advanced installation method, you can set the following Ansible variables in your inventory file, as required:

```
cli_docker_additional_registries=<registry_hostname>
cli_docker_insecure_registries=<registry_hostname>
cli_docker_blocked_registries=<registry_hostname>
```

For the quick installation method, you can export the following environment variables on each target host:

```
# export OO_INSTALL_ADDITIONAL_REGISTRIES=<registry_hostname>
# export OO_INSTALL_INSECURE_REGISTRIES=<registry_hostname>
```

Blocked Docker registries cannot currently be specified using the quick installation method.

The configuration of additional, insecure, and blocked Docker registries occurs at the beginning of the installation process to ensure that these settings are applied before attempting to pull any of the required images.

2.4.3. Starting and Stopping Containers

The installation process creates relevant **systemd** units which can be used to start, stop, and poll services using normal **systemctl** commands. For containerized installations, these unit names match those of an RPM installation, with the exception of the **etcd** service which is named **etcd_container**.

This change is necessary as currently RHEL Atomic Host ships with the **etcd** package installed as part of the operating system, so a containerized version is used for the OpenShift Container Platform installation instead. The installation process disables the default **etcd** service. The **etcd** package is slated to be removed from RHEL Atomic Host in the future.

2.4.4. File Paths

All OpenShift configuration files are placed in the same locations during containerized installation as RPM based installations and will survive **os-tree** upgrades.

However, [the default image stream and template files](#) are installed at **/etc/origin/examples/** for containerized installations rather than the standard **/usr/share/openshift/examples/**, because that directory is read-only on RHEL Atomic Host.

2.4.5. Storage Requirements

RHEL Atomic Host installations normally have a very small root file system. However, the **etcd**, master, and node containers persist data in the **/var/lib/** directory. Ensure that you have enough space on the root file system before installing OpenShift Container Platform; see the [System Requirements](#) section for details.

2.4.6. Open vSwitch SDN Initialization

OpenShift Container Platform SDN initialization requires that the Docker bridge be reconfigured and that Docker is restarted. This complicates the situation when the node is running within a container. When using the Open vSwitch (OVS) SDN, you will see the node start, reconfigure Docker, restart Docker (which restarts all containers), and finally start successfully.

In this case, the node service may fail to start and be restarted a few times because the master services are also restarted along with Docker. The current implementation uses a workaround which relies on setting the **Restart=always** parameter in the Docker based **systemd** units.

2.5. QUICK INSTALLATION

2.5.1. Overview

The *quick installation* method allows you to use an interactive CLI utility, the **atomic-openshift-installer** command, to install OpenShift Container Platform across a set of hosts. This installer can deploy OpenShift Container Platform components on targeted hosts by either installing RPMs or running containerized services.

This installation method is provided to make the installation experience easier by [interactively gathering the data](#) needed to run on each host. The installer is a self-contained wrapper intended for usage on a Red Hat Enterprise Linux (RHEL) 7 system. While RHEL Atomic Host is supported for running containerized OpenShift Container Platform services, the installer is [provided by an RPM](#) not available by default in RHEL Atomic Host, and must therefore be run from a RHEL 7 system. The host initiating the installation does not need to be intended for inclusion in the OpenShift Container Platform cluster, but it can be.

In addition to running [interactive installations](#) from scratch, the **atomic-openshift-installer** command can also be run or re-run using a predefined installation configuration file. This file can be used with the installer to:

- run an [unattended installation](#),
- [add nodes](#) to an existing cluster,
- [upgrade your cluster](#), or
- [reinstall](#) the OpenShift Container Platform cluster completely.

Alternatively, you can use the [advanced installation](#) method for more complex environments.



NOTE

To install OpenShift Container Platform as a stand-alone registry, see [Installing a Stand-alone Registry](#).

2.5.2. Before You Begin

The installer allows you to install OpenShift Container Platform [master](#) and [node](#) components on a defined set of hosts.



NOTE

By default, any hosts you designate as masters during the installation process are automatically also configured as nodes so that the masters are configured as part of the [OpenShift Container Platform SDN](#). The node component on the masters, however, are marked unschedulable, which blocks pods from being scheduled on it. After the installation, you can [mark them schedulable](#) if you want.

Before installing OpenShift Container Platform, you must first [satisfy the prerequisites](#) on your hosts, which includes verifying system and environment requirements and properly installing and configuring Docker. You must also be prepared to provide or validate the following information for each of your targeted hosts during the course of the installation:

- User name on the target host that should run the Ansible-based installation (can be root or non-root)
- Host name
- Whether to install components for master, node, or both
- Whether to use the RPM or containerized method
- Internal and external IP addresses

If you are interested in installing OpenShift Container Platform using the containerized method (optional for RHEL but required for RHEL Atomic Host), see [RPM vs Containerized](#) to ensure that you understand the differences between these methods, then return to this topic to continue.

After following the instructions in the [Prerequisites](#) topic and deciding between the RPM and containerized methods, you can continue to running an [interactive](#) or [unattended](#) installation.

2.5.3. Running an Interactive Installation



NOTE

Ensure you have read through [Before You Begin](#).

You can start the interactive installation by running:

```
$ atomic-openshift-installer install
```

Then follow the on-screen instructions to install a new OpenShift Container Platform cluster.

After it has finished, ensure that you back up the `~/.config/openshift/installer.cfg.yml` [installation configuration file](#) that is created, as it is required if you later want to re-run the installation, add hosts to the cluster, or [upgrade your cluster](#). Then, [verify the installation](#).

2.5.4. Defining an Installation Configuration File

The installer can use a predefined installation configuration file, which contains information about your installation, individual hosts, and cluster. When running an [interactive installation](#), an installation configuration file based on your answers is created for you in `~/.config/openshift/installer.cfg.yml`. The file is created if you are instructed to exit the installation to manually modify the configuration or when the installation completes. You can also create the configuration file manually from scratch to perform an [unattended installation](#).

Example 2.1. Installation Configuration File Specification

```
version: v2 ❶
variant: openshift-enterprise ❷
variant_version: 3.3 ❸
ansible_log_path: /tmp/ansible.log ❹
deployment:
  ansible_ssh_user: root ❺
  hosts: ❻
  - ip: 10.0.0.1 ❼
    hostname: master-private.example.com ❸
    public_ip: 24.222.0.1 ❹
    public_hostname: master.example.com ❿
    roles: ❶
      - master
      - node
    containerized: true ❷
    connect_to: 24.222.0.1 ❸
  - ip: 10.0.0.2
    hostname: node1-private.example.com
    public_ip: 24.222.0.2
    public_hostname: node1.example.com
    node_labels: {'region': 'infra'} ❶
    roles:
      - node
    connect_to: 10.0.0.2
  - ip: 10.0.0.3
```

```

hostname: node2-private.example.com
public_ip: 24.222.0.3
public_hostname: node2.example.com
roles:
  - node
connect_to: 10.0.0.3
roles: 15
  master:
    <variable_name1>: "<value1>" 16
    <variable_name2>: "<value2>"
  node:
    <variable_name1>: "<value1>" 17

```

- 1 The version of this installation configuration file. As of OpenShift Container Platform 3.3, the only valid version here is **v2**.
- 2 The OpenShift Container Platform variant to install. For OpenShift Container Platform, set this to **openshift-enterprise**.
- 3 A valid version of your selected variant: **3.3**, **3.2**, or **3.1**. If not specified, this defaults to the latest version for the specified variant.
- 4 Defines where the Ansible logs are stored. By default, this is the **/tmp/ansible.log** file.
- 5 Defines which user Ansible uses to SSH in to remote systems for gathering facts and for the installation. By default, this is the root user, but you can set it to any user that has **sudo** privileges.
- 6 Defines a list of the hosts onto which you want to install the OpenShift Container Platform master and node components.
- 7 8 Required. Allows the installer to connect to the system and gather facts before proceeding with the install.
- 9 10 Required for unattended installations. If these details are not specified, then this information is pulled from the facts gathered by the installer, and you are asked to confirm the details. If undefined for an unattended installation, the installation fails.
- 11 Determines the type of services that are installed. Specified as a list.
- 12 If set to **true**, containerized OpenShift Container Platform services are run on target master and node hosts instead of installed using RPM packages. If set to **false** or unset, the default RPM method is used. RHEL Atomic Host requires the containerized method, and is automatically selected for you based on the detection of the **/run/ostree-booted** file. See [RPM vs Containerized](#) for more details.
- 13 The IP address that Ansible attempts to connect to when installing, upgrading, or uninstalling the systems. If the configuration file was auto-generated, then this is the value you first enter for the host during that interactive install process.
- 14 Node labels can optionally be set per-host.
- 15 Defines a dictionary of roles across the deployment.
- 16 17 Any ansible variables that should only be applied to hosts assigned a role can be defined. For examples, see [Configuring Ansible](#).

2.5.5. Running an Unattended Installation



NOTE

Ensure you have read through the [Before You Begin](#).

Unattended installations allow you to define your hosts and cluster configuration in an [installation configuration file](#) before running the installer so that you do not have to go through all of the [interactive installation](#) questions and answers. It also allows you to resume an interactive installation you may have left unfinished, and quickly get back to where you left off.

To run an unattended installation, first define an [installation configuration file](#) at `~/.config/openshift/installer.cfg.yml`. Then, run the installer with the `-u` flag:

```
$ atomic-openshift-installer -u install
```

By default in interactive or unattended mode, the installer uses the configuration file located at `~/.config/openshift/installer.cfg.yml` if the file exists. If it does not exist, attempting to start an unattended installation fails.

Alternatively, you can specify a different location for the configuration file using the `-c` option, but doing so will require you to specify the file location every time you run the installation:

```
$ atomic-openshift-installer -u -c </path/to/file> install
```

After the unattended installation finishes, ensure that you back up the `~/.config/openshift/installer.cfg.yml` file that was used, as it is required if you later want to re-run the installation, add hosts to the cluster, or [upgrade your cluster](#). Then, [verify the installation](#).

2.5.6. Verifying the Installation

After the installation completes:

1. Verify that the master is started and nodes are registered and reporting in **Ready** status. **On the master host**, run the following as root:

```
# oc get nodes
```

NAME	STATUS	AGE
master.example.com	Ready,SchedulingDisabled	165d
node1.example.com	Ready	165d
node2.example.com	Ready	165d

2. To verify that the web console is installed correctly, use the master host name and the console port number to access the console with a web browser.

For example, for a master host with a hostname of `master.openshift.com` and using the default port of **8443**, the web console would be found at:

```
https://master.openshift.com:8443/console
```

3. Now that the install has been verified, run the following command on each master and node host to add the **atomic-openshift** packages back to the list of yum excludes on the host:

```
# atomic-openshift-excluder exclude
```

Then, see [What's Next](#) for the next steps on configuring your OpenShift Container Platform cluster.

2.5.7. Uninstalling OpenShift Container Platform

You can uninstall OpenShift Container Platform on all hosts in your cluster using the installer by running:

```
$ atomic-openshift-installer uninstall
```

See the [advanced installation method](#) for more options.

2.5.8. What's Next?

Now that you have a working OpenShift Container Platform instance, you can:

- [Configure authentication](#); by default, authentication is set to [Deny All](#).
- Configure the automatically-deployed [integrated Docker registry](#).
- Configure the automatically-deployed [router](#).

2.6. ADVANCED INSTALLATION

2.6.1. Overview

A reference configuration implemented using [Ansible](#) playbooks is available as the *advanced installation* method for installing a OpenShift Container Platform cluster. Familiarity with Ansible is assumed, however you can use this configuration as a reference to create your own implementation using the configuration management tool of your choosing.

While RHEL Atomic Host is supported for running containerized OpenShift Container Platform services, the advanced installation method utilizes Ansible, which is not available in RHEL Atomic Host, and must therefore be run from a RHEL 7 system. The host initiating the installation does not need to be intended for inclusion in the OpenShift Container Platform cluster, but it can be.

Alternatively, you can use the [quick installation](#) method if you prefer an interactive installation experience.



NOTE

To install OpenShift Container Platform as a stand-alone registry, see [Installing a Stand-alone Registry](#).

2.6.2. Before You Begin

Before installing OpenShift Container Platform, you must first see the [Prerequisites](#) and [Host Preparation](#) topics to prepare your hosts. This includes verifying system and environment requirements per component type and properly installing and configuring Docker. It also includes installing Ansible

version 2.2.0 or later, as the advanced installation method is based on Ansible playbooks and as such requires directly invoking Ansible.

If you are interested in installing OpenShift Container Platform using the containerized method (optional for RHEL but required for RHEL Atomic Host), see [RPM vs Containerized](#) to ensure that you understand the differences between these methods, then return to this topic to continue.

After following the instructions in the [Prerequisites](#) topic and deciding between the RPM and containerized methods, you can continue in this topic to [Configuring Ansible](#).

2.6.3. Configuring Ansible

The `/etc/ansible/hosts` file is Ansible's inventory file for the playbook to use during the installation. The inventory file describes the configuration for your OpenShift Container Platform cluster. You must replace the default contents of the file with your desired configuration.

The following sections describe commonly-used variables to set in your inventory file during an advanced installation, followed by example inventory files you can use as a starting point for your installation. The examples describe various environment topographies, including [using multiple masters for high availability](#). You can choose an example that matches your requirements, modify it to match your own environment, and use it as your inventory file when [running the advanced installation](#).

2.6.3.1. Configuring Host Variables

To assign environment variables to hosts during the Ansible installation, indicate the desired variables in the `/etc/ansible/hosts` file after the host entry in the `[masters]` or `[nodes]` sections. For example:

```
[masters]
ec2-52-6-179-239.compute-1.amazonaws.com openshift_public_hostname=ose3-
master.public.example.com
```

The following table describes variables for use with the Ansible installer that can be assigned to individual host entries:

Table 2.9. Host Variables

Variable	Purpose
<code>openshift_hostname</code>	This variable overrides the internal cluster host name for the system. Use this when the system's default IP address does not resolve to the system host name.
<code>openshift_public_hostname</code>	This variable overrides the system's public host name. Use this for cloud installations, or for hosts on networks using a network address translation (NAT).
<code>openshift_ip</code>	This variable overrides the cluster internal IP address for the system. Use this when using an interface that is not configured with the default route. This variable can also be used for etcd.

Variable	Purpose
openshift_public_ip	This variable overrides the system's public IP address. Use this for cloud installations, or for hosts on networks using a network address translation (NAT).
containerized	If set to true , containerized OpenShift Container Platform services are run on target master and node hosts instead of installed using RPM packages. If set to false or unset, the default RPM method is used. RHEL Atomic Host requires the containerized method, and is automatically selected for you based on the detection of the <code>/run/ostree-booted</code> file. See RPM vs Containerized for more details. Containerized installations are supported starting in OpenShift Container Platform 3.1.1.
openshift_node_labels	This variable adds labels to nodes during installation. See Configuring Node Host Labels for more details.
openshift_node_kubelet_args	This variable is used to configure kubeletArguments on nodes, such as arguments used in container and image garbage collection , and to specify resources per node . kubeletArguments are key value pairs that are passed directly to the Kubelet that match the Kubelet's command line arguments . kubeletArguments are not migrated or validated and may become invalid if used. These values override other settings in node configuration which may cause invalid configurations. Example usage: <code>{'image-gc-high-threshold': ['90'],'image-gc-low-threshold': ['80']}</code> .
openshift_docker_options	This variable configures additional Docker options within <code>/etc/sysconfig/docker</code> , such as options used in Managing Container Logs . Example usage: <code>--log-driver json-file --log-opt max-size=1M --log-opt max-file=3</code> .
openshift_schedulable	This variable configures whether the host is marked as a schedulable node, meaning that it is available for placement of new pods. See Configuring Schedulability on Masters .

Image Version Policy

Images require a version number policy in order to maintain updates. See the [Image Version Tag Policy](#) section in the Architecture Guide for more information.

2.6.3.2. Configuring Cluster Variables

To assign environment variables during the Ansible install that apply more globally to your OpenShift Container Platform cluster overall, indicate the desired variables in the `/etc/ansible/hosts` file on separate, single lines within the `[OSEv3:vars]` section. For example:

```
[OSEv3:vars]

openshift_master_identity_providers=[{'name': 'htpasswd_auth',
'login': 'true', 'challenge': 'true',
'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/htpasswd'}]

openshift_master_default_subdomain=apps.test.example.com
```

The following table describes variables for use with the Ansible installer that can be assigned cluster-wide:

Table 2.10. Cluster Variables

Variable	Purpose
ansible_ssh_user	This variable sets the SSH user for the installer to use and defaults to root . This user should allow SSH-based authentication without requiring a password . If using SSH key-based authentication, then the key should be managed by an SSH agent.
ansible_become	If ansible_ssh_user is not root , this variable must be set to true and the user must be configured for passwordless sudo .
debug_level	<p>This variable sets which INFO messages are logged to the systemd-journald.service. Set one of the following:</p> <ul style="list-style-type: none"> • 0 to log errors and warnings only • 2 to log normal information (This is the default level.) • 4 to log debugging-level information • 6 to log API-level debugging information (request / response) • 8 to log body-level API debugging information <p>For more information on debug log levels, see Configuring Logging Levels.</p>

Variable	Purpose
containerized	If set to true , containerized OpenShift Container Platform services are run on all target master and node hosts in the cluster instead of installed using RPM packages. If set to false or unset, the default RPM method is used. RHEL Atomic Host requires the containerized method, and is automatically selected for you based on the detection of the <i>/run/ostree-booted</i> file. See RPM vs Containerized for more details. Containerized installations are supported starting in OpenShift Container Platform 3.1.1.
openshift_master_cluster_hostname	This variable overrides the host name for the cluster, which defaults to the host name of the master.
openshift_master_cluster_public_host name	<p>This variable overrides the public host name for the cluster, which defaults to the host name of the master. If you use an external load balancer, specify the address of the external load balancer.</p> <p>For example:</p> <pre>---- openshift_master_cluster_public_hostname=openshift-ansible.public.example.com ----</pre>
openshift_master_cluster_method	Optional. This variable defines the HA method when deploying multiple masters. Supports the native method. See Multiple Masters for more information.
openshift_rolling_restart_mode	This variable enables rolling restarts of HA masters (i.e., masters are taken down one at a time) when running the upgrade playbook directly . It defaults to services , which allows rolling restarts of services on the masters. It can instead be set to system , which enables rolling, full system restarts and also works for single master clusters.
os_sdn_network_plugin_name	This variable configures which OpenShift Container Platform SDN plug-in to use for the pod network, which defaults to redhat/openshift-ovs-subnet for the standard SDN plug-in. Set the variable to redhat/openshift-ovs-multitenant to use the multitenant plug-in.
openshift_master_identity_providers	This variable overrides the identity provider , which defaults to Deny All .

Variable	Purpose
<code>openshift_master_named_certificates</code>	These variables are used to configure custom certificates which are deployed as part of the installation. See Configuring Custom Certificates for more information.
<code>openshift_master_overwrite_named_certificates</code>	
<code>openshift_master_session_name</code>	These variables override defaults for session options in the OAuth configuration. See Configuring Session Options for more information.
<code>openshift_master_session_max_seconds</code>	
<code>openshift_master_session_auth_secrets</code>	
<code>openshift_master_session_encryption_secrets</code>	
<code>openshift_portal_net</code>	This variable configures the subnet in which services will be created within the OpenShift Container Platform SDN . This network block should be private and must not conflict with any existing network blocks in your infrastructure to which pods, nodes, or the master may require access to, or the installation will fail. Defaults to 172.30.0.0/16 , and cannot be re-configured after deployment. If changing from the default, avoid 172.17.0.0/16 , which the docker0 network bridge uses by default, or modify the docker0 network.
<code>openshift_master_default_subdomain</code>	This variable overrides the default subdomain to use for exposed routes .
<code>openshift_node_proxy_mode</code>	This variable specifies the service proxy mode to use: either iptables for the default, pure- iptables implementation, or userspace for the user space proxy.
<code>openshift_hosted_router_selector</code>	Default node selector for automatically deploying router pods. See Configuring Node Host Labels for details.
<code>openshift_registry_selector</code>	Default node selector for automatically deploying registry pods. See Configuring Node Host Labels for details.
<code>osm_default_node_selector</code>	This variable overrides the node selector that projects will use by default when placing pods.

Variable	Purpose
<code>osm_cluster_network_cidr</code>	This variable overrides the SDN cluster network CIDR block. This is the network from which pod IPs are assigned. This network block should be a private block and must not conflict with existing network blocks in your infrastructure to which pods, nodes, or the master may require access. Defaults to 10.128.0.0/14 and cannot be arbitrarily re-configured after deployment, although certain changes to it can be made in the SDN master configuration .
<code>osm_host_subnet_length</code>	This variable specifies the size of the per host subnet allocated for pod IPs by OpenShift Container Platform SDN . Defaults to 9 which means that a subnet of size /23 is allocated to each host; for example, given the default 10.128.0.0/14 cluster network, this will allocate 10.128.0.0/23, 10.128.2.0/23, 10.128.4.0/23, and so on. This cannot be re-configured after deployment.
<code>openshift_docker_additional_registries</code>	OpenShift Container Platform adds the specified additional registry or registries to the Docker configuration.
<code>openshift_docker_insecure_registries</code>	OpenShift Container Platform adds the specified additional insecure registry or registries to the Docker configuration.
<code>openshift_docker_blocked_registries</code>	OpenShift Container Platform adds the specified blocked registry or registries to the Docker configuration.
<code>openshift_hosted_metrics_public_url</code>	This variable sets the host name for integration with the metrics console. The default is https://hawkular-metrics.{{openshift_master_default_subdomain}}/hawkular/metrics . If you alter this variable, ensure the host name is accessible via your router.

2.6.3.3. Configuring a Registry Location

If you are using an image registry other than the default at `registry.access.redhat.com`, specify the desired registry within the `/etc/ansible/hosts` file.

```
oreg_url=example.com/openshift3/ose-${component}:${version}
openshift_examples_modify_imagestreams=true
```

Table 2.11. Registry Variables

Variable	Purpose
oreg_url	Set to the alternate image location. Necessary if you are not using the default registry at registry.access.redhat.com .
openshift_examples_modify_imagestreams	Set to true if pointing to a registry other than the default. Modifies the image stream location to the value of oreg_url .

2.6.3.4. Configuring Global Proxy Options

If your hosts require use of a HTTP or HTTPS proxy in order to connect to external hosts, there are many components that must be configured to use the proxy, including masters, Docker, and builds. Node services only connect to the master API requiring no external access and therefore do not need to be configured to use a proxy.

In order to simplify this configuration, the following Ansible variables can be specified at a cluster or host level to apply these settings uniformly across your environment.



NOTE

See [Configuring Global Build Defaults and Overrides](#) for more information on how the proxy environment is defined for builds.

Table 2.12. Cluster Proxy Variables

Variable	Purpose
openshift_http_proxy	This variable specifies the HTTP_PROXY environment variable for masters and the Docker daemon.
openshift_https_proxy	This variable specifies the HTTPS_PROXY environment variable for masters and the Docker daemon.
openshift_no_proxy	This variable is used to set the NO_PROXY environment variable for masters and the Docker daemon. This value should be set to a comma separated list of host names or wildcard host names that should not use the defined proxy. This list will be augmented with the list of all defined OpenShift Container Platform host names by default.
openshift_generate_no_proxy_hosts	This boolean variable specifies whether or not the names of all defined OpenShift hosts and *.cluster.local should be automatically appended to the NO_PROXY list. Defaults to true ; set it to false to override this option.

Variable	Purpose
openshift_builddefaults_http_proxy	This variable defines the HTTP_PROXY environment variable inserted into builds using the BuildDefaults admission controller. If openshift_http_proxy is set, this variable will inherit that value; you only need to set this if you want your builds to use a different value.
openshift_builddefaults_https_proxy	This variable defines the HTTPS_PROXY environment variable inserted into builds using the BuildDefaults admission controller. If openshift_https_proxy is set, this variable will inherit that value; you only need to set this if you want your builds to use a different value.
openshift_builddefaults_no_proxy	This variable defines the NO_PROXY environment variable inserted into builds using the BuildDefaults admission controller. If openshift_no_proxy is set, this variable will inherit that value; you only need to set this if you want your builds to use a different value.
openshift_builddefaults_git_http_proxy	This variable defines the HTTP proxy used by git clone operations during a build, defined using the BuildDefaults admission controller. If openshift_builddefaults_http_proxy is set, this variable will inherit that value; you only need to set this if you want your git clone operations to use a different value.
openshift_builddefaults_git_https_proxy	This variable defines the HTTPS proxy used by git clone operations during a build, defined using the BuildDefaults admission controller. If openshift_builddefaults_https_proxy is set, this variable will inherit that value; you only need to set this if you want your git clone operations to use a different value.

2.6.3.5. Configuring Schedulability on Masters

Any hosts you designate as masters during the installation process should also be configured as nodes so that the masters are configured as part of the [OpenShift SDN](#). You must do so by adding entries for these hosts to the **[nodes]** section:

```
[nodes]
master.example.com
```

In order to ensure that your masters are not burdened with running pods, they are automatically marked unschedulable by default by the installer, meaning that new pods cannot be placed on the hosts. This is the same as setting the **openshift_schedulable=false** host variable.

You can manually set a master host to schedulable during installation using the **openshift_schedulable=true** host variable, though this is not recommended in production environments:

```
[nodes]
master.example.com openshift_schedulable=true
```

If you want to change the schedulability of a host post-installation, see [Marking Nodes as Unschedulable or Schedulable](#).

2.6.3.6. Configuring Node Host Labels

You can assign [labels](#) to node hosts during the Ansible install by configuring the `/etc/ansible/hosts` file. Labels are useful for determining the placement of pods onto nodes using the [scheduler](#). Other than **region=infra** (discussed in [Configuring Dedicated Infrastructure Nodes](#)), the actual label names and values are arbitrary and can be assigned however you see fit per your cluster's requirements.

To assign labels to a node host during an Ansible install, use the **openshift_node_labels** variable with the desired labels added to the desired node host entry in the **[nodes]** section. In the following example, labels are set for a region called **primary** and a zone called **east**:

```
[nodes]
node1.example.com openshift_node_labels="{ 'region': 'primary', 'zone': 'east' }"
```

2.6.3.6.1. Configuring Dedicated Infrastructure Nodes

The **openshift_router_selector** and **openshift_registry_selector** Ansible settings determine the label selectors used when placing registry and router pods. They are set to **region=infra** by default:

```
# default selectors for router and registry services
# openshift_router_selector='region=infra'
# openshift_registry_selector='region=infra'
```

The default router and registry will be automatically deployed during installation if nodes exist in the **[nodes]** section that match the selector settings. For example:

```
[nodes]
infra-node1.example.com openshift_node_labels="{ 'region': 'infra', 'zone': 'default' }"
```



IMPORTANT

The registry and router are only able to run on node hosts with the **region=infra** label. Ensure that at least one node host in your OpenShift Container Platform environment has the **region=infra** label.

It is recommended for production environments that you maintain dedicated infrastructure nodes where the registry and router pods can run separately from pods used for user applications.

As described in [Configuring Schedulability on Masters](#), master hosts are marked unschedulable by

default. If you label a master host with **region=infra** and have no other dedicated infrastructure nodes, you must also explicitly mark these master hosts as schedulable. Otherwise, the registry and router pods cannot be placed anywhere:

```
[nodes]
master.example.com openshift_node_labels="{ 'region': 'infra', 'zone':
'default' }" openshift_schedulable=true
```

2.6.3.7. Configuring Session Options

[Session options](#) in the OAuth configuration are configurable in the inventory file. By default, Ansible populates a **sessionSecretsFile** with generated authentication and encryption secrets so that sessions generated by one master can be decoded by the others. The default location is **/etc/origin/master/session-secrets.yaml**, and this file will only be re-created if deleted on all masters.

You can set the session name and maximum number of seconds with **openshift_master_session_name** and **openshift_master_session_max_seconds**:

```
openshift_master_session_name=ssn
openshift_master_session_max_seconds=3600
```

If provided, **openshift_master_session_auth_secrets** and **openshift_master_encryption_secrets** must be equal length.

For **openshift_master_session_auth_secrets**, used to authenticate sessions using HMAC, it is recommended to use secrets with 32 or 64 bytes:

```
openshift_master_session_auth_secrets=[ 'DONT+USE+THIS+SECRET+b4NV+pmZNSO' ]
```

For **openshift_master_encryption_secrets**, used to encrypt sessions, secrets must be 16, 24, or 32 characters long, to select AES-128, AES-192, or AES-256:

```
openshift_master_session_encryption_secrets=
[ 'DONT+USE+THIS+SECRET+b4NV+pmZNSO' ]
```

2.6.3.8. Configuring Custom Certificates

[Custom serving certificates](#) for the public host names of the OpenShift Container Platform API and [web console](#) can be deployed during an advanced installation and are configurable in the inventory file.



NOTE

Custom certificates should only be configured for the host name associated with the **publicMasterURL** which can be set using **openshift_master_cluster_public_hostname**. Using a custom serving certificate for the host name associated with the **masterURL** (**openshift_master_cluster_hostname**) will result in TLS errors as infrastructure components will attempt to contact the master API using the internal **masterURL** host.

Certificate and key file paths can be configured using the **openshift_master_named_certificates** cluster variable:

■

```
openshift_master_named_certificates=[{"certfile": "/path/to/custom1.crt",
"keyfile": "/path/to/custom1.key"}]
```

File paths must be local to the system where Ansible will be run. Certificates are copied to master hosts and are deployed within the ***/etc/origin/master/named_certificates/*** directory.

Ansible detects a certificate's **Common Name** and **Subject Alternative Names**. Detected names can be overridden by providing the "names" key when setting

openshift_master_named_certificates:

```
openshift_master_named_certificates=[{"certfile": "/path/to/custom1.crt",
"keyfile": "/path/to/custom1.key", "names": ["public-master-host.com"]}]
```

Certificates configured using **openshift_master_named_certificates** are cached on masters, meaning that each additional Ansible run with a different set of certificates results in all previously deployed certificates remaining in place on master hosts and within the master configuration file.

If you would like **openshift_master_named_certificates** to be overwritten with the provided value (or no value), specify the **openshift_master_overwrite_named_certificates** cluster variable:

```
openshift_master_overwrite_named_certificates=true
```

For a more complete example, consider the following cluster variables in an inventory file:

```
openshift_master_cluster_method=native
openshift_master_cluster_hostname=lb-internal.openshift.com
openshift_master_cluster_public_hostname=custom.openshift.com
```

To overwrite the certificates on a subsequent Ansible run, you could set the following:

```
openshift_master_named_certificates=[{"certfile":
"/root/STAR.openshift.com.crt", "keyfile": "/root/STAR.openshift.com.key",
"names": ["custom.openshift.com"]}]
openshift_master_overwrite_named_certificates=true
```

2.6.3.9. Configuring Deployment Type

Various defaults used throughout the playbooks and roles in this repository are set based on the deployment type configuration (usually defined in an Ansible hosts file).

Ensure the **deployment_type** parameter in your inventory file is set to **openshift-enterprise**.

2.6.4. Configuring Cluster Metrics

Cluster metrics are not set to automatically deploy by default. Set the following to enable cluster metrics when using the advanced install:

```
[OSEv3:vars]

openshift_hosted_metrics_deploy=true
```

**IMPORTANT**

In accordance with upstream Kubernetes rules, metrics can be collected only on the default interface of **eth0**.

2.6.4.1. Metrics Storage

The **openshift_hosted_metrics_storage_kind** variable must be set in order to use persistent storage. If **openshift_hosted_metrics_storage_kind** is not set, then cluster metrics data is stored in an **EmptyDir** volume, which will be deleted when the Cassandra pod terminates.

There are three options for enabling cluster metrics storage when using the advanced install:

Option A - NFS Host Group

When the following variables are set, an NFS volume is created during an advanced install with path **<nfs_directory>/<volume_name>** on the host within the [nfs] host group. For example, the volume path using these options would be **/exports/metrics**:

```
[OSEv3:vars]

openshift_hosted_metrics_storage_kind=nfs
openshift_hosted_metrics_storage_access_modes=['ReadWriteOnce']
openshift_hosted_metrics_storage_nfs_directory=/exports
openshift_hosted_metrics_storage_nfs_options='*(rw,root_squash)'
openshift_hosted_metrics_storage_volume_name=metrics
openshift_hosted_metrics_storage_volume_size=10Gi
```

Option B - External NFS Host

To use an external NFS volume, one must already exist with a path of **<nfs_directory>/<volume_name>** on the storage host.

```
[OSEv3:vars]

openshift_hosted_metrics_storage_kind=nfs
openshift_hosted_metrics_storage_access_modes=['ReadWriteOnce']
openshift_hosted_metrics_storage_host=nfs.example.com
openshift_hosted_metrics_storage_nfs_directory=/exports
openshift_hosted_metrics_storage_volume_name=metrics
openshift_hosted_metrics_storage_volume_size=10Gi
```

The remote volume path using the following options would be **nfs.example.com:/exports/metrics**.

Option C - Dynamic


Use the following variable if your OpenShift Container Platform environment supports dynamic volume provisioning for your cloud platform:

```
[OSEv3:vars]

#openshift_hosted_metrics_storage_kind=dynamic
```

2.6.5. Single Master Examples

You can configure an environment with a single master and multiple nodes, and either a single embedded **etcd** or multiple external **etcd** hosts.



NOTE

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

Single Master and Multiple Nodes

The following table describes an example environment for a single **master** (with embedded **etcd**) and two **nodes**:

Host Name	Infrastructure Component to Install
master.example.com	Master and node
node1.example.com	Node
node2.example.com	

You can see these example hosts present in the **[masters]** and **[nodes]** sections of the following example inventory file:

Example 2.2. Single Master and Multiple Nodes Inventory File

```
# Create an OSEv3 group that contains the masters and nodes groups
[OSEv3:children]
masters
nodes

# Set variables common for all OSEv3 hosts
[OSEv3:vars]
# SSH user, this user should allow ssh based auth without requiring a
password
ansible_ssh_user=root

# If ansible_ssh_user is not root, ansible_become must be set to true
#ansible_become=true

deployment_type=openshift-enterprise

# uncomment the following to enable htpasswd authentication; defaults to
DenyAllPasswordIdentityProvider
#openshift_master_identity_providers=[{'name': 'htpasswd_auth', 'login':
'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/htpasswd'}]

# host group for masters
[masters]
master.example.com

# host group for nodes, includes region info
[nodes]
```

```

master.example.com
node1.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'west' }"
infra-node1.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"
infra-node2.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"

```

To use this example, modify the file to match your environment and specifications, and save it as */etc/ansible/hosts*.

Single Master, Multiple etcd, and Multiple Nodes

The following table describes an example environment for a single [master](#), three [etcd](#) hosts, and two [nodes](#):

Host Name	Infrastructure Component to Install
master.example.com	Master and node
etcd1.example.com	etcd
etcd2.example.com	
etcd3.example.com	
node1.example.com	Node
node2.example.com	



NOTE

When specifying multiple **etcd** hosts, external **etcd** is installed and configured. Clustering of OpenShift Container Platform's embedded **etcd** is not supported.

You can see these example hosts present in the **[masters]**, **[nodes]**, and **[etcd]** sections of the following example inventory file:

Example 2.3. Single Master, Multiple etcd, and Multiple Nodes Inventory File

```

# Create an OSEv3 group that contains the masters, nodes, and etcd
groups
[OSEv3:children]
masters
nodes
etcd

# Set variables common for all OSEv3 hosts

```

```

[OSEv3:vars]
ansible_ssh_user=root
deployment_type=openshift-enterprise

# uncomment the following to enable htpasswd authentication; defaults to
# DenyAllPasswordIdentityProvider
#openshift_master_identity_providers=[{'name': 'htpasswd_auth', 'login':
# 'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
# 'filename': '/etc/origin/master/htpasswd'}]

# host group for masters
[masters]
master.example.com

# host group for etcd
[etcd]
etcd1.example.com
etcd2.example.com
etcd3.example.com

# host group for nodes, includes region info
[nodes]
master.example.com
node1.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'west' }"
infra-node1.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"
infra-node2.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"

```

To use this example, modify the file to match your environment and specifications, and save it as */etc/ansible/hosts*.

2.6.6. Multiple Masters Examples

You can configure an environment with multiple masters, multiple **etcd** hosts, and multiple nodes. Configuring [multiple masters for high availability](#) (HA) ensures that the cluster has no single point of failure.



NOTE

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

When configuring multiple masters, the advanced installation supports the following high availability (HA) method:

native	Leverages the native HA master capabilities built into OpenShift Container Platform and can be combined with any load balancing solution. If a host is defined in the [lb] section of the inventory file, Ansible installs and configures HAProxy automatically as the load balancing solution. If no host is defined, it is assumed you have pre-configured a load balancing solution of your choice to balance the master API (port 8443) on all master hosts.
---------------	---

For your pre-configured load balancing solution, you must have:

- A pre-created load balancer VIP configured for SSL passthrough.
- A domain name for VIP registered in DNS.
 - The domain name will become the value of both **openshift_master_cluster_public_hostname** and **openshift_master_cluster_hostname** in the OpenShift Container Platform installer.

See [External Load Balancer Integrations](#) for more information.



NOTE

For more on the high availability master architecture, see [Kubernetes Infrastructure](#).

Note the following when using the **native** HA method:

- The advanced installation method does not currently support multiple HAProxy load balancers in an active-passive setup. See the [Load Balancer Administration documentation](#) for post-installation amendments.
- In a HAProxy setup, controller manager servers run as standalone processes. They elect their active leader with a lease stored in **etcd**. The lease expires after 30 seconds by default. If a failure happens on an active controller server, it will take up to this number of seconds to elect another leader. The interval can be configured with the **osm_controller_lease_ttl** variable.

To configure multiple masters, refer to the following section.

Multiple Masters with Multiple etcd, and Using Native HA

The following describes an example environment for three [masters](#), one HAProxy load balancer, three [etcd](#) hosts, and two [nodes](#) using the **native** HA method:

Host Name	Infrastructure Component to Install
master1.example.com	Master (clustered using native HA) and node
master2.example.com	
master3.example.com	
lb.example.com	HAProxy to load balance API master endpoints

Host Name	Infrastructure Component to Install
etcd1.example.com	etcd
etcd2.example.com	
etcd3.example.com	
node1.example.com	Node
node2.example.com	



NOTE

When specifying multiple **etcd** hosts, external **etcd** is installed and configured. Clustering of OpenShift Container Platform's embedded **etcd** is not supported.

You can see these example hosts present in the **[masters]**, **[etcd]**, **[lb]**, and **[nodes]** sections of the following example inventory file:

Example 2.4. Multiple Masters Using HAProxy Inventory File

```
# Create an OSEv3 group that contains the master, nodes, etcd, and lb
groups.
# The lb group lets Ansible configure HAProxy as the load balancing
solution.
# Comment lb out if your load balancer is pre-configured.
[OSEv3:children]
masters
nodes
etcd
lb

# Set variables common for all OSEv3 hosts
[OSEv3:vars]
ansible_ssh_user=root
deployment_type=openshift-enterprise

# Uncomment the following to enable htpasswd authentication; defaults to
# DenyAllPasswordIdentityProvider.
#openshift_master_identity_providers=[{'name': 'htpasswd_auth', 'login':
'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/htpasswd'}]

# Native high availability cluster method with optional load balancer.
# If no lb group is defined installer assumes that a load balancer has
# been preconfigured. For installation the value of
# openshift_master_cluster_hostname must resolve to the load balancer
# or to one or all of the masters defined in the inventory if no load
# balancer is present.
openshift_master_cluster_method=native
openshift_master_cluster_hostname=openshift-internal.example.com
```

```

openshift_master_cluster_public_hostname=openshift-cluster.example.com

# apply updated node defaults
openshift_node_kubelet_args={'pods-per-core': ['10'], 'max-pods':
['250'], 'image-gc-high-threshold': ['90'], 'image-gc-low-threshold':
['80']}

# override the default controller lease ttl
#osm_controller_lease_ttl=30

# enable ntp on masters to ensure proper failover
openshift_clock_enabled=true

# host group for masters
[masters]
master1.example.com
master2.example.com
master3.example.com

# host group for etcd
[etcd]
etcd1.example.com
etcd2.example.com
etcd3.example.com

# Specify load balancer host
[lb]
lb.example.com

# host group for nodes, includes region info
[nodes]
master[1:3].example.com
node1.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'west' }"
infra-node1.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"
infra-node2.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"

```

To use this example, modify the file to match your environment and specifications, and save it as ***/etc/ansible/hosts***.

Multiple Masters with Master and etcd on the Same Host, and Using Native HA

The following describes an example environment for three **masters** with **etcd** on each host, one HAProxy load balancer, and two **nodes** using the **native** HA method:

Host Name	Infrastructure Component to Install
master1.example.com	Master (clustered using native HA) and node with etcd on each host

Host Name	Infrastructure Component to Install
master2.example.com	
master3.example.com	
lb.example.com	HAProxy to load balance API master endpoints
node1.example.com	Node
node2.example.com	

You can see these example hosts present in the **[masters]**, **[etcd]**, **[lb]**, and **[nodes]** sections of the following example inventory file:

```
# Create an OSEv3 group that contains the master, nodes, etcd, and lb
groups.
# The lb group lets Ansible configure HAProxy as the load balancing
solution.
# Comment lb out if your load balancer is pre-configured.
[OSEv3:children]
masters
nodes
etcd
lb

# Set variables common for all OSEv3 hosts
[OSEv3:vars]
ansible_ssh_user=root
deployment_type=openshift-enterprise

# Uncomment the following to enable htpasswd authentication; defaults to
# DenyAllPasswordIdentityProvider.
#openshift_master_identity_providers=[{'name': 'htpasswd_auth', 'login':
'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/htpasswd'}]

# Native high availability cluster method with optional load balancer.
# If no lb group is defined installer assumes that a load balancer has
# been preconfigured. For installation the value of
# openshift_master_cluster_hostname must resolve to the load balancer
# or to one or all of the masters defined in the inventory if no load
# balancer is present.
openshift_master_cluster_method=native
openshift_master_cluster_hostname=openshift-internal.example.com
openshift_master_cluster_public_hostname=openshift-cluster.example.com

# override the default controller lease ttl
#osm_controller_lease_ttl=30

# host group for masters
[masters]
```

```

master1.example.com
master2.example.com
master3.example.com

# host group for etcd
[etcd]
master1.example.com
master2.example.com
master3.example.com

# Specify load balancer host
[lb]
lb.example.com

# host group for nodes, includes region info
[nodes]
master[1:3].example.com
node1.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'west' }"
infra-node1.example.com openshift_node_labels="{ 'region': 'infra', 'zone':
'default' }"
infra-node2.example.com openshift_node_labels="{ 'region': 'infra', 'zone':
'default' }"

```

To use this example, modify the file to match your environment and specifications, and save it as ***/etc/ansible/hosts***.

2.6.7. Running the Advanced Installation

After you have [configured Ansible](#) by defining an inventory file in ***/etc/ansible/hosts***, you can run the advanced installation using the following playbook:

```

# ansible-playbook /usr/share/ansible/openshift-
ansible/playbooks/byo/config.yml

```

If for any reason the installation fails, before re-running the installer, see [Known Issues](#) to check for any specific instructions or workarounds.



WARNING

The installer caches playbook configuration values for 10 minutes, by default. If you change any system, network, or inventory configuration, and then re-run the installer within that 10 minute period, the new values are not used, and the previous values are used instead. You can delete the contents of the cache, which is defined by the ***fact_caching_connection*** value in the ***/etc/ansible/ansible.cfg*** file.



NOTE

Due to a known issue, after running the installation, if NFS volumes are provisioned for any component, the following directories might be created whether their components are being deployed to NFS volumes or not:

- `/exports/logging-es`
- `/exports/logging-es-ops/`
- `/exports/metrics/`
- `/exports/prometheus`
- `/exports/prometheus-alertbuffer/`
- `/exports/prometheus-alertmanager/`

You can delete these directories after installation, as needed.

2.6.8. Verifying the Installation

After the installation completes:

1. Verify that the master is started and nodes are registered and reporting in **Ready** status. **On the master host**, run the following as root:

```
# oc get nodes
```

NAME	STATUS	AGE
master.example.com	Ready, SchedulingDisabled	165d
node1.example.com	Ready	165d
node2.example.com	Ready	165d

2. To verify that the web console is installed correctly, use the master host name and the console port number to access the console with a web browser.
For example, for a master host with a hostname of **master.openshift.com** and using the default port of **8443**, the web console would be found at:

```
https://master.openshift.com:8443/console
```

3. Now that the install has been verified, run the following command on each master and node host to add the **atomic-openshift** packages back to the list of yum excludes on the host:

```
# atomic-openshift-excluder exclude
```



NOTE

The default port for the console is **8443**. If this was changed during the installation, the port can be found at **openshift_master_console_port** in the `/etc/ansible/hosts` file.

Multiple etcd Hosts

If you installed multiple **etcd** hosts:

1. First, verify that the **etcd** package, which provides the **etcdctl** command, is installed:

```
# yum install etcd
```

2. On a etcd host, verify the **etcd** cluster health, substituting for the FQDNs of your **etcd** hosts in the following:

```
# etcdctl -C \
https://etcd1.example.com:2379,https://etcd2.example.com:2379,https://etcd3.example.com:2379 \
--ca-file=/etc/origin/master/master.etcd-ca.crt \
--cert-file=/etc/origin/master/master.etcd-client.crt \
--key-file=/etc/origin/master/master.etcd-client.key cluster-health
```

3. Also verify the member list is correct:

```
# etcdctl -C \
https://etcd1.example.com:2379,https://etcd2.example.com:2379,https://etcd3.example.com:2379 \
--ca-file=/etc/origin/master/master.etcd-ca.crt \
--cert-file=/etc/origin/master/master.etcd-client.crt \
--key-file=/etc/origin/master/master.etcd-client.key member list
```

Multiple Masters Using HAProxy

If you installed multiple masters using HAProxy as a load balancer, browse to the following URL according to your **[lb]** section definition and check HAProxy's status:

```
http://<lb_hostname>:9000
```

You can verify your installation by consulting the [HAProxy Configuration documentation](#).

2.6.9. Optionally Securing Builds

Running **docker build** is a privileged process, so the container has more access to the node than might be considered acceptable in some multi-tenant environments. If you do not trust your users, you can use a more secure option at the time of installation. Disable Docker builds on the cluster and require that users build images outside of the cluster. See [Securing Builds by Strategy](#) for more information on this optional process.

2.6.10. Uninstalling OpenShift Container Platform

You can uninstall OpenShift Container Platform hosts in your cluster by running the ***uninstall.yml*** playbook. This playbook deletes OpenShift Container Platform content installed by Ansible, including:

- Configuration
- Containers
- Default templates and image streams

- Images
- RPM packages

The playbook will delete content for any hosts defined in the inventory file that you specify when running the playbook. If you want to uninstall OpenShift Container Platform across all hosts in your cluster, run the playbook using the inventory file you used when installing OpenShift Container Platform initially or ran most recently:

```
# ansible-playbook [-i /path/to/file] \
    /usr/share/ansible/openshift-ansible/playbooks/adhoc/uninstall.yml
```

2.6.10.1. Uninstalling Nodes

You can also uninstall node components from specific hosts using the ***uninstall.yml*** playbook while leaving the remaining hosts and cluster alone:



WARNING

This method should only be used when attempting to uninstall specific node hosts and not for specific masters or etcd hosts, which would require further configuration changes within the cluster.

1. First follow the steps in [Deleting Nodes](#) to remove the node object from the cluster, then continue with the remaining steps in this procedure.
2. Create a different inventory file that only references those hosts. For example, to only delete content from one node:

```
[OSEv3:children]
nodes ❶

[OSEv3:vars]
ansible_ssh_user=root
deployment_type=openshift-enterprise

[nodes]
node3.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'west' }" ❷
```

- ❶ Only include the sections that pertain to the hosts you are interested in uninstalling.
- ❷ Only include hosts that you want to uninstall.

3. Specify that new inventory file using the ***-i*** option when running the ***uninstall.yml*** playbook:

```
# ansible-playbook -i /path/to/new/file \
    /usr/share/ansible/openshift-
    ansible/playbooks/adhoc/uninstall.yml
```

■

When the playbook completes, all OpenShift Container Platform content should be removed from any specified hosts.

2.6.11. Known Issues

- On failover in multiple master clusters, it is possible for the controller manager to overcorrect, which causes the system to run more pods than what was intended. However, this is a transient event and the system does correct itself over time. See <https://github.com/kubernetes/kubernetes/issues/10030> for details.
- On failure of the Ansible installer, you must start from a clean operating system installation. If you are using virtual machines, start from a fresh image. If you are using bare metal machines, see [Uninstalling OpenShift Container Platform](#) for instructions.

2.6.12. What's Next?

Now that you have a working OpenShift Container Platform instance, you can:

- [Configure authentication](#); by default, authentication is set to [Deny All](#).
- Deploy an [integrated Docker registry](#).
- Deploy a [router](#).

2.7. DISCONNECTED INSTALLATION

2.7.1. Overview

Frequently, portions of a datacenter may not have access to the Internet, even via proxy servers. Installing OpenShift Container Platform in these environments is considered a disconnected installation.

An OpenShift Container Platform disconnected installation differs from a regular installation in two primary ways:

- The OpenShift Container Platform software channels and repositories are not available via Red Hat's content distribution network.
- OpenShift Container Platform uses several containerized components. Normally, these images are pulled directly from Red Hat's Docker registry. In a disconnected environment, this is not possible.

A disconnected installation ensures the OpenShift Container Platform software is made available to the relevant servers, then follows the same installation process as a standard connected installation. This topic additionally details how to manually download the container images and transport them onto the relevant servers.

Once installed, in order to use OpenShift Container Platform, you will need source code in a source control repository (for example, Git). This topic assumes that an internal Git repository is available that can host source code and this repository is accessible from the OpenShift Container Platform nodes. Installing the source control repository is outside the scope of this document.

Also, when building applications in OpenShift Container Platform, your build may have some external dependencies, such as a Maven Repository or Gem files for Ruby applications. For this reason, and because they might require certain tags, many of the Quickstart templates offered by OpenShift

Container Platform may not work on a disconnected environment. However, while Red Hat container images try to reach out to external repositories by default, you can configure OpenShift Container Platform to use your own internal repositories. For the purposes of this document, we assume that such internal repositories already exist and are accessible from the OpenShift Container Platform nodes hosts. Installing such repositories is outside the scope of this document.



NOTE

You can also have a [Red Hat Satellite](#) server that provides access to Red Hat content via an intranet or LAN. For environments with Satellite, you can synchronize the OpenShift Container Platform software onto the Satellite for use with the OpenShift Container Platform servers.

[Red Hat Satellite 6.1](#) also introduces the ability to act as a Docker registry, and it can be used to host the OpenShift Container Platform containerized components. Doing so is outside of the scope of this document.

2.7.2. Prerequisites

This document assumes that you understand [OpenShift Container Platform's overall architecture](#) and that you have already planned out what the topology of your environment will look like.

2.7.3. Required Software and Components

In order to pull down the required software repositories and container images, you will need a Red Hat Enterprise Linux (RHEL) 7 server with access to the Internet and at least 100GB of additional free space. All steps in this section should be performed on the Internet-connected server as the root system user.

2.7.3.1. Syncing Repositories

Before you sync with the required repositories, you may need to import the appropriate GPG key:

```
# rpm --import /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
```

If the key is not imported, the indicated package is deleted after syncing the repository.

To sync the required repositories:

1. Register the server with the Red Hat Customer Portal. You must use the login and password associated with the account that has access to the OpenShift Container Platform subscriptions:

```
# subscription-manager register
```

2. Attach to a subscription that provides OpenShift Container Platform channels. You can find the list of available subscriptions using:

```
# subscription-manager list --available --matches '*OpenShift*'
```

Then, find the pool ID for the subscription that provides OpenShift Container Platform, and attach it:

```
# subscription-manager attach --pool=<pool_id>
# subscription-manager repos --disable=""
# subscription-manager repos \
```

```
--enable="rhel-7-server-rpms" \
--enable="rhel-7-server-extras-rpms" \
--enable="rhel-7-server-ose-3.3-rpms"
```

3. The **yum-utils** command provides the **reposync** utility, which lets you mirror yum repositories, and **createrepo** can create a usable **yum** repository from a directory:

```
# yum -y install yum-utils createrepo docker git
```

You will need up to 110GB of free space in order to sync the software. Depending on how restrictive your organization's policies are, you could re-connect this server to the disconnected LAN and use it as the repository server. You could use USB-connected storage and transport the software to another server that will act as the repository server. This topic covers these options.

4. Make a path to where you want to sync the software (either locally or on your USB or other device):

```
# mkdir -p </path/to/repos>
```

5. Sync the packages and create the repository for each of them. You will need to modify the command for the appropriate path you created above:

```
# for repo in \
rhel-7-server-rpms rhel-7-server-extras-rpms \
rhel-7-server-ose-3.3-rpms
do
    reposync --gpgcheck -lm --repoid=${repo} --
download_path=/path/to/repos
    createrepo -v </path/to/repos/>${repo} -o </path/to/repos/>${repo}
done
```

2.7.3.2. Syncing Images

To sync the container images:

1. Start the Docker daemon:

```
# systemctl start docker
```

2. Pull all of the required OpenShift Container Platform containerized components. Replace **<tag>** with **v3.3.1.25** for the latest version.

```
# docker pull registry.access.redhat.com/openshift3/ose-haproxy-
router:<tag>
# docker pull registry.access.redhat.com/openshift3/ose-deployer:
<tag>
# docker pull registry.access.redhat.com/openshift3/ose-recycler:
<tag>
# docker pull registry.access.redhat.com/openshift3/ose-sti-builder:
<tag>
# docker pull registry.access.redhat.com/openshift3/ose-docker-
builder:<tag>
# docker pull registry.access.redhat.com/openshift3/ose-pod:<tag>
```

```
# docker pull docker.io/openshift/hello-openshift:latest
# docker pull registry.access.redhat.com/openshift3/ose-docker-
registry:<tag>
```

3. Pull all of the required OpenShift Container Platform containerized components for the additional centralized log aggregation and metrics aggregation components. Replace **<tag>** with **3.3.1** for the latest version.

```
# docker pull registry.access.redhat.com/openshift3/logging-
deployer:<tag>
# docker pull registry.access.redhat.com/openshift3/logging-
elasticsearch:<tag>
# docker pull registry.access.redhat.com/openshift3/logging-kibana:
<tag>
# docker pull registry.access.redhat.com/openshift3/logging-fluentd:
<tag>
# docker pull registry.access.redhat.com/openshift3/logging-curator:
<tag>
# docker pull registry.access.redhat.com/openshift3/logging-auth-
proxy:<tag>
# docker pull registry.access.redhat.com/openshift3/metrics-
deployer:<tag>
# docker pull registry.access.redhat.com/openshift3/metrics-
hawkular-metrics:<tag>
# docker pull registry.access.redhat.com/openshift3/metrics-
cassandra:<tag>
# docker pull registry.access.redhat.com/openshift3/metrics-
heapster:<tag>
```

4. Pull the Red Hat-certified [Source-to-Image \(S2I\)](#) builder images that you intend to use in your OpenShift environment. You can pull the following images:

- jboss-eap70-openshift
- jboss-amq-62
- jboss-datagrid65-openshift
- jboss-decisionserver62-openshift
- jboss-eap64-openshift
- jboss-eap70-openshift
- jboss-webserver30-tomcat7-openshift
- jboss-webserver30-tomcat8-openshift
- mongodb
- mysql
- nodejs
- perl
- php

- postgresql
- python
- redhat-sso70-openshift
- ruby

Make sure to indicate the correct tag specifying the desired version number. For example, to pull both the previous and latest version of the Tomcat image:

```
# docker pull \
registry.access.redhat.com/jboss-webserver-3/webserver30-tomcat7-
openshift:latest
# docker pull \
registry.access.redhat.com/jboss-webserver-3/webserver30-tomcat7-
openshift:1.1
```

5. If you are using a stand-alone registry or plan to enable the registry console with the integrated registry, you must pull the **registry-console** image.
Replace **<tag>** with **3.3** for the latest version.

```
# docker pull registry.access.redhat.com/openshift3/registry-
console:<tag>
```

2.7.3.3. Preparing Images for Export

Container images can be exported from a system by first saving them to a tarball and then transporting them:

1. Make and change into a repository home directory:

```
# mkdir </path/to/repos/images>
# cd </path/to/repos/images>
```

2. Export the OpenShift Container Platform containerized components:

```
# docker save -o ose3-images.tar \
registry.access.redhat.com/openshift3/ose-haproxy-router \
registry.access.redhat.com/openshift3/ose-deployer \
registry.access.redhat.com/openshift3/ose-recycler \
registry.access.redhat.com/openshift3/ose-sti-builder \
registry.access.redhat.com/openshift3/ose-docker-builder \
registry.access.redhat.com/openshift3/ose-pod \
docker.io/openshift/hello-openshift \
registry.access.redhat.com/openshift3/ose-docker-registry
```

3. If you synchronized the metrics and log aggregation images, export:

```
# docker save -o ose3-logging-metrics-images.tar \
registry.access.redhat.com/openshift3/logging-deployer \
registry.access.redhat.com/openshift3/logging-elasticsearch \
registry.access.redhat.com/openshift3/logging-kibana \
registry.access.redhat.com/openshift3/logging-fluentd \
registry.access.redhat.com/openshift3/logging-auth-proxy \
```



```
registry.access.redhat.com/openshift3/metrics-deployer \
registry.access.redhat.com/openshift3/metrics-hawkular-metrics \
registry.access.redhat.com/openshift3/metrics-cassandra \
registry.access.redhat.com/openshift3/metrics-heapster
```

4. Export the S2I builder images that you synced in the previous section. For example, if you synced only the Tomcat image:

```
# docker save -o ose3-builder-images.tar \
  registry.access.redhat.com/jboss-webserver-3/webserver30-
tomcat7-openshift:latest \
  registry.access.redhat.com/jboss-webserver-3/webserver30-
tomcat7-openshift:1.1
```

2.7.4. Repository Server

During the installation (and for later updates, should you so choose), you will need a webserver to host the repositories. RHEL 7 can provide the Apache webserver.

Option 1: Re-configuring as a Web server

If you can re-connect the server where you synchronized the software and images to your LAN, then you can simply install Apache on the server:

```
# yum install httpd
```

Skip to [Placing the Software](#).

Option 2: Building a Repository Server

If you need to build a separate server to act as the repository server, install a new RHEL 7 system with at least 110GB of space. On this repository server during the installation, make sure you select the **Basic Web Server** option.

2.7.4.1. Placing the Software

1. If necessary, attach the external storage, and then copy the repository files into Apache's root folder. Note that the below copy step (**cp -a**) should be substituted with move (**mv**) if you are repurposing the server you used to sync:

```
# cp -a /path/to/repos /var/www/html/
# chmod -R +r /var/www/html/repos
# restorecon -vR /var/www/html
```

2. Add the firewall rules:

```
# firewall-cmd --permanent --add-service=http
# firewall-cmd --reload
```

3. Enable and start Apache for the changes to take effect:

```
# systemctl enable httpd
# systemctl start httpd
```

2.7.5. OpenShift Container Platform Systems

2.7.5.1. Building Your Hosts

At this point you can perform the initial creation of the hosts that will be part of the OpenShift Container Platform environment. It is recommended to use the latest version of RHEL 7 and to perform a minimal installation. You will also want to pay attention to the other [OpenShift Container Platform-specific prerequisites](#).

Once the hosts are initially built, the repositories can be set up.

2.7.5.2. Connecting the Repositories

On all of the relevant systems that will need OpenShift Container Platform software components, create the required repository definitions. Place the following text in the `/etc/yum.repos.d/ose.repo` file, replacing `<server_IP>` with the IP or host name of the Apache server hosting the software repositories:

```
[rhel-7-server-rpms]
name=rhel-7-server-rpms
baseurl=http://<server_IP>/repos/rhel-7-server-rpms
enabled=1
gpgcheck=0
[rhel-7-server-extras-rpms]
name=rhel-7-server-extras-rpms
baseurl=http://<server_IP>/repos/rhel-7-server-extras-rpms
enabled=1
gpgcheck=0
[rhel-7-server-ose-3.3-rpms]
name=rhel-7-server-ose-3.3-rpms
baseurl=http://<server_IP>/repos/rhel-7-server-ose-3.3-rpms
enabled=1
gpgcheck=0
```

2.7.5.3. Host Preparation

At this point, the systems are ready to continue to be prepared [following the OpenShift Container Platform documentation](#).

Skip the section titled **Host Registration** and start with **Installing Base Packages**.

2.7.6. Installing OpenShift Container Platform

2.7.6.1. Importing OpenShift Container Platform Containerized Components

To import the relevant components, securely copy the images from the connected host to the individual OpenShift Container Platform hosts:

```
# scp /var/www/html/repos/images/ose3-images.tar
root@<openshift_host_name>:
# ssh root@<openshift_host_name> "docker load -i ose3-images.tar"
```

If you prefer, you could use **wget** on each OpenShift Container Platform host to fetch the tar file, and then perform the Docker import command locally. Perform the same steps for the metrics and logging images, if you synchronized them.

On the host that will act as an OpenShift Container Platform master, copy and import the builder images:

```
# scp /var/www/html/images/ose3-builder-images.tar
root@<openshift_master_host_name>:
# ssh root@<openshift_master_host_name> "docker load -i ose3-builder-
images.tar"
```

2.7.6.2. Running the OpenShift Container Platform Installer

You can now choose to follow the [quick](#) or [advanced](#) OpenShift Container Platform installation instructions in the documentation.

2.7.6.3. Creating the Internal Docker Registry

You now need to [create the internal Docker registry](#).

If you want to [install a stand-alone registry](#), you must [pull the registry-console container image](#) and set **deployment_subtype=registry** in the inventory file.

2.7.7. Post-Installation Changes

In one of the previous steps, the S2I images were imported into the Docker daemon running on one of the OpenShift Container Platform master hosts. In a connected installation, these images would be pulled from Red Hat's registry on demand. Since the Internet is not available to do this, the images must be made available in another Docker registry.

OpenShift Container Platform provides an internal registry for storing the images that are built as a result of the S2I process, but it can also be used to hold the S2I builder images. The following steps assume you did not customize the service IP subnet (172.30.0.0/16) or the Docker registry port (5000).

2.7.7.1. Re-tagging S2I Builder Images

1. On the master host where you imported the S2I builder images, obtain the service address of your Docker registry that you installed on the master:

```
# export REGISTRY=$(oc get service docker-registry -t
'{{.spec.clusterIP}}{{"\n"}}')
```

2. Next, tag all of the builder images that you synced and exported before pushing them into the OpenShift Container Platform Docker registry. For example, if you synced and exported only the Tomcat image:

```
# docker tag \
registry.access.redhat.com/jboss-webserver-3/webserver30-tomcat7-
openshift:1.1 \
$REGISTRY:5000/openshift/webserver30-tomcat7-openshift:1.1
# docker tag \
registry.access.redhat.com/jboss-webserver-3/webserver30-tomcat7-
openshift:latest \
$REGISTRY:5000/openshift/webserver30-tomcat7-openshift:1.2
```

```
# docker tag \
registry.access.redhat.com/jboss-webserver-3/webserver30-tomcat7-
openshift:latest \
$REGISTRY:5000/openshift/webserver30-tomcat7-openshift:latest
```

2.7.7.2. Configuring a Registry Location

If you are using an image registry other than the default at **registry.access.redhat.com**, specify the desired registry within the **/etc/ansible/hosts** file.

```
oreg_url=example.com/openshift3/ose-${component}:${version}
openshift_examples_modify_imagestreams=true
```

Depending on your registry, you may need to configure:

```
openshift_docker_additional_registries=example.com
openshift_docker_insecure_registries=example.com
```

Table 2.13. Registry Variables

Variable	Purpose
oreg_url	Set to the alternate image location. Necessary if you are not using the default registry at registry.access.redhat.com .
openshift_examples_modify_imagestreams	Set to true if pointing to a registry other than the default. Modifies the image stream location to the value of oreg_url .
openshift_docker_additional_registries	Set openshift_docker_additional_registries to add its value in the add_registry line in /etc/sysconfig/docker . With add_registry , you can add your own registry to be used for Docker search and Docker pull. Use the add_registry option to list a set of registries, each prepended with --add-registry flag. The first registry added will be the first registry searched. For example, add_registry=--add-registry registry.access.redhat.com --add-registry example.com .

Variable	Purpose
openshift_docker_insecure_registries	Set openshift_docker_insecure_registries to add its value in the insecure_registry line in /etc/sysconfig/docker . If you have a registry secured with HTTPS but do not have proper certificates distributed, you can tell Docker not to look for full authorization by adding the registry to the insecure_registry line and uncommenting it. For example, insecure_registry= insecure-registry example.com .

2.7.7.3. Creating an Administrative User

Pushing the container images into OpenShift Container Platform's Docker registry requires a user with **cluster-admin** privileges. Because the default OpenShift Container Platform system administrator does not have a standard authorization token, they cannot be used to log in to the Docker registry.

To create an administrative user:

1. Create a new user account in the authentication system you are using with OpenShift Container Platform. For example, if you are using local **htpasswd**-based authentication:

```
# htpasswd -b /etc/openshift/openshift-passwd <admin_username>
<password>
```

2. The external authentication system now has a user account, but a user must log in to OpenShift Container Platform before an account is created in the internal database. Log in to OpenShift Container Platform for this account to be created. This assumes you are using the self-signed certificates generated by OpenShift Container Platform during the installation:

```
# oc login --certificate-authority=/etc/origin/master/ca.crt \
-u <admin_username> https://<openshift_master_host>:8443
```

3. Get the user's authentication token:

```
# MYTOKEN=$(oc whoami -t)
# echo $MYTOKEN
iwo7hc4Xi1D2K0LL4V1055ExH2V1PmLD-W2-J0d6Fko
```

2.7.7.4. Modifying the Security Policies

1. Using **oc login** switches to the new user. Switch back to the OpenShift Container Platform system administrator in order to make policy changes:

```
# oc login -u system:admin
```

2. In order to push images into the OpenShift Container Platform Docker registry, an account must have the **image-builder** security role. Add this to your OpenShift Container Platform administrative user:

```
# oadm policy add-role-to-user system:image-builder <admin_username>
```

3. Next, add the administrative role to the user in the **openshift** project. This allows the administrative user to edit the **openshift** project, and, in this case, push the container images:

```
# oadm policy add-role-to-user admin <admin_username> -n openshift
```

2.7.7.5. Editing the Image Stream Definitions

The **openshift** project is where all of the image streams for builder images are created by the installer. They are loaded by the installer from the **/usr/share/openshift/examples** directory. Change all of the definitions by deleting the image streams which had been loaded into OpenShift Container Platform's database, then re-create them:

1. Delete the existing image streams:

```
# oc delete is -n openshift --all
```

2. Make a backup of the files in **/usr/share/openshift/examples/** if you desire. Next, edit the file **image-streams-rhel7.json** in the **/usr/share/openshift/examples/image-streams** folder. You will find an image stream section for each of the builder images. Edit the **spec** stanza to point to your internal Docker registry.

For example, change:

```
"spec": {
  "dockerImageRepository":
    "registry.access.redhat.com/rhsc1/mongodb-26-rhel7",
```

to:

```
"spec": {
  "dockerImageRepository": "172.30.69.44:5000/openshift/mongodb-26-
rhel7",
```

In the above, the repository name was changed from **rhsc1** to **openshift**. You will need to ensure the change, regardless of whether the repository is **rhsc1**, **openshift3**, or another directory. Every definition should have the following format:

```
<registry_ip>:5000/openshift/<image_name>
```

Repeat this change for every image stream in the file. Ensure you use the correct IP address that you determined earlier. When you are finished, save and exit. Repeat the same process for the JBoss image streams in the **/usr/share/openshift/examples/xpaas-streams/jboss-image-streams.json** file.

3. Load the updated image stream definitions:

```
# oc create -f /usr/share/openshift/examples/image-streams/image-
streams-rhel7.json -n openshift
# oc create -f /usr/share/openshift/examples/xpaas-streams/jboss-
image-streams.json -n openshift
```

2.7.7.6. Loading the Container Images

At this point the system is ready to load the container images.

1. Log in to the Docker registry using the token and registry service IP obtained earlier:

```
# docker login -u adminuser -e mailto:adminuser@abc.com \
  -p $MYTOKEN $REGISTRY:5000
```

2. Push the Docker images:

```
# docker push $REGISTRY:5000/openshift/webserver30-tomcat7-
openshift:1.1
# docker push $REGISTRY:5000/openshift/webserver30-tomcat7-
openshift:1.2
# docker push $REGISTRY:5000/openshift/webserver30-tomcat7-
openshift:latest
```

3. Verify that all the image streams now have the tags populated:

```
# oc get imagestreams -n openshift
NAME                                     DOCKER REPO
TAGS                                     UPDATED
jboss-webserver30-tomcat7-openshift    $REGISTRY/jboss-webserver-
3/webserver30-jboss-tomcat7-openshift  1.1,1.1-2,1.1-6 + 2 more...
2 weeks ago
...
```

2.7.8. Installing a Router

At this point, the OpenShift Container Platform environment is almost ready for use. It is likely that you will want to [install and configure a router](#).

2.8. INSTALLING A STAND-ALONE DEPLOYMENT OF OPENSIFT CONTAINER REGISTRY

2.8.1. About OpenShift Container Registry

OpenShift Container Platform is a fully-featured enterprise solution that includes an integrated container registry called [OpenShift Container Registry](#) (OCR). Alternatively, instead of deploying OpenShift Container Platform as a full PaaS environment for developers, you can install OCR as a stand-alone container registry to run on-premise or in the cloud.

When installing a stand-alone deployment of OCR, a cluster of masters and nodes is still installed, similar to a typical OpenShift Container Platform installation. Then, the container registry is deployed to run on the cluster. This stand-alone deployment option is useful for administrators that want a container registry, but do not require the full OpenShift Container Platform environment that includes the developer-focused web console and application build and deployment tools.

**NOTE**

OCR should not be confused with the upstream project [Atomic Registry](#), which is a different implementation using a non-Kubernetes deployment method that leverages **systemd** and local configuration files to manage services.

OCR provides the following capabilities:

- A user-focused [registry web console](#).
- [Secured traffic](#) by default, served via TLS.
- Global [identity provider authentication](#).
- A [project namespace](#) model to enable teams to collaborate through [role-based access control \(RBAC\)](#) authorization.
- A [Kubernetes-based cluster](#) to manage services.
- An image abstraction called [image streams](#) to enhance image management.

Administrators may want to deploy a stand-alone OCR to manage a registry separately that supports multiple OpenShift Container Platform clusters. A stand-alone OCR also enables administrators to separate their registry to satisfy their own security or compliance requirements.

2.8.2. Minimum Hardware Requirements

Installing a stand-alone OCR has the following hardware requirements:

- Physical or virtual system, or an instance running on a public or private IaaS.
- Base OS: RHEL 7.3 with the "Minimal" installation option and the latest packages from the RHEL 7 Extras channel, or RHEL Atomic Host 7.3.2 or later. RHEL 7.2 is also supported using Docker 1.12 and its dependencies.
- NetworkManager 1.0 or later
- 2 vCPU.
- Minimum 16 GB RAM.
- Minimum 15 GB hard disk space for the file system containing **/var/**.
- An additional minimum 15 GB unallocated space to be used for Docker's storage back end; see [Configuring Docker Storage](#) for details.

**IMPORTANT**

OpenShift Container Platform only supports servers with x86_64 architecture.

**NOTE**

Meeting the **/var/** file system sizing requirements in RHEL Atomic Host requires making changes to the default configuration. See [Managing Storage in Red Hat Enterprise Linux Atomic Host](#) for instructions on configuring this during or after installation.

2.8.3. Supported System Topologies

The following system topologies are supported for stand-alone OCR:

All-in-one	A single host that includes the master, node, etcd, and registry components.
Multiple Masters (Highly-Available)	Three hosts with all components included on each (master, node, etcd, and registry), with the masters configured for native high-availability.

2.8.4. Host Preparation

Before installing stand-alone OCR, all of the same steps detailed in the [Host Preparation](#) topic for installing a full OpenShift Container Platform PaaS must be performed. This includes registering and subscribing the host(s) to the proper repositories, installing or updating certain packages, and setting up Docker and its storage requirements.

Follow the steps in the [Host Preparation](#) topic, then continue to [Installation Methods](#).

2.8.5. Installation Methods

To install a stand-alone registry, use either of the standard installation methods (quick or advanced) used to install any variant of OpenShift Container Platform.

2.8.5.1. Quick Installation for Stand-alone OpenShift Container Registry

When using the quick installation method to install stand-alone OCR, start the interactive installation by running:

```
$ atomic-openshift-installer install
```

Then follow the on-screen instructions to install a new registry. The installation questions will be largely the same as if you were installing a full OpenShift Container Platform PaaS, but when you reach the following screen:

```
Which variant would you like to install?
```

```
(1) OpenShift Container Platform 3.3
(2) Registry 3.3
```

Be sure to choose **2** to follow the registry installation path.



NOTE

For further usage details on the quick installer in general, see the full topic at [Quick Installation](#).

2.8.5.2. Advanced Installation for Stand-alone OpenShift Container Registry

When using the advanced installation method to install stand-alone OCR, use the same steps for

installing a full OpenShift Container Platform PaaS using Ansible described in the full [Advanced Installation](#) topic. The main difference is that you must set **deployment_subtype=registry** in the inventory file within the **[OSEv3:vars]** section for the playbooks to follow the registry installation path.

See the following example inventory files for the different supported system topologies:

All-in-one Stand-alone OpenShift Container Registry Inventory File

```
# Create an OSEv3 group that contains the masters and nodes groups
[OSEv3:children]
masters
nodes

# Set variables common for all OSEv3 hosts
[OSEv3:vars]
# SSH user, this user should allow ssh based auth without requiring a
password
ansible_ssh_user=root

openshift_master_default_subdomain=apps.test.example.com

# If ansible_ssh_user is not root, ansible_become must be set to true
#ansible_become=true

deployment_type=openshift-enterprise
deployment_subtype=registry ❶

# uncomment the following to enable htpasswd authentication; defaults to
DenyAllPasswordIdentityProvider
#openshift_master_identity_providers=[{'name': 'htpasswd_auth', 'login':
'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/htpasswd'}]

# host group for masters
[masters]
registry.example.com

# host group for nodes, includes region info
[nodes]
registry.example.com openshift_node_labels="{ 'region': 'infra', 'zone':
'default' }" openshift_schedulable=true ❷
```

- ❶ Set **deployment_subtype=registry** to ensure installation of stand-alone OCR and not a full OpenShift Container Platform environment.
- ❷ Set **openshift_schedulable=true** on the node entry to make the single node schedulable for pod placement.

Multiple Masters (Highly-Available) Stand-alone OpenShift Container Registry Inventory File

```
# Create an OSEv3 group that contains the master, nodes, etcd, and lb
groups.
# The lb group lets Ansible configure HAProxy as the load balancing
```

```

solution.
# Comment lb out if your load balancer is pre-configured.
[OSEv3:children]
masters
nodes
etcd
lb

# Set variables common for all OSEv3 hosts
[OSEv3:vars]
ansible_ssh_user=root
deployment_type=openshift-enterprise
deployment_subtype=registry ❶

openshift_master_default_subdomain=apps.test.example.com

# Uncomment the following to enable httpasswd authentication; defaults to
# DenyAllPasswordIdentityProvider.
#openshift_master_identity_providers=[{'name': 'httpasswd_auth', 'login':
'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/httpasswd'}]

# Native high availability cluster method with optional load balancer.
# If no lb group is defined installer assumes that a load balancer has
# been preconfigured. For installation the value of
# openshift_master_cluster_hostname must resolve to the load balancer
# or to one or all of the masters defined in the inventory if no load
# balancer is present.
openshift_master_cluster_method=native
openshift_master_cluster_hostname=openshift-internal.example.com
openshift_master_cluster_public_hostname=openshift-cluster.example.com

# apply updated node defaults
openshift_node_kubelet_args={'pods-per-core': ['10'], 'max-pods': ['250'],
'image-gc-high-threshold': ['90'], 'image-gc-low-threshold': ['80']}

# override the default controller lease ttl
#osm_controller_lease_ttl=30

# enable ntp on masters to ensure proper failover
openshift_clock_enabled=true

# host group for masters
[masters]
master1.example.com
master2.example.com
master3.example.com

# host group for etcd
[etcd]
etcd1.example.com
etcd2.example.com
etcd3.example.com

# Specify load balancer host
[lb]

```


```
lb.example.com
```

```
# host group for nodes, includes region info
[nodes]
master[1:3].example.com openshift_node_labels="{ 'region': 'infra', 'zone':
'default' }" openshift_schedulable=true
node1.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary', 'zone':
'west' }"
```

- 1 Set **deployment_subtype=registry** to ensure installation of stand-alone OCR and not a full OpenShift Container Platform environment.

After you have configured Ansible by defining an inventory file in */etc/ansible/hosts*, you can run the advanced installation using the following playbook:

```
# ansible-playbook /usr/share/ansible/openshift-
ansible/playbooks/byo/config.yml
```



NOTE

For more detailed usage information on the advanced installation method, including a comprehensive list of available Ansible variables, see the full topic at [Advanced Installation](#).

CHAPTER 3. SETTING UP THE REGISTRY

3.1. REGISTRY OVERVIEW

3.1.1. About the Registry

OpenShift Container Platform can build [container images](#) from your source code, deploy them, and manage their lifecycle. To enable this, OpenShift Container Platform provides an internal, [integrated Docker registry](#) that can be deployed in your OpenShift Container Platform environment to locally manage images.

3.1.2. Integrated or Stand-alone Registries

During an initial installation of a full OpenShift Container Platform cluster, it is likely that the registry was deployed automatically during the installation process. If it was not, or if you want to further customize the configuration of your registry, see [Deploying a Registry on Existing Clusters](#).

While it can be deployed to run as an integrated part of your full OpenShift Container Platform cluster, the OpenShift Container Platform registry can alternatively be installed separately as a stand-alone container image registry.

To install a stand-alone registry, follow [Installing a Stand-alone Registry](#). This installation path deploys an all-in-one cluster running a registry and specialized web console.

3.2. DEPLOYING A REGISTRY ON EXISTING CLUSTERS

3.2.1. Overview

If the integrated registry was not previously deployed automatically during the initial installation of your OpenShift Container Platform cluster, or if it is no longer running successfully and you need to redeploy it on your existing cluster, see the following sections for options on deploying a new registry.



NOTE

This topic is not required if you installed a [stand-alone registry](#).

3.2.2. Deploying the Registry

To deploy the integrated Docker registry, use the **oadm registry** command as a user with cluster administrator privileges. For example:

```
$ oadm registry --config=/etc/origin/master/admin.kubeconfig \ ❶
    --service-account=registry \ ❷
    --images='registry.access.redhat.com/openshift3/ose-
    ${component}:${version}' ❸
```

❶ **--config** is the path to the [CLI configuration file](#) for the [cluster administrator](#).

❷ **--service-account** is the service account used to run the registry's pod.

❸ Required to pull the correct image for OpenShift Container Platform.

This creates a service and a deployment configuration, both called **docker-registry**. Once deployed successfully, a pod is created with a name similar to **docker-registry-1-cpty9**.

To see a full list of options that you can specify when creating the registry:

```
$ oadm registry --help
```

3.2.3. Deploying the Registry as a DaemonSet

Use the **oadm registry** command to deploy the registry as a **DaemonSet** with the **--daemonset** option.

Daemonsets ensure that when nodes are created, they contain copies of a specified pod. When the nodes are removed, the pods are garbage collected.

For more information on **DaemonSets**, see [Using Daemonsets](#).

3.2.4. Registry Compute Resources

By default, the registry is created with no settings for [compute resource requests or limits](#). For production, it is highly recommended that the deployment configuration for the registry be updated to set resource requests and limits for the registry pod. Otherwise, the registry pod will be considered a **BestEffort** pod.

See [Compute Resources](#) for more information on configuring requests and limits.

3.2.5. Storage for the Registry

The registry stores container images and metadata. If you simply deploy a pod with the registry, it uses an ephemeral volume that is destroyed if the pod exits. Any images anyone has built or pushed into the registry would disappear.

This section lists the supported [registry storage drivers](#).


The following list includes storage drivers that need to be configured in the registry's configuration file:

- [Filesystem](#). Filesystem is the default and does not need to be configured.
- [S3](#). Learn more about [CloudFront configuration](#).
- [OpenStack Swift](#)
- [Google Cloud Storage \(GCS\)](#)
- [Microsoft Azure](#)
- [Aliyun OSS](#)

[General registry storage configuration options](#) are supported.

The following storage options need to be configured through the [filesystem driver](#):

- [Backing Docker Registry with GlusterFS Storage](#)
- [Ceph Rados Block Device](#)

**NOTE**


For more information on supported persistent storage drivers, see [Configuring Persistent Storage](#) and [Persistent Storage Examples](#).

3.2.5.1. Production Use

For production use, attach a remote volume or [define and use the persistent storage method of your choice](#).

For example, to use an existing persistent volume claim:

```
$ oc volume deploymentconfigs/docker-registry --add --name=registry-
storage -t pvc \
    --claim-name=<pvc_name> --overwrite
```

**NOTE**

See [Known Issues](#) if using a scaled registry with a shared NFS volume.

3.2.5.1.1. Use Amazon S3 as a Storage Back-end

There is also an option to use Amazon Simple Storage Service storage with the internal Docker registry. It is a secure cloud storage manageable through [AWS Management Console](#). To use it, the registry's configuration file must be manually edited and mounted to the registry pod. However, before you start with the configuration, look at upstream's [recommended steps](#).

Take a [default YAML configuration file](#) as a base and replace the **filesystem** entry in the **storage** section with **s3** entry such as below. The resulting storage section may look like this:

```
storage:
  cache:
    layerinfo: inmemory
  delete:
    enabled: true
  s3:
    accesskey: awsaccesskey ❶
    secretkey: awssecretkey ❷
    region: us-west-1
    regionendpoint: http://myobjects.local
    bucket: bucketname
    encrypt: true
    keyid: mykeyid
    secure: true
    v4auth: false
    chunksize: 5242880
    rootdirectory: /s3/object/name/prefix
```

❶ Replace with your Amazon access key.

❷ Replace with your Amazon secret key.

All of the **s3** configuration options are documented in upstream's [driver reference documentation](#).

[Overriding the registry configuration](#) will take you through the additional steps on mounting the configuration file into pod.




WARNING

When the registry runs on the S3 storage back-end, there are [reported issues](#).

3.2.5.2. Non-Production Use

For non-production use, you can use the `--mount-host=<path>` option to specify a directory for the registry to use for persistent storage. The registry volume is then created as a host-mount at the specified `<path>`.



IMPORTANT

The `--mount-host` option mounts a directory from the node on which the registry container lives. If you scale up the **docker-registry** deployment configuration, it is possible that your registry pods and containers will run on different nodes, which can result in two or more registry containers, each with its own local storage. This will lead to unpredictable behavior, as subsequent requests to pull the same image repeatedly may not always succeed, depending on which container the request ultimately goes to.

The `--mount-host` option requires that the registry container run in privileged mode. This is automatically enabled when you specify `--mount-host`. However, not all pods are allowed to run [privileged containers](#) by default. If you still want to use this option, create the registry and specify that it use the **registry** service account that was created during installation:

```
$ oadm registry --service-account=registry \
  --config=/etc/origin/master/admin.kubeconfig \
  --images='registry.access.redhat.com/openshift3/ose-
  ${component}:${version}' \
  --mount-host=<path>
```



IMPORTANT

The Docker registry pod runs as user **1001**. This user must be able to write to the host directory. You may need to change directory ownership to user ID **1001** with this command:

```
$ sudo chown 1001:root <path>
```

3.2.6. Enabling the Registry Console


OpenShift Container Platform provides a web-based interface to the integrated registry. This registry console is an optional component for browsing and managing images. It is deployed as a stateless service running as a pod.

**NOTE**

If you installed OpenShift Container Platform as a [stand-alone registry](#), the registry console is already deployed and secured automatically during installation.

**IMPORTANT**

If Cockpit is already running, you'll need to shut it down before proceeding in order to avoid a port conflict (9090 by default) with the registry console.

3.2.6.1. Deploying the Registry Console**IMPORTANT**

You must first have [exposed the registry](#).

1. Create a passthrough route in the **default** project. You will need this when creating the registry console application in the next step.

```
$ oc create route passthrough --service registry-console \
  --port registry-console \
  -n default
```

2. Deploy the registry console application. Replace **<openshift_oauth_url>** with the URL of the OpenShift Container Platform OAuth provider, which is typically the master.

```
$ oc new-app -n default --template=registry-console -p \
  OPENSIFT_OAUTH_PROVIDER_URL="https://<openshift_oauth_url>:8443",RE
  GISTRY_HOST=$(oc get route docker-registry -n default --template='{{
  .spec.host }}'),COCKPIT_KUBE_URL=$(oc get route registry-console -n
  default --template='https://{{ .spec.host }}')
```

**NOTE**

If the redirection URL is wrong when you are trying to log in to the registry console, check your OAuth client with **oc get oauthclients**.

1. Finally, use a web browser to view the console using the route URI.

3.2.6.2. Securing the Registry Console

By default, the registry console generates self-signed TLS certificates if deployed manually per the steps in [Deploying the Registry Console](#). See [Troubleshooting the Registry Console](#) for more information.

Use the following steps to add your organization's signed certificates as a secret volume. This assumes your certificates are available on the **oc** client host.

1. Create a **.cert** file containing the certificate and key. Format the file with:
 - One or more **BEGIN CERTIFICATE** blocks for the server certificate and the intermediate certificate authorities

- A block containing a **BEGIN PRIVATE KEY** or similar for the key. The key must not be encrypted

For example:

```
-----BEGIN CERTIFICATE-----
MIIDUzCCAjugAwIBAgIJAPXW+CuNYS6QMA0GCSqGSIb3DQEBCwUAMD8xKTAnBgNV
BAoMIGI00GE2NGNkNmMwNTQ1YThhZTgx0TEzZDE5YmJjMmRjMRIwEAYDVQQDDAIs
...
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
MIIDUzCCAjugAwIBAgIJAPXW+CuNYS6QMA0GCSqGSIb3DQEBCwUAMD8xKTAnBgNV
BAoMIGI00GE2NGNkNmMwNTQ1YThhZTgx0TEzZDE5YmJjMmRjMRIwEAYDVQQDDAIs
...
-----END CERTIFICATE-----
-----BEGIN PRIVATE KEY-----
MIIEvgIBADANBgkqhkiG9w0BAQEFAASCBAkwggSkAgEAAoIBAQCyoJ5garOYw0sm
8TBCDSqQ/H1awGMzDYdB11xuHHsxYS2VepPMzMzryHR137I4dGFLhvdTvJUH81US
...
-----END PRIVATE KEY-----
```

- The secured registry should contain the following Subject Alternative Names (SAN) list:

- Two service hostnames.

For example:

```
docker-registry.default.svc.cluster.local
docker-registry.default.svc
```

- Service IP address.

For example:

```
172.30.124.220
```

Use the following command to get the Docker registry service IP address:

```
oc get service docker-registry --
template='{{.spec.clusterIP}}'
```

- Public hostname.

For example:

```
docker-registry-default.apps.example.com
```

Use the following command to get the Docker registry public hostname:

```
oc get route docker-registry --template '{{.spec.host}}'
```

For example, the server certificate should contain SAN details similar to the following:

```
X509v3 Subject Alternative Name:
          DNS:docker-registry-public.openshift.com,
          DNS:docker-registry.default.svc, DNS:docker-
```

```
registry.default.svc.cluster.local, DNS:172.30.2.98, IP
Address:172.30.2.98
```

The registry console loads a certificate from the `/etc/cockpit/ws-certs.d` directory. It uses the last file with a `.cert` extension in alphabetical order. Therefore, the `.cert` file should contain at least two PEM blocks formatted in the OpenSSL style.

If no certificate is found, a self-signed certificate is created using the `openssl` command and stored in the `0-self-signed.cert` file.

2. Create the secret:

```
$ oc secrets new console-secret \
/path/to/console.cert
```

3. Add the secrets to the **registry-console** deployment configuration:

```
$ oc volume dc/registry-console --add --type=secret \
--secret-name=console-secret -m /etc/cockpit/ws-certs.d
```

This triggers a new deployment of the registry console to include your signed certificates.

3.2.6.3. Troubleshooting the Registry Console

3.2.6.3.1. Debug Mode

The registry console debug mode is enabled using an environment variable. The following command redeploys the registry console in debug mode:

```
$ oc set env dc registry-console G_MESSAGES_DEBUG=cockpit-ws,cockpit-
wrapper
```

Enabling debug mode allows more verbose logging to appear in the registry console's pod logs.

3.2.6.3.2. Display SSL Certificate Path

To check which certificate the registry console is using, a command can be run from inside the console pod.

1. List the pods in the **default** project and find the registry console's pod name:

```
$ oc get pods -n default
```

NAME	READY	STATUS	RESTARTS	AGE
registry-console-1-rssrw	1/1	Running	0	1d

2. Using the pod name from the previous command, get the certificate path that the **cockpit-ws** process is using. This example shows the console using the auto-generated certificate:

```
$ oc exec registry-console-1-rssrw remotectl certificate
certificate: /etc/cockpit/ws-certs.d/0-self-signed.cert
```

3.3. ACCESSING THE REGISTRY

3.3.1. Viewing Logs

To view the logs for the Docker registry, use the **oc logs** command with the deployment config:

```
$ oc logs dc/docker-registry
2015-05-01T19:48:36.300593110Z time="2015-05-01T19:48:36Z" level=info
msg="version=v2.0.0+unknown"
2015-05-01T19:48:36.303294724Z time="2015-05-01T19:48:36Z" level=info
msg="redis not configured" instance.id=9ed6c43d-23ee-453f-9a4b-
031fea646002
2015-05-01T19:48:36.303422845Z time="2015-05-01T19:48:36Z" level=info
msg="using inmemory layerinfo cache" instance.id=9ed6c43d-23ee-453f-9a4b-
031fea646002
2015-05-01T19:48:36.303433991Z time="2015-05-01T19:48:36Z" level=info
msg="Using OpenShift Auth handler"
2015-05-01T19:48:36.303439084Z time="2015-05-01T19:48:36Z" level=info
msg="listening on :5000" instance.id=9ed6c43d-23ee-453f-9a4b-031fea646002
```

3.3.2. File Storage

Tag and image metadata is stored in OpenShift Container Platform, but the registry stores layer and signature data in a volume that is mounted into the registry container at **/registry**. As **oc exec** does not work on privileged containers, to view a registry's contents you must manually SSH into the node housing the registry pod's container, then run **docker exec** on the container itself:

1. List the current pods to find the pod name of your Docker registry:

```
# oc get pods
```

Then, use **oc describe** to find the host name for the node running the container:

```
# oc describe pod <pod_name>
```

2. Log into the desired node:

```
# ssh node.example.com
```

3. List the running containers on the node host and identify the container ID for the Docker registry:

```
# docker ps | grep ose-docker-registry
```

4. List the registry contents using the **docker exec** command:

```
# docker exec -it 4c01db0b339c find /registry
/registry/docker
/registry/docker/registry
/registry/docker/registry/v2
/registry/docker/registry/v2/blobs 1
/registry/docker/registry/v2/blobs/sha256
/registry/docker/registry/v2/blobs/sha256/ed
/registry/docker/registry/v2/blobs/sha256/ed/ede17b139a271d6b1331ca3
d83c648c24f92cece5f89d95ac6c34ce751111810
/registry/docker/registry/v2/blobs/sha256/ed/ede17b139a271d6b1331ca3
```

```

d83c648c24f92cece5f89d95ac6c34ce751111810/data 2
/registry/docker/registry/v2/blobs/sha256/a3
/registry/docker/registry/v2/blobs/sha256/a3/a3ed95caeb02ffe68cdd9fd
84406680ae93d633cb16422d00e8a7c22955b46d4
/registry/docker/registry/v2/blobs/sha256/a3/a3ed95caeb02ffe68cdd9fd
84406680ae93d633cb16422d00e8a7c22955b46d4/data
/registry/docker/registry/v2/blobs/sha256/f7
/registry/docker/registry/v2/blobs/sha256/f7/f72a00a23f01987b42cb26f
259582bb33502bdb0fcf5011e03c60577c4284845
/registry/docker/registry/v2/blobs/sha256/f7/f72a00a23f01987b42cb26f
259582bb33502bdb0fcf5011e03c60577c4284845/data
/registry/docker/registry/v2/repositories 3
/registry/docker/registry/v2/repositories/p1
/registry/docker/registry/v2/repositories/p1/pause 4
/registry/docker/registry/v2/repositories/p1/pause/_manifests
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons/sha256
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons/sha256/e9a2ac6418981897b399d3709f1b4a6d2723cd38a4909215ce2752a5c
068b1cf
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons/sha256/e9a2ac6418981897b399d3709f1b4a6d2723cd38a4909215ce2752a5c
068b1cf/signatures 5
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons/sha256/e9a2ac6418981897b399d3709f1b4a6d2723cd38a4909215ce2752a5c
068b1cf/signatures/sha256
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons/sha256/e9a2ac6418981897b399d3709f1b4a6d2723cd38a4909215ce2752a5c
068b1cf/signatures/sha256/ede17b139a271d6b1331ca3d83c648c24f92cece5f
89d95ac6c34ce751111810
/registry/docker/registry/v2/repositories/p1/pause/_manifests/revisi
ons/sha256/e9a2ac6418981897b399d3709f1b4a6d2723cd38a4909215ce2752a5c
068b1cf/signatures/sha256/ede17b139a271d6b1331ca3d83c648c24f92cece5f
89d95ac6c34ce751111810/link 6
/registry/docker/registry/v2/repositories/p1/pause/_uploads 7
/registry/docker/registry/v2/repositories/p1/pause/_layers 8
/registry/docker/registry/v2/repositories/p1/pause/_layers/sha256
/registry/docker/registry/v2/repositories/p1/pause/_layers/sha256/a3
ed95caeb02ffe68cdd9fd84406680ae93d633cb16422d00e8a7c22955b46d4
/registry/docker/registry/v2/repositories/p1/pause/_layers/sha256/a3
ed95caeb02ffe68cdd9fd84406680ae93d633cb16422d00e8a7c22955b46d4/link
9
/registry/docker/registry/v2/repositories/p1/pause/_layers/sha256/f7
2a00a23f01987b42cb26f259582bb33502bdb0fcf5011e03c60577c4284845
/registry/docker/registry/v2/repositories/p1/pause/_layers/sha256/f7
2a00a23f01987b42cb26f259582bb33502bdb0fcf5011e03c60577c4284845/link

```

- 1 This directory stores all layers and signatures as blobs.
- 2 This file contains the blob's contents.
- 3 This directory stores all the image repositories.

- 4 This directory is for a single image repository **p1/pause**.
- 5 This directory contains signatures for a particular image manifest revision.
- 6 This file contains a reference back to a blob (which contains the signature data).
- 7 This directory contains any layers that are currently being uploaded and staged for the given repository.
- 8 This directory contains links to all the layers this repository references.
- 9 This file contains a reference to a specific layer that has been linked into this repository via an image.

3.3.3. Accessing the Registry Directly

For advanced usage, you can access the registry directly to invoke **docker** commands. This allows you to push images to or pull them from the integrated registry directly using operations like **docker push** or **docker pull**. To do so, you must be logged in to the registry using the **docker login** command. The operations you can perform depend on your user permissions, as described in the following sections.

3.3.3.1. User Prerequisites

To access the registry directly, the user that you use must satisfy the following, depending on your intended usage:

- For any direct access, you must have a [regular user](#), if one does not already exist, for your preferred [identity provider](#). A regular user can generate an access token required for logging in to the registry. [System users](#), such as **system:admin**, cannot obtain access tokens and, therefore, cannot access the registry directly.

For example, if you are using **HTPASSWD** authentication, you can create one using the following command:

```
# htpasswd /etc/origin/openshift-htpasswd <user_name>
```

- The user must have the **system:registry** role. To add this role:

```
# oadm policy add-role-to-user system:registry <user_name>
```

- Have the **admin** role for the project associated with the Docker operation. For example, if accessing images in the global **openshift** project:

```
$ oadm policy add-role-to-user admin <user_name> -n openshift
```

- For writing or pushing images, for example when using the **docker push** command, the user must have the **system:image-builder** role. To add this role:

```
$ oadm policy add-role-to-user system:image-builder <user_name>
```

For more information on user permissions, see [Managing Role Bindings](#).

3.3.3.2. Logging in to the Registry



NOTE

Ensure your user satisfies the [prerequisites](#) for accessing the registry directly.

To log in to the registry directly:

1. Ensure you are logged in to OpenShift Container Platform as a **regular user**:

```
$ oc login
```

2. Get your access token:

```
$ oc whoami -t
```

3. Log in to the Docker registry:

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

3.3.3.3. Pushing and Pulling Images

After [logging in to the registry](#), you can perform **docker pull** and **docker push** operations against your registry.



IMPORTANT

You can pull arbitrary images, but if you have the **system:registry** role added, you can only push images to the registry in your project.

In the following examples, we use:

Component	Value
<registry_ip>	172.30.124.220
<port>	5000
<project>	openshift
<image>	busybox
<tag>	omitted (defaults to latest)

1. Pull an arbitrary image:

```
$ docker pull docker.io/busybox
```

2. Tag the new image with the form `<registry_ip>:<port>/<project>/<image>`. The project name **must** appear in this [pull specification](#) for OpenShift Container Platform to correctly place and later access the image in the registry.

```
$ docker tag docker.io/busybox 172.30.124.220:5000/openshift/busybox
```



NOTE

Your regular user must have the **system:image-builder** role for the specified project, which allows the user to write or push an image. Otherwise, the **docker push** in the next step will fail. To test, you can [create a new project](#) to push the **busybox** image.

3. Push the newly-tagged image to your registry:

```
$ docker push 172.30.124.220:5000/openshift/busybox
...
cf2616975b4a: Image successfully pushed
Digest:
sha256:3662dd821983bc4326bee12caec61367e7fb6f6a3ee547cbaff98f77403cab55
```

3.4. SECURING AND EXPOSING THE REGISTRY

3.4.1. Securing the Registry

Optionally, you can secure the registry so that it serves traffic via TLS:

1. [Deploy the registry](#).
2. Fetch the service IP and port of the registry:

```
$ oc get svc/docker-registry
NAME                                LABELS
SELECTOR                            IP(S)                                PORT(S)
docker-registry                    docker-registry=default              docker-
registry=default                    172.30.124.220                      5000/TCP
```

3. You can use an existing server certificate, or create a key and server certificate valid for specified IPs and host names, signed by a specified CA. To create a server certificate for the registry service IP and the **docker-registry.default.svc.cluster.local** host name:

```
$ oadm ca create-server-cert \
  --signer-cert=/etc/origin/master/ca.crt \
  --signer-key=/etc/origin/master/ca.key \
  --signer-serial=/etc/origin/master/ca.serial.txt \
  --hostnames='docker-registry.default.svc.cluster.local,docker-
registry.default.svc,172.30.124.220' \
  --cert=/etc/secrets/registry.crt \
  --key=/etc/secrets/registry.key
```


If the router will be [exposed externally](#), add the public route host name in the **--hostnames** flag:

```
--hostnames='mydocker-registry.example.com,docker-registry.default.svc.cluster.local,172.30.124.220 \
```

See [Redeploying Registry and Router Certificates](#) for additional details on updating the default certificate so that the route is externally accessible.



NOTE

The **oadm ca create-server-cert** command generates a certificate that is valid for two years. This can be altered with the **--expire-days** option, but for security reasons, it is recommended to not make it greater than this value.

4. Create the secret for the registry certificates:

```
$ oc secrets new registry-secret \
  /etc/secrets/registry.crt \
  /etc/secrets/registry.key
```

5. Add the secret to the registry pod's service accounts (including the **default** service account):

```
$ oc secrets link registry registry-secret
$ oc secrets link default registry-secret
```



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.

6. Add the secret volume to the registry deployment configuration:

```
$ oc volume dc/docker-registry --add --type=secret \
  --secret-name=registry-secret -m /etc/secrets
```

7. Enable TLS by adding the following environment variables to the registry deployment configuration:

```
$ oc set env dc/docker-registry \
  REGISTRY_HTTP_TLS_CERTIFICATE=/etc/secrets/registry.crt \
  REGISTRY_HTTP_TLS_KEY=/etc/secrets/registry.key
```

See more details on [overriding registry options](#).

8. Update the scheme used for the registry's liveness probe from HTTP to HTTPS:

```
$ oc patch dc/docker-registry -p '{"spec": {"template": {"spec": {"containers": [{
```

```
"name": "registry",
"livenessProbe": {"httpGet": {"scheme": "HTTPS"}}
}}}}}'
```

9. If your registry was initially deployed on OpenShift Container Platform 3.2 or later, update the scheme used for the registry's readiness probe from HTTP to HTTPS:

```
$ oc patch dc/docker-registry -p '{"spec": {"template": {"spec":
{"containers": [{
  "name": "registry",
  "readinessProbe": {"httpGet": {"scheme": "HTTPS"}}
}]}}}'
```

10. Validate the registry is running in TLS mode. Wait until the latest **docker-registry** deployment completes and verify the Docker logs for the registry container. You should find an entry for **listening on :5000, tls**.

```
$ oc logs dc/docker-registry | grep tls
time="2015-05-27T05:05:53Z" level=info msg="listening on :5000, tls"
instance.id=deeba528-c478-41f5-b751-dc48e4935fc2
```

11. Copy the CA certificate to the Docker certificates directory. This must be done on all nodes in the cluster:

```
$ dcertsdir=/etc/docker/certs.d
$ destdir_addr=$dcertsdir/172.30.124.220:5000
$ destdir_name=$dcertsdir/docker-
registry.default.svc.cluster.local:5000

$ sudo mkdir -p $destdir_addr $destdir_name
$ sudo cp ca.crt $destdir_addr 1
$ sudo cp ca.crt $destdir_name
```

1 The **ca.crt** file is a copy of **/etc/origin/master/ca.crt** on the master.

12. When using authentication, some versions of **docker** also require you to configure your cluster to trust the certificate at the OS level.

- a. Copy the certificate:

```
$ cp /etc/origin/master/ca.crt /etc/pki/ca-
trust/source/anchors/myregistrydomain.com.crt
```

- b. Run:

```
$ update-ca-trust enable
```

13. Remove the **--insecure-registry** option only for this particular registry in the **/etc/sysconfig/docker** file. Then, reload the daemon and restart the **docker** service to reflect this configuration change:

```
$ sudo systemctl daemon-reload
$ sudo systemctl restart docker
```

14. Validate the **docker** client connection. Running **docker push** to the registry or **docker pull** from the registry should succeed. Make sure you have [logged into the registry](#).

```
$ docker tag|push <registry/image> <internal_registry/project/image>
```

For example:

```
$ docker pull busybox
$ docker tag docker.io/busybox 172.30.124.220:5000/openshift/busybox
$ docker push 172.30.124.220:5000/openshift/busybox
...
cf2616975b4a: Image successfully pushed
Digest:
sha256:3662dd821983bc4326bee12caec61367e7fb6f6a3ee547cbaff98f77403ca
b55
```

3.4.2. Exposing the Registry

To expose your internal registry externally, it is recommended that you run a [secure registry](#). To expose the registry you must first have [deployed a router](#).

1. [Deploy the registry](#).
2. [Secure the registry](#).
3. [Deploy a router](#).
4. Create a [passthrough](#) route via the **oc create route passthrough** command, specifying the registry as the route's service. By default, the name of the created route is the same as the service name.

For example:

```
$ oc get svc
NAME                                CLUSTER_IP          EXTERNAL_IP  PORT(S)
SELECTOR                            AGE
docker-registry                    172.30.69.167       <none>       5000/TCP
docker-registry=default            4h
kubernetes                         172.30.0.1          <none>
443/TCP, 53/UDP, 53/TCP            <none>              4h
router                             172.30.172.132      <none>       80/TCP
router=router                      4h

$ oc create route passthrough \
  --service=docker-registry \ 1
  --hostname=<host>
route "docker-registry" created 2
```

- 1 Specify the registry as the route's service.
- 2 The route name is identical to the service name.


```
$ oc get route/docker-registry -o yaml
apiVersion: v1
```

```

kind: Route
metadata:
  name: docker-registry
spec:
  host: <host> ❶
  to:
    kind: Service
    name: docker-registry ❷
  tls:
    termination: passthrough ❸

```

- ❶ The host for your route. You must be able to resolve this name externally via DNS to the router's IP address.
- ❷ The service name for your registry.
- ❸ Specify this route as a passthrough route.

**NOTE**

Passthrough is currently the only type of route supported for exposing the secure registry.

5. Next, you must trust the certificates being used for the registry on your host system. The certificates referenced were created when you secured your registry.

```

$ sudo mkdir -p /etc/docker/certs.d/<host>
$ sudo cp <ca certificate file> /etc/docker/certs.d/<host>
$ sudo systemctl restart docker

```

6. [Log in to the registry](#) using the information from securing the registry. However, this time point to the host name used in the route rather than your service IP. You should now be able to tag and push images using the route host.

```

$ oc get imagestreams -n test
NAME          DOCKER REPO    TAGS      UPDATED

$ docker pull busybox
$ docker tag busybox <host>/test/busybox
$ docker push <host>/test/busybox
The push refers to a repository [<host>/test/busybox] (len: 1)
8c2e06607696: Image already exists
6ce2e90b0bc7: Image successfully pushed
cf2616975b4a: Image successfully pushed
Digest:
sha256:6c7e676d76921031532d7d9c0394d0da7c2906f4cb4c049904c4031147d8c
a31

$ docker pull <host>/test/busybox
latest: Pulling from <host>/test/busybox
cf2616975b4a: Already exists
6ce2e90b0bc7: Already exists
8c2e06607696: Already exists
Digest:

```

```
sha256:6c7e676d76921031532d7d9c0394d0da7c2906f4cb4c049904c4031147d8c
a31
```

```
Status: Image is up to date for <host>/test/busybox:latest
```

```
$ oc get imagestreams -n test
```

NAME	DOCKER REPO	TAGS	UPDATED
busybox	172.30.11.215:5000/test/busybox	latest	2 seconds ago



NOTE

Your image streams will have the IP address and port of the registry service, not the route name and port. See **oc get imagestreams** for details.



NOTE

In the **<host>/test/busybox** example above, **test** refers to the project name.

3.4.2.1. Exposing a Secure Registry

Instead of logging in to the registry from within the OpenShift Container Platform cluster, you can gain external access to it by first [securing the registry](#) and then [exposing the registry](#). This allows you to log in to the registry from outside the cluster using the route address, and to tag and push images using the route host.

When [logging in](#) to the secured and exposed registry, make sure you specify the registry in the login command. For example:

```
docker login -e user@company.com -u f83j5h6 -p Ju1PeM47R0B92Lk3AZp-
bWJSck2F7aGCiZ66aFGZrs2 registry.example.com:80
```

3.4.2.2. Exposing a Non-Secure Registry

Instead of securing the registry in order to expose the registry, you can simply expose a non-secure registry for non-production OpenShift Container Platform environments. This allows you to have an external route to the registry without using SSL certificates.



WARNING

Only non-production environments should expose a non-secure registry to external access.

To expose a non-secure registry:

1. Expose the registry:

```
# oc expose service docker-registry --hostname=<hostname> -n default
```

This creates the following JSON file:

```

apiVersion: v1
kind: Route
metadata:
  creationTimestamp: null
  labels:
    docker-registry: default
  name: docker-registry
spec:
  host: registry.example.com
  port:
    targetPort: "5000"
  to:
    kind: Service
    name: docker-registry
status: {}

```

2. Verify that the route has been created successfully:

```

# oc get route
NAME                                HOST/PORT                                PATH      SERVICE
LABELS                                INSECURE POLICY    TLS TERMINATION
docker-registry    registry.example.com                                docker-registry
docker-registry=default

```

3. Check the health of the registry:

```
$ curl -v http://registry.example.com/healthz
```


Expect an HTTP 200/OK message.

After exposing the registry, update your `/etc/sysconfig/docker` file by adding the port number to the **OPTIONS** entry. For example:

```

OPTIONS='--selinux-enabled --insecure-registry=172.30.0.0/16 --
insecure-registry registry.example.com:80'

```



IMPORTANT

The above options should be added on the client from which you are trying to log in.

Also, ensure that Docker is running on the client.

When [logging in](#) to the non-secured and exposed registry, make sure you specify the registry in the login command. For example:

```

docker login -e user@company.com -u f83j5h6 -p Ju1PeM47R0B92Lk3AZp-
bWJSck2F7aGCiZ66aFGZrs2 registry.example.com:80

```

3.5. EXTENDED REGISTRY CONFIGURATION

3.5.1. Maintaining the Registry IP Address

OpenShift Container Platform refers to the integrated registry by its service IP address, so if you decide to delete and recreate the **docker-registry** service, you can ensure a completely transparent transition by arranging to re-use the old IP address in the new service. If a new IP address cannot be avoided, you can minimize cluster disruption by rebooting only the masters.

Re-using the Address

To re-use the IP address, you must save the IP address of the old **docker-registry** service prior to deleting it, and arrange to replace the newly assigned IP address with the saved one in the new **docker-registry** service.

1. Make a note of the **clusterIP** for the service:

```
$ oc get svc/docker-registry -o yaml | grep clusterIP:
```

2. Delete the service:

```
$ oc delete svc/docker-registry dc/docker-registry
```

3. Create the registry definition in **registry.yaml**, replacing **<options>** with, for example, those used in step 3 of the instructions in the [Non-Production Use](#) section:

```
$ oadm registry <options> -o yaml > registry.yaml
```

4. Edit **registry.yaml**, find the **Service** there, and change its **clusterIP** to the address noted in step 1.
5. Create the registry using the modified **registry.yaml**:

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

Rebooting the Masters

If you are unable to re-use the IP address, any operation that uses a [pull specification](#) that includes the old IP address will fail. To minimize cluster disruption, you must reboot the masters:

```
# systemctl restart atomic-openshift-master
```

This ensures that the old registry URL, which includes the old IP address, is cleared from the cache.



NOTE

We recommend against rebooting the entire cluster because that incurs unnecessary downtime for pods and does not actually clear the cache.

3.5.2. Whitelisting Docker Registries

You can specify a whitelist of docker registries, allowing you to curate a set of images and templates that are available for download by OpenShift Container Platform users. This curated set can be placed in one or more docker registries, and then added to the whitelist. When using a whitelist, only the specified

registries are accessible within OpenShift Container Platform, and all other registries are denied access by default.

To configure a whitelist:

1. Edit the `/etc/sysconfig/docker` file to block all registries:

```
BLOCK_REGISTRY='--block-registry=all'
```

You may need to uncomment the **BLOCK_REGISTRY** line.

2. In the same file, add registries to which you want to allow access:

```
ADD_REGISTRY='--add-registry=<registry1> --add-registry=<registry2>'
```

Example 3.1. Allowing Access to Registries

```
ADD_REGISTRY='--add-registry=registry.access.redhat.com'
```

This example would restrict access to images available on the [Red Hat Customer Portal](#).

Once the whitelist is configured, if a user tries to pull from a docker registry that is not on the whitelist, they will receive an error message stating that this registry is not allowed.

3.5.3. Overriding the Registry Configuration

You can override the integrated registry's default configuration, found by default at `/config.yml` in a running registry's container, with your own [custom configuration](#).



NOTE

Upstream configuration options in this file may also be overridden using environment variables. The [middleware section](#) is an exception as there are just a few options that can be overridden using environment variables. [Learn how to override specific configuration options.](#)

To enable management of the registry configuration file directly and deploy an updated configuration, mount the configuration file as a [secret volume](#):

1. [Deploy the registry.](#)
2. Edit the registry configuration file locally as needed. The initial YAML file deployed on the registry is provided below. [Review supported options.](#)

Example 3.2. Registry Configuration File

```
version: 0.1
log:
  level: debug
http:
  addr: :5000
storage:
  cache:
```



```

    blobdescriptor: inmemory
    filesystem:
      rootdirectory: /registry
    delete:
      enabled: true
  auth:
    openshift:
      realm: openshift
  middleware:
    registry:
      - name: openshift
    repository:
      - name: openshift
      options:
        acceptschema2: false
        pullthrough: true
        enforcequota: false
        projectcachettl: 1m
        blobrepositorycachettl: 10m
  storage:
    - name: openshift

```

3. Create a new secret called **registry-config** from your custom registry configuration file you edited locally:

```

$ oc secrets new registry-config config.yml=
</path/to/custom/registry/config.yml>

```

4. Add the **registry-config** secret as a volume to the registry's deployment configuration to mount the custom configuration file at **/etc/docker/registry/**:

```

$ oc volume dc/docker-registry --add --type=secret \
  --secret-name=registry-config -m /etc/docker/registry/

```

5. Update the registry to reference the configuration path from the previous step by adding the following environment variable to the registry's deployment configuration:

```

$ oc set env dc/docker-registry \
  REGISTRY_CONFIGURATION_PATH=/etc/docker/registry/config.yml

```

This may be performed as an iterative process to achieve the desired configuration. For example, during troubleshooting, the configuration may be temporarily updated to put it in **debug** mode.

To update an existing configuration:



WARNING

This procedure will overwrite the currently deployed registry configuration.

1. Edit the local registry configuration file, **config.yml**.

2. Delete the **registry-config** secret:

```
$ oc delete secret registry-config
```

3. Recreate the secret to reference the updated configuration file:

```
$ oc secrets new registry-config config.yml=  
</path/to/custom/registry/config.yml>
```

4. Redeploy the registry to read the updated configuration:

```
$ oc deploy docker-registry --latest
```

TIP

Maintain configuration files in a source control repository.

3.5.4. Registry Configuration Reference

There are many configuration options available in the upstream [Docker Distribution](#) library. Not all [configuration options](#) are supported or enabled. Use this section as a reference when [overriding the registry configuration](#).



NOTE

Upstream configuration options in this file may also be overridden using environment variables. However, the [middleware section](#) may **not** be overridden using environment variables. [Learn how to override specific configuration options](#).

3.5.4.1. Log

[Upstream options](#) are supported.

Example:

```
log:  
  level: debug  
  formatter: text  
  fields:  
    service: registry  
    environment: staging
```

3.5.4.2. Hooks

Mail hooks are not supported.

3.5.4.3. Storage

This section lists the supported [registry storage drivers](#).

The following list includes storage drivers that need to be configured in the registry's configuration file:

- [Filesystem](#). Filesystem is the default and does not need to be configured.
- [S3](#). Learn more about [CloudFront configuration](#).
- [OpenStack Swift](#)
- [Google Cloud Storage \(GCS\)](#)
- [Microsoft Azure](#)
- [Aliyun OSS](#)

[General registry storage configuration options](#) are supported.

The following storage options need to be configured through the [filesystem driver](#):

- [Backing Docker Registry with GlusterFS Storage](#)
- [Ceph Rados Block Device](#)



NOTE

For more information on supported persistent storage drivers, see [Configuring Persistent Storage](#) and [Persistent Storage Examples](#).

Example 3.3. General Storage Configuration Options

```
storage:
  delete:
    enabled: true 1
  redirect:
    disable: false
  cache:
    blobdescriptor: inmemory
  maintenance:
    uploadpurging:
      enabled: true
      age: 168h
      interval: 24h
      dryrun: false
  readonly:
    enabled: false
```

1 This entry is **mandatory** for image pruning to work properly.

3.5.4.4. Auth

Auth options should not be altered. The **openshift** extension is the only supported option.

```
auth:
  openshift:
```

```
realm: openshift
```

3.5.4.5. Middleware

The **repository** middleware extension allows to configure OpenShift Container Platform middleware responsible for interaction with OpenShift Container Platform and image proxying.

```
middleware:
  registry:
    - name: openshift ❶
  repository:
    - name: openshift ❷
      options:
        acceptschema2: false ❸
        pullthrough: true ❹
        enforcequota: false ❺
        projectcachettl: 1m ❻
        blobrepositorycachettl: 10m ❼
  storage:
    - name: openshift ❽
```

❶ ❷ ❽ These entries are mandatory. Their presence ensures required components get loaded. These values shouldn't be changed.

❸ Allow to store [manifest v2 schema 2](#) during a push to the registry. See [below](#) for more details.

❹ Let the registry act as a proxy for remote blobs. See [below](#) for more details.

❺ Prevent blob uploads exceeding size limit defined in targeted project.

❻ An expiration timeout for limits cached in the registry. The lower the value, the less time will it take for the limit changes to propagate to the registry. However, the registry will query limits from the server more frequently and, as a consequence, pushes will be slower.

❼ An expiration timeout for remembered associations between blob and repository. The higher the value, the higher probability of fast lookup and more efficient registry operation. On the other hand, memory usage will raise as well as a risk of serving image layer to user, who is no longer authorized to access it.

3.5.4.5.1. CloudFront Middleware

The [CloudFront middleware extension](#) can be added to support AWS, CloudFront CDN storage provider. CloudFront middleware speeds up distribution of image content internationally. The blobs are distributed to several edge locations around the world. The client is always directed to the edge with the lowest latency.



NOTE

The [CloudFront middleware extension](#) can be only used with [S3](#) storage. It is utilized only during blob serving. Therefore, only blob downloads can be speeded up, not uploads.

The following is an example of minimal configuration of S3 storage driver with a CloudFront middleware:

```
version: 0.1
log:
  level: debug
http:
  addr: :5000
storage:
  cache:
    blobdescriptor: inmemory
  delete:
    enabled: true
s3: ❶
  accesskey: BJKMSZBRESWJQXRWMAEQ
  secretkey: 5ah5I91SNXbeoUXXDasFtadRq0dy62Jzln0W1goS
  region: us-east-1
  bucket: docker.myregistry.com
auth:
  openshift:
    realm: openshift
middleware:
  registry:
    - name: openshift
  repository:
    - name: openshift
  storage:
    - name: cloudfront ❷
      options:
        baseurl: https://jrpbyn0k5k88bi.cloudfront.net/ ❸
        privatekey: /etc/docker/cloudfront-ABCDEFGHJKLMNOPQRST.pem ❹
        keypairid: ABCDEFGHJKLMNOPQRST ❺
    - name: openshift
```

- ❶ The S3 storage must be configured the same way regardless of CloudFront middleware.
- ❷ The CloudFront storage middleware needs to be listed before OpenShift middleware.
- ❸ The CloudFront base URL. In the AWS management console, this is listed as **Domain Name** of CloudFront distribution.
- ❹ The location of your AWS private key on the filesystem. This must be not confused with Amazon EC2 key pair. Please refer to [AWS documentation](#) on creating CloudFront key pairs for your trusted signers. The file needs to be mounted as a secret [secret](#) into the registry pod.
- ❺ The ID of your Cloudfront key pair.

3.5.4.5.2. Overriding Middleware Configuration Options

The **middleware** section cannot be overridden using environment variables. There are a few exceptions, however. For example:

```
middleware:
  repository:
    - name: openshift
```

```
options:
  acceptschema2: false 1
  enforcequota: false 2
  projectcachettl: 1m 3
  blobrepositorycachettl: 10m 4
```

- 1 A configuration option that can be overridden by the boolean environment variable **REGISTRY_MIDDLEWARE_REPOSITORY_OPENSHIFT_ACCEPTSCHEMA2**, which allows for the ability to accept manifest v2 schema 2 on manifest put requests.
- 2 A configuration option that can be overridden by the boolean environment variable **REGISTRY_MIDDLEWARE_REPOSITORY_OPENSHIFT_ENFORCEQUOTA**, which allows the ability to turn quota enforcement on or off. By default, quota enforcement is off. It overrides OpenShift Container Platform middleware configuration option. Recognized values are **true** and **false**.
- 3 A configuration option that can be overridden by the environment variable **REGISTRY_MIDDLEWARE_REPOSITORY_OPENSHIFT_PROJECTCACHETTTL**, specifying an eviction timeout for project quota objects. It takes a valid time duration string (for example, **2m**). If empty, you get the default timeout. If zero (**0m**), caching is disabled.
- 4 A configuration option that can be overridden by the environment variable **REGISTRY_MIDDLEWARE_REPOSITORY_OPENSHIFT_BLOBREPOSITORYCACHETTTL**, specifying an eviction timeout for associations between blob and containing repository. The format of the value is the same as in **projectcachettl** case.

3.5.4.5.3. Image Pullthrough

If enabled, the registry will attempt to fetch requested blob from a remote registry unless the blob exists locally. The remote candidates are calculated from **DockerImage** entries stored in status of the [image stream](#), a client pulls from. All the unique remote registry references in such entries will be tried in turn until the blob is found. The blob, served this way, will not be stored in the registry.

You must ensure that your registry has appropriate certificates to trust any external registries you do a pullthrough against. The certificates need to be placed in the **/etc/pki/tls/certs** directory on the pod. You can mount the certificates using a [configuration map](#) or [secret](#). Note that the entire **/etc/pki/tls/certs** directory must be replaced. You must include the new certificates and replace the system certificates in your secret or configuration map that you mount.

This feature is on by default. However, it can be disabled using a [configuration option](#).

3.5.4.5.4. Manifest V2 Schema 2 Support

Each image has a manifest describing its blobs, instructions for running it and additional metadata. The manifest is versioned which have different structure and fields as it evolves over time. The same image can be represented by multiple manifest versions. Each version will have different digest though.

The registry currently supports [manifest v2 schema 1](#) (**schema 1**) and [manifest v2 schema 2](#) (**schema 2**). The former is being obsoleted but will be supported for an extended amount of time.

You should be wary of compatibility issues with various Docker clients:

- Docker clients of version 1.9 or older support only **schema 1**. Any manifest this client pulls or pushes will be of this legacy schema.

- Docker clients of version 1.10 support both **schema 1** and **schema 2**. And by default, it will push the latter to the registry if it supports newer schema.

The registry, storing an image with **schema 1** will always return it unchanged to the client. **Schema 2** will be transferred unchanged only to newer Docker client. For the older one, it will be converted on-the-fly to **schema 1**.

This has significant consequences. For example an image pushed to the registry by a newer Docker client cannot be pulled by the older Docker by its digest. That's because the stored image's manifest is of **schema 2** and its digest can be used to pull only this version of manifest.

For this reason, the registry is configured by default not to store **schema 2**. This ensures that any docker client will be able to pull from the registry any image pushed there regardless of client's version.

Once you're confident that all the registry clients support **schema 2**, you'll be safe to enable its support in the registry. See the [middleware configuration reference](#) above for particular option.

3.5.4.6. Reporting

Reporting is unsupported.

3.5.4.7. HTTP

[Upstream options](#) are supported. [Learn how to alter these settings via environment variables](#). Only the **tls** section should be altered. For example:

```
http:
  addr: :5000
  tls:
    certificate: /etc/secrets/registry.crt
    key: /etc/secrets/registry.key
```

3.5.4.8. Notifications

[Upstream options](#) are supported. The [REST API Reference](#) provides more comprehensive integration options.

Example:

```
notifications:
  endpoints:
    - name: registry
      disabled: false
      url: https://url:port/path
      headers:
        Accept:
          - text/plain
      timeout: 500
      threshold: 5
      backoff: 1000
```

3.5.4.9. Redis

Redis is not supported.

3.5.4.10. Health

[Upstream options](#) are supported. The registry deployment configuration provides an integrated health check at `/healthz`.

3.5.4.11. Proxy

Proxy configuration should not be enabled. This functionality is provided by the [OpenShift Container Platform repository middleware extension](#), `pullthrough: true`.

3.6. KNOWN ISSUES

3.6.1. Overview

The following are the known issues when deploying or using the integrated registry.

3.6.2. Image Push Errors with Scaled Registry Using Shared NFS Volume

When using a scaled registry with a shared NFS volume, you may see one of the following errors during the push of an image:

- **digest invalid: provided digest did not match uploaded content**
- **blob upload unknown**
- **blob upload invalid**

These errors are returned by an internal registry service when Docker attempts to push the image. Its cause originates in the synchronization of file attributes across nodes. Factors such as NFS client side caching, network latency, and layer size can all contribute to potential errors that might occur when pushing an image using the default round-robin load balancing configuration.

You can perform the following steps to minimize the probability of such a failure:


1. Ensure that the **sessionAffinity** of your **docker-registry** service is set to **ClientIP**:

```
$ oc get svc/docker-registry --template='{{.spec.sessionAffinity}}'
```

This should return **ClientIP**, which is the default in recent OpenShift Container Platform versions. If not, change it:

```
$ oc get -o yaml svc/docker-registry | \
    sed 's/\(sessionAffinity:\s*\).*\/\1ClientIP/' | \
    oc replace -f -
```

2. Ensure that the NFS export line of your registry volume on your NFS server has the **no_wdelay** options listed. The **no_wdelay** option prevents the server from delaying writes, which greatly improves read-after-write consistency, a requirement of the registry.



NOTE

The guidelines for NFS are recommended to help you get started. You may switch off from NFS when moving to production.

3.6.3. Pull of Internally Managed Image Fails with "not found" Error

This error occurs when the pulled image is pushed to an image stream different from the one it is being pulled from. This is caused by re-tagging a built image into an arbitrary image stream:

```
$ oc tag srcimagestream:latest anyproject/pullimagestream:latest
```

And subsequently pulling from it, using an image reference such as:

```
internal.registry.url:5000/anyproject/pullimagestream:latest
```

During a manual Docker pull, this will produce a similar error:

```
Error: image anyproject/pullimagestream:latest not found
```

To prevent this, avoid the tagging of internally managed images completely, or re-push the built image to the desired namespace [manually](#).

3.6.4. Image Push Fails with "500 Internal Server Error" on S3 Storage

There are problems reported happening when the registry runs on S3 storage back-end. Pushing to a Docker registry occasionally fails with the following error:

```
Received unexpected HTTP status: 500 Internal Server Error
```

To debug this, you need to [view the registry logs](#). In there, look for similar error messages occurring at the time of the failed push:

```
time="2016-03-30T15:01:21.22287816-04:00" level=error msg="unknown error
completing upload: driver.Error{DriverName:\"s3\", Enclosed:(*url.Error)
(0xc20901cea0)}" http.request.method=PUT
...
time="2016-03-30T15:01:21.493067808-04:00" level=error msg="response
completed with error" err.code=UNKNOWN err.detail="s3: Put
https://s3.amazonaws.com/oso-tsi-
docker/registry/docker/registry/v2/blobs/sha256/ab/abe5af443833d60cf672e2a
c57589410dddec060ed725d3e676f1865af63d2e2/data: EOF" err.message="unknown
error" http.request.method=PUT
...
time="2016-04-02T07:01:46.056520049-04:00" level=error msg="error putting
into main store: s3: The request signature we calculated does not match
the signature you provided. Check your key and signing method."
http.request.method=PUT
atest
```

If you see such errors, contact your Amazon S3 support. There may be a problem in your region or with your particular bucket.

3.6.5. Image Pruning Fails

If you encounter the following error when pruning images:

```
BLOB
sha256:49638d540b2b62f3b01c388e9d8134c55493b1fa659ed84e97cb59b87a6b8e6c
error deleting blob
```

And your [registry log](#) contains the following information:

```
error deleting blob
\"sha256:49638d540b2b62f3b01c388e9d8134c55493b1fa659ed84e97cb59b87a6b8e6c\
\": operation unsupported
```

It means that your [custom configuration file](#) lacks mandatory entries in the [storage section](#), namely **storage.delete.enabled** set to **true**. Add them, re-deploy the registry, and repeat your image pruning operation.

CHAPTER 4. SETTING UP A ROUTER

4.1. ROUTER OVERVIEW

4.1.1. About Routers

The OpenShift Container Platform [router](#) is the ingress point for all external traffic destined for [services](#) in your OpenShift Container Platform installation. OpenShift Container Platform provides and supports the following two router plug-ins:

- The [HAProxy template router](#) is the default plug-in. It uses the **openshift3/ose-haproxy-router** image to run an HAProxy instance alongside the template router plug-in inside a container on OpenShift Container Platform. It currently supports HTTP(S) traffic and TLS-enabled traffic via SNI. The router's container listens on the host network interface, unlike most containers that listen only on private IPs. The router proxies external requests for route names to the IPs of actual pods identified by the service associated with the route.
- The [F5 router](#) integrates with an existing **F5 BIG-IP®** system in your environment to synchronize routes. **F5 BIG-IP®** version 11.4 or newer is required in order to have the F5 iControl REST API.



NOTE

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

4.1.2. Router Service Account

Before deploying an OpenShift Container Platform cluster, you must have a service account for the router. Starting in OpenShift Container Platform 3.1, a router [service account](#) is automatically created during a quick or advanced installation (previously, this required manual creation). This service account has permissions to a [security context constraint](#) (SCC) that allows it to specify host ports.


Use of labels (e.g., to define [router shards](#)) requires **cluster-reader** permission.

```
$ oadm policy add-cluster-role-to-user \
    cluster-reader \
    system:serviceaccount:default:router
```

4.2. USING THE DEFAULT HAPROXY ROUTER


4.2.1. Overview

The **oadm router** command is provided with the administrator CLI to simplify the tasks of setting up routers in a new installation. If you followed the [quick installation](#), then a default router was automatically created for you. The **oadm router** command creates the service and deployment configuration objects. Just about every form of communication between OpenShift Container Platform components is secured by TLS and uses various certificates and authentication methods. Use the **--service-account** option to specify the service account the router will use to contact the master.



IMPORTANT

Routers directly attach to port 80 and 443 on all interfaces on a host. Restrict routers to hosts where port 80/443 is available and not being consumed by another service, and set this using node selectors and the [scheduler configuration](#). As an example, you can achieve this by dedicating infrastructure nodes to run services such as routers.




IMPORTANT

It is recommended to use separate distinct **openshift-router** service account with your router. This can be provided using the **--service-account** flag to the **oadm router** command.

```
$ oadm router --dry-run --service-account=router 1
```

1 **--service-account** is the name of a [service account](#) for the **openshift-router**.



IMPORTANT

Router pods created using **oadm router** have default resource requests that a node must satisfy for the router pod to be deployed. In an effort to increase the reliability of infrastructure components, the default resource requests are used to increase the QoS tier of the router pods above pods without resource requests. The default values represent the observed minimum resources required for a basic router to be deployed and can be edited in the routers deployment configuration and you may want to increase them based on the load of the router.

4.2.2. Creating a Router

The [quick installation](#) process automatically creates a default router. If the router does not exist, run the following to create a router:

```
$ oadm router <router_name> --replicas=<number> --service-account=router
```

You can also [use router shards](#) to ensure that the router is filtered to specific namespaces or routes, or [set any environment variables](#) after router creation.

4.2.3. Other Basic Router Commands

Checking the Default Router

The default router service account, named **router**, is automatically created during quick and advanced installations. To verify that this account already exists:

```
$ oadm router --dry-run --service-account=router
```

Viewing the Default Router

To see what the default router would look like if created:

```
$ oadm router -o yaml --service-account=router
```

Deploying the Router to a Labeled Node

To deploy the router to any node(s) that match a specified [node label](#):

```
$ oadm router <router_name> --replicas=<number> --selector=<label> \
  --service-account=router
```

For example, if you want to create a router named **router** and have it placed on a node labeled with **region=infra**:

```
$ oadm router router --replicas=1 --selector='region=infra' \
  --service-account=router
```

During [advanced installation](#), the **openshift_hosted_router_selector** and **openshift_registry_selector** Ansible settings are set to **region=infra** by default. The default router and registry will only be automatically deployed if a node exists that matches the **region=infra** label.

Multiple instances are created on different hosts according to the [scheduler policy](#).

Using a Different Router Image

To use a different router image and view the router configuration that would be used:

```
$ oadm router <router_name> -o <format> --images=<image> \
  --service-account=router
```

For example:

```
$ oadm router region-west -o yaml --images=myrepo/somerouter:mytag \
  --service-account=router
```

4.2.4. Filtering Routes to Specific Routers

Using the **ROUTE_LABELS** environment variable, you can filter routes so that they are used only by specific routers.

For example, if you have multiple routers, and 100 routes, you can attach labels to the routes so that a portion of them are handled by one router, whereas the rest are handled by another.

1. After [creating a router](#), use the **ROUTE_LABELS** environment variable to tag the router:

```
$ oc env dc/<router=name> ROUTE_LABELS="key=value"
```

2. Add the label to the desired routes:

```
oc label route <route=name> key=value
```

3. To verify that the label has been attached to the route, check the route configuration:

```
$ oc describe dc/<route_name>
```

4.2.5. Highly-Available Routers

You can [set up a highly-available router](#) on your OpenShift Container Platform cluster using IP failover.

4.2.6. Customizing the Router Service Ports

You can customize the service ports that a template router binds to by setting the environment variables **ROUTER_SERVICE_HTTP_PORT** and **ROUTER_SERVICE_HTTPS_PORT**. This can be done by creating a template router, then editing its deployment configuration.

The following example creates a router deployment with **0** replicas and customizes the router service HTTP and HTTPS ports, then scales it appropriately (to **1** replica).

```
$ oadm router --replicas=0 --ports='10080:10080,10443:10443' 1
$ oc set env dc/router ROUTER_SERVICE_HTTP_PORT=10080 \
    ROUTER_SERVICE_HTTPS_PORT=10443
$ oc scale dc/router --replicas=1
```

- 1** Ensures exposed ports are appropriately set for routers that use the container networking mode `--host-network=false`.



IMPORTANT

If you do customize the template router service ports, you will also need to ensure that the nodes where the router pods run have those custom ports opened in the firewall (either via Ansible or **iptables**, or any other custom method that you use via **firewall-cmd**).

The following is an example using **iptables** to open the custom router service ports.

```
$ iptables -A INPUT -p tcp --dport 10080 -j ACCEPT
$ iptables -A INPUT -p tcp --dport 10443 -j ACCEPT
```

4.2.7. Working With Multiple Routers

An administrator can create multiple routers with the same definition to serve the same set of routes. By having different groups of routers with different namespace or route selectors, they can vary the routes that the router serves.

Multiple routers can be grouped to distribute routing load in the cluster and separate tenants to different routers or [shards](#). Each router or shard in the group handles routes based on the selectors in the router. An administrator can create shards over the whole cluster using **ROUTE_LABELS**. A user can create shards over a namespace (project) by using **NAMESPACE_LABELS**.

4.2.8. Adding a Node Selector to a Deployment Configuration

Making specific routers deploy on specific nodes requires two steps:

1. Add a [label](#) to the desired node:

```
$ oc label node 10.254.254.28 "router=first"
```

2. Add a node selector to the router deployment configuration:

```
$ oc edit dc <deploymentConfigName>
```

Add the `template.spec.nodeSelector` field with a key and value corresponding to the label:

```
...
template:
  metadata:
    creationTimestamp: null
    labels:
      router: router1
  spec:
    nodeSelector:
      router: "first"
...
```

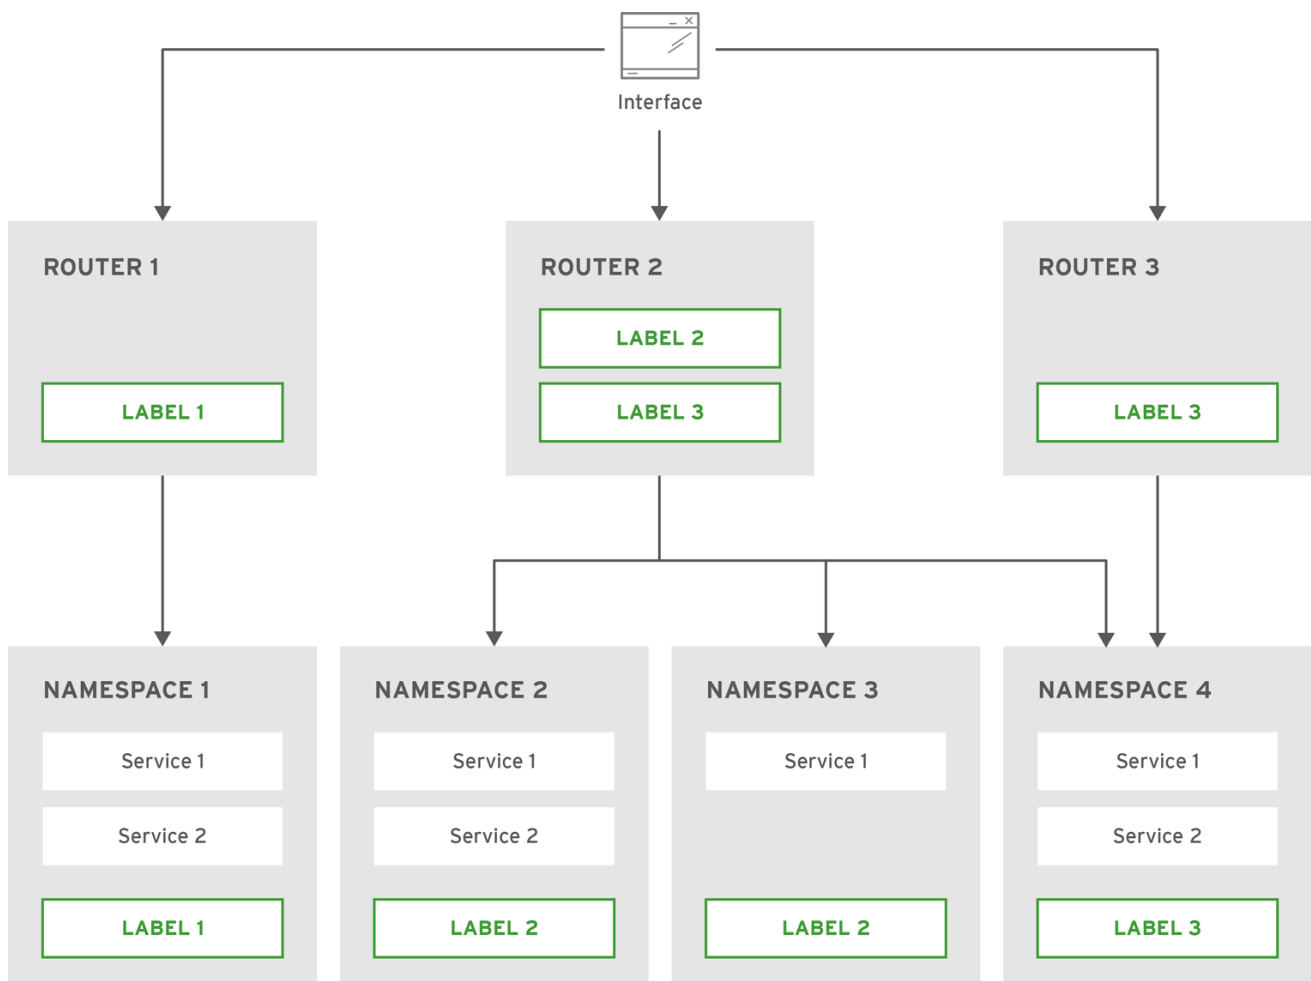
- 1** The key and value are **router** and **first**, respectively, corresponding to the **router=first** label.

4.2.9. Using Router Shards

The access controls are based on the service account that the router is run with.

Using **NAMESPACE_LABELS** and/or **ROUTE_LABELS**, a router can filter out the namespaces and/or routes that it should service. This enables you to partition routes amongst multiple router deployments effectively distributing the set of routes.

Figure 4.1. Router Sharding Based on Namespace Labels



OPENSIFT_415490_0217

Example: A router deployment **finops-router** is run with route selector **NAMESPACE_LABELS="name in (finance, ops)"** and a router deployment **dev-router** is run with route selector **NAMESPACE_LABELS="name=dev"**.

If all routes are in the three namespaces **finance**, **ops** or **dev**, then this could effectively distribute your routes across two router deployments.

In the above scenario, sharding becomes a special case of partitioning with no overlapping sets. Routes are divided amongst multiple router shards.

The criteria for route selection governs how the routes are distributed. It is possible to have routes that overlap across multiple router deployments.

Example: In addition to the **finops-router** and **dev-router** in the example above, you also have **devops-router**, which is run with a route selector **NAMESPACE_LABELS="name in (dev, ops)"**.

The routes in namespaces **dev** or **ops** now are serviced by two different router deployments. This becomes a case in which you have partitioned the routes with an overlapping set.

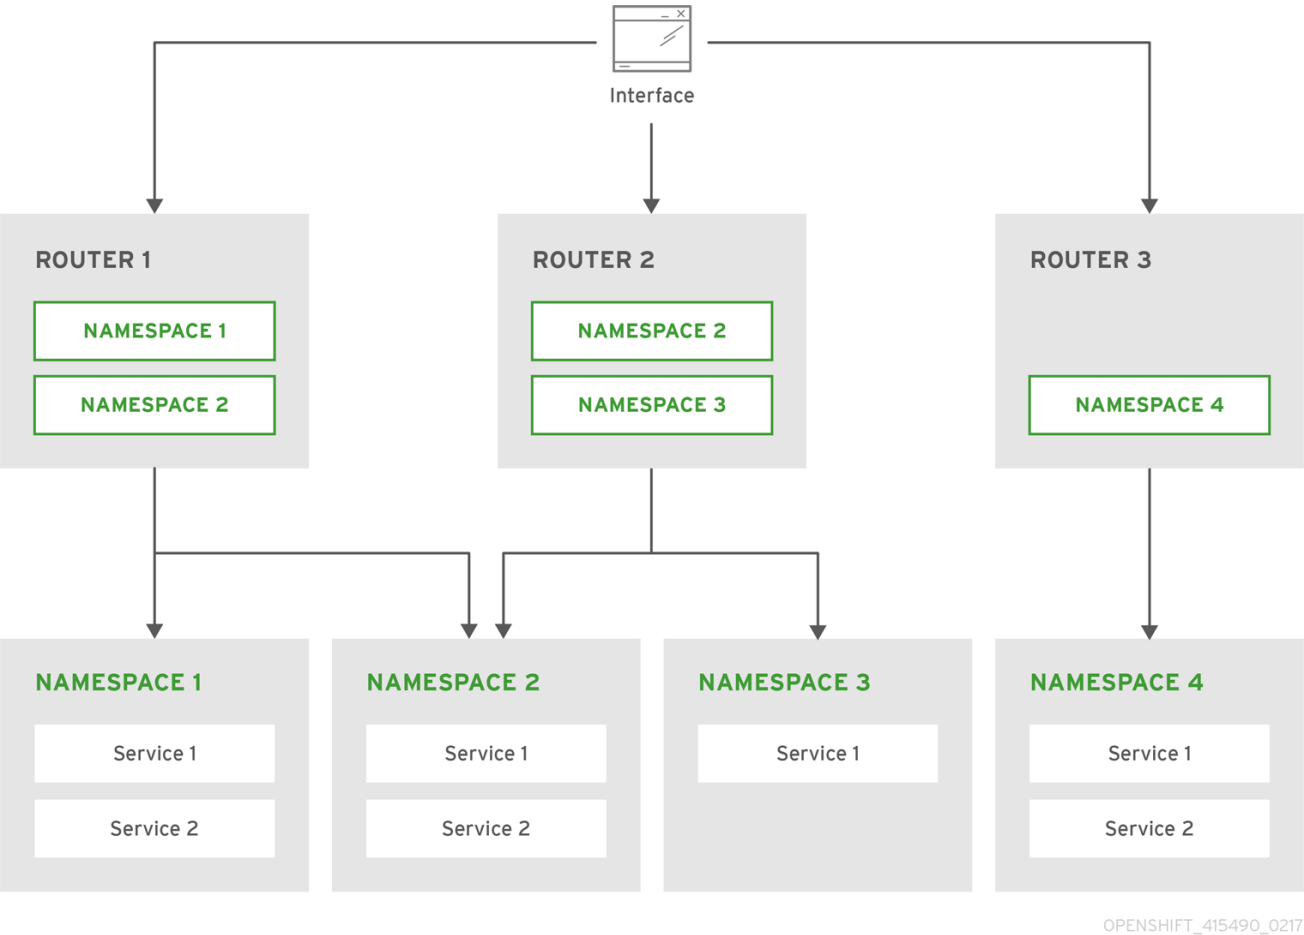
In addition, this enables you to create more complex routing rules, allowing the diversion of high priority traffic to the dedicated **finops-router**, but sending the lower priority ones to the **devops-router**.

NAMESPACE_LABELS allows filtering of the projects to service and selecting all the routes from those projects, but you may want to partition routes based on other criteria in the routes themselves. The **ROUTE_LABELS** selector allows you to slice-and-dice the routes themselves.

Example: A router deployment **prod-router** is run with route selector **ROUTE_LABELS="mydeployment=prod"** and a router deployment **devtest-router** is run with route selector **ROUTE_LABELS="mydeployment in (dev, test)"**.

The example assumes you have all the routes you want to be serviced tagged with a label **"mydeployment=<tag>"**.

Figure 4.2. Router Sharding Based on Namespace Names



4.2.9.1. Creating Router Shards

Router sharding lets you select how routes are distributed among a set of routers.

Router sharding is [based on labels](#); you set labels on the routes in the pool, and express the desired subset of those routes for the router to serve with a selection expression via the `oc set env` command.

First, ensure that service account associated with the router has the [cluster reader](#) permission.

The rest of this section describes an extended example. Suppose there are 26 routes, named **a** — **z**, in the pool, with various labels:

Possible labels on routes in the pool

sla=high	geo=east	hw=modest	dept=finance
sla=medium	geo=west	hw=strong	dept=dev
sla=low			dept=ops

These labels express the concepts: service level agreement, geographical location, hardware requirements, and department. The routes in the pool can have at most one label from each column. Some routes may have other labels, entirely, or none at all.

Name(s)	SLA	Geo	HW	Dept	Other Labels
a	high	east	modest	finance	type=static

Name(s)	SLA	Geo	HW	Dept	Other Labels
b		west	strong		type=dynamic
c, d, e	low		modest		type=static
g—k	medium		strong	dev	
l—s	high		modest	ops	
t—z		west			type=dynamic

Here is a convenience script **mkshard** that illustrates how **oadm router**, **oc set env**, and **oc scale** work together to make a router shard.

```
#!/bin/bash
# Usage: mkshard ID SELECTION-EXPRESSION
id=$1
sel="$2"
router=router-shard-$id
oadm router $router --replicas=0
dc=dc/router-shard-$id
oc set env $dc ROUTE_LABELS="$sel"
oc scale $dc --replicas=3
```

- ❶ The created router has name **router-shard-*<id>***.
- ❷ Specify no scaling for now.
- ❸ The deployment configuration for the router.
- ❹ Set the selection expression using **oc set env**. The selection expression is the value of the **ROUTE_LABELS** environment variable.
- ❺ Scale it up.

Running **mkshard** several times creates several routers:

Router	Selection Expression	Routes
router-shard-1	sla=high	a, l—s
router-shard-2	geo=west	b, t—z
router-shard-3	dept=dev	g—k

4.2.9.2. Modifying Router Shards

Because a router shard is a construct [based on labels](#), you can modify either the labels (via `oc label`) or the selection expression.

This section extends the example started in the [Creating Router Shards](#) section, demonstrating how to change the selection expression.

Here is a convenience script **modshard** that modifies an existing router to use a new selection expression:

```
#!/bin/bash
# Usage: modshard ID SELECTION-EXPRESSION...
id=$1
shift
router=router-shard-$id
dc=dc/$router
oc scale $dc --replicas=0
oc set env $dc "$@"
oc scale $dc --replicas=3
```

- 1 The modified router has name **router-shard-`<id>`**.
- 2 The deployment configuration where the modifications occur.
- 3 Scale it down.
- 4 Set the new selection expression using `oc set env`. Unlike **mkshard** from the [Creating Router Shards](#) section, the selection expression specified as the non-**ID** arguments to **modshard** must include the environment variable name as well as its value.
- 5 Scale it back up.



NOTE

In **modshard**, the `oc scale` commands are not necessary if the [deployment strategy](#) for **router-shard-`<id>`** is **Rolling**.

For example, to expand the department for **router-shard-3** to include **ops** as well as **dev**:

```
$ modshard 3 ROUTE_LABELS='dept in (dev, ops)'
```

The result is that **router-shard-3** now selects routes **g — s** (the combined sets of **g — k** and **l — s**).

This example takes into account that there are only three departments in this example scenario, and specifies a department to leave out of the shard, thus achieving the same result as the preceding example:

```
$ modshard 3 ROUTE_LABELS='dept != finance'
```

This example specifies shows three comma-separated qualities, and results in only route **b** being selected:

```
$ modshard 3 ROUTE_LABELS='hw=strong,type=dynamic,geo=west'
```

Similarly to **ROUTE_LABELS**, which involve a route's labels, you can select routes based on the labels of the route's namespace labels, with the **NAMESPACE_LABELS** environment variable. This example modifies **router-shard-3** to serve routes whose namespace has the label **frequency=weekly**:

```
$ modshard 3 NAMESPACE_LABELS='frequency=weekly'
```

The last example combines **ROUTE_LABELS** and **NAMESPACE_LABELS** to select routes with label **sla=low** and whose namespace has the label **frequency=weekly**:

```
$ modshard 3 \
  NAMESPACE_LABELS='frequency=weekly' \
  ROUTE_LABELS='sla=low'
```

4.2.9.3. Using Namespace Router Shards

The routes for a project can be handled by a selected router by using **NAMESPACE_LABELS**. The router is given a selector for a **NAMESPACE_LABELS** label and the project that wants to use the router applies the **NAMESPACE_LABELS** label to its namespace.

First, ensure that service account associated with the router has the **cluster reader** permission. This permits the router to read the labels that are applied to the namespaces.

Now create and label the router:

```
$ oadm router ... --service-account=router
$ oc set env dc/router NAMESPACE_LABELS="router=r1"
```

Because the router has a selector for a namespace, the router will handle routes for that namespace. So, for example:

```
$ oc label namespace default "router=r1"
```

Now create routes in the default namespace, and the route is available in the default router:

```
$ oc create -f route1.yaml
```

Now create a new project (namespace) and create a route, route2.

```
$ oc new-project p1
$ oc create -f route2.yaml
```

And notice the route is not available in your router. Now label namespace p1 with "router=r1"

```
$ oc label namespace p1 "router=r1"
```

Which makes the route available to the router.

Note that removing the label from the namespace won't have immediate effect (as we don't see the updates in the router), so if you redeploy/start a new router pod, you should see the unlabelled effects.

```
$ oc scale dc/router --replicas=0 && oc scale dc/router --replicas=1
```

4.2.10. Customizing the Default Routing Subdomain

You can customize the suffix used as the default routing subdomain for your environment by modifying the [master configuration file](#) (the `/etc/origin/master/master-config.yaml` file by default). Routes that do not specify a host name would have one generated using this default routing subdomain.

The following example shows how you can set the configured suffix to **v3.openshift.test**:

```
routingConfig:
  subdomain: v3.openshift.test
```



NOTE

This change requires a restart of the master if it is running.

With the OpenShift Container Platform master(s) running the above configuration, the [generated host name](#) for the example of a route named **no-route-hostname** without a host name added to a namespace **mynamespace** would be:

```
no-route-hostname-mynamespace.v3.openshift.test
```

4.2.11. Forcing Route Host Names to a Custom Routing Subdomain

If an administrator wants to restrict all routes to a specific routing subdomain, they can pass the **--force-subdomain** option to the **oadm router** command. This forces the router to override any host names specified in a route and generate one based on the template provided to the **--force-subdomain** option.

The following example runs a router, which overrides the route host names using a custom subdomain template `${name}-${namespace}.apps.example.com`.

```
$ oadm router --force-subdomain='${name}-${namespace}.apps.example.com'
```

4.2.12. Using Wildcard Certificates

A TLS-enabled route that does not include a certificate uses the router's default certificate instead. In most cases, this certificate should be provided by a trusted certificate authority, but for convenience you can use the OpenShift Container Platform CA to create the certificate. For example:


```
$ CA=/etc/origin/master
$ oadm ca create-server-cert --signer-cert=$CA/ca.crt \
  --signer-key=$CA/ca.key --signer-serial=$CA/ca.serial.txt \
  --hostnames='*.cloudapps.example.com' \
  --cert=cloudapps.crt --key=cloudapps.key
```

The router expects the certificate and key to be in PEM format in a single file:

```
$ cat cloudapps.crt cloudapps.key $CA/ca.crt > cloudapps.router.pem
```

From there you can use the **--default-cert** flag:

```
$ oadm router --default-cert=cloudapps.router.pem --service-account=router
```



NOTE

Browsers only consider wildcards valid for subdomains one level deep. So in this example, the certificate would be valid for *a.cloudapps.example.com* but not for *a.b.cloudapps.example.com*.

4.2.13. Manually Redeploy Certificates

To manually redeploy the router certificates:

1. Check to see if a secret containing the default router certificate was added to the router:

```
$ oc volumes dc/router
deploymentconfigs/router
secret/router-certs as server-certificate
mounted at /etc/pki/tls/private
```

If the certificate is added, skip the following step and overwrite the secret.

2. Make sure that you have a default certificate directory set for the following variable **DEFAULT_CERTIFICATE_DIR**:

```
$ oc env dc/router --list
DEFAULT_CERTIFICATE_DIR=/etc/pki/tls/private
```

If not, create the directory using the following command:

```
$ oc env dc/router DEFAULT_CERTIFICATE_DIR=/etc/pki/tls/private
```

3. Export the certificate to PEM format:

```
$ cat custom-router.key custom-router.crt custom-ca.crt > custom-
router.crt
```

4. Overwrite or create a router certificate secret:

If the certificate secret was added to the router, overwrite the secret. If not, create a new secret.

To overwrite the secret, run the following command:

```
$ oc secrets new router-certs tls.crt=custom-router.crt
tls.key=custom-router.key -o json --type='kubernetes.io/tls' --
confirm | oc replace -f -
```

To create a new secret, run the following commands:

```
$ oc secrets new router-certs tls.crt=custom-router.crt
tls.key=custom-router.key --type='kubernetes.io/tls' --confirm
```

```
$ oc volume dc/router --add --mount-path=/etc/pki/tls/private --
secret-name='router-certs' --name router-certs
```

5. Deploy the router.

```
$ oc deploy router --latest
```

4.2.14. Using Secured Routes

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>
```

Here is an example of how to use a secure edge terminated route with TLS termination occurring on the router before traffic is proxied to the destination. The secure edge terminated route specifies the TLS certificate and key information. The TLS certificate is served by the router front end.

First, start up a router instance:

```
# oadm router --replicas=1 --service-account=router
```

Next, create a private key, csr and certificate for our edge secured route. The instructions on how to do that would be specific to your certificate authority and provider. For a simple self-signed certificate for a domain named **www.example.test**, see the example shown below:

```
# sudo openssl genrsa -out example-test.key 2048
#
# sudo openssl req -new -key example-test.key -out example-test.csr \
-subj "/C=US/ST=CA/L=Mountain View/O=OS3/OU=Eng/CN=www.example.test"
#
# sudo openssl x509 -req -days 366 -in example-test.csr \
-signkey example-test.key -out example-test.crt
```

Generate a route using the above certificate and key.

```
$ oc create route edge --service=my-service \
--hostname=www.example.test \
--key=example-test.key --cert=example-test.crt
route "my-service" created
```

Look at its definition.

```
$ oc get route/my-service -o yaml
apiVersion: v1
kind: Route
metadata:
  name: my-service
spec:
  host: www.example.test
  to:
    kind: Service
```

```

    name: my-service
  tls:
    termination: edge
    key: |
      -----BEGIN PRIVATE KEY-----
      [...]
      -----END PRIVATE KEY-----
    certificate: |
      -----BEGIN CERTIFICATE-----
      [...]
      -----END CERTIFICATE-----

```

Make sure your DNS entry for **www.example.test** points to your router instance(s) and the route to your domain should be available. The example below uses curl along with a local resolver to simulate the DNS lookup:

```

# routerip="4.1.1.1" # replace with IP address of one of your router
instances.
# curl -k --resolve www.example.test:443:$routerip
https://www.example.test/

```

4.2.15. Using the Container Network Stack

The OpenShift Container Platform router runs inside a container and the default behavior is to use the network stack of the host (i.e., the node where the router container runs). This default behavior benefits performance because network traffic from remote clients does not need to take multiple hops through user space to reach the target service and container.

Additionally, this default behavior enables the router to get the actual source IP address of the remote connection rather than getting the node's IP address. This is useful for defining ingress rules based on the originating IP, supporting sticky sessions, and monitoring traffic, among other uses.

This host network behavior is controlled by the **--host-network** router command line option, and the default behaviour is the equivalent of using **--host-network=true**. If you wish to run the router with the container network stack, use the **--host-network=false** option when creating the router. For example:

```
$ oadm router --service-account=router --host-network=false
```

Internally, this means the router container must publish the 80 and 443 ports in order for the external network to communicate with the router.



NOTE

Running with the container network stack means that the router sees the source IP address of a connection to be the NATed IP address of the node, rather than the actual remote IP address.


**NOTE**

On OpenShift Container Platform clusters using [multi-tenant network isolation](#), routers on a non-default namespace with the `--host-network=false` option will load all routes in the cluster, but routes across the namespaces will not be reachable due to network isolation. With the `--host-network=true` option, routes bypass the container network and it can access any pod in the cluster. If isolation is needed in this case, then do not add routes across the namespaces.

4.2.16. Exposing Router Metrics

Using the `--metrics-image` and `--expose-metrics` options, you can configure the OpenShift Container Platform router to run a sidecar container that exposes or publishes router metrics for consumption by external metrics collection and aggregation systems (e.g. Prometheus, statsd).

Depending on your router implementation, the image is appropriately set up and the metrics sidecar container is started when the router is deployed. For example, the HAProxy-based router implementation defaults to using the `prom/haproxy-exporter` image to run as a sidecar container, which can then be used as a metrics datasource by the Prometheus server.

**NOTE**

The `--metrics-image` option overrides the defaults for HAProxy-based router implementations and, in the case of custom implementations, enables the image to use for a custom metrics exporter or publisher.

1. Grab the HAProxy Prometheus exporter image from the Docker registry:

```
$ sudo docker pull prom/haproxy-exporter
```

2. Create the OpenShift Container Platform router:

```
$ oadm router --service-account=router --expose-metrics
```

Or, optionally, use the `--metrics-image` option to override the HAProxy defaults:

```
$ oadm router --service-account=router --expose-metrics \
  --metrics-image=prom/haproxy-exporter
```

3. Once the haproxy-exporter containers (and your HAProxy router) have started, point Prometheus to the sidecar container on port 9101 on the node where the haproxy-exporter container is running:

```
$ haproxy_exporter_ip="<enter-ip-address-or-hostname>"
$ cat > haproxy-scraper.yml <<CFGE0F
---
global:
  scrape_interval: "60s"
  scrape_timeout: "10s"
  # external_labels:
  #   source: openshift-router

scrape_configs:
```

```

- job_name: "haproxy"
  target_groups:
    - targets:
      - "${haproxy_exporter_ip}:9101"
CFGEOF

$ # And start prometheus as you would normally using the above
  config file.
$ echo " - Example: prometheus -config.file=haproxy-scraper.yml "
$ echo " or you can start it as a container on
{product-title}!!

$ echo " - Once the prometheus server is up, view the {product-
title} HAProxy "
$ echo " router metrics at:
http://<ip>:9090/consales/haproxy.html "
```

4.2.17. Preventing Connection Failures During Restarts

If you connect to the router while the proxy is reloading, there is a small chance that your connection will end up in the wrong network queue and be dropped. The issue is being addressed. In the meantime, it is possible to work around the problem by installing **iptables** rules to prevent connections during the reload window. However, doing so means that the router needs to run with elevated privilege so that it can manipulate **iptables** on the host. It also means that connections that happen during the reload are temporarily ignored and must retransmit their connection start, lengthening the time it takes to connect, but preventing connection failure.

To prevent this, configure the router to use **iptables** by changing the service account, and setting an environment variable on the router.

Use a Privileged SCC

When creating the router, allow it to use the privileged SCC. This gives the router user the ability to create containers with root privileges on the nodes:

```
$ oadm policy add-scc-to-user privileged -z router
```

Patch the Router Deployment Configuration to Create a Privileged Container

You can now create privileged containers. Next, configure the router deployment configuration to use the privilege so that the router can set the iptables rules it needs. This patch changes the router deployment configuration so that the container that is created runs as privileged (and therefore gets correct capabilities) and run as root:

```
$ oc patch dc router -p '{"spec":{"template":{"spec":{"containers":
[{"name":"router","securityContext":
{"privileged":true}]},"securityContext":{"runAsUser": 0}}}}}'
```

Configure the Router to Use iptables

Set the option on the router deployment configuration:

```
$ oc set env dc/router -c router DROP_SYN_DURING_RESTART=true
```

If you used a non-default name for the router, you must change **dc/router** accordingly.

4.2.18. Protecting Against DDoS Attacks

Add **timeout http-request** to the default HAProxy router image to protect the deployment against distributed denial-of-service (DDoS) attacks (for example, slowloris):

```
# and the haproxy stats socket is available at /var/run/haproxy.stats
global
  stats socket ./haproxy.stats level admin

defaults
  option http-server-close
  mode http
  timeout http-request 5s
  timeout connect 5s
  timeout server 10s
  timeout client 30s
```

- 1** **timeout http-request** is set up to 5 seconds. HAProxy gives a client 5 seconds *to send its whole HTTP request. Otherwise, HAProxy shuts the connection with *an error.

Also, when the environment variable **ROUTER_SLOWLORIS_TIMEOUT** is set, it limits the amount of time a client has to send the whole HTTP request. Otherwise, HAProxy will shut down the connection.

Setting the environment variable allows information to be captured as part of the router's deployment configuration and does not require manual modification of the template, whereas manually adding the HAProxy setting requires you to rebuild the router pod and maintain your router template file.

Using annotations implements basic DDoS protections in the HAProxy template router, including the ability to limit the:

- number of concurrent TCP connections
- rate at which a client can request TCP connections
- rate at which HTTP requests can be made

These are enabled on a per route basis because applications can have extremely different traffic patterns.

Table 4.1. HAProxy Template Router Settings

Setting	Description
haproxy.router.openshift.io/rate-limit-connections	Enables the settings be configured (when set to true , for example).
haproxy.router.openshift.io/rate-limit-connections.concurrent-tcp	The number of concurrent TCP connections that can be made by the same IP address on this route.
haproxy.router.openshift.io/rate-limit-connections.rate-tcp	The number of TCP connections that can be opened by a client IP.

Setting	Description
<code>haproxy.router.openshift.io/rate-limit-connections.rate-http</code>	The number of HTTP requests that a client IP can make in a 3-second period.

4.3. DEPLOYING A CUSTOMIZED HAPROXY ROUTER

4.3.1. Overview

The default HAProxy router is intended to satisfy the needs of most users. However, it does not expose all of the capability of HAProxy. Therefore, users may need to modify the router for their own needs.

You may need to implement new features within the application back-ends, or modify the current operation. The router plug-in provides all the facilities necessary to make this customization.

The router pod uses a template file to create the needed HAProxy configuration file. The template file is a [golang template](#). When processing the template, the router has access to OpenShift Container Platform information, including the router's deployment configuration, the set of admitted routes, and some helper functions.

When the router pod starts, and every time it reloads, it creates an HAProxy configuration file, and then it starts HAProxy. The [HAProxy configuration manual](#) describes all of the features of HAProxy and how to construct a valid configuration file.

A [configMap](#) can be used to add the new template to the router pod. With this approach, the router deployment configuration is modified to mount the **configMap** as a volume in the router pod. The **TEMPLATE_FILE** environment variable is set to the full path name of the template file in the router pod.

Alternatively, you can build a custom router image and use it when deploying some or all of your routers. There is no need for all routers to run the same image. To do this, modify the **haproxy-template.config** file, and [rebuild](#) the router image. The new image is pushed to the cluster's Docker repository, and the router's deployment configuration **image:** field is updated with the new name. When the cluster is updated, the image needs to be rebuilt and pushed.

In either case, the router pod starts with the template file.

4.3.2. Obtaining the Router Configuration Template

The HAProxy template file is fairly large and complex. For some changes, it may be easier to modify the existing template rather than writing a complete replacement. You can obtain a **haproxy-config.template** file from a running router by running this on master, referencing the router pod:

```
# oc get po
NAME                                READY    STATUS    RESTARTS   AGE
router-2-40fc3                     1/1      Running   0           11d
# oc rsh router-2-40fc3 cat haproxy-config.template > haproxy-
config.template
# oc rsh router-2-40fc3 cat haproxy.config > haproxy.config
```

Alternatively, you can log onto the node that is running the router:

```
# docker run --rm --interactive=true --tty --entrypoint=cat \
    registry.access.redhat.com/openshift3/ose-haproxy-router:$version
haproxy-config.template
```

The image name is from **docker images**.

Save this content to a file for use as the basis of your customized template. The saved **haproxy.config** shows what is actually running.

4.3.3. Modifying the Router Configuration Template

4.3.3.1. Background

The template is based on the [golang template](#). It can reference any of the environment variables in the router's deployment configuration, any configuration information that is described below, and router provided helper functions.

The structure of the template file mirrors the resulting HAProxy configuration file. As the template is processed, anything not surrounded by `{{" something "}}` is directly copied to the configuration file. Passages that are surrounded by `{{" something "}}` are evaluated. The resulting text, if any, is copied to the configuration file.

4.3.3.2. Go Template Actions

The **define** action names the file that will contain the processed template.

```
{{define "/var/lib/haproxy/conf/haproxy.config"}}pipeline{{end}}
```

Table 4.2. Template Router Functions

Function	Meaning
endpointsForAlias(alias ServiceAliasConfig, svc ServiceUnit) []Endpoint	Returns the list of valid endpoints.
env(variable, default string) string	Tries to get the named environment variable from the pod. Returns the second argument if the variable cannot be read or is empty.
matchPattern(pattern, s string) bool	The first argument is a string that contains the regular expression, the second argument is the variable to test. Returns a Boolean value indicating whether the regular expression provided as the first argument matches the string provided as the second argument.
isInteger(s string) bool	Determines if a given variable is an integer.
matchValues(s string, allowedValues ...string) bool	Compares a given string to a list of allowed strings.

Function	Meaning
genSubdomainWildcardRegex(hostname, path string, exactPath bool) string	Generates a regular expression matching the subdomain for hosts (and paths) with a wildcard policy.
generateRouteRegex(hostname, path string, wildcard bool) string	Generates a regular expression matching the route hosts (and paths). The first argument is the host name, the second is the path, and the third is a wildcard Boolean.
genCertificateHostName(hostname string, wildcard bool) string	Generates host name to use for serving/matching certificates. First argument is the host name and the second is the wildcard Boolean.

These functions are provided by the HAProxy template router plug-in.

4.3.3.3. Router Provided Information

This section reviews the OpenShift Container Platform information that the router makes available to the template. The router configuration parameters are the set of data that the HAProxy router plug-in is given. The fields are accessed by **(dot) .Fieldname**.

The tables below the Router Configuration Parameters expand on the definitions of the various fields. In particular, **.State** has the set of admitted routes.

Table 4.3. Router Configuration Parameters

Field	Type	Description
WorkingDir	string	The directory that files will be written to, defaults to <i>/var/lib/containers/router</i>
State	map[string] (ServiceAliasConfig) `	The routes.
ServiceUnits	map[string]ServiceUnit	The service lookup.
DefaultCertificate	string	Full path name to the default certificate in pem format.
PeerEndpoints	`[]Endpoint	Peers.
StatsUser	string	User name to expose stats with (if the template supports it).
StatsPassword	string	Password to expose stats with (if the template supports it).

Field	Type	Description
StatsPort	int	Port to expose stats with (if the template supports it).
BindPorts	bool	Whether the router should bind the default ports.

Table 4.4. Router ServiceAliasConfig (A Route)

Field	Type	Description
Name	string	The user-specified name of the route.
Namespace	string	The namespace of the route.
Host	string	The host name. For example, www.example.com .
Path	string	Optional path. For example, www.example.com/myservice where myservice is the path.
TLSTermination	routeapi.TLSTerminationType	The termination policy for this back-end; drives the mapping files and router configuration.
Certificates	map[string]Certificate	Certificates used for securing this back-end. Keyed by the certificate ID.
Status	ServiceAliasConfigStatus	Indicates the status of configuration that needs to be persisted.
PreferPort	string	Indicates the port the user wants to expose. If empty, a port will be selected for the service.
InsecureEdgeTerminationPolicy	routeapi.InsecureEdgeTerminationPolicyType	Indicates desired behavior for insecure connections to an edge-terminated route: none (or disable), allow , or redirect .

Field	Type	Description
RoutingKeyName	string	Hash of the route + namespace name used to obscure the cookie ID.
IsWildcard	bool	Indicates this service unit needing wildcard support.
Annotations	map[string]string	Annotations attached to this route.
ServiceUnitNames	map[string]int32	Collection of services that support this route, keyed by service name and valued on the weight attached to it with respect to other entries in the map.
ActiveServiceUnits	int	Count of the ServiceUnitNames with a non-zero weight.

The **ServiceAliasConfig** is a route for a service. Uniquely identified by host + path. The default template iterates over routes using `{{range $cfgIdx, $cfg := .State }}`. Within such a `{{range}}` block, the template can refer to any field of the current **ServiceAliasConfig** using `$cfg.Field`.

Table 4.5. Router ServiceUnit

Field	Type	Description
Name	string	Name corresponds to a service name + namespace. Uniquely identifies the ServiceUnit .
EndpointTable	[]Endpoint	Endpoints that back the service. This translates into a final back-end implementation for routers.

ServiceUnit is an encapsulation of a service, the endpoints that back that service, and the routes that point to the service. This is the data that drives the creation of the router configuration files

Table 4.6. Router Endpoint

Field	Type
ID	string
IP	string

Field	Type
Port	string
TargetName	string
PortName	string
IdHash	string
NoHealthCheck	bool

Endpoint is an internal representation of a Kubernetes endpoint.

Table 4.7. Router Certificate, ServiceAliasConfigStatus

Field	Type	Description
Certificate	string	Represents a public/private key pair. It is identified by an ID, which will become the file name. A CA certificate will not have a PrivateKey set.
ServiceAliasConfigStatus	string	Indicates that the necessary files for this configuration have been persisted to disk. Valid values: "saved", "".

Table 4.8. Router Certificate Type

Field	Type	Description
ID	string	
Contents	string	The certificate.
PrivateKey	string	The private key.

Table 4.9. Router TLSTerminationType

Field	Type	Description
TLSTerminationType	string	Dictates where the secure communication will stop.

Field	Type	Description
InsecureEdgeTerminationPolicyType	string	Indicates the desired behavior for insecure connections to a route. While each router may make its own decisions on which ports to expose, this is normally port 80.

TLSTerminationType and **InsecureEdgeTerminationPolicyType** dictate where the secure communication will stop.

Table 4.10. Router TLSTerminationType Values

Constant	Value	Meaning
TLSTerminationEdge	edge	Terminate encryption at the edge router.
TLSTerminationPassthrough	passthrough	Terminate encryption at the destination, the destination is responsible for decrypting traffic.
TLSTerminationReencrypt	reencrypt	Terminate encryption at the edge router and re-encrypt it with a new certificate supplied by the destination.

Table 4.11. Router InsecureEdgeTerminationPolicyType Values

Type	Meaning
Allow	Traffic is sent to the server on the insecure port (default).
Disable	No traffic is allowed on the insecure port.
Redirect	Clients are redirected to the secure port.

None ("") is the same as **Disable**.

4.3.3.4. Annotations

Each route can have annotations attached. Each annotation is just a name and a value.

```
apiVersion: v1
kind: Route
metadata:
```

```

    annotations:
      haproxy.router.openshift.io/timeout: 5500ms
  [...]

```

The name can be anything that does not conflict with existing Annotations. The value is any string. The string can have multiple tokens separated by a space. For example, **aa bb cc**. The template uses `{{index}}` to extract the value of an annotation. For example:

```

{{ $balanceAlgo := index $cfg.Annotations
  "haproxy.router.openshift.io/balance" }}

```

This is an example of how this could be used for mutual client authorization.

```

{{ with $cnList := index $cfg.Annotations "whiteListCertCommonName" }}
  {{ if ne $cnList "" }}
    acl test ssl_c_s_dn(CN) -m str {{ $cnList }}
    http-request deny if !test
  {{ end }}
{{ end }}

```

Then, you can handle the white-listed CNs with this command.

```

$ oc annotate route <route-name> --overwrite whiteListCertCommonName="CN1
CN2 CN3"

```

4.3.3.5. Environment Variables

The template can use any environment variables that exist in the router pod. The environment variables can be set in the deployment configuration. New environment variables can be added.

They are referenced by the **env** function:

```

{{env "ROUTER_MAX_CONNECTIONS" "20000"}}

```

The first string is the variable, and the second string is the default when the variable is missing or **nil**. When **ROUTER_MAX_CONNECTIONS** is not set or is **nil**, 20000 is used. Environment variables are a map where the key is the environment variable name and the content is the value of the variable.

See [Route-specific Environment variables](#) for more information.

4.3.3.6. Example Usage

Here is a simple template based on the HAProxy template file.

Start with a comment:

```

{{/*
  Here is a small example of how to work with templates
  taken from the HAProxy template file.
*/}}

```

The template can create any number of output files. Use a `define` construct to create an output file. The file name is specified as an argument to `define`, and everything inside the `define` block up to the matching `end` is written as the contents of that file.

```
{{ define "/var/lib/haproxy/conf/haproxy.config" }}
global
{{ end }}
```

The above will copy **global** to the **/var/lib/haproxy/conf/haproxy.config** file, and then close the file.

Set up logging based on environment variables.

```
{{ with (env "ROUTER_SYSLOG_ADDRESS" "") }}
  log {{.}} {{env "ROUTER_LOG_FACILITY" "local1"}} {{env
"ROUTER_LOG_LEVEL" "warning"}}
{{ end }}
```

The **env** function extracts the value for the environment variable. If the environment variable is not defined or **nil**, the second argument is returned.

The **with** construct sets the value of **.** (dot) within the **with** block to whatever value is provided as an argument to **with**. The **with** action tests **Dot** for **nil**. If not **nil**, the clause is processed up to the **end**. In the above, assume **ROUTER_SYSLOG_ADDRESS** contains **/var/log/msg**, **ROUTER_LOG_FACILITY** is not defined, and **ROUTER_LOG_LEVEL** contains **info**. The following will be copied to the output file:

```
log /var/log/msg local1 info
```

Each admitted route ends up generating lines in the configuration file. Use **range** to go through the admitted routes:

```
{{ range $cfgIdx, $cfg := .State }}
  backend be_http_{{ $cfgIdx }}
{{ end }}
```

.State is a map of **ServiceAliasConfig**, where the key is the route name. **range** steps through the map and, for each pass, it sets **\$cfgIdx** with the **key**, and sets **\$cfg** to point to the **ServiceAliasConfig** that describes the route. If there are two routes named **myroute** and **hisroute**, the above will copy the following to the output file:

```
backend be_http_myroute
backend be_http_hisroute
```

Route Annotations, **\$cfg.Annotations**, is also a map with the annotation name as the key and the content string as the value. The route can have as many annotations as desired and the use is defined by the template author. The user codes the annotation into the route and the template author customized the HAProxy template to handle the annotation.

The common usage is to index the annotation to get the value.

```
{{ $balanceAlgo := index $cfg.Annotations
"haproxy.router.openshift.io/balance" }}
```

The index extracts the value for the given annotation, if any. Therefore, ``$balanceAlgo` will contain the string associated with the annotation or `nil`. As above, you can test for a non-`nil` string and act on it with the `with` construct.

```
{{ with $balanceAlgo }}
  balance $balanceAlgo
{{ end }}
```

Here when `$balanceAlgo` is not `nil`, `balance $balanceAlgo` is copied to the output file.

In a second example, you want to set a server timeout based on a timeout value set in an annotation.

```
$value := index $cfg.Annotations "haproxy.router.openshift.io/timeout"
```

The `$value` can now be evaluated to make sure it contains a properly constructed string. The `matchPattern` function accepts a regular expression and returns `true` if the argument satisfies the expression.

```
matchPattern "[1-9][0-9]*(us|ms|s|m|h|d)?" $value
```

This would accept `5000ms` but not `7y`. The results can be used in a test.

```
{{if (matchPattern "[1-9][0-9]*(us|ms|s|m|h|d)?" $value) }}
  timeout server {{$value}}
{{ end }}
```

It can also be used to match tokens:

```
matchPattern "roundrobin|leastconn|source" $balanceAlgo
```

Alternatively `matchValues` can be used to match tokens:

```
matchValues $balanceAlgo "roundrobin" "leastconn" "source"
```

4.3.4. Using a ConfigMap to Replace the Router Configuration Template

You can use a [ConfigMap](#) to customize the router instance without rebuilding the router image. The *haproxy-config.template*, *reload-haproxy*, and other scripts can be modified as well as creating and modifying router environment variables.

1. Copy the *haproxy-config.template* that you want to modify as [described above](#). Modify it as desired.
2. Create a ConfigMap:

```
$ oc create configmap customrouter --from-file=haproxy-
config.template
```

The `customrouter` ConfigMap now contains a copy of the modified *haproxy-config.template* file.

3. Modify the router deployment configuration to mount the ConfigMap as a file and point the **TEMPLATE_FILE** environment variable to it. This can be done via **oc set env** and **oc volume** commands, or alternatively by editing the router deployment configuration.

Using oc commands

```
$ oc set env dc/router \
    TEMPLATE_FILE=/var/lib/haproxy/conf/custom/haproxy-
config.template
$ oc volume dc/router --add --overwrite \
    --name=config-volume \
    --mount-path=/var/lib/haproxy/conf/custom \
    --source='{"configMap": { "name": "customrouter"}}'
```

Editing the Router Deployment Configuration

Use **oc edit dc router** to edit the router deployment configuration with a text editor.

```
...
    - name: STATS_USERNAME
      value: admin
    - name: TEMPLATE_FILE 1
      value: /var/lib/haproxy/conf/custom/haproxy-
config.template
    image: openshift/origin-haproxy-routerp
...
    terminationMessagePath: /dev/termination-log
    volumeMounts: 2
    - mountPath: /var/lib/haproxy/conf/custom
      name: config-volume
    dnsPolicy: ClusterFirst
...
    terminationGracePeriodSeconds: 30
    volumes: 3
    - configMap:
      name: customrouter
      name: config-volume
    test: false
...
```

- 1** In the **spec.container.env** field, add the **TEMPLATE_FILE** environment variable to point to the mounted **haproxy-config.template** file.
- 2** Add the **spec.container.volumeMounts** field to create the mount point.
- 3** Add a new **spec.volumes** field to mention the ConfigMap.

Save the changes and exit the editor. This restarts the router.

4.3.5. Using Stick Tables

The following example customization can be used in a [highly-available routing setup](#) to use stick-tables that synchronize between peers.

Adding a Peer Section

In order to synchronize stick-tables amongst peers you must define a peers section in your HAProxy configuration. This section determines how HAProxy will identify and connect to peers. The plug-in provides data to the template under the **.PeerEndpoints** variable to allow you to easily identify members of the router service. You may add a peer section to the *haproxy-config.template* file inside the router image by adding:

```
{{ if (len .PeerEndpoints) gt 0 }}
peers openshift_peers
  {{ range $endpointID, $endpoint := .PeerEndpoints }}
    peer {{$endpoint.TargetName}} {{$endpoint.IP}}:1937
  {{ end }}
{{ end }}
```

Changing the Reload Script

When using stick-tables, you have the option of telling HAProxy what it should consider the name of the local host in the peer section. When creating endpoints, the plug-in attempts to set the **TargetName** to the value of the endpoint's **TargetRef.Name**. If **TargetRef** is not set, it will set the **TargetName** to the IP address. The **TargetRef.Name** corresponds with the Kubernetes host name, therefore you can add the **-L** option to the **reload-haproxy** script to identify the local host in the peer section.

```
peer_name=$HOSTNAME ❶

if [ -n "$old_pid" ]; then
  /usr/sbin/haproxy -f $config_file -p $pid_file -L $peer_name -sf
  $old_pid
else
  /usr/sbin/haproxy -f $config_file -p $pid_file -L $peer_name
fi
```

❶ Must match an endpoint target name that is used in the peer section.

Modifying Back Ends

Finally, to use the stick-tables within back ends, you can modify the HAProxy configuration to use the stick-tables and peer set. The following is an example of changing the existing back end for TCP connections to use stick-tables:

```
        {{ if eq $cfg.TLS Termination "passthrough" }}
backend be_tcp_{{$cfgIdx}}
  balance leastconn
  timeout check 5000ms
  stick-table type ip size 1m expire 5m{{ if (len $.PeerEndpoints) gt 0 }}
peers openshift_peers {{ end }}
  stick on src
        {{ range $endpointID, $endpoint :=
$serviceUnit.EndpointTable }}
  server {{$endpointID}} {{$endpoint.IP}}:{{$endpoint.Port}} check inter
5000ms
        {{ end }}
{{ end }}
```

After this modification, you can [rebuild your router](#).

4.3.6. Rebuilding Your Router

In order to rebuild the router, you need copies of several files that are present on a running router. Make a work directory and copy the files from the router:

```
# mkdir -p myrouter/conf
# cd myrouter
# oc get po
NAME                                READY    STATUS    RESTARTS   AGE
router-2-40fc3                      1/1      Running   0           11d
# oc rsh router-2-40fc3 cat haproxy-config.template > conf/haproxy-
config.template
# oc rsh router-2-40fc3 cat error-page-503.http > conf/error-page-503.http
# oc rsh router-2-40fc3 cat default_pub_keys.pem >
conf/default_pub_keys.pem
# oc rsh router-2-40fc3 cat ../Dockerfile > Dockerfile
# oc rsh router-2-40fc3 cat ../reload-haproxy > reload-haproxy
```

You can edit or replace any of these files. However, ***conf/haproxy-config.template*** and ***reload-haproxy*** are the most likely to be modified.

After updating the files:

```
# docker build -t openshift/origin-haproxy-router-myversion .
# docker tag openshift/origin-haproxy-router-myversion
172.30.243.98:5000/openshift/haproxy-router-myversion 1
# docker push 172.30.243.98:5000/openshift/origin-haproxy-router-pc:latest
2
```

- 1** Tag the version with the repository. In this case the repository is **172.30.243.98:5000**.
- 2** Push the tagged version to the repository. It may be necessary to **docker login** to the repository first.

To use the new router, edit the router deployment configuration either by changing the **image:** string or by adding the **--images=<repo>/<image>:<tag>** flag to the **oadm router** command.

When debugging the changes, it is helpful to set **imagePullPolicy: Always** in the deployment configuration to force an image pull on each pod creation. When debugging is complete, you can change it back to **imagePullPolicy: IfNotPresent** to avoid the pull on each pod start.

4.4. CONFIGURING THE HAPROXY ROUTER TO USE THE PROXY PROTOCOL

4.4.1. Overview

By default, the HAProxy router expects incoming connections to unsecure, edge, and re-encrypt routes to use HTTP. However, you can configure the router to expect incoming requests by using [the PROXY protocol](#) instead. This topic describes how to configure the HAProxy router and an external load balancer to use the PROXY protocol.

4.4.2. Why Use the PROXY Protocol?

When an intermediary service such as a proxy server or load balancer forwards an HTTP request, it appends the source address of the connection to the request's "Forwarded" header in order to provide this information to subsequent intermediaries and to the back-end service to which the request is ultimately forwarded. However, if the connection is encrypted, intermediaries cannot modify the "Forwarded" header. In this case, the HTTP header will not accurately communicate the original source address when the request is forwarded.

To solve this problem, some load balancers encapsulate HTTP requests using the PROXY protocol as an alternative to simply forwarding HTTP. Encapsulation enables the load balancer to add information to the request without modifying the forwarded request itself. In particular, this means that the load balancer can communicate the source address even when forwarding an encrypted connection.

The HAProxy router can be configured to accept the PROXY protocol and decapsulate the HTTP request. Because the router terminates encryption for edge and re-encrypt routes, the router can then update the "Forwarded" HTTP header (and related HTTP headers) in the request, appending any source address that is communicated using the PROXY protocol.



WARNING

The PROXY protocol and HTTP are incompatible and cannot be mixed. If you use a load balancer in front of the router, both must use either the PROXY protocol or HTTP. Configuring one to use one protocol and the other to use the other protocol will cause routing to fail.

4.4.3. Using the PROXY Protocol

By default, the HAProxy router does not use the PROXY protocol. The router can be configured using the **ROUTER_USE_PROXY_PROTOCOL** environment variable to expect the PROXY protocol for incoming connections:

Enable the PROXY Protocol

```
$ oc env dc/router ROUTER_USE_PROXY_PROTOCOL=true
```

Set the variable to any value other than **true** or **TRUE** to disable the PROXY protocol:

Disable the PROXY Protocol

```
$ oc env dc/router ROUTER_USE_PROXY_PROTOCOL=false
```

If you enable the PROXY protocol in the router, you must configure your load balancer in front of the router to use the PROXY protocol as well. Following is an example of configuring Amazon's Elastic Load Balancer (ELB) service to use the PROXY protocol. This example assumes that ELB is forwarding ports 80 (HTTP), 443 (HTTPS), and 5000 (for the image registry) to the router running on one or more EC2 instances.

Configure Amazon ELB to Use the PROXY Protocol

1. To simplify subsequent steps, first set some shell variables:

```
$ lb='infra-lb' ❶
$ instances=( 'i-079b4096c654f563c' ) ❷
$ secgroups=( 'sg-e1760186' ) ❸
$ subnets=( 'subnet-cf57c596' ) ❹
```

- ❶ The name of your ELB.
- ❷ The instance or instances on which the router is running.
- ❸ The security group or groups for this ELB.
- ❹ The subnet or subnets for this ELB.

2. Next, create the ELB with the appropriate listeners, security groups, and subnets.



NOTE

You must configure all listeners to use the TCP protocol, not the HTTP protocol.

```
$ aws elb create-load-balancer --load-balancer-name "$lb" \
    --listeners \
    'Protocol=TCP,LoadBalancerPort=80,InstanceProtocol=TCP,InstancePort=80' \
    'Protocol=TCP,LoadBalancerPort=443,InstanceProtocol=TCP,InstancePort=443' \
    'Protocol=TCP,LoadBalancerPort=5000,InstanceProtocol=TCP,InstancePort=5000' \
    --security-groups $secgroups \
    --subnets $subnets
{
  "DNSName": "infra-lb-2006263232.us-east-1.elb.amazonaws.com"
}
```

3. Register your router instance or instances with the ELB:

```
$ aws elb register-instances-with-load-balancer --load-balancer-name "$lb" \
    --instances $instances
{
  "Instances": [
    {
      "InstanceId": "i-079b4096c654f563c"
    }
  ]
}
```

4. Configure the ELB's health check:

```
$ aws elb configure-health-check --load-balancer-name "$lb" \
  --health-check
  'Target=HTTP:1936/healthz,Interval=30,UnhealthyThreshold=2,HealthyTh
  reshold=2,Timeout=5'
{
  "HealthCheck": {
    "HealthyThreshold": 2,
    "Interval": 30,
    "Target": "HTTP:1936/healthz",
    "Timeout": 5,
    "UnhealthyThreshold": 2
  }
}
```

- Finally, create a load-balancer policy with the **ProxyProtocol** attribute enabled, and configure it on the ELB's TCP ports 80 and 443:

```
$ aws elb create-load-balancer-policy --load-balancer-name "$lb" \
  --policy-name "${lb}-ProxyProtocol-policy" \
  --policy-type-name 'ProxyProtocolPolicyType' \
  --policy-attributes
  'AttributeName=ProxyProtocol,AttributeValue=true'
$ for port in 80 443
do
  aws elb set-load-balancer-policies-for-backend-server \
    --load-balancer-name "$lb" \
    --instance-port "$port" \
    --policy-names "${lb}-ProxyProtocol-policy"
done
```

Verify the Configuration

You can examine the load balancer as follows to verify that the configuration is correct:

```
$ aws elb describe-load-balancers --load-balancer-name "$lb" |
jq '.LoadBalancerDescriptions| [0].ListenerDescriptions'
[
  [
    {
      "Listener": {
        "InstancePort": 80,
        "LoadBalancerPort": 80,
        "Protocol": "TCP",
        "InstanceProtocol": "TCP"
      },
      "PolicyNames": ["infra-lb-ProxyProtocol-policy"] ❶
    },
    {
      "Listener": {
        "InstancePort": 443,
        "LoadBalancerPort": 443,
        "Protocol": "TCP",
        "InstanceProtocol": "TCP"
      },
      "PolicyNames": ["infra-lb-ProxyProtocol-policy"] ❷
    }
  ]
]
```

```

    },
    {
      "Listener": {
        "InstancePort": 5000,
        "LoadBalancerPort": 5000,
        "Protocol": "TCP",
        "InstanceProtocol": "TCP"
      },
      "PolicyNames": [] ❸
    }
  ]
]

```

- ❶ The listener for TCP port 80 should have the policy for using the PROXY protocol.
- ❷ The listener for TCP port 443 should have the same policy.
- ❸ The listener for TCP port 5000 should **not** have the policy.

Alternatively, if you already have an ELB configured, but it is not configured to use the PROXY protocol, you will need to change the existing listener for TCP port 80 to use the TCP protocol instead of HTTP (TCP port 443 should already be using the TCP protocol):

```

$ aws elb delete-load-balancer-listeners --load-balancer-name "$lb" \
  --load-balancer-ports 80
$ aws elb create-load-balancer-listeners --load-balancer-name "$lb" \
  --listeners
  'Protocol=TCP,LoadBalancerPort=80,InstanceProtocol=TCP,InstancePort=80'

```

Verify the Protocol Updates

Verify that the protocol has been updated as follows:

```

$ aws elb describe-load-balancers --load-balancer-name "$lb" |
  jq '[.LoadBalancerDescriptions[]|.ListenerDescriptions]'
[
  [
    {
      "Listener": {
        "InstancePort": 443,
        "LoadBalancerPort": 443,
        "Protocol": "TCP",
        "InstanceProtocol": "TCP"
      },
      "PolicyNames": []
    },
    {
      "Listener": {
        "InstancePort": 5000,
        "LoadBalancerPort": 5000,
        "Protocol": "TCP",
        "InstanceProtocol": "TCP"
      },
      "PolicyNames": []
    }
  ],
  [
    {
      "Listener": {
        "InstancePort": 5000,
        "LoadBalancerPort": 5000,
        "Protocol": "TCP",
        "InstanceProtocol": "TCP"
      },
      "PolicyNames": []
    }
  ]
]

```

```
{
  "Listener": {
    "InstancePort": 80,
    "LoadBalancerPort": 80,
    "Protocol": "TCP", 1
    "InstanceProtocol": "TCP"
  },
  "PolicyNames": []
}
```

- 1 All listeners, including the listener for TCP port 80, should be using the TCP protocol.

Then, create a load-balancer policy and add it to the ELB as described in Step 5 above.

4.5. USING THE F5 ROUTER PLUG-IN

4.5.1. Overview



NOTE

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

The F5 router plug-in is provided as a container image and run as a pod, just like the [default HAProxy router](#). Deploying the F5 router is done similarly as well, using the **oadm router** command but providing additional flags (or environment variables) to specify the following parameters for the **F5 BIG-IP®** host:

Flag	Description
--type=f5-router	Specifies that an F5 router should be launched (the default --type is haproxy-router).
--external-host	Specifies the F5 BIG-IP® host's management interface's host name or IP address.
--external-host-username	Specifies the F5 BIG-IP® user name (typically admin).
--external-host-password	Specifies the F5 BIG-IP® password.
--external-host-http-vserver	Specifies the name of the F5 virtual server for HTTP connections.

Flag	Description
--external-host-https-vserver	Specifies the name of the F5 virtual server for HTTPS connections.
--external-host-private-key	Specifies the path to the SSH private key file for the F5 BIG-IP® host. Required to upload and delete key and certificate files for routes.
--external-host-insecure	A Boolean flag that indicates that the F5 router should skip strict certificate verification with the F5 BIG-IP® host.

As with the HAProxy router, the **oadm router** command creates the service and deployment configuration objects, and thus the replication controllers and pod(s) in which the F5 router itself runs. The replication controller restarts the F5 router in case of crashes. Because the F5 router is only watching routes and endpoints and configuring **F5 BIG-IP®** accordingly, running the F5 router in this way along with an appropriately configured **F5 BIG-IP®** deployment should satisfy high-availability requirements.

4.5.2. Deploying the F5 Router

The F5 router must be run in privileged mode because route certificates get copied using **scp**:

```
$ oadm policy remove-scc-from-user hostnetwork -z router
$ oadm policy add-scc-to-user privileged -z router
```

To deploy the F5 router:

1. First, [establish a tunnel using a ramp node](#), which allows for the routing of traffic to pods through the [OpenShift Container Platform SDN](#).
2. Run the **oadm router** command with the [appropriate flags](#). For example:

```
$ oadm router \
  --type=f5-router \
  --external-host=10.0.0.2 \
  --external-host-username=admin \
  --external-host-password=mypassword \
  --external-host-http-vserver=ose-vserver \
  --external-host-https-vserver=https-ose-vserver \
  --external-host-private-key=/path/to/key \
  --service-account=router
```

CHAPTER 5. UPGRADING A CLUSTER

5.1. OVERVIEW

When new versions of OpenShift Container Platform are released, you can upgrade your existing cluster to apply the latest enhancements and bug fixes. This includes upgrading from previous minor versions, such as release 3.2 to 3.3, and applying asynchronous errata updates within a minor version (3.3.z releases). See the [OpenShift Container Platform 3.3 Release Notes](#) to review the latest changes.



NOTE

Due to the [core architectural changes](#) between the major versions, OpenShift Enterprise 2 environments cannot be upgraded to OpenShift Container Platform 3 and require a fresh installation.

Unless noted otherwise, node and masters within a major version are forward and backward compatible [across one minor version](#), so upgrading your cluster should go smoothly. However, you should not run mismatched versions longer than necessary to upgrade the entire cluster.

5.1.1. In-place or Blue-Green Upgrades

There are two methods for performing OpenShift Container Platform cluster upgrades. You can either do in-place upgrades (automated or manual), or upgrade using a blue-green deployment method.

In-place Upgrades

With in-place upgrades, the cluster upgrade is performed on all hosts in a single cluster. If you installed using the [quick](#) or [advanced installation](#) and the `~/.config/openshift/installer.cfg.yml` or inventory file that was used is available, you can perform an [automated in-place upgrade](#). Alternatively, you can [upgrade in-place manually](#).

Blue-green Deployments

With [blue-green deployments](#), you can reduce downtime caused while upgrading an environment by creating a parallel environment on which the new deployment can be installed. If a problem is detected, and after the new deployment is verified, traffic can be switched over with the option to rollback.

5.2. PERFORMING AUTOMATED IN-PLACE CLUSTER UPGRADES


5.2.1. Overview

If you installed using the [advanced installation](#) and the inventory file that was used is available, you can use the upgrade playbook to automate the OpenShift cluster upgrade process. If you installed using the [quick installation](#) method and a `~/.config/openshift/installer.cfg.yml` file is available, you can use the installer to perform the automated upgrade.

The automated upgrade performs the following steps for you:

- Applies the latest configuration.
- Upgrades and restart master services.
- Upgrades and restart node services.


- Applies the latest cluster policies.
- Updates the default router if one exists.
- Updates the default registry if one exists.
- Updates default image streams and InstantApp templates.



IMPORTANT

Ensure that you have met all [prerequisites](#) before proceeding with an upgrade. Failure to do so can result in a failed upgrade.

5.2.2. Preparing for an Automated Upgrade



IMPORTANT

Before upgrading your cluster to OpenShift Container Platform 3.3, the cluster must be already upgraded to the [latest asynchronous release of version 3.2](#). Cluster upgrades cannot span more than one minor version at a time, so if your cluster is at version 3.0 or 3.1, you must first upgrade incrementally (e.g., 3.0 to 3.1, then 3.1 or 3.2).

To prepare for an automated upgrade:

1. If you are upgrading from version 3.2 to 3.3, manually disable the 3.2 channel and enable the 3.3 channel on each master and node host:

```
# subscription-manager repos --disable="rhel-7-server-ose-3.2-rpms" \
  --enable="rhel-7-server-ose-3.3-rpms" \
  --enable="rhel-7-server-rpms" \
  --enable="rhel-7-server-extras-rpms"
# yum clean all
```

2. For any upgrade path, always ensure that you have the latest version of the **atomic-openshift-utils** package, which should also update the **openshift-ansible-*** packages:

```
# yum update atomic-openshift-utils
```

3. Install or update to the following latest available ***-excluder** packages on each RHEL 7 system, which helps ensure your systems stay on the correct versions of **atomic-openshift** and **docker** packages when you are not trying to upgrade, according to the OpenShift Container Platform version:

```
# yum install atomic-openshift-excluder atomic-openshift-docker-excluder
```

These packages add entries to the **exclude** directive in the host's **/etc/yum.conf** file.

4. You must be logged in as a cluster administrative user on the master host for the upgrade to succeed:

```
$ oc login
```


After satisfying these steps, there are two methods for running the automated upgrade:

- [Using the installer](#)
- [Running the upgrade playbook directly](#)

Choose and follow one of these methods.

5.2.3. Using the Installer to Upgrade

If you installed OpenShift Container Platform using the [quick installation](#) method, you should have an installation configuration file located at `~/.config/openshift/installer.cfg.yml`. The installer requires this file to start an upgrade.

The installer supports upgrading between minor versions of OpenShift Container Platform (one minor version at a time, e.g., 3.2 to 3.3) as well as between [asynchronous errata updates](#) within a minor version (e.g., 3.3.z).

If you have an older format installation configuration file in `~/.config/openshift/installer.cfg.yml` from an installation of a previous cluster version, the installer will attempt to upgrade the file to the new supported format. If you do not have an installation configuration file of any format, you can [create one manually](#).

To start an upgrade with the quick installer:

1. Satisfy the steps in [Preparing for an Automated Upgrade](#) to ensure you are using the latest upgrade playbooks.
2. Run the installer with the **upgrade** subcommand:

```
# atomic-openshift-installer upgrade
```

3. Then, follow the on-screen instructions to upgrade to the latest release.
4. After all master and node upgrades have completed, a recommendation will be printed to reboot all hosts.
5. After rebooting, if there are no additional features enabled, you can [verify the upgrade](#). Otherwise, the next step depends on what additional features have you previously enabled.

Feature	Next Step
Aggregated Logging	Upgrade the EFK logging stack.
Cluster Metrics	Upgrade cluster metrics.

5.2.4. Running the Upgrade Playbook Directly

You can run the automated upgrade playbook using Ansible directly, similar to the advanced installation method, if you have an inventory file.

The same **v3_3** upgrade playbook can be used to upgrade either of the following to the latest 3.3 release:

- [Existing OpenShift Container Platform 3.2 clusters](#)

- [Existing OpenShift Container Platform 3.3 clusters](#)

5.2.4.1. Upgrading to OpenShift Container Platform 3.3

To run an upgrade from OpenShift Container Platform 3.2 to 3.3:

1. Satisfy the steps in [Preparing for an Automated Upgrade](#) to ensure you are using the latest upgrade playbooks.
2. Ensure the **deployment_type** parameter in your inventory file is set to **openshift-enterprise**.
3. If you have multiple masters configured and want to enable rolling, full system restarts of the hosts, you can set the **openshift_rolling_restart_mode** parameter in your inventory file to **system**. Otherwise, the default value **services** performs rolling service restarts on HA masters, but does not reboot the systems. See [Configuring Cluster Variables](#) for details.
4. Run the **v3_3** upgrade playbook. If your inventory file is located somewhere other than the default **/etc/ansible/hosts**, add the **-i** flag to specify the location. If you previously used the **atomic-openshift-installer** command to run your installation, you can check **~/.config/openshift/hosts** (previously located at **~/.config/openshift/.ansible/hosts**) for the last inventory file that was used, if needed.

```
# ansible-playbook [-i </path/to/inventory/file>] \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
    cluster/upgrades/v3_3/upgrade.yml
```

5. After all master and node upgrades have completed, a recommendation will be printed to reboot all hosts.
6. After rebooting, if there are no additional features enabled, you can [verify the upgrade](#). Otherwise, the next step depends on what additional features have you previously enabled.

Feature	Next Step
Aggregated Logging	Upgrade the EFK logging stack.
Cluster Metrics	Upgrade cluster metrics.

5.2.4.2. Upgrading to OpenShift Container Platform 3.3 Asynchronous Releases

To apply [asynchronous errata updates](#) to an existing OpenShift Container Platform 3.3 cluster:

1. Satisfy the steps in [Preparing for an Automated Upgrade](#) to ensure you are using the latest upgrade playbooks.
2. Run the **v3_3** upgrade playbook (the same playbook that is used for [upgrading from OpenShift Container Platform 3.2 to 3.3](#)). If your inventory file is located somewhere other than the default **/etc/ansible/hosts**, add the **-i** flag to specify the location. If you previously used the **atomic-openshift-installer** command to run your installation, you can check **~/.config/openshift/hosts** (previously located at **~/.config/openshift/.ansible/hosts**) for the last inventory file that was used, if needed.

```
# ansible-playbook [-i </path/to/inventory/file>] \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
    cluster/upgrades/v3_3/upgrade.yml
```

- After all master and node upgrades have completed, a recommendation will be printed to reboot all hosts.
- After rebooting, if there are no additional features enabled, you can [verify the upgrade](#). Otherwise, the next step depends on what additional features have you previously enabled.

Feature	Next Step
Aggregated Logging	Upgrade the EFK logging stack.
Cluster Metrics	Upgrade cluster metrics.

5.2.5. Upgrading the EFK Logging Stack

If you have previously [deployed the EFK logging stack](#) and want to upgrade to the latest logging component images, the steps must be performed manually as shown in [Manual Upgrades](#).

5.2.6. Upgrading Cluster Metrics

If you have previously [deployed cluster metrics](#), you must manually [update](#) to the latest metric components.

5.2.7. Verifying the Upgrade

To verify the upgrade:

- First check that all nodes are marked as **Ready**:

```
# oc get nodes
NAME                                STATUS              AGE
master.example.com                 Ready,SchedulingDisabled 165d
node1.example.com                  Ready                165d
node2.example.com                  Ready                165d
```

- Then, verify that you are running the expected versions of the **docker-registry** and **router** images, if deployed. Replace **<tag>** with **v3.3.1.25** for the latest version.

```
# oc get -n default dc/docker-registry -o json | grep \"image\"
  \"image\": \"openshift3/ose-docker-registry:<tag>\",
# oc get -n default dc/router -o json | grep \"image\"
  \"image\": \"openshift3/ose-haproxy-router:<tag>\",
```

- After upgrading, you can use the diagnostics tool on the master to look for common issues:

```
# oadm diagnostics
...
[Note] Summary of diagnostics execution:
```

[Note] Completed with no errors or warnings seen.

5.3. PERFORMING MANUAL IN-PLACE CLUSTER UPGRADES

5.3.1. Overview

As an alternative to performing an [automated upgrade](#), you can manually upgrade your OpenShift cluster. To manually upgrade without disruption, it is important to upgrade each component as documented in this topic.

Before you begin your upgrade, familiarize yourself now with the entire procedure. [Specific releases may require additional steps](#) to be performed at key points before or during the standard upgrade process.



IMPORTANT

Ensure that you have met all [prerequisites](#) before proceeding with an upgrade. Failure to do so can result in a failed upgrade.

5.3.2. Preparing for a Manual Upgrade



NOTE

Before upgrading your cluster to OpenShift Container Platform 3.3, the cluster must be already upgraded to the [latest asynchronous release of version 3.2](#). Cluster upgrades cannot span more than one minor version at a time, so if your cluster is at version 3.0 or 3.1, you must first upgrade incrementally (e.g., 3.0 to 3.1, then 3.1 or 3.2).

To prepare for a manual upgrade, follow these steps:

1. If you are upgrading from version 3.2 to 3.3, manually disable the 3.2 channel and enable the 3.3 channel on each host:

```
# subscription-manager repos --disable="rhel-7-server-ose-3.2-rpms" \
--enable="rhel-7-server-ose-3.3-rpms" \
--enable="rhel-7-server-extras-rpms"
```

On RHEL 7 systems, also clear the **yum** cache:

```
# yum clean all
```

2. Install or update to the latest available version of the **atomic-openshift-utils** package on each RHEL 7 system, which provides files that will be used in later sections:

```
# yum install atomic-openshift-utils
```

3. Install or update to the following latest available ***-excluder** packages on each RHEL 7 system, which helps ensure your systems stay on the correct versions of **atomic-openshift** and **docker** packages when you are not trying to upgrade, according to the OpenShift Container Platform version:

```
# yum install atomic-openshift-excluder atomic-openshift-docker-
excluder
```

These packages add entries to the **exclude** directive in the host's */etc/yum.conf* file.

4. Create an **etcd** backup on each master. The **etcd** package is required, even if using embedded etcd, for access to the **etcdctl** command to make the backup. The package is installed by default for RHEL Atomic Host 7 systems. If the master is a RHEL 7 system, ensure the package is installed:

```
# yum install etcd
```

Then, create the backup:

```
# ETCD_DATA_DIR=/var/lib/origin/openshift.local.etcd ❶
# etcdctl backup \
  --data-dir $ETCD_DATA_DIR \
  --backup-dir $ETCD_DATA_DIR.bak.<date> ❷
```

- ❶ This directory is for embedded etcd. If you use a separate etcd cluster, use */var/lib/etcd* instead.
- ❷ Use the date of the backup, or some unique identifier, for **<date>**. The command will not make a backup if the **--backup-dir** location already exists.

5. For any upgrade path, ensure that you are running the latest kernel on each RHEL 7 system:

```
# yum update kernel
```

5.3.3. Upgrading Master Components

Upgrade your master hosts first:

1. Run the following command on each master to remove the **atomic-openshift** packages from the list of yum excludes on the host:

```
# atomic-openshift-excluder unexclude
```

2. Upgrade the **atomic-openshift** packages or related images.

- a. For masters using the RPM-based method on a RHEL 7 system, upgrade all installed **atomic-openshift** packages:

```
# yum upgrade atomic-openshift\*
```

- b. For masters using the containerized method on a RHEL 7 or RHEL Atomic Host 7 system, set the **IMAGE_VERSION** parameter to the version you are upgrading to in the following files:

- */etc/sysconfig/atomic-openshift-master* (single master clusters only)
- */etc/sysconfig/atomic-openshift-master-controllers* (multi-master clusters only)
- */etc/sysconfig/atomic-openshift-master-api* (multi-master clusters only)

- **/etc/sysconfig/atomic-openshift-node**
- **/etc/sysconfig/atomic-openshift-openvswitch**

For example:

```
IMAGE_VERSION=<tag>
```

Replace **<tag>** with **v3.3** for the latest version.

3. In OpenShift Container Platform 3.3, protocol buffers are used by default for internal communications between node, masters, and controllers. To configure this, the following stanzas must be altered in the **/etc/origin/master-config.yaml** file on each master:

```
masterClients:
  externalKubernetesClientConnectionOverrides:
    acceptContentTypes:
      application/vnd.kubernetes.protobuf,application/json
    contentType: application/vnd.kubernetes.protobuf
    burst: 400
    qps: 200
  externalKubernetesKubeConfig: ""
  openshiftLoopbackClientConnectionOverrides:
    acceptContentTypes:
      application/vnd.kubernetes.protobuf,application/json
    contentType: application/vnd.kubernetes.protobuf
    burst: 600
    qps: 300
  openshiftLoopbackKubeConfig: openshift-master.kubeconfig
```

For more information on protocol buffers, see <https://developers.google.com/protocol-buffers/docs/overview>.

4. In OpenShift Container Platform 3.3, the **kubernetesMasterConfig.admissionConfig.pluginConfig** parameter in the **/etc/origin/master-config.yaml** file is being deprecated. If you are upgrading from version 3.2 to 3.3 and this parameter is in use, see [General Admission Rules](#) for guidance on moving and merging into **admissionConfig.pluginConfig**.
5. Restart the master service(s) on each master and review logs to ensure they restart successfully.
For single master clusters:

```
# systemctl restart atomic-openshift-master
# journalctl -r -u atomic-openshift-master
```

For multi-master clusters:

```
# systemctl restart atomic-openshift-master-controllers
# systemctl restart atomic-openshift-master-api
# journalctl -r -u atomic-openshift-master-controllers
# journalctl -r -u atomic-openshift-master-api
```

6. Because masters also have node components running on them in order to be configured as part of the OpenShift SDN, restart the **atomic-openshift-node** and **openvswitch** services:

```
# systemctl restart atomic-openshift-node
# systemctl restart openvswitch
# journalctl -r -u openvswitch
# journalctl -r -u atomic-openshift-node
```

7. Run the following command on each master to add the **atomic-openshift** packages back to the list of yum excludes on the host:

```
# atomic-openshift-excluder exclude
```

Upgrade any external etcd hosts using the RPM-based method on a RHEL 7 system:

1. Upgrade the **etcd** package:

```
# yum update etcd
```

2. Restart the **etcd** service and review the logs to ensure it restarts successfully:

```
# systemctl restart etcd
# journalctl -r -u etcd
```

NOTE

During the cluster upgrade, it can sometimes be useful to take a master out of rotation since some DNS client libraries will not properly to the other masters for cluster DNS. In addition to stopping the master and controller services, you can remove the EndPoint from the Kubernetes service's **subsets.addresses**.

```
$ oc edit ep/kubernetes -n default
```

When the master is restarted, the Kubernetes service will be automatically updated.

5.3.4. Updating Policy Definitions

After a cluster upgrade, the recommended [default cluster roles](#) may be updated. To check if an update is recommended for your environment, you can run:

```
# oadm policy reconcile-cluster-roles
```



WARNING

If you have customized default cluster roles and want to ensure a role reconciliation does not modify those customized roles, annotate them with **openshift.io/reconcile-protect** set to **true**. Doing so means you are responsible for manually updating those roles with any new or required permissions during upgrades.

This command outputs a list of roles that are out of date and their new proposed values. For example:

```
# oadm policy reconcile-cluster-roles
apiVersion: v1
items:
- apiVersion: v1
  kind: ClusterRole
  metadata:
    creationTimestamp: null
    name: admin
  rules:
  - attributeRestrictions: null
    resources:
    - builds/custom
  ...
```



NOTE

Your output will vary based on the OpenShift version and any local customizations you have made. Review the proposed policy carefully.

You can either modify this output to re-apply any local policy changes you have made, or you can automatically apply the new policy using the following process:

1. Reconcile the cluster roles:

```
# oadm policy reconcile-cluster-roles \
  --additive-only=true \
  --confirm
```

2. Reconcile the cluster role bindings:

```
# oadm policy reconcile-cluster-role-bindings \
  --exclude-groups=system:authenticated \
  --exclude-groups=system:authenticated:oauth \
  --exclude-groups=system:unauthenticated \
  --exclude-users=system:anonymous \
  --additive-only=true \
  --confirm
```

3. Reconcile security context constraints:

```
# oadm policy reconcile-sccs \
  --additive-only=true \
  --confirm
```

5.3.5. Upgrading Nodes

After upgrading your masters, you can upgrade your nodes. When restarting the **atomic-openshift-node** service, there will be a brief disruption of outbound network connectivity from running pods to services while the [service proxy](#) is restarted. The length of this disruption should be very short and scales based on the number of services in the entire cluster.

**NOTE**

[Blue-green deployments](#) are another proven approach to reducing downtime caused while updating an environment.

One at a time for each node that is not also a master, you must disable scheduling and evacuate its pods to other nodes, then upgrade packages and restart services.


1. Run the following command on each node to remove the **atomic-openshift** packages from the list of yum excludes on the host:

```
# atomic-openshift-excluder unexclude
```

2. As a user with **cluster-admin** privileges, disable scheduling for the node:

```
# oadm manage-node <node> --schedulable=false
```

3. Evacuate pods on the node to other nodes:

**IMPORTANT**

The **--force** option deletes any pods that are not backed by a replication controller.

```
# oadm manage-node <node> --evacuate --force
```

4. Upgrade the node component packages or related images.
 - a. For nodes using the RPM-based method on a RHEL 7 system, upgrade all installed **atomic-openshift** packages:

```
# yum upgrade atomic-openshift\*
```

- b. For nodes using the containerized method on a RHEL 7 or RHEL Atomic Host 7 system, set the **IMAGE_VERSION** parameter in the **/etc/sysconfig/atomic-openshift-node** and **/etc/sysconfig/opensvswitch** files to the version you are upgrading to. For example:

```
IMAGE_VERSION=<tag>
```

Replace **<tag>** with **v3.3** for the latest version.

5. In OpenShift Container Platform 3.3, protocol buffers are used by default for internal communications between node, masters, and controllers. To configure this, the following stanzas must be altered in the **/etc/origin/node-config.yaml** file on each node:

```
masterClientConnectionOverrides:
  acceptContentTypes:
  application/vnd.kubernetes.protobuf,application/json
  contentType: application/vnd.kubernetes.protobuf
  burst: 200
  qps: 100
```

For more information on protocol buffers, see <https://developers.google.com/protocol-buffers/docs/overview>.

- Restart the **atomic-openshift-node** and **openvswitch** services and review the logs to ensure they restart successfully:

```
# systemctl restart atomic-openshift-node
# systemctl restart openvswitch
# journalctl -r -u atomic-openshift-node
# journalctl -r -u openvswitch
```

- Re-enable scheduling for the node:

```
# oadm manage-node <node> --schedulable
```

- Run the following command on the node to add the **atomic-openshift** packages back to the list of yum excludes on the host:

```
# atomic-openshift-excluder exclude
```

- Repeat these steps on the next node, and continue repeating these steps until all nodes have been upgraded.

- After all nodes have been upgraded, as a user with **cluster-admin** privileges, verify that all nodes are showing as **Ready**:

```
# oc get nodes
NAME                                STATUS                                AGE
master.example.com                 Ready,SchedulingDisabled            165d
node1.example.com                  Ready                                165d
node2.example.com                  Ready                                165d
```

5.3.6. Upgrading the Router

If you have previously [deployed a router](#), the router deployment configuration must be upgraded to apply updates contained in the router image. To upgrade your router without disrupting services, you must have previously deployed a [highly-available routing service](#).



NOTE

If you previously customized your HAProxy routing template, then, depending on the changes, additional steps may be required due to changes in the routing data structure starting in OpenShift Container Platform 3.3. See [Routing Data Structure Changes](#) in the OpenShift Container Platform 3.3 Release Notes for details.

Edit your router's deployment configuration. For example, if it has the default **router** name:

```
# oc edit dc/router
```

Apply the following changes:

```
...
spec:
```

```

template:
  spec:
    containers:
      - env:
          ...
          image: registry.access.redhat.com/openshift3/ose-haproxy-router:
<tag> 1
          imagePullPolicy: IfNotPresent
          ...

```

- 1** Adjust **<tag>** to match the version you are upgrading to (use **v3.3** for the latest version).

You should see one router pod updated and then the next.

5.3.7. Upgrading the Registry

The registry must also be upgraded for changes to take effect in the registry image. If you have used a **PersistentVolumeClaim** or a host mount point, you may restart the registry without losing the contents of your registry. [Storage for the Registry](#) details how to configure persistent storage for the registry.

Edit your registry's deployment configuration:

```
# oc edit dc/docker-registry
```

Apply the following changes:

```

...
spec:
  template:
    spec:
      containers:
        - env:
            ...
            image: registry.access.redhat.com/openshift3/ose-docker-registry:
<tag> 1
            imagePullPolicy: IfNotPresent
            ...

```


- 1** Adjust **<tag>** to match the version you are upgrading to (use **v3.3** for the latest version).



IMPORTANT

Images that are being pushed or pulled from the internal registry at the time of upgrade will fail and should be restarted automatically. This will not disrupt pods that are already running.

5.3.7.1. Updating Custom Registry Configuration Files

**NOTE**

You may safely skip this part if you do not use a custom registry configuration file.

The internal Docker registry version 3.3.0 and higher requires following entries in the [middleware section](#) of the configuration file:

```
middleware:
  registry:
    - name: openshift
  repository:
    - name: openshift
  storage:
    - name: openshift
```

1. Edit your custom configuration file, adding the missing entries.
2. [Deploy your updated configuration.](#)
3. Append the **--overwrite** flag to **oc volume dc/docker-registry --add** to replace a volume mount of your previous secret.
4. You can safely remove the old secret.

5.3.7.2. Enforcing Quota in the Registry

Quota must be enforced to prevent layer blobs that exceed the size limit from being written to the registry's storage. This can be achieved via a [configuration file](#):

```
...
middleware:
  repository:
    - name: openshift
      options:
        enforcequota: true
...
```

Alternatively, use the **REGISTRY_MIDDLEWARE_REPOSITORY_OPENSIFT_ENFORCEQUOTA** environment variable, which is set to **true** for the new registry deployments by default. Existing deployments need to be modified using:

```
# oc set env dc/docker-registry
REGISTRY_MIDDLEWARE_REPOSITORY_OPENSIFT_ENFORCEQUOTA=true
```

5.3.8. Updating the Default Image Streams and Templates

By default, the [quick](#) and [advanced installation](#) methods automatically create default image streams, InstantApp templates, and database service templates in the **openshift** project, which is a default project to which all users have view access. These objects were created during installation from the JSON files located under the **/usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/** directory.



NOTE

Because RHEL Atomic Host 7 cannot use **yum** to update packages, the following steps must take place on a RHEL 7 system.

1. Update the packages that provide the example JSON files. On a subscribed Red Hat Enterprise Linux 7 system where you can run the CLI as a user with **cluster-admin** permissions, install or update to the latest version of the **atomic-openshift-utils** package, which should also update the **openshift-ansible-** packages:

```
# yum update atomic-openshift-utils
```

The **openshift-ansible-roles** package provides the latest example JSON files.

2. After a manual upgrade, get the latest templates from **openshift-ansible-roles**:

```
rpm -ql openshift-ansible-roles | grep examples | grep v1.3
```

In this example, **/usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/image-streams/image-streams-rhel7.json** is the latest file that you want in the latest **openshift-ansible-roles** package.

/usr/share/openshift/examples/image-streams/image-streams-rhel7.json is not owned by a package, but is updated by Ansible. If you are upgrading outside of Ansible, you need to get the latest .json files on the system where you are running **oc**, which can run anywhere that has access to the master.

3. Install **atomic-openshift-utils** and its dependencies to install the new content into **/usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/**:

```
$ oc create -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/image-streams/image-streams-rhel7.json
$ oc create -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/image-streams/dotnet_imagestreams.json
$ oc replace -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/image-streams/image-streams-rhel7.json
$ oc replace -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/image-streams/dotnet_imagestreams.json
```

4. Update the templates:

```
$ oc create -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/quickstart-templates/
$ oc create -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/db-templates/
$ oc create -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/infrastructure-templates/
$ oc create -n openshift -f /usr/share/ansible/openshift-ansible/roles/openshift_examples/files/examples/v1.3/xpaas-
```

```

templates/
$ oc create -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/xpaas-streams/
$ oc replace -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/
$ oc replace -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/db-templates/
$ oc replace -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/infrastructure-
templates/
$ oc replace -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/xpaas-
templates/
$ oc replace -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/xpaas-streams/

```

Errors are generated for items that already exist. This is expected behavior:

```

# oc create -n openshift -f /usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/
Error from server: error when creating
"/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/cakephp-mysql.json": templates "cakephp-mysql-example"
already exists
Error from server: error when creating
"/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/cakephp.json": templates "cakephp-example" already exists
Error from server: error when creating
"/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/dancer-mysql.json": templates "dancer-mysql-example"
already exists
Error from server: error when creating
"/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/dancer.json": templates "dancer-example" already exists
Error from server: error when creating
"/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates/django-postgresql.json": templates "django-psql-example"
already exists

```

Now, content can be updated. Without running the automated upgrade playbooks, the content is not updated in **/usr/share/openshift/**.

5.3.9. Importing the Latest Images

After [updating the default image streams](#), you may also want to ensure that the images within those streams are updated. For each image stream in the default **openshift** project, you can run:

```
# oc import-image -n openshift <imagestream>
```

For example, get the list of all image streams in the default **openshift** project:


```
# oc get is -n openshift
NAME          DOCKER REPO
TAGS          UPDATED
mongodb       registry.access.redhat.com/openshift3/mongodb-24-rhel7
2.4,latest,v3.1.1.6    16 hours ago
mysql         registry.access.redhat.com/openshift3/mysql-55-rhel7
5.5,latest,v3.1.1.6    16 hours ago
nodejs        registry.access.redhat.com/openshift3/nodejs-010-rhel7
0.10,latest,v3.1.1.6    16 hours ago
...
```

Update each image stream one at a time:

```
# oc import-image -n openshift nodejs
The import completed successfully.

Name:      nodejs
Created:   10 seconds ago
Labels:    <none>
Annotations: openshift.io/image.dockerRepositoryCheck=2016-07-
05T19:20:30Z
Docker Pull Spec: 172.30.204.22:5000/openshift/nodejs

Tag Spec          Created    PullSpec          Image
latest 4          9 seconds ago registry.access.redhat.com/rhsc1/nodejs-4-
rhel7:latest
570ad8ed927fd5c2c9554ef4d9534cef808dfa05df31ec491c0969c3bd372b05
4 registry.access.redhat.com/rhsc1/nodejs-4-rhel7:latest 9 seconds ago
<same>
570ad8ed927fd5c2c9554ef4d9534cef808dfa05df31ec491c0969c3bd372b05
0.10 registry.access.redhat.com/openshift3/nodejs-010-rhel7:latest 9
seconds ago <same>
a1ef33be788a28ec2bdd48a9a5d174ebcfbe11c8e986d2996b77f5bccaaa4774
```



IMPORTANT

In order to update your S2I-based applications, you must manually trigger a new build of those applications after importing the new images using **oc start-build <app-name>**.

5.3.10. Upgrading the EFK Logging Stack

Use the following to upgrade an [already-deployed EFK logging stack](#).



NOTE

The following steps apply when upgrading to OpenShift Container Platform 3.3+.

1. Ensure you are working in the project where the EFK stack was previously deployed. For example, if the project is named **logging**:

```
$ oc project logging
```

2. Recreate the deployer templates for service accounts and running the deployer:

```
$ oc apply -n openshift -f \
  usr/share/ansible/openshift-
  ansible/roles/openshift_hosted_templates/files/v1.3/enterprise/metri
  cs-deployer.yaml
```

3. Generate any missing service accounts and roles:

```
$ oc process logging-deployer-account-template | oc apply -f -
```

4. Ensure that the cluster role **oauth-editor** is assigned to the **logging-deployer** service account:

```
$ oadm policy add-cluster-role-to-user oauth-editor \
  system:serviceaccount:logging:logging-deployer
```

5. In preparation for running the deployer, ensure that you have the configurations for your current deployment in the [logging-deployer ConfigMap](#).




IMPORTANT

Ensure that your image version is the latest version, not the currently installed version.

6. Run the deployer with the parameter in **upgrade** mode:

```
$ oc new-app logging-deployer-template -p MODE=upgrade
```

Running the deployer in this mode handles scaling down the components to minimize loss of logs, patching configurations, generating missing secrets and keys, and scaling the components back up to their previous replica count.




IMPORTANT

Due to the privileges needed to label and unlabel a node for controlling the deployment of Fluentd pods, the deployer does delete the **logging-fluentd** Daemonset and recreates it from the **logging-fluentd-template** template.

5.3.11. Upgrading Cluster Metrics

After upgrading an [already-deployed Cluster Metrics install](#), you must update to a newer version of the metrics components.

- The update process stops all the metrics containers, updates the metrics configuration files, and redeploys the newer components.
- It does not change the metrics route.
- It does not delete the metrics persistent volume claim. Metrics stored to persistent volumes before the update are available after the update completes.

**IMPORTANT**

The update deletes all non-persisted metric values and overwrites local changes to the metrics configurations. For example, the number of instances in a replica set is not saved.

To update, follow the same steps as when the metrics components were [first deployed](#), using the [correct template](#), except this time, specify the **MODE=refresh** option:

```
$ oc new-app -f metrics-deployer.yaml \
  -p HAWKULAR_METRICS_HOSTNAME=hm.example.com,MODE=refresh 1
```

1 In the original deployment command, there was no **MODE=refresh**.

**NOTE**

During the update, the metrics components do not run. Because of this, they cannot collect data and a gap normally appears in the graphs.

5.3.12. Additional Manual Steps Per Release

Some OpenShift Container Platform releases may have additional instructions specific to that release that must be performed to fully apply the updates across the cluster. This section will be updated over time as new asynchronous updates are released for OpenShift Container Platform 3.3.

As of the latest 3.3 release (v3.3), there are no 3.3 asynchronous releases that require additional instructions during upgrade. See the [OpenShift Container Platform 3.3 Release Notes](#) to review the latest release notes.

5.3.13. Verifying the Upgrade

To verify the upgrade, first check that all nodes are marked as **Ready**:

```
# oc get nodes
NAME                                STATUS                                AGE
master.example.com                  Ready,SchedulingDisabled             165d
node1.example.com                   Ready                                165d
node2.example.com                   Ready                                165d
```

Then, verify that you are running the expected versions of the **docker-registry** and **router** images, if deployed. Replace **<tag>** with **v3.3** for the latest version.


```
# oc get -n default dc/docker-registry -o json | grep "\"image\""
"image": "openshift3/ose-docker-registry:<tag>",
# oc get -n default dc/router -o json | grep "\"image\""
"image": "openshift3/ose-haproxy-router:<tag>",
```

After upgrading, you can use the diagnostics tool on the master to look for common issues:

```
# oadm diagnostics
...
[Note] Summary of diagnostics execution:
[Note] Completed with no errors or warnings seen.
```

5.4. BLUE-GREEN DEPLOYMENTS

5.4.1. Overview



NOTE

This topic serves as an alternative node upgrade method to the approach in [Manual In-place Upgrades](#).

Blue-green deployments are a proven approach to reducing downtime caused while upgrading an environment. This is done by creating a parallel environment on which the new deployment can be installed. If a problem is detected, and after the new deployment is verified, traffic can be switched over with the option to rollback.

While blue-green is a valid strategy for deploying just about any software, there are always trade-offs. Not all environments have the same uptime requirements or the resources to properly perform blue-green deployments. In an OpenShift Container Platform environment, the most suitable candidate for blue-green deployments are the nodes. All user processes run on these systems and even critical pieces of OpenShift Container Platform infrastructure are self-hosted there. Uptime is most important for these workloads and the additional complexity of blue-green deployments can be justified. The exact implementation of this approach varies based on your requirements. Often the main challenge is having the excess capacity to facilitate such an approach.

Another lesser challenge is that the administrator must temporarily share the Red Hat software entitlements between the blue-green deployments or provide access to the installation content by means of a system such as Red Hat Satellite. This can be accomplished by sharing the consumer ID from the previous host.


5.4.2. Preparing for Upgrade

1. On the old host:

```
# subscription-manager identity | grep system
system identity: 6699375b-06db-48c4-941e-689efd6ce3aa
```

2. On the new host:

```
# subscription-manager register --consumerid=6699375b-06db-48c4-941e-689efd6ce3aa
```



IMPORTANT

After a successful deployment, remember to unregister the old host with **subscription-manager clean** to prevent the environment from being out of compliance.

3. [After the master and etcd servers have been upgraded](#), you must ensure that your current production nodes are labeled either blue or green. In this example, the current installation will be blue and the new environment will be green. On each production node in your current installation:

```
$ oc label --all nodes color=blue
```

In the case of nodes requiring the uptime guarantees of a blue-green deployment, the **-l** flag can be used to match a subset of the environment using a selector.

4. Create the new green environment for any nodes that are to be replaced by [adding an equal number of new nodes](#) to the existing cluster. Ansible can apply the **color=green** label using the **openshift_node_labels** variable for each node.
5. In order to delay workload scheduling until the nodes are [healthy](#), be sure to set the **openshift_schedulable=false** variable. After the green nodes are in **Ready** state, they can be made schedulable.
6. Blue nodes are disabled so that no new pods are run on them:

```
# oadm manage-node --schedulable=true --selector=color=green
# oadm manage-node --schedulable=false --selector=color=blue
```

A common practice is to scale the registry and router pods until they are migrated to the green nodes. For these pods, a *canary* deployment approach is commonly used. Scaling them up will make them immediately active on the new nodes. Pointing the deployment configuration to the new image initiates a rolling update. However, because of node anti-affinity, and the fact that the blue nodes are still unschedulable, the deployments to the old nodes will fail. At this point, the registry and router deployments can be scaled down to the original number of pods. At any given point, the original number of pods is still available so no capacity is lost.

5.4.3. Warming the New Nodes

In order for pods to be migrated from the blue environment to the green, the images must be pulled. Network latency and load on the registry can cause delays if there is not sufficient capacity built in to the environment. Often, the best way to minimize impact to the running system is to trigger new pod deployments that will land on the new nodes. Accomplish this by importing new image streams.

A major release of OpenShift Container Platform is the motivation for a blue-green deployment. At that time, new image streams become available for users of Source-to-Image (S2I). Upon import, any builds or deployments configured with **ImageChangeTriggers** are automatically created.



NOTE

To continue with the upgrade process, [update the default image streams and templates](#) and [import the latest images](#).

It is important to realize that this process can trigger a large number of builds. The good news is that the builds are performed on the green nodes and, therefore, do not impact any traffic on the blue deployment.

To monitor build progress across all namespaces (projects) in the cluster:

```
$ oc get events -w --all-namespaces
```

In large environments, builds rarely completely stop. However, you should see a large increase and decrease caused by the administrative import.

Another benefit of triggering the builds is that it does a fairly good job of fetching the majority of the ancillary images to all nodes such as the various build images, the pod infrastructure image, and deployers. Everything else can be moved over using node evacuation and will proceed more quickly as a

result.

5.4.4. Node Evacuation

For larger deployments, it is possible to have other labels that help determine how evacuation can be coordinated. The most conservative approach for avoiding downtime is to evacuate one node at a time. If services are composed of pods using zone anti-affinity, then an entire zone can be evacuated at once. It is important to ensure that the storage volumes used are available in the new zone as this detail can vary among cloud providers.

In OpenShift Enterprise 3.2 and later, a node evacuation is triggered whenever the service is stopped. Achieve manual evacuation and deletion of all blue nodes at once by:

```
# oadm manage-node --selector=color=blue --evacuate
# oc delete node --selector=color=blue
```

5.5. OPERATING SYSTEM UPDATES AND UPGRADES

5.5.1. Updating and Upgrading the Operating System

Updating or upgrading your operating system (OS), by either changing OS versions or updating the system software, can impact the OpenShift Container Platform software running on those machines. In particular, these updates can affect the **iptables** rules or **ovs** flows that OpenShift Container Platform requires to operate.

Use the following to safely upgrade the OS on a host:

1. Ensure the host is unschedulable, meaning that no new pods will be placed onto the host:

```
$ oadm manage-node <node_name> --schedulable=false
```

2. Migrate the pods from the host:

```
$ oadm drain <node_name> --force --delete-local-data --ignore-daemonsets
```

3. Install or update the ***-excluder** packages on each host with the following. This ensures the hosts stay on the correct versions of OpenShift Container Platform, as per the **atomic-openshift** and **docker** packages, instead of the most current versions:

```
# yum install atomic-openshift-excluder atomic-openshift-docker-excluder
```

This adds entries to the **exclude** directive in the host's **/etc/yum.conf** file.

4. Update or upgrade the host packages, and reboot the host. A reboot ensures that the host is running the newest versions, and means that the **docker** and OpenShift Container Platform processes have been restarted, which will force them to check that all of the rules in other services are correct.

However, instead of rebooting a node host, you can restart the services that are affected, or preserve the **iptables** state. Both processes are described in the [OpenShift Container Platform IPtables](#) topic. The **ovs** flow rules do not need to be saved, but restarting the OpenShift Container Platform node software will fix the flow rules.

5. Configure the host to be schedulable again:

```
$ oadm manage-node <node_name> --schedulable=true
```

CHAPTER 6. DOWNGRADING OPENSIFT

6.1. OVERVIEW

Following an OpenShift Container Platform [upgrade](#), it may be desirable in extreme cases to downgrade your cluster to a previous version. The following sections outline the required steps for each system in a cluster to perform such a downgrade for the OpenShift Container Platform 3.3 to 3.2 downgrade path.



WARNING

These steps are currently only supported for [RPM-based installations](#) of OpenShift Container Platform and assumes downtime of the entire cluster.

6.2. VERIFYING BACKUPS

The Ansible playbook used during the [upgrade process](#) should have created a backup of the **master-config.yaml** file and the etcd data directory. Ensure these exist on your masters and etcd members:

```
/etc/origin/master/master-config.yaml.<timestamp>
/var/lib/origin/etcd-backup-<timestamp>
```

Also, back up the **node-config.yaml** file on each node (including masters, which have the node component on them) with a timestamp:

```
/etc/origin/node/node-config.yaml.<timestamp>
```

If you use a separate etcd cluster instead of a single embedded etcd instance, the backup is likely created on all etcd members, though only one is required for the recovery process. You can run a separate etcd instance that is co-located with your master nodes.

The RPM downgrade process in a later step should create **.rpmsave** backups of the following files, but it may be a good idea to keep a separate copy regardless:

```
/etc/sysconfig/atomic-openshift-master
/etc/etcd/etcd.conf 1
```

1 Only required if using a separate etcd cluster.

6.3. SHUTTING DOWN THE CLUSTER

On all masters, nodes, and etcd members, if you use a separate etcd cluster that runs on different nodes, ensure the relevant services are stopped.

On the master in a single master cluster:

```
# systemctl stop atomic-openshift-master
```

On each master in a multi-master cluster:

```
# systemctl stop atomic-openshift-master-api
# systemctl stop atomic-openshift-master-controllers
```

On all master and node hosts:

```
# systemctl stop atomic-openshift-node
```

On any etcd hosts for a separate etcd cluster:

```
# systemctl stop etcd
```

6.4. REMOVING RPMS

1. The ***-excluder** packages add entries to the exclude directive in the host's **/etc/yum.conf** file when installed. Run the following command on each host to remove the **atomic-openshift-*** and **docker** packages from the exclude list:

```
# atomic-openshift-excluder unexclude
# atomic-openshift-docker-excluder unexclude
```

2. On all masters, nodes, and etcd members, if you use a separate etcd cluster that runs on different nodes, remove the following packages:

```
# yum remove atomic-openshift \
  atomic-openshift-clients \
  atomic-openshift-node \
  atomic-openshift-master \
  openvswitch \
  atomic-openshift-sdn-ovs \
  tuned-profiles-atomic-openshift-node \
  atomic-openshift-excluder \
  atomic-openshift-docker-excluder
```

3. If you use a separate etcd cluster, also remove the **etcd** package:

```
# yum remove etcd
```

If using the embedded etcd, leave the **etcd** package installed. It is required for running the **etcdctl** command to issue operations in later steps.

6.5. DOWNGRADING DOCKER

OpenShift Container Platform 3.2 requires Docker 1.9.1 and also supports Docker 1.10.3.

Downgrade to Docker 1.9.1 on each host using the following steps:

1. Remove all local containers and images on the host. Any pods backed by a replication controller will be recreated.

**WARNING**

The following commands are destructive and should be used with caution.

Delete all containers:

```
# docker rm $(docker ps -a -q)
```

Delete all images:

```
# docker rmi $(docker images -q)
```

2. Use **yum swap** (instead of **yum downgrade**) to install Docker 1.9.1:

```
# yum swap docker-* docker-*1.9.1
# sed -i 's/--storage-opt dm.use_deferred_deletion=true//'
/etc/sysconfig/docker-storage
# systemctl restart docker
```

3. You should now have Docker 1.9.1 installed and running on the host. Verify with the following:

```
# docker version
Client:
 Version:      1.9.1-e17
 API version:  1.20
 Package Version: docker-1.9.1-10.el7.x86_64
 [...]

# systemctl status docker
• docker.service - Docker Application Container Engine
   Loaded: loaded (/usr/lib/systemd/system/docker.service; enabled;
 vendor preset: disabled)
   Active: active (running) since Wed 2017-06-21 15:44:20 EDT; 30min
 ago
   [...]

```

6.6. REINSTALLING RPMS

1. Disable the OpenShift Container Platform 3.3 repositories, and re-enable the 3.2 repositories:

```
# subscription-manager repos \
  --disable=rhel-7-server-ose-3.3-rpms \
  --enable=rhel-7-server-ose-3.2-rpms
```

2. On each master, install the following packages:

```
# yum install atomic-openshift \
  atomic-openshift-clients \
```



```
atomic-openshift-node \
atomic-openshift-master \
openvswitch \
atomic-openshift-sdn-ovs \
tuned-profiles-atomic-openshift-node \
atomic-openshift-excluder \
atomic-openshift-docker-excluder
```

3. On each node, install the following packages:

```
# yum install atomic-openshift \
atomic-openshift-node \
openvswitch \
atomic-openshift-sdn-ovs \
tuned-profiles-atomic-openshift-node \
atomic-openshift-excluder \
atomic-openshift-docker-excluder
```

4. If you use a separate etcd cluster, install the following package on each etcd member:

```
# yum install etcd
```

6.7. BRINGING OPENSIFT CONTAINER PLATFORM SERVICES BACK ONLINE

See [Backup and Restore](#).

6.8. VERIFYING THE DOWNGRADE

1. To verify the downgrade, first check that all nodes are marked as **Ready**:

```
# oc get nodes
NAME                                STATUS                                AGE
master.example.com                  Ready,SchedulingDisabled             165d
node1.example.com                   Ready                                165d
node2.example.com                   Ready                                165d
```

2. Then, verify that you are running the expected versions of the **docker-registry** and **router** images, if deployed:

```
# oc get -n default dc/docker-registry -o json | grep "\"image\""
"image": "openshift3/ose-docker-registry:v3.2.1.34-5",
# oc get -n default dc/router -o json | grep "\"image\""
"image": "openshift3/ose-haproxy-router:v3.2.1.34-5",
```

3. You can use the [diagnostics tool](#) on the master to look for common issues and provide suggestions:

```
# oc adm diagnostics
...
[Note] Summary of diagnostics execution:
[Note] Completed with no errors or warnings seen.
```

CHAPTER 7. MASTER AND NODE CONFIGURATION

7.1. OVERVIEW

The **openshift start** command is used to launch OpenShift Container Platform servers. The command and its subcommands (**master** to launch a [master server](#) and **node** to launch a [node server](#)) all take a limited set of arguments that are sufficient for launching servers in a development or experimental environment.

However, these arguments are insufficient to describe and control the full set of configuration and security options that are necessary in a production environment. To provide those options, it is necessary to use the dedicated master and node configuration files.

[Master configuration files](#) and [node configuration files](#) are fully specified with no default values. Therefore, any empty value indicates that you want to start up with an empty value for that parameter. This makes it easy to reason about exactly what your configuration is, but it also makes it difficult to remember all of the options to specify. To make this easier, the configuration files can be created with the **--write-config** option and then used with the **--config** option.

7.2. MASTER CONFIGURATION FILES

This section reviews parameters mentioned in the **master-config.yaml** file.

You can [create a new master configuration file](#) to see the valid options for your installed version of OpenShift Container Platform.

7.2.1. Admission Control Configuration

Table 7.1. Admission Control Configuration Parameters

Parameter Name	Description
AdmissionConfig	Contains admission control plug-in configuration.
APIServerArguments	Key-value pairs that will be passed directly to the Kube API server that match the API servers' command line arguments. These are not migrated, but if you reference a value that does not exist the server will not start. These values may override other settings in KubernetesMasterConfig , which may cause invalid configurations.
ControllerArguments	Key-value pairs that will be passed directly to the Kube controller manager that match the controller manager's command line arguments. These are not migrated, but if you reference a value that does not exist the server will not start. These values may override other settings in KubernetesMasterConfig , which may cause invalid configurations.
DefaultAdmissionConfig	Used to enable or disable various admission plug-ins. When this type is present as the configuration object under pluginConfig and if the admission plug-in supports it, this will cause an off by default admission plug-in to be enabled.

Parameter Name	Description
PluginConfig	Allows specifying a configuration file per admission control plug-in.
PluginOrderOverride	A list of admission control plug-in names that will be installed on the master. Order is significant. If empty, a default list of plug-ins is used.
SchedulerArguments	Key-value pairs that will be passed directly to the Kube scheduler that match the scheduler's command line arguments. These are not migrated, but if you reference a value that does not exist the server will not start. These values may override other settings in KubernetesMasterConfig , which may cause invalid configurations.

7.2.2. Asset Configuration

Table 7.2. Asset Configuration Parameters

Parameter Name	Description
AssetConfig	Holds the necessary configuration options for serving assets.
DisabledFeatures	A list of features that should not be started. You will likely want to set this as null . It is very unlikely that anyone will want to manually disable features and that is not encouraged.
Extensions	Files to serve from the asset server file system under a subcontext.
ExtensionDevelopment	When set to true , tells the asset server to reload extension scripts and stylesheets for every request rather than only at startup. It lets you develop extensions without having to restart the server for every change.
ExtensionProperties	Key- (string) and value- (string) pairs that will be injected into the console under the global variable OPENSIFT_EXTENSION_PROPERTIES .
ExtensionScripts	File paths on the asset server files to load as scripts when the web console loads.
ExtensionStylesheets	File paths on the asset server files to load as style sheets when the web console loads.
LoggingPublicURL	The public endpoint for logging (optional).
LogoutURL	An optional, absolute URL to redirect web browsers to after logging out of the web console. If not specified, the built-in logout page is shown.
MasterPublicURL	How the web console can access the OpenShift Container Platform server.

Parameter Name	Description
MetricsPublicURL	The public endpoint for metrics (optional).
PublicURL	URL of the the asset server.

7.2.3. Authentication and Authorization Configuration

Table 7.3. Authentication and Authorization Parameters

Parameter Name	Description
authConfig	Holds authentication and authorization configuration options.
AuthenticationCacheSize	Indicates how many authentication results should be cached. If 0, the default cache size is used.
AuthorizationCacheTTL	Indicates how long an authorization result should be cached. It takes a valid time duration string (e.g. "5m"). If empty, you get the default timeout. If zero (e.g. "0m"), caching is disabled.

7.2.4. Controller Configuration

Table 7.4. Controller Configuration Parameters

Parameter Name	Description
Controllers	List of the controllers that should be started. If set to none , no controllers will start automatically. The default value is * which will start all controllers. When using *, you may exclude controllers by prepending a - in front of their name. No other values are recognized at this time.
ControllerLeaseTTL	Enables controller election, instructing the master to attempt to acquire a lease before controllers start and renewing it within a number of seconds defined by this value. Setting this value non-negative forces pauseControllers=true . This value defaults off (0, or omitted) and controller election can be disabled with -1.
PauseControllers	Instructs the master to not automatically start controllers, but instead to wait until a notification to the server is received before launching them.

7.2.5. etcd Configuration

Table 7.5. etcd Configuration Parameters

Parameter Name	Description
Address	The advertised host:port for client connections to etcd.
etcdClientInfo	Contains information about how to connect to etcd.
etcdConfig	Holds the necessary configuration options for connecting with an etcd database.
etcdStorageConfig	Contains information about how API resources are stored in etcd. These values are only relevant when etcd is the backing store for the cluster.
KubernetesStoragePrefix	The path within etcd that the Kubernetes resources will be rooted under. This value, if changed, will mean existing objects in etcd will no longer be located. The default value is kubernetes.io .
KubernetesStorageVersion	The API version that Kubernetes resources in etcd should be serialized to. This value should not be advanced until all clients in the cluster that read from etcd have code that allows them to read the new version.
OpenShiftStoragePrefix	The path within etcd that the OpenShift Container Platform resources will be rooted under. This value, if changed, will mean existing objects in etcd will no longer be located. The default value is openshift.io .
OpenShiftStorageVersion	API version that OS resources in etcd should be serialized to. This value should not be advanced until all clients in the cluster that read from etcd have code that allows them to read the new version.
PeerAddress	The advertised host:port for peer connections to etcd .
PeerServingInfo	Describes how to start serving the etcd peer.
ServingInfo	Describes how to start serving the etcd master.
StorageDir	The path to the etcd storage directory.

7.2.6. Grant Configuration

Table 7.6. Grant Configuration Parameters

Parameter Name	Description
GrantConfig	Describes how to handle grants.
GrantHandlerAuto	Auto-approves client authorization grant requests.
GrantHandlerDeny	Auto-denies client authorization grant requests.

Parameter Name	Description
GrantHandlerPrompt	Prompts the user to approve new client authorization grant requests.
Method	<p>Determines the default strategy to use when an OAuth client requests a grant. This method will be used only if the specific OAuth client does not provide a strategy of their own. Valid grant handling methods are:</p> <ul style="list-style-type: none"> • auto: always approves grant requests, useful for trusted clients • prompt: prompts the end user for approval of grant requests, useful for third-party clients • deny: always denies grant requests, useful for black-listed clients

7.2.7. Image Configuration

Table 7.7. Image Configuration Parameters

Parameter Name	Description
DisableScheduledImport	Allows scheduled background import of images to be disabled.
Format	The format of the name to be built for the system component.
ImageConfig	Holds options that describe how to build image names for system components.
ImagePolicyConfig	Controls limits and behavior for importing images.
Latest	Determines if the latest tag will be pulled from the registry.
MaxImagesBulkImportedPerRepository	Controls the number of images that are imported when a user does a bulk import of a Docker repository. This number defaults to 5 to prevent users from importing large numbers of images accidentally. Set -1 for no limit.
MaxScheduledImageImportsPerMinute	The maximum number of scheduled image streams that will be imported in the background per minute. The default value is 60.
ScheduledImageImportMinimumIntervalSeconds	The minimum number of seconds that can elapse between when image streams scheduled for background import are checked against the upstream repository. The default value is 15 minutes.

7.2.8. Kubernetes Master Configuration

Table 7.8. Kubernetes Master Configuration Parameters

Parameter Name	Description
APILevels	A list of API levels that should be enabled on startup, v1 as examples.
DisabledAPIGroupVersions	A map of groups to the versions (or *) that should be disabled.
KubeletClientInfo	Contains information about how to connect to kubelets.
KubernetesMasterConfig	Holds the necessary configuration options for the Kubernetes master.
MasterCount	The number of expected masters that should be running. This value defaults to 1 and may be set to a positive integer, or if set to -1, indicates this is part of a cluster.
MasterIP	The public IP address of Kubernetes resources. If empty, the first result from net.InterfaceAddrs will be used.
MasterKubeConfig	File name for the <i>.kubeconfig</i> file that describes how to connect this node to the master.
ServicesNodePortRange	The range to use for assigning service public ports on a host.
ServicesSubnet	The subnet to use for assigning service IPs.
StaticNodeNames	The list of nodes that are statically known.

7.2.9. Network Configuration

Table 7.9. Network Configuration Parameters

Parameter Name	Description
ClusterNetworkCIDR	The CIDR string to specify the global overlay network's L3 space.
externalIPNetworkCIDRs	Controls what values are acceptable for the service external IP field. If empty, no externalIP may be set. It may contain a list of CIDRs which are checked for access. If a CIDR is prefixed with !, IPs in that CIDR will be rejected. Rejections will be applied first, then the IP checked against one of the allowed CIDRs. You should ensure this range does not overlap with your nodes, pods, or service CIDRs for security reasons.
HostSubnetLength	The number of bits to allocate to each host's subnet. For example, 8 would mean a /24 network on the host.

Parameter Name	Description
ingressIPNetworkCIDR	Controls the range to assign ingress IPs from for services of type LoadBalancer on bare metal. If empty, ingress IPs will not be assigned. It may contain a single CIDR that will be allocated from. For security reasons, you should ensure that this range does not overlap with the CIDRs reserved for external IPs, nodes, pods, or services.
NetworkConfig	Provides network options for the node.
NetworkPluginName	The name of the network plug-in to use.
ServiceNetwork	The CIDR string to specify the service networks.

7.2.10. OAuth Authentication Configuration

Table 7.10. OAuth Configuration Parameters

Parameter Name	Description
AlwaysShowProviderSelection	Forces the provider selection page to render even when there is only a single provider.
AssetPublicURL	Used for building valid client redirect URLs for external access.
Error	A path to a file containing a go template used to render error pages during the authentication or grant flow. If unspecified, the default error page is used.
IdentityProviders	Ordered list of ways for a user to identify themselves.
Login	A path to a file containing a go template used to render the login page. If unspecified, the default login page is used.
MasterCA	CA for verifying the TLS connection back to the MasterURL .
MasterPublicURL	Used for building valid client redirect URLs for external access.
MasterURL	Used for making server-to-server calls to exchange authorization codes for access tokens.
OAuthConfig	Holds the necessary configuration options for OAuth authentication.
OAuthTemplates	Allows for customization of pages like the login page.
ProviderSelection	A path to a file containing a go template used to render the provider selection page. If unspecified, the default provider selection page is used.

Parameter Name	Description
SessionConfig	Holds information about configuring sessions.
Templates	Allows you to customize pages like the login page.
TokenConfig	Contains options for authorization and access tokens.

7.2.11. Project Configuration

Table 7.11. Project Configuration Parameters

Parameter Name	Description
DefaultNodeSelector	Holds default project node label selector.
ProjectConfig	Holds information about project creation and defaults.
ProjectRequestMessage	The string presented to a user if they are unable to request a project via the project request API endpoint.
ProjectRequestTemplate	The template to use for creating projects in response to projectrequest . It is in the format namespace/template and it is optional. If it is not specified, a default template is used.

7.2.12. Scheduler Configuration

Table 7.12. Scheduler Configuration Parameters

Parameter Name	Description
SchedulerConfigFile	Points to a file that describes how to set up the scheduler. If empty, you get the default scheduling rules

7.2.13. Security Allocator Configuration

Table 7.13. Security Allocator Parameters

Parameter Name	Description
----------------	-------------

Parameter Name	Description
MCSAllocatorRange	Defines the range of MCS categories that will be assigned to namespaces. The format is <prefix>/<numberOfLabels>[,<maxCategory>] . The default is s0/2 and will allocate from c0 to c1023, which means a total of 535k labels are available (1024 choose 2 ~ 535k). If this value is changed after startup, new projects may receive labels that are already allocated to other projects. Prefix may be any valid SELinux set of terms (including user, role, and type), although leaving them as the default will allow the server to set them automatically.
SecurityAllocator	Controls the automatic allocation of UIDs and MCS labels to a project. If nil, allocation is disabled.
UIDAllocatorRange	Defines the total set of Unix user IDs (UIDs) that will be allocated to projects automatically, and the size of the block each namespace gets. For example, 1000-1999/10 will allocate ten UIDs per namespace, and will be able to allocate up to 100 blocks before running out of space. The default is to allocate from 1 billion to 2 billion in 10k blocks (which is the expected size of the ranges container images will use once user namespaces are started).

7.2.14. Service Account Configuration

Table 7.14. Service Account Configuration Parameters

Parameter Name	Description
LimitSecretReferences	Controls whether or not to allow a service account to reference any secret in a namespace without explicitly referencing them.
ManagedNames	A list of service account names that will be auto-created in every namespace. If no names are specified, the ServiceAccountsController will not be started.
MasterCA	The CA for verifying the TLS connection back to the master. The service account controller will automatically inject the contents of this file into pods so they can verify connections to the master.
PrivateKeyFile	A file containing a PEM-encoded private RSA key, used to sign service account tokens. If no private key is specified, the service account TokensController will not be started.

Parameter Name	Description
PublicKeyFiles	A list of files, each containing a PEM-encoded public RSA key. If any file contains a private key, the public portion of the key is used. The list of public keys is used to verify presented service account tokens. Each key is tried in order until the list is exhausted or verification succeeds. If no keys are specified, no service account authentication will be available.
ServiceAccountConfig	Holds the necessary configuration options for a service account.

7.2.15. Serving Information Configuration

Table 7.15. Serving Information Configuration Parameters

Parameter Name	Description
AllowRecursiveQueries	Allows the DNS server on the master to answer queries recursively. Note that open resolvers can be used for DNS amplification attacks and the master DNS should not be made accessible to public networks.
BindAddress	The ip:port to serve on.
BindNetwork	Controls limits and behavior for importing images.
CertFile	A file containing a PEM-encoded certificate.
CertInfo	TLS cert information for serving secure traffic.
ClientCA	The certificate bundle for all the signers that you recognize for incoming client certificates.
dnsConfig	Holds the necessary configuration options for DNS.
DNSDomain	Holds the domain suffix.
DNSIP	Holds the IP.
KeyFile	A file containing a PEM-encoded private key for the certificate specified by CertFile .
MasterClientConnection Overrides	Provides overrides to the client connection used to connect to the master.
MaxRequestsInFlight	The number of concurrent requests allowed to the server. If zero, no limit.
NamedCertificates	A list of certificates to use to secure requests to specific host names.

Parameter Name	Description
RequestTimeoutSecond	The number of seconds before requests are timed out. The default is 60 minutes. If -1, there is no limit on requests.
ServingInfo	The HTTP serving information for the assets.

7.2.16. Volume Configuration

Table 7.16. Volume Configuration Parameters

Parameter Name	Description
DynamicProvisioningEnabled	A boolean to enable or disable dynamic provisioning. Default is true .
FSGroup	Can be specified to enable a quota on local storage use per unique FSGroup ID. At present this is only implemented for emptyDir volumes, and if the underlying volumeDirectory is on an XFS filesystem.
LocalQuota	Contains options for controlling local volume quota on the node.
MasterVolumeConfig	Contains options for configuring volume plug-ins in the master node.
NodeVolumeConfig	Contains options for configuring volumes on the node.
VolumeConfig	Contains options for configuring volumes on the node.
VolumeDirectory	The directory that volumes are stored under.

7.2.17. Audit Configuration

Audit provides a security-relevant chronological set of records documenting the sequence of activities that have affected system by individual users, administrators, or other components of the system.

Audit works at the API server level, logging all requests coming to the server. Each audit log contains two entries:

1. The request line containing:
 - a. A Unique ID allowing to match the response line (see #2)
 - b. The source IP of the request
 - c. The HTTP method being invoked
 - d. The original user invoking the operation
 - e. The impersonated user for the operation

- f. The namespace of the request or <none>
 - g. The URI as requested
2. The response line containing:
 - a. The the unique ID from #1
 - b. The response code

Example output for user **admin** asking for a list of pods:

```
AUDIT: id="5c3b8227-4af9-4322-8a71-542231c3887b" ip="127.0.0.1"
method="GET" user="admin" as "<self>" namespace="default"
uri="/api/v1/namespaces/default/pods"
AUDIT: id="5c3b8227-4af9-4322-8a71-542231c3887b" response="200"
```

The **openshift_master_audit_config** variable enables API service auditing. It takes an array of the following options:

Table 7.17. Audit Configuration Parameters

Parameter Name	Description
enabled	A boolean to enable or disable audit logs. Default is false .
auditFilePath	File path where the requests should be logged to. If not set, logs are printed to master logs.
maximumFileRetentionDays	Specifies maximum number of days to retain old audit log files based on the time stamp encoded in their filename.
maximumRetainedFiles	Specifies the maximum number of old audit log files to retain.
maximumFileSizeMegabytes	Specifies maximum size in megabytes of the log file before it gets rotated. Defaults to 100MB.

Example Audit Configuration

```
auditConfig:
  auditFilePath: "/var/log/audit-ocp.log"
  enabled: true
  maximumFileRetentionDays: 10
  maximumFileSizeMegabytes: 10
  maximumRetainedFiles: 10
```

7.3. NODE CONFIGURATION FILES

The following **node-config.yaml** file is a sample node configuration file that was generated with the default values as of writing. You can [create a new node configuration file](#) to see the valid options for your installed version of OpenShift Container Platform.

Example 7.1. Sample Node Configuration File

```

allowDisabledDocker: false
apiVersion: v1
authConfig:
  authenticationCacheSize: 1000
  authenticationCacheTTL: 5m
  authorizationCacheSize: 1000
  authorizationCacheTTL: 5m
dnsDomain: cluster.local
dnsIP: 10.0.2.15 ❶
dockerConfig:
  execHandlerName: native
imageConfig:
  format: openshift/origin-${component}:${version}
  latest: false
iptablesSyncPeriod: 5s
kind: NodeConfig
masterKubeConfig: node.kubeconfig
networkConfig:
  mtu: 1450
  networkPluginName: ""
nodeIP: ""
nodeName: node1.example.com
podManifestConfig: ❷
  path: "/path/to/pod-manifest-file" ❸
  fileCheckIntervalSeconds: 30 ❹
proxyArguments:
  proxy-mode:
    - iptables ❺
volumeConfig:
  localQuota:
    perFSGroup: null ❻
servingInfo:
  bindAddress: 0.0.0.0:10250
  bindNetwork: tcp4
  certFile: server.crt
  clientCA: node-client-ca.crt
  keyFile: server.key
  namedCertificates: null
volumeDirectory: /root/openshift.local.volumes

```

- ❶ Configures an IP address to be prepended to a pod's */etc/resolv.conf* by adding the address here.
- ❷ Allows pods to be placed directly on certain set of nodes, or on all nodes without going through the scheduler. You can then use pods to perform the same administrative tasks and support the same services on each node.
- ❸ Specifies the path for the [pod manifest file](#) or directory. If it is a directory, then it is expected to contain one or more manifest files. This is used by the Kubelet to create pods on the node.
- ❹ This is the interval (in seconds) for checking the manifest file for new data. The interval must be a positive value.

- 5 The [service proxy implementation](#) to use.
- 6 Preliminary support for local emptyDir volume quotas, set this value to a resource quantity representing the desired quota per FSGroup, per node. (i.e. 1Gi, 512Mi, etc) Currently requires that the **volumeDirectory** be on an XFS filesystem mounted with the 'gquota' option, and the matching security context constraint's fsGroup type set to 'MustRunAs'.

7.3.1. Pod and Node Configuration

Table 7.18. Pod and Node Configuration Parameters

Parameter Name	Description
NodeConfig	The fully specified configuration starting an OpenShift Container Platform node.
NodeIP	Node may have multiple IPs, so this specifies the IP to use for pod traffic routing. If not specified, network parse/lookup on the nodeName is performed and the first non-loopback address is used.
NodeName	The value used to identify this particular node in the cluster. If possible, this should be your fully qualified hostname. If you are describing a set of static nodes to the master, this value must match one of the values in the list.
PodEvictionTimeout	Controls grace period for deleting pods on failed nodes. It takes valid time duration string. If empty, you get the default pod eviction timeout.
ProxyClientInfo	Specifies the client cert/key to use when proxying to pods.

7.3.2. Docker Configuration

Table 7.19. Docker Configuration Parameters

Parameter Name	Description
AllowDisabledDocker	If true, the kubelet will ignore errors from Docker. This means that a node can start on a machine that does not have docker started.
DockerConfig	Holds Docker related configuration options
ExecHandlerName	The handler to use for executing commands in Docker containers.

7.3.3. Parallel Image Pulls with Docker 1.9+

If you are using Docker 1.9+, you may want to consider enabling parallel image pulling, as the default is to pull images one at a time.

**NOTE**

There is a potential issue with data corruption prior to Docker 1.9. However, starting with 1.9, the corruption issue is resolved and it is safe to switch to parallel pulls.

```
kubeletArguments:
  serialize-image-pulls:
    - "false" 1
```

- 1** Change to true to disable parallel pulls. (This is the default config)

7.4. PASSWORDS AND OTHER SENSITIVE DATA

For some [authentication configurations](#), an LDAP **bindPassword** or OAuth **clientSecret** value is required. Instead of specifying these values directly in the master configuration file, these values may be provided as environment variables, external files, or in encrypted files.

Environment Variable Example

```
...
bindPassword:
  env: BIND_PASSWORD_ENV_VAR_NAME
```

External File Example

```
...
bindPassword:
  file: bindPassword.txt
```

Encrypted External File Example

```
...
bindPassword:
  file: bindPassword.encrypted
  keyFile: bindPassword.key
```

To create the encrypted file and key file for the above example:

```
$ oadm ca encrypt --genkey=bindPassword.key --out=bindPassword.encrypted
> Data to encrypt: B1ndPass0rd!
```

**WARNING**

Encrypted data is only as secure as the decrypting key. Care should be taken to limit filesystem permissions and access to the key file.

7.5. CREATING NEW CONFIGURATION FILES

When defining an OpenShift Container Platform configuration from scratch, start by creating new configuration files.

For master host configuration files, use the **openshift start** command with the **--write-config** option to write the configuration files. For node hosts, use the **oadm create-node-config** command to write the configuration files.

The following commands write the relevant launch configuration file(s), certificate files, and any other necessary files to the specified **--write-config** or **--node-dir** directory.

To create configuration files for an all-in-one server (a master and a node on the same host) in the specified directory:

```
$ openshift start --write-config=/openshift.local.config
```

To create a [master configuration file](#) and other required files in the specified directory:

```
$ openshift start master --write-config=/openshift.local.config/master
```

To create a [node configuration file](#) and other related files in the specified directory:

```
$ oadm create-node-config \
  --node-dir=/openshift.local.config/node-<node_hostname> \
  --node=<node_hostname> \
  --hostnames=<node_hostname>,<ip_address> \
  --certificate-authority="/path/to/ca.crt" \
  --signer-cert="/path/to/ca.crt" \
  --signer-key="/path/to/ca.key" \
  --signer-serial="/path/to/ca.serial.txt" \
  --node-client-certificate-authority="/path/to/ca.crt"
```

When creating node configuration files, the **--hostnames** option accepts a comma-delimited list of every host name or IP address you want server certificates to be valid for.

7.6. LAUNCHING SERVERS USING CONFIGURATION FILES

Once you have modified the master and/or node configuration files to your specifications, you can use them when launching servers by specifying them as an argument. Keep in mind that if you specify a configuration file, none of the other command line options you pass are respected.

To launch an all-in-one server using a master configuration and a node configuration file:

```
$ openshift start --master-config=/openshift.local.config/master/master-
config.yaml --node-config=/openshift.local.config/node-
<node_hostname>/node-config.yaml
```

To launch a master server using a master configuration file:

```
$ openshift start master --config=/openshift.local.config/master/master-
config.yaml
```

To launch a node server using a node configuration file:

```
$ openshift start node --config=/openshift.local.config/node-
<node_hostname>/node-config.yaml
```

7.7. CONFIGURING LOGGING LEVELS

OpenShift Container Platform uses the **systemd-journald.service** to collect log messages for debugging, using five log message severities. The logging levels are based on Kubernetes logging conventions, as follows:

Table 7.20. Log Level Options

Option	Description
0	Errors and warnings only
2	Normal information
4	Debugging-level information
6	API-level debugging information (request / response)
8	Body-level API debugging information

You can control which INFO messages are logged by setting the `loglevel` option in the in **`/etc/sysconfig/atomic-openshift-node`** or **`/etc/sysconfig/atomic-openshift-master`** file. Configuring the logs to collect all messages can lead to large logs that are difficult to interpret and can take up excessive space. Collecting all messages should only be used in debug situations.



NOTE

Messages with FATAL, ERROR, WARNING and some INFO severities appear in the logs regardless of the log configuration.

You can view logs for the master or the node system using the following command:

```
# journalctl -r -u <journal_name>
```

Use the **`-r`** option to show the newest entries first.

For example:

```
# journalctl -r -u atomic-openshift-master.service
# journalctl -r -u atomic-openshift-node.service
```

To change the logging level:

1. Edit the **`/etc/sysconfig/atomic-openshift-master`** file for the master or **`/etc/sysconfig/atomic-openshift-node`** file for the nodes.

2. Enter a value from the **Log Level Options** table above in the **OPTIONS=- -loglevel=** field.
For example:

```
OPTIONS=- -loglevel=4
```

3. Restart the master or node host as appropriate:

```
# systemctl restart atomic-openshift-master
# systemctl restart atomic-openshift-node
```

After the restart, all new log messages will conform to the new setting. Older messages do not change.



NOTE

The default log level can be set using the Advanced Install. For more information, see [Cluster Variables](#).

The following examples are excerpts from a master **journal** log at various log levels. Timestamps and system information have been removed from these examples.

Excerpt of `journalctl -u atomic-openshift-master.service` output at `loglevel=0`

```
fit failure on node: PodFitsHostPorts
E0222 factory.go:361] Error scheduling default router-1
I0222 event.go:211] Event(api.ObjectReference
I0222 start_master.go:644] Started Origin Controllers
I0222 start_master.go:623] Started Kubernetes Controllers
I0222 endpoints_controller.go:283] Waiting for pods controller
I0222 event.go:211] Event(api.ObjectReference{Kind:"Node"
W0222 nodecontroller.go:671] Missing timestamp for Node
W0222 nodecontroller.go:671] Missing timestamp for Node
```

Excerpt of `journalctl -u atomic-openshift-master.service` output at `loglevel=2`

```
E0222 factory.go:361] Error scheduling default router-1
Always TerminationGracePeriodSeconds:0xc20d819b78 ActiveDeadlineSeconds:
<nil>
I0222 scheduler.go:117] Failed to schedule: &{TypeMeta:
I0222 event.go:211] Event(api.ObjectReference{Kind:
I0222 replication_controller.go:434] Too few "default"/"router"
I0222 start_master.go:644] Started Origin Controllers
I0222 subnets.go:27] Found existing HostSubnet
I0222 subnets.go:27] Found existing HostSubnet
I0222 common.go:79] Initializing single-tenant plugin
I0222 common.go:54] Starting with configured hostname
I0222 start_master.go:623] Started Kubernetes Controllers
I0222 plugins.go:291] Loaded volume plugin "kubernetes.io/nf"
I0222 plugins.go:291] Loaded volume plugin "kubernetes.io/ho"
I0222 endpoints_controller.go:283] Waiting for pods controller"
I0222 endpoints_controller.go:283] Waiting for pods controller"
W0222 nodecontroller.go:671] Missing timestamp for Node
I0222 nodecontroller.go:604] Recording Registered Node
I0222 nodecontroller.go:416] NodeController observed
```

Excerpt of journalctl -u atomic-openshift-master.service output at loglevel=4

```

controller_utils.go:592] Ignoring inactive pod default/router-1-0ww1g in
state Failed, deletion time <nil>
I0222 replication_controller.go:497] Finished syncing controller
"default/docker-registry-2" (206.507µs)
I0222 controller_utils.go:592] Ignoring inactive pod default/router-1-
zyi9y in state Failed, deletion time <nil>
I0222 replication_controller.go:497] Finished syncing controller
"default/docker-registry-1" (176.748µs)
I0222 controller_utils.go:160] Controller default/router-1 either never
recorded expectations, or the ttl expired.
I0222 controller_utils.go:160] Controller default/docker-registry-2 either
never recorded expectations, or the ttl expired.
I0222 nodecontroller.go:709] Node dell-r430-20.gsslab.pnq.redhat.com
ReadyCondition updated. Updating timestamp.
I0222 factory.go:474] Backing off 1m0s for pod &{backoff:600000000000
lastUpdate:{sec:636233 loc:0x59a6560} reqInFlight:1}
; retrying
fit failure on node (ibm-x3650m4-01-vm-02.lab.eng.bos.redhat.com):
PodFitsHostPorts
E0222 factory.go:361] Error scheduling default router-1-d7svd: pod
(router-1-d7svd) failed to fit in any node
b089d8 ActiveDeadlineSeconds:<nil> DNSPolicy:ClusterFirst
NodeSelector:map[region:infra] ServiceAccountName:router NodeName:
SecurityContext:0xc20ecd00c0 ImagePullSec
}] Spec:{Volumes:[{Name:router-token-h6q3n VolumeSource:{HostPath:<nil>
EmptyDir:<nil> GCEPersistentDisk:<nil> AWSElasticBlockStore:<nil> GitRepo:
<nil> Secret:0xc20eb
I0222 factory.go:474] Backing off 16s for pod &{backoff:320000000000
lastUpdate:{sec:63623353879 nsec:506358920 loc:0x59a6560} reqInFlight:1}
b118b8 ActiveDeadlineSeconds:<nil> DNSPolicy:ClusterFirst
NodeSelector:map[region:infra] ServiceAccountName:router NodeName:

```

Excerpt of journalctl -u atomic-openshift-master.service output at loglevel=8

```

I0222 round_tripper.go:271] Request Headers:
I0222 round_tripper.go:264] GET https://dell-r430-20.gsslab
I0222 round_tripper.go:274] User-Agent: openshift/v3.2.
I0222 round_tripper.go:274] Accept: application/json, *
I0222 round_tripper.go:271] Request Headers:
I0222 round_tripper.go:264] GET https://dell-r430-20.gsslab
I0222 request.go:870] Response Body: {"kind":"ClusterRole", "
I0222 round_tripper.go:295] Content-Type: application/j
I0222 round_tripper.go:295] Cache-Control: no-store
I0222 round_tripper.go:295] Content-Length: 898
I0222 round_tripper.go:295] Date: Wed, 22 Feb 2017 09:3
I0222 round_tripper.go:292] Response Headers:
I0222 round_tripper.go:289] Response Status: 200 OK in 7 mi
I0222 configgetter.go:127] using watch cache storage (capacity=1000)
I0222 controller_utils.go:592] Ignoring inactive pod

```

CHAPTER 8. ADDING HOSTS TO AN EXISTING CLUSTER

8.1. OVERVIEW

Depending on how your OpenShift Container Platform cluster was installed, you can add new hosts (either nodes or masters) to your installation by using the install tool for quick installations, or by using the *scaleup.yml* playbook for advanced installations.

8.2. ADDING HOSTS USING THE QUICK INSTALLER TOOL


If you used the quick install tool to install your OpenShift Container Platform cluster, you can use the quick install tool to add a new node host to your existing cluster, or to reinstall the cluster entirely.



NOTE

Currently, you can not use the quick installer tool to add new master hosts. You must use the [advanced installation](#) method to do so.

If you used the installer in either [interactive](#) or [unattended](#) mode, you can re-run the installation as long as you have an [installation configuration file](#) at `~/config/openshift/installer.cfg.yml` (or specify a different location with the `-c` option).



IMPORTANT

The recommended maximum number of nodes is 1000.

To add nodes to your installation:

1. Re-run the installer with the **install** subcommand in interactive or unattended mode:

```
$ atomic-openshift-installer [-u] [-c </path/to/file>] install
```

2. The installer detects your current environment and allows you to either add an additional node or re-perform a clean install:

```
Gathering information from hosts...
Installed environment detected.
By default the installer only adds new nodes to an installed
environment.
Do you want to (1) only add additional nodes or (2) perform a clean
install?:
```


Choose (1) and follow the on-screen instructions to complete your desired task.

8.3. ADDING HOSTS USING THE ADVANCED INSTALL

If you installed using the advanced install, you can add new hosts to your cluster by running the *scaleup.yml* playbook. This playbook queries the master, generates and distributes new certificates for the new hosts, then runs the configuration playbooks on the new hosts only. Before running the *scaleup.yml* playbook, complete all prerequisite [host preparation](#) steps.

This process is similar to re-running the installer in the [quick installation method to add nodes](#), however you have more configuration options available when using the advanced method and when running the playbooks directly.

You must have an existing inventory file (for example, `/etc/ansible/hosts`) that is representative of your current cluster configuration in order to run the `scaleup.yml` playbook. If you previously used the `atomic-openshift-installer` command to run your installation, you can check `~/.config/openshift/hosts` (previously located at `~/.config/openshift/ansible/hosts`) for the last inventory file that the installer generated, and use or modify that as needed as your inventory file. You must then specify the file location with `-i` when calling `ansible-playbook` later.



IMPORTANT

The recommended maximum number of nodes is 1000.

To add a host to an existing cluster:

1. Ensure you have the latest playbooks by updating the `atomic-openshift-utils` package:

```
# yum update atomic-openshift-utils
```

2. Edit your `/etc/ansible/hosts` file and add `new_<host_type>` to the `[OSEv3:children]` section: For example, to add a new node host, add `new_nodes`:

```
[OSEv3:children]
masters
nodes
new_nodes
```

To add new master hosts, add `new_masters`.

3. Create a `[new_<host_type>]` section much like an existing section, specifying host information for any new hosts you want to add. For example, when adding a new node:

```
[nodes]
master[1:3].example.com
node1.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'west' }"
infra-node1.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"
infra-node2.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"

[new_nodes]
node3.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'west' }"
```

See [Configuring Host Variables](#) for more options.

When adding new masters, hosts added to the `[new_masters]` section must also be added to the `[new_nodes]` section. This ensures the new master host is part of the OpenShift SDN.


```
[masters]
master[1:2].example.com

[new_masters]
master3.example.com

[nodes]
master[1:2].example.com
node1.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'east' }"
node2.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'west' }"
infra-node1.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"
infra-node2.example.com openshift_node_labels="{ 'region': 'infra',
'zone': 'default' }"

[new_nodes]
master3.example.com
```

Masters are also automatically marked as unschedulable for pod placement by the installer.



IMPORTANT

If you label a master host with the **region=infra** label and have no other dedicated infrastructure nodes, you must also explicitly mark the host as schedulable by adding **openshift_schedulable=true** to the entry. Otherwise, the registry and router pods cannot be placed anywhere.

4. Run the **scaleup.yml** playbook. If your inventory file is located somewhere other than the default of **/etc/ansible/hosts**, specify the location with the **-i option**.

For additional nodes:

```
# ansible-playbook [-i /path/to/file] \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
    node/scaleup.yml
```

For additional masters:

```
# ansible-playbook [-i /path/to/file] \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
    master/scaleup.yml
```

5. After the playbook completes successfully, [verify the installation](#).
6. Finally, move any hosts you had defined in the **[new_<host_type>]** section into their appropriate section (but leave the **[new_<host_type>]** section definition itself in place) so that subsequent runs using this inventory file are aware of the nodes but do not handle them as new nodes. For example, when adding new nodes:

```
[nodes]
master[1:3].example.com
node1.example.com openshift_node_labels="{ 'region': 'primary',
'zone': 'east' }"
```

```
node2.example.com openshift_node_labels="{ 'region': 'primary',  
'zone': 'west' }"  
node3.example.com openshift_node_labels="{ 'region': 'primary',  
'zone': 'west' }"  
infra-node1.example.com openshift_node_labels="{ 'region': 'infra',  
'zone': 'default' }"  
infra-node2.example.com openshift_node_labels="{ 'region': 'infra',  
'zone': 'default' }"  
  
[new_nodes]
```


CHAPTER 9. LOADING THE DEFAULT IMAGE STREAMS AND TEMPLATES

9.1. OVERVIEW

Your OpenShift Container Platform installation includes useful sets of Red Hat-provided [image streams](#) and [templates](#) to make it easy for developers to create new applications. By default, the [quick](#) and [advanced installation](#) methods automatically create these sets in the **openshift** project, which is a default global project to which all users have view access.

9.2. OFFERINGS BY SUBSCRIPTION TYPE

Depending on the active subscriptions on your Red Hat account, the following sets of image streams and templates are provided and supported by Red Hat. Contact your Red Hat sales representative for further subscription details.

9.2.1. OpenShift Container Platform Subscription

The core set of image streams and templates are provided and supported with an active *OpenShift Container Platform subscription*. This includes the following technologies:

Type	Technology
Languages & Frameworks	<ul style="list-style-type: none">• .NET Core• Node.js• Perl• PHP• Python• Ruby
Databases	<ul style="list-style-type: none">• MongoDB• MySQL• PostgreSQL
Middleware Services	<ul style="list-style-type: none">• Red Hat JBoss Web Server (Tomcat)• Red Hat Single Sign-on
Other Services	<ul style="list-style-type: none">• Jenkins

9.2.2. xPaaS Middleware Add-on Subscriptions

Support for xPaaS middleware images are provided by *xPaaS Middleware add-on subscriptions*, which are separate subscriptions for each xPaaS product. If the relevant subscription is active on your account, image streams and templates are provided and supported for the following technologies:

Type	Technology
Middleware Services	<ul style="list-style-type: none"> • Red Hat JBoss A-MQ • Red Hat JBoss BPM Suite Intelligent Process Server • Red Hat JBoss BRMS Decision Server • Red Hat JBoss Data Grid • Red Hat JBoss EAP • Red Hat JBoss Fuse Integration Services

9.3. BEFORE YOU BEGIN

Before you consider performing the tasks in this topic, confirm if these image streams and templates are already registered in your OpenShift Container Platform cluster by doing one of the following:

- Log into the web console and click **Add to Project**.
- List them for the **openshift** project using the CLI:

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

If the default image streams and templates are ever removed or changed, you can follow this topic to create the default objects yourself. Otherwise, the following instructions are not necessary.

9.4. PREREQUISITES

Before you can create the default image streams and templates:

- The [integrated Docker registry](#) service must be deployed in your OpenShift Container Platform installation.
- You must be able to run the **oc create** command with [cluster-admin privileges](#), because they operate on the default **openshiftproject**.
- You must have installed the **atomic-openshift-utils** RPM package. See [Software Prerequisites](#) for instructions.
- Define shell variables for the directories containing image streams and templates. This significantly shortens the commands in the following sections. To do this:

```
$ IMAGESTREAMDIR="/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/image-streams";
```

```
\
    XPAASSTREAMDIR="/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/xpaas-streams";
\
    XPAASTEMPLATES="/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/xpaas-
templates"; \
    DBTEMPLATES="/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/db-templates";
\
    QSTEMPLATES="/usr/share/ansible/openshift-
ansible/roles/openshift_examples/files/examples/v1.3/quickstart-
templates"
```

9.5. CREATING IMAGE STREAMS FOR OPENSIFT CONTAINER PLATFORM IMAGES

If your node hosts are subscribed using Red Hat Subscription Manager and you want to use the core set of image streams that used Red Hat Enterprise Linux (RHEL) 7 based images:

```
$ oc create -f $IMAGESTREAMDIR/image-streams-rhel7.json -n openshift
```

Alternatively, to create the core set of image streams that use the CentOS 7 based images:

```
$ oc create -f $IMAGESTREAMDIR/image-streams-centos7.json -n openshift
```

Creating both the CentOS and RHEL sets of image streams is not possible, because they use the same names. To have both sets of image streams available to users, either create one set in a different project, or edit one of the files and modify the image stream names to make them unique.

9.6. CREATING IMAGE STREAMS FOR XPAAS MIDDLEWARE IMAGES

The xPaaS Middleware image streams provide images for **JBoss EAP**, **JBoss JWS**, **JBoss A-MQ**, **JBoss Fuse Integration Services**, **Decision Server**, and **JBoss Data Grid**. They can be used to build applications for those platforms using the provided templates.

To create the xPaaS Middleware set of image streams:

```
$ oc create -f $XPAASSTREAMDIR/jboss-image-streams.json -n openshift
```



NOTE

Access to the images referenced by these image streams requires the relevant xPaaS Middleware subscriptions.

9.7. CREATING DATABASE SERVICE TEMPLATES

The database service templates make it easy to run a database image which can be utilized by other components. For each database ([MongoDB](#), [MySQL](#), and [PostgreSQL](#)), two templates are defined.

One template uses ephemeral storage in the container which means data stored will be lost if the container is restarted, for example if the pod moves. This template should be used for demonstration purposes only.

The other template defines a persistent volume for storage, however it requires your OpenShift Container Platform installation to have [persistent volumes](#) configured.

To create the core set of database templates:

```
$ oc create -f $DBTEMPLATES -n openshift
```


After creating the templates, users are able to easily instantiate the various templates, giving them quick access to a database deployment.

9.8. CREATING INSTANT APP AND QUICKSTART TEMPLATES

The Instant App and Quickstart templates define a full set of objects for a running application. These include:

- [Build configurations](#) to build the application from source located in a GitHub public repository
- [Deployment configurations](#) to deploy the application image after it is built.
- [Services](#) to provide load balancing for the application [pods](#).
- [Routes](#) to provide external access to the application.

Some of the templates also define a database deployment and service so the application can perform database operations.



NOTE

The templates which define a database use 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.

Using these templates, users are able to easily instantiate full applications using the various language images provided with OpenShift Container Platform. They can also customize the template parameters during instantiation so that it builds source from their own repository rather than the sample repository, so this provides a simple starting point for building new applications.

To create the core Instant App and Quickstart templates:

```
$ oc create -f $QSTEMPLATES -n openshift
```

There is also a set of templates for creating applications using various xPaaS Middleware products (**JBoss EAP**, **JBoss JWS**, **JBoss A-MQ**, **JBoss Fuse Integration Services**, **Decision Server**, and **JBoss Data Grid**), which can be registered by running:

```
$ oc create -f $XPAASTEMPLATES -n openshift
```

**NOTE**

The xPaaS Middleware templates require the [xPaaS Middleware image streams](#), which in turn require the relevant xPaaS Middleware subscriptions.

**NOTE**

The templates which define a database use 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.

9.9. WHAT'S NEXT?

With these artifacts created, developers can now [log into the web console](#) and follow the flow for [creating from a template](#). Any of the database or application templates can be selected to create a running database service or application in the current project. Note that some of the application templates define their own database services as well.

The example applications are all built out of GitHub repositories which are referenced in the templates by default, as seen in the **SOURCE_REPOSITORY_URL** parameter value. Those repositories can be forked, and the fork can be provided as the **SOURCE_REPOSITORY_URL** parameter value when creating from the templates. This allows developers to experiment with creating their own applications.

You can direct your developers to the [Using the Instant App and Quickstart Templates](#) section in the Developer Guide for these instructions.

CHAPTER 10. CONFIGURING CUSTOM CERTIFICATES

10.1. OVERVIEW

Administrators can configure custom serving certificates for the public host names of the OpenShift Container Platform API and [web console](#). This can be done during an [advanced installation](#) or configured after installation.

10.2. CONFIGURING CUSTOM CERTIFICATES WITH ANSIBLE

During [advanced installations](#), custom certificates can be configured using the `openshift_master_named_certificates` and `openshift_master_overwrite_named_certificates` parameters, which are configurable in the inventory file. More details are available about [configuring custom certificates with Ansible](#).

Example Custom Certificate Configuration with Ansible

```
# Configure custom named certificates
# NOTE: openshift_master_named_certificates is cached on masters and is an
# additive fact, meaning that each run with a different set of
# certificates
# will add the newly provided certificates to the cached set of
# certificates.
#
# An optional CA may be specified for each named certificate. CAs will
# be added to the OpenShift CA bundle which allows for the named
# certificate to be served for internal cluster communication.
#
# If you would like openshift_master_named_certificates to be overwritten
# with
# the provided value, specify
# openshift_master_overwrite_named_certificates.
# openshift_master_overwrite_named_certificates=true
#
# Provide local certificate paths which will be deployed to masters
# openshift_master_named_certificates=[{"certfile":
# "/path/on/host/to/custom1.crt", "keyfile": "/path/on/host/to/custom1.key",
# "cafile": "/path/on/host/to/custom-ca1.crt"}]
#
# Detected names may be overridden by specifying the "names" key
# openshift_master_named_certificates=[{"certfile":
# "/path/on/host/to/custom1.crt", "keyfile": "/path/on/host/to/custom1.key",
# "names": ["public-master-host.com"], "cafile": "/path/on/host/to/custom-
# ca1.crt"}]
```

10.3. CONFIGURING CUSTOM CERTIFICATES

In the [master configuration file](#) you can list the `namedCertificates` section in the `assetConfig.servingInfo` section so the custom certificate serves up for the web console, and in the `servingInfo` section so the custom certificate serves up for the CLI and other API calls. Multiple certificates can be configured this way and each certificate may be associated with multiple host names or wildcards.

A default certificate must be configured in the **servingInfo.certFile** and **servingInfo.keyFile** configuration sections in addition to **namedCertificates**.



NOTE

The **namedCertificates** section should only be configured for the host name associated with the **masterPublicURL**, **assetConfig.publicURL**, and **oauthConfig.assetPublicURL** settings. Using a custom serving certificate for the host name associated with the **masterURL** will result in TLS errors as infrastructure components will attempt to contact the master API using the internal **masterURL** host.

Custom Certificates Configuration

```
servingInfo:
  ...
  namedCertificates:
  - certFile: custom.crt
    keyFile: custom.key
    names:
    - "customhost.com"
    - "api.customhost.com"
    - "console.customhost.com"
  ...
```

Relative paths are resolved relative to the master configuration file. Restart the server to pick up the configuration changes.

For the master API or web console, wildcard names are accepted.

CHAPTER 11. REDEPLOYING CERTIFICATES

11.1. OVERVIEW

OpenShift Container Platform uses certificates to provide secure connections for the following components:

- masters (API server and controllers)
- etcd
- nodes
- registry
- router

You can use Ansible playbooks provided with the installer to automate checking expiration dates for cluster certificates. Playbooks are also provided to automate backing up and redeploying these certificates, which can fix common certificate errors.

Possible use cases for redeploying certificates include:

- The installer detected the wrong host names and the issue was identified too late.
- The certificates are expired and you need to update them.
- You have a new CA and would like to create certificates using it instead.

11.2. CHECKING CERTIFICATE EXPIRATIONS

You can use the installer to warn you about any certificates expiring within a configurable window of days and notify you about any certificates that have already expired. Certificate expiry playbooks use the Ansible role **openshift_certificate_expiry**.

Certificates examined by the role include:

- Master and node service certificates
- Router and registry service certificates from etcd secrets
- Master, node, router, registry, and *kubeconfig* files for **cluster-admin** users
- etcd certificates (including embedded)

11.2.1. Role Variables

The **openshift_certificate_expiry** role uses the following variables:

Table 11.1. Core Variables

Variable Name	Default Value	Description
<code>openshift_certificate_expiry_config_base</code>	<code>/etc/origin</code>	Base OpenShift Container Platform configuration directory.
<code>openshift_certificate_expiry_warning_days</code>	<code>30</code>	Flag certificates that will expire in this many days from now.
<code>openshift_certificate_expiry_show_all</code>	<code>no</code>	Include healthy (non-expired and non-warning) certificates in results.

Table 11.2. Optional Variables

Variable Name	Default Value	Description
<code>openshift_certificate_expiry_generate_html_report</code>	<code>no</code>	Generate an HTML report of the expiry check results.
<code>openshift_certificate_expiry_html_report_path</code>	<code>/tmp/cert-expiry-report.html</code>	The full path for saving the HTML report.
<code>openshift_certificate_expiry_save_json_results</code>	<code>no</code>	Save expiry check results as a JSON file.
<code>openshift_certificate_expiry_json_results_path</code>	<code>/tmp/cert-expiry-report.json</code>	The full path for saving the JSON report.

11.2.2. Running Certificate Expiration Playbooks

The OpenShift Container Platform installer provides a set of example certificate expiration playbooks, using different sets of configuration for the `openshift_certificate_expiry` role.

These playbooks must be used with an [inventory file](#) that is representative of the cluster. For best results, run **ansible-playbook** with the `-v` option.

Using the ***easy-mode.yaml*** example playbook, you can try the role out before tweaking it to your specifications as needed. This playbook:

- Produces JSON and stylized HTML reports in */tmp/*.
- Sets the warning window very large, so you will almost always get results back.
- Includes all certificates (healthy or not) in the results.

easy-mode.yaml Playbook

```
- name: Check cert expirys
  hosts: nodes:masters:etcd
  become: yes
```

```
gather_facts: no
vars:
  openshift_certificate_expiry_warning_days: 1500
  openshift_certificate_expiry_save_json_results: yes
  openshift_certificate_expiry_generate_html_report: yes
  openshift_certificate_expiry_show_all: yes
roles:
  - role: openshift_certificate_expiry
```

To run the ***easy-mode.yaml*** playbook:

```
$ ansible-playbook -v -i <inventory_file> \
  /usr/share/ansible/openshift-
  ansible/playbooks/certificate_expiry/easy-mode.yaml
```

Other Example Playbooks

The other example playbooks are also available to run directly out of the ***/usr/share/ansible/openshift-ansible/playbooks/certificate_expiry/*** directory.

Table 11.3. Other Example Playbooks

File Name	Usage
<i>default.yaml</i>	Produces the default behavior of the <i>openshift_certificate_expiry</i> role.
<i>html_and_json_default_paths.yaml</i>	Generates HTML and JSON artifacts in their default paths.
<i>longer_warning_period.yaml</i>	Changes the expiration warning window to 1500 days.
<i>longer-warning-period-json-results.yaml</i>	Changes the expiration warning window to 1500 days and saves the results as a JSON file.

To run any of these example playbooks:

```
$ ansible-playbook -v -i <inventory_file> \
  /usr/share/ansible/openshift-
  ansible/playbooks/certificate_expiry/<playbook>
```

11.2.3. Output Formats

As noted above, there are two ways to format your check report. In JSON format for machine parsing, or as a stylized HTML page for easy skimming.

HTML Report

An example of an HTML report is provided with the installer. You can open the following file in your browser to view it:

/usr/share/ansible/openshift-ansible/roles/openshift_certificate_expiry/examples/cert-expiry-report.html

JSON Report

There are two top-level keys in the saved JSON results: **data** and **summary**.

The **data** key is a hash where the keys are the names of each host examined and the values are the check results for the certificates identified on each respective host.

The **summary** key is a hash that summarizes the total number of certificates:

- examined on the entire cluster
- that are OK
- expiring within the configured warning window
- already expired

For an example of the full JSON report, see */usr/share/ansible/openshift-ansible/roles/openshift_certificate_expiry/examples/cert-expiry-report.json*.

The summary from the JSON data can be easily checked for warnings or expirations using a variety of command-line tools. For example, using **grep** you can look for the word **summary** and print out the two lines after the match (**-A2**):

```
$ grep -A2 summary /tmp/cert-expiry-report.json
    "summary": {
        "warning": 16,
        "expired": 0
```

If available, the **jq** tool can also be used to pick out specific values. The first two examples below show how to select just one value, either **warning** or **expired**. The third example shows how to select both values at once:

```
$ jq '.summary.warning' /tmp/cert-expiry-report.json
16

$ jq '.summary.expired' /tmp/cert-expiry-report.json
0

$ jq '.summary.warning, .summary.expired' /tmp/cert-expiry-report.json
16
0
```

11.3. REDEPLOYING CERTIFICATES

Use the following playbooks to redeploy master, etcd, node, registry, and router certificates on all relevant hosts. You can redeploy all of them at once using the current CA, redeploy certificates for specific components only, or redeploy a newly generated or custom CA on its own.

Just like the certificate expiry playbooks, these playbooks must be run with an [inventory file](#) that is representative of the cluster.

In particular, the inventory must specify or override all host names and IP addresses set via the following variables such that they match the current cluster configuration:

- **openshift_hostname**

- `openshift_public_hostname`
- `openshift_ip`
- `openshift_public_ip`
- `openshift_master_cluster_hostname`
- `openshift_master_cluster_public_hostname`

11.3.1. Redeploying All Certificates Using the Current OpenShift Container Platform and etcd CA

The ***redeploy-certificates.yml*** playbook does *not* regenerate the OpenShift Container Platform CA certificate. New master, etcd, node, registry, and router certificates are created using the current CA certificate to sign new certificates.

This also includes serial restarts of:

- etcd
- master services
- node services

To redeploy master, etcd, and node certificates using the current OpenShift Container Platform CA, run this playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
    cluster/redeploy-certificates.yml
```

11.3.2. Redeploying a New or Custom OpenShift Container Platform CA

The ***redeploy-openshift-ca.yml*** playbook redeployes the OpenShift Container Platform CA certificate by generating a new CA certificate and distributing an updated bundle to all components including client ***kubeconfig*** files and the node's database of trusted CAs (the CA-trust).

This also includes serial restarts of:

- master services
- node services
- docker

Additionally, you can specify a [custom CA certificate](#) when redeploying certificates instead of relying on a CA generated by OpenShift Container Platform.

When the master services are restarted, the registry and routers can continue to communicate with the master without being redeployed because the master's serving certificate is the same, and the CA the registry and routers have are still valid.

To redeploy a newly generated or custom CA:

1. If you want to use a custom CA, set the following variable in your inventory file:

```
# Configure custom ca certificate
# NOTE: CA certificate will not be replaced with existing clusters.
# This option may only be specified when creating a new cluster or
# when redeploying cluster certificates with the redeploy-
# certificates
# playbook.
openshift_master_ca_certificate={'certfile': '</path/to/ca.crt>',
'keyfile': '</path/to/ca.key>'}
```

If you do not set the above, then the current CA will be regenerated in the next step.

2. Run the ***redeploy-openshift-ca.yml*** playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
  /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
  cluster/redeploy-openshift-ca.yml
```

With the new OpenShift Container Platform CA in place, you can then use the [redeploy-certificates.yml](#) [playbook](#) at your discretion whenever you want to redeploy certificates signed by the new CA on all components.

11.3.3. Redeploying a New etcd CA

The ***redeploy-etcd-ca.yml*** playbook redeploys the etcd CA certificate by generating a new CA certificate and distributing an updated bundle to all etcd peers and master clients.

This also includes serial restarts of:

- etcd
- master services

To redeploy a newly generated etcd CA:

1. Run the ***redeploy-etcd-ca.yml*** playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
  /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
  cluster/redeploy-etcd-ca.yml
```


With the new etcd CA in place, you can then use the [redeploy-etcd-certificates.yml](#) [playbook](#) at your discretion whenever you want to redeploy certificates signed by the new etcd CA on etcd peers and master clients. Alternatively, you can use the [redeploy-certificates.yml](#) [playbook](#) to redeploy certificates for OpenShift Container Platform components in addition to etcd peers and master clients.

11.3.4. Redeploying Master Certificates Only

The ***redeploy-master-certificates.yml*** playbook only redeploys master certificates. This also includes serial restarts of master services.

To redeploy master certificates, run this playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
cluster/redeploy-master-certificates.yml
```



NOTE

After running this playbook, you need to regenerate any [service signing certificate or key pairs](#) by deleting existing secrets that contain service serving certificates or removing and re-adding annotations to appropriate services.

11.3.5. Redeploying etcd Certificates Only

The ***redeploy-etcd-certificates.yml*** playbook only redeploys etcd certificates including master client certificates.

This also include serial restarts of:

- etcd
- master services.

To redeploy etcd certificates, run this playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
cluster/redeploy-etcd-certificates.yml
```

11.3.6. Redeploying Node Certificates Only

The ***redeploy-node-certificates.yml*** playbook only redeploys node certificates. This also include serial restarts of node services.

To redeploy node certificates, run this playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
cluster/redeploy-node-certificates.yml
```

11.3.7. Redeploying Registry or Router Certificates Only

The ***redeploy-registry-certificates.yml*** and ***redeploy-router-certificates.yml*** playbooks replace installer-created certificates for the registry and router. If custom certificates are in use for these components, see [Redeploying Custom Registry or Router Certificates](#) to replace them manually.

11.3.7.1. Redeploying Registry Certificates Only

To redeploy registry certificates, run the following playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
    /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
cluster/redeploy-registry-certificates.yml
```

11.3.7.2. Redeploying Router Certificates Only

To redeploy router certificates, run the following playbook, specifying your inventory file:

```
$ ansible-playbook -i <inventory_file> \
  /usr/share/ansible/openshift-ansible/playbooks/byo/openshift-
  cluster/redeploy-router-certificates.yml
```

11.3.8. Redeploying Custom Registry or Router Certificates

When nodes are evacuated due to a redeployed CA, registry and router pods are restarted. If the registry and router certificates were not also redeployed with the new CA, this can cause outages because they cannot reach the masters using their old certificates.

The playbooks for redeploying certificates cannot redeploy custom registry or router certificates, so to address this issue, you can manually redeploy the registry and router certificates.

11.3.8.1. Redeploying Registry Certificates Manually

To redeploy registry certificates manually, you must add new registry certificates to a secret named **registry-certificates**, then redeploy the registry:

1. Switch to the **default** project for the remainder of these steps:

```
$ oc project default
```

2. If your registry was initially created on OpenShift Container Platform 3.1 or earlier, it may still be using environment variables to store certificates (which has been deprecated in favor of using secrets).

- a. Run the following and look for the **OPENSIFT_CA_DATA**, **OPENSIFT_CERT_DATA**, **OPENSIFT_KEY_DATA** environment variables:

```
$ oc env dc/docker-registry --list
```

- b. If they do not exist, skip this step. If they do, create the following **ClusterRoleBinding**:

```
$ cat <<EOF |
apiVersion: v1
groupNames: null
kind: ClusterRoleBinding
metadata:
  creationTimestamp: null
  name: registry-registry-role
roleRef:
  kind: ClusterRole
  name: system:registry
subjects:
- kind: ServiceAccount
  name: registry
  namespace: default
userNames:
```

```
- system:serviceaccount:default:registry
EOF
oc create -f -
```

Then, run the following to remove the environment variables:

```
$ oc env dc/docker-registry OPENSIFT_CA_DATA-
OPENSIFT_CERT_DATA- OPENSIFT_KEY_DATA- OPENSIFT_MASTER-
```

3. Set the following environment variables locally to make later commands less complex:

```
$ REGISTRY_IP=`oc get service docker-registry -o
jsonpath='{.spec.clusterIP}'`
$ REGISTRY_HOSTNAME=`oc get route/docker-registry -o
jsonpath='{.spec.host}'`
```

4. Create new registry certificates:

```
$ oc adm ca create-server-cert \
  --signer-cert=/etc/origin/master/ca.crt \
  --signer-key=/etc/origin/master/ca.key \
  --hostnames=$REGISTRY_IP,docker-
registry.default.svc.cluster.local,$REGISTRY_HOSTNAME \
  --cert=/etc/origin/master/registry.crt \
  --key=/etc/origin/master/registry.key \
  --signer-serial=/etc/origin/master/ca.serial.txt
```

5. Update the **registry-certificates** secret with the new registry certificates:

```
$ oc secret new registry-certificates \
  /etc/origin/master/registry.crt \
  /etc/origin/master/registry.key \
  -o json | oc replace -f -
```

6. Redeploy the registry:

```
$ oc deploy dc/docker-registry --latest
```

11.3.8.2. Redeploying Router Certificates Manually

When routers are initially deployed, an annotation is added to the router's service that automatically creates a [service serving certificate secret](#).

To redeploy router certificates manually, that service serving certificate can be triggered to be recreated by deleting the secret, removing and re-adding annotations to the **router** service, then redeploying the router:

1. Switch to the **default** project for the remainder of these steps:

```
$ oc project default
```


2. If your router was initially created on OpenShift Container Platform 3.1 or earlier, it might still use environment variables to store certificates, which has been deprecated in favor of using service serving certificate secret.

- a. Run the following command and look for the **OPENSIFT_CA_DATA**, **OPENSIFT_CERT_DATA**, **OPENSIFT_KEY_DATA** environment variables:

```
$ oc env dc/router --list
```

- b. If those variables exist, create the following **ClusterRoleBinding**:

```
$ cat <<EOF |
apiVersion: v1
groupNames: null
kind: ClusterRoleBinding
metadata:
  creationTimestamp: null
  name: router-router-role
roleRef:
  kind: ClusterRole
  name: system:router
subjects:
- kind: ServiceAccount
  name: router
  namespace: default
userNames:
- system:serviceaccount:default:router
EOF
oc create -f -
```

- c. If those variables exist, run the following command to remove them:

```
$ oc env dc/router OPENSIFT_CA_DATA- OPENSIFT_CERT_DATA-
OPENSIFT_KEY_DATA- OPENSIFT_MASTER-
```

3. Obtain a certificate.

- If you use an external Certificate Authority (CA) to sign your certificates, create a new certificate and provide it to OpenShift Container Platform by following your internal processes.
- If you use the internal OpenShift Container Platform CA to sign certificates, run the following commands:



IMPORTANT

The following commands generate a certificate that is internally signed. It will be trusted by only clients that trust the OpenShift Container Platform CA.

```
$ cd /root
$ mkdir cert ; cd cert
$ oc adm ca create-server-cert \
  --signer-cert=/etc/origin/master/ca.crt \
  --signer-key=/etc/origin/master/ca.key \
```

```
--signer-serial=/etc/origin/master/ca.serial.txt \
--hostnames='*.hostnames.for.the.certificate' \
--cert=router.crt \
--key=router.key \
```

These commands generate the following files:

- o A new certificate named **router.crt**.
- o A copy of the signing CA certificate chain, **/etc/origin/master/ca.crt**. This chain can contain more than one certificate if you use intermediate CAs.
- o A corresponding private key named **router.key**.

4. Create a new file that concatenates the generated certificates:

```
$ cat router.crt /etc/origin/master/ca.crt router.key > router.pem
```

5. Before you generate a new secret, back up the current one:

```
$ oc export secret router-certs > ~/old-router-certs-secret.yaml
```

6. Create a new secret to hold the new certificate and key, and replace the contents of the existing secret:

```
$ oc create secret tls router-certs --cert=router.pem \ 1
--key=router.key -o json --dry-run | \
oc replace -f -
```

1 **router.pem** is the file that contains the concatenation of the certificates that you generated.

7. Remove the following annotations from the **router** service:

```
$ oc annotate service router \
service.alpha.openshift.io/serving-cert-secret-name- \
service.alpha.openshift.io/serving-cert-signed-by-
```

8. Re-add the annotations:

```
$ oc annotate service router \
service.alpha.openshift.io/serving-cert-secret-name=router-certs
```

9. Redeploy the router:


```
$ oc deploy dc/router --latest
```

CHAPTER 12. CONFIGURING AUTHENTICATION AND USER AGENT

12.1. OVERVIEW

The OpenShift Container Platform [master](#) includes a built-in [OAuth server](#). Developers and administrators obtain [OAuth access tokens](#) to authenticate themselves to the API.

As an administrator, you can configure OAuth using the [master configuration file](#) to specify an [identity provider](#). It is a best practice to configure your identity provider during [advanced installation](#), but you can configure it after installation.



NOTE

OpenShift Container Platform user names containing `/`, `:`, and `%` are not supported.

If you installed OpenShift Container Platform using the [Quick Installation](#) or [Advanced Installation](#) method, the [Deny All](#) identity provider is used by default, which denies access for all user names and passwords. To allow access, you must choose a different identity provider and configure the master configuration file appropriately (located at `/etc/origin/master/master-config.yaml` by default).

When you run a master without a configuration file, the [Allow All](#) identity provider is used by default, which allows any non-empty user name and password to log in. This is useful for testing purposes. To use other identity providers, or to modify any [token](#), [grant](#), or [session options](#), you must run the master from a configuration file.



NOTE

[Roles](#) need to be assigned to administer the setup with an external user.



NOTE

After making changes to an identity provider you must restart the master service for the changes to take effect:

```
# systemctl restart atomic-openshift-master
```

12.2. IDENTITY PROVIDER PARAMETERS

There are four parameters common to all identity providers:

Parameter	Description
name	The provider name is prefixed to provider user names to form an identity name.

Parameter	Description
challenge	<p>When true, unauthenticated token requests from non-web clients (like the CLI) are sent a WWW-Authenticate challenge header. Not supported by all identity providers.</p> <p>To prevent cross-site request forgery (CSRF) attacks against browser clients Basic authentication challenges are only sent if a X-CSRF-Token header is present on the request. Clients that expect to receive Basic WWW-Authenticate challenges should set this header to a non-empty value.</p>
login	<p>When true, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider. Not supported by all identity providers.</p> <p>If you want users to be sent to a branded page before being redirected to the identity provider's login, then set oauthConfig → alwaysShowProviderSelection: true in the master configuration file. This provider selection page can be customized.</p>
mappingMethod	<p>Defines how new identities are mapped to users when they log in. Enter one of the following values:</p> <p>claim</p> <p>The default value. Provisions a user with the identity's preferred user name. Fails if a user with that user name is already mapped to another identity.</p> <p>lookup</p> <p>Looks up an existing identity, user identity mapping, and user, but does not automatically provision users or identities. This allows cluster administrators to set up identities and users manually, or using an external process.</p> <p>generate</p> <p>Provisions a user with the identity's preferred user name. If a user with the preferred user name is already mapped to an existing identity, a unique user name is generated. For example, myuser2. This method should not be used in combination with external processes that require exact matches between OpenShift Container Platform user names and identity provider user names, such as LDAP group sync.</p> <p>add</p> <p>Provisions a user with the identity's preferred user name. If a user with that user name already exists, the identity is mapped to the existing user, adding to any existing identity mappings for the user. Required when multiple identity providers are configured that identify the same set of users and map to the same user names.</p>

**NOTE**

When adding or changing identity providers, you can map identities from the new provider to existing users by setting the **mappingMethod** parameter to **add**.

12.3. CONFIGURING IDENTITY PROVIDERS

OpenShift Container Platform supports configuring only a single identity provider. However, you can extend the basic authentication for more complex configurations such as [LDAP failover](#).

You can use these parameters to define the identity provider during installation or after installation.

12.3.1. Configuring identity providers with Ansible

For initial [advanced installations](#), the [Deny All](#) identity provider is configured by default, though it can be [overridden during installation](#) using the `openshift_master_identity_providers` parameter, which is configurable in the inventory file. [Session options in the OAuth configuration](#) are also configurable in the inventory file.

Example 12.1. Example identity provider configuration with Ansible

```
# httpasswd auth
openshift_master_identity_providers=[{'name': 'htpasswd_auth', 'login':
'true', 'challenge': 'true', 'kind': 'HTPasswdPasswordIdentityProvider',
'filename': '/etc/origin/master/htpasswd'}]
# Defining htpasswd users
#openshift_master_htpasswd_users={'user1': '<pre-hashed password>',
'user2': '<pre-hashed password>'}
# or
#openshift_master_htpasswd_file=<path to local pre-generated htpasswd
file>

# Allow all auth
#openshift_master_identity_providers=[{'name': 'allow_all', 'login':
'true', 'challenge': 'true', 'kind':
'AllowAllPasswordIdentityProvider'}]

# LDAP auth
#openshift_master_identity_providers=[{'name': 'my_ldap_provider',
'challenge': 'true', 'login': 'true', 'kind':
'LDAPPasswordIdentityProvider', 'attributes': {'id': ['dn'], 'email':
['mail'], 'name': ['cn'], 'preferredUsername': ['uid']}, 'bindDN': '',
'bindPassword': '', 'ca': '', 'insecure': 'false', 'url':
'ldap://ldap.example.com:389/ou=users,dc=example,dc=com?uid'}]
# Configuring the ldap ca certificate ❶
#openshift_master_ldap_ca=<ca text>
# or
#openshift_master_ldap_ca_file=<path to local ca file to use>

# Available variables for configuring certificates for other identity
providers:
#openshift_master_openid_ca
#openshift_master_openid_ca_file
#openshift_master_request_header_ca
#openshift_master_request_header_ca_file
```

❶ If you specify your CA certificate location in the `openshift_master_identity_providers` parameter, do not specify a certificate value in the `openshift_master_ldap_ca` parameter or path in the `openshift_master_ldap_ca_file` parameter.

12.3.2. Configuring identity providers in the master configuration file

You can configure the master host for authentication using your desired identity provider by modifying the [master configuration file](#).

Example 12.2. Example identity provider configuration in the master configuration file

```
...
oauthConfig:
  identityProviders:
  - name: htpasswd_auth
    challenge: true
    login: true
    mappingMethod: "claim"
...
```

When set to the default **claim** value, OAuth will fail if the identity is mapped to a previously-existing user name.

Parameter	Description
claim	The default value. Provisions a user with the identity's preferred user name. Fails if a user with that user name is already mapped to another identity.
lookup	Looks up an existing identity, user identity mapping, and user, but does not automatically provision users or identities. This allows cluster administrators to set up identities and users manually, or using an external process.
generate	Provisions a user with the identity's preferred user name. If a user with the preferred user name is already mapped to an existing identity, a unique user name is generated. For example, myuser2 . This method should not be used in combination with external processes that require exact matches between OpenShift Container Platform user names and identity provider user names, such as LDAP group sync.
add	Provisions a user with the identity's preferred user name. If a user with that user name already exists, the identity is mapped to the existing user, adding to any existing identity mappings for the user. Required when multiple identity providers are configured that identify the same set of users and map to the same user names.

12.3.3. Allow all

Set **AllowAllPasswordIdentityProvider** in the **identityProviders** stanza to allow any non-empty user name and password to log in.

Example 12.3. Master Configuration Using AllowAllPasswordIdentityProvider

```
oauthConfig:
...
identityProviders:
- name: my_allow_provider ❶
  challenge: true ❷
  login: true ❸
  mappingMethod: claim ❹
```

```

provider:
  apiVersion: v1
  kind: AllowAllPasswordIdentityProvider

```

- 1 This provider name is prefixed to provider user names to form an identity name.
- 2 When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider.
- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).

12.3.4. Deny all

Set **DenyAllPasswordIdentityProvider** in the **identityProviders** stanza to deny access for all user names and passwords.

Example 12.4. Master Configuration Using DenyAllPasswordIdentityProvider

```

oauthConfig:
  ...
  identityProviders:
  - name: my_deny_provider 1
    challenge: true 2
    login: true 3
    mappingMethod: claim 4
    provider:
      apiVersion: v1
      kind: DenyAllPasswordIdentityProvider

```

- 1 This provider name is prefixed to provider user names to form an identity name.
- 2 When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider.
- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).

12.3.5. HTTPasswd

Set **HTTPasswdPasswordIdentityProvider** in the **identityProviders** stanza to validate user names and passwords against a flat file generated using [httpasswd](#).

**NOTE**

The **htpasswd** utility is in the **httpd-tools** package:

```
# yum install httpd-tools
```

OpenShift Container Platform supports the Bcrypt, SHA-1, and MD5 cryptographic hash functions, and MD5 is the default for **htpasswd**. Plaintext, encrypted text, and other hash functions are not currently supported.

The flat file is reread if its modification time changes, without requiring a server restart.

To use the **htpasswd** command:

- To create a flat file with a user name and hashed password, run:

```
$ htpasswd -c </path/to/users.htpasswd> <user_name>
```

Then, enter and confirm a clear-text password for the user. The command generates a hashed version of the password.

For example:

```
htpasswd -c users.htpasswd user1
New password:
Re-type new password:
Adding password for user user1
```

**NOTE**

You can include the **-b** option to supply the password on the command line:

```
$ htpasswd -c -b <user_name> <password>
```

For example:

```
$ htpasswd -c -b file user1 MyPassword!
Adding password for user user1
```

- To add or update a login to the file, run:

```
$ htpasswd </path/to/users.htpasswd> <user_name>
```

- To remove a login from the file, run:

```
$ htpasswd -D </path/to/users.htpasswd> <user_name>
```

Example 12.5. Master Configuration Using HTTPasswdPasswordIdentityProvider

```
oauthConfig:
  ...
```



```
identityProviders:
- name: my_htpasswd_provider ❶
  challenge: true ❷
  login: true ❸
  mappingMethod: claim ❹
  provider:
    apiVersion: v1
    kind: HTTPasswdPasswordIdentityProvider
    file: /path/to/users.htpasswd ❺
```

- ❶ This provider name is prefixed to provider user names to form an identity name.
- ❷ When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider.
- ❸ When **true**, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider.
- ❹ Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- ❺ File generated using [htpasswd](#).

12.3.6. Keystone

Set **KeystonePasswordIdentityProvider** in the **identityProviders** stanza to validate user names and passwords against an OpenStack Keystone v3 server. This enables shared authentication with an OpenStack server configured to store users in an internal Keystone database.

Example 12.6. Master Configuration Using KeystonePasswordIdentityProvider

```
oauthConfig:
...
identityProviders:
- name: my_keystone_provider ❶
  challenge: true ❷
  login: true ❸
  mappingMethod: claim ❹
  provider:
    apiVersion: v1
    kind: KeystonePasswordIdentityProvider
    domainName: default ❺
    url: http://keystone.example.com:5000 ❻
    ca: ca.pem ❼
    certFile: keystone.pem ❽
    keyFile: keystonekey.pem ❾
```

- ❶ This provider name is prefixed to provider user names to form an identity name.
- ❷ When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider.

- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 Keystone domain name. In Keystone, usernames are domain-specific. Only a single domain is supported.
- 6 The URL to use to connect to the Keystone server (required).
- 7 Optional: Certificate bundle to use to validate server certificates for the configured URL.
- 8 Optional: Client certificate to present when making requests to the configured URL.
- 9 Key for the client certificate. Required if **certFile** is specified.

12.3.7. LDAP authentication

Set **LDAPPasswordIdentityProvider** in the **identityProviders** stanza to validate user names and passwords against an LDAPv3 server, using simple bind authentication.



NOTE

If you require failover for your LDAP server, instead of following these steps, extend the basic authentication method by [configuring SSSD for LDAP failover](#).

During authentication, the LDAP directory is searched for an entry that matches the provided user name. If a single unique match is found, a simple bind is attempted using the distinguished name (DN) of the entry plus the provided password.

These are the steps taken:

1. Generate a search filter by combining the attribute and filter in the configured **url** with the user-provided user name.
2. Search the directory using the generated filter. If the search does not return exactly one entry, deny access.
3. Attempt to bind to the LDAP server using the DN of the entry retrieved from the search, and the user-provided password.
4. If the bind is unsuccessful, deny access.
5. If the bind is successful, build an identity using the configured attributes as the identity, email address, display name, and preferred user name.

The configured **url** is an RFC 2255 URL, which specifies the LDAP host and search parameters to use. The syntax of the URL is:

```
ldap://host:port/basedn?attribute?scope?filter
```

For the above example:

URL Component	Description
ldap	For regular LDAP, use the string ldap . For secure LDAP (LDAPS), use ldaps instead.
host:port	The name and port of the LDAP server. Defaults to localhost:389 for ldap and localhost:636 for LDAPS.
basedn	The DN of the branch of the directory where all searches should start from. At the very least, this must be the top of your directory tree, but it could also specify a subtree in the directory.
attribute	The attribute to search for. Although RFC 2255 allows a comma-separated list of attributes, only the first attribute will be used, no matter how many are provided. If no attributes are provided, the default is to use uid . It is recommended to choose an attribute that will be unique across all entries in the subtree you will be using.
scope	The scope of the search. Can be either either one or sub . If the scope is not provided, the default is to use a scope of sub .
filter	A valid LDAP search filter. If not provided, defaults to (objectClass=*)

When doing searches, the attribute, filter, and provided user name are combined to create a search filter that looks like:

```
(<filter>(<attribute>=<username>))
```

For example, consider a URL of:

```
ldap://ldap.example.com/o=Acme?cn?sub?(enabled=true)
```

When a client attempts to connect using a user name of **bob**, the resulting search filter will be **(&(enabled=true)(cn=bob))**.

If the LDAP directory requires authentication to search, specify a **bindDN** and **bindPassword** to use to perform the entry search.

Master Configuration Using LDAPPasswordIdentityProvider

```
oauthConfig:
  ...
  identityProviders:
    - name: "my_ldap_provider" ①
      challenge: true ②
      login: true ③
      mappingMethod: claim ④
      provider:
        apiVersion: v1
        kind: LDAPPasswordIdentityProvider
        attributes:
```

```

id: 5
- dn
email: 6
- mail
name: 7
- cn
preferredUsername: 8
- uid
bindDN: "" 9
bindPassword: "" 10
ca: my-ldap-ca-bundle.crt 11
insecure: false 12
url: "ldap://ldap.example.com/ou=users,dc=acme,dc=com?uid" 13

```

- 1 This provider name is prefixed to the returned user ID to form an identity name.
- 2 When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider.
- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 List of attributes to use as the identity. First non-empty attribute is used. At least one attribute is required. If none of the listed attribute have a value, authentication fails.
- 6 List of attributes to use as the email address. First non-empty attribute is used.
- 7 List of attributes to use as the display name. First non-empty attribute is used.
- 8 List of attributes to use as the preferred user name when provisioning a user for this identity. First non-empty attribute is used.
- 9 Optional DN to use to bind during the search phase.
- 10 Optional password to use to bind during the search phase. This value may also be provided in an [environment variable](#), [external file](#), or [encrypted file](#).
- 11 Certificate bundle to use to validate server certificates for the configured URL. If empty, system trusted roots are used. Only applies if **insecure: false**.
- 12 When **true**, no TLS connection is made to the server. When **false**, **ldaps://** URLs connect using TLS, and **ldap://** URLs are upgraded to TLS.
- 13 An RFC 2255 URL which specifies the LDAP host and search parameters to use, [as described above](#).



NOTE

To whitelist users for an LDAP integration, use the **lookup** mapping method. Before a login from LDAP would be allowed, a cluster administrator must create an identity and user object for each LDAP user.

12.3.8. Basic authentication (remote)

Basic Authentication is a generic backend integration mechanism that allows users to log in to OpenShift Container Platform with credentials validated against a remote identity provider.

Because basic authentication is generic, you can use this identity provider for advanced authentication configurations. You can configure [LDAP failover](#) or use the [containerized basic authentication](#) repository as a starting point for another advanced remote basic authentication configuration.

CAUTION

Basic authentication must use an HTTPS connection to the remote server to prevent potential snooping of the user ID and password and man-in-the-middle attacks.

With **BasicAuthPasswordIdentityProvider** configured, users send their user name and password to OpenShift Container Platform, which then validates those credentials against a remote server by making a server-to-server request, passing the credentials as a Basic Auth header. This requires users to send their credentials to OpenShift Container Platform during login.

Set **BasicAuthPasswordIdentityProvider** in the **identityProviders** stanza to validate user names and passwords against a remote server using a server-to-server Basic authentication request. User names and passwords are validated against a remote URL that is protected by Basic authentication and returns JSON.

A **401** response indicates failed authentication.

A non-**200** status, or the presence of a non-empty "error" key, indicates an error:

```
{"error": "Error message"}
```

A **200** status with a **sub** (subject) key indicates success:

```
{"sub": "userid"} 1
```

1 The subject must be unique to the authenticated user and must not be able to be modified.

A successful response may optionally provide additional data, such as:

- A display name using the **name** key. For example:

```
{"sub": "userid", "name": "User Name", ...}
```

- An email address using the **email** key. For example:

```
{"sub": "userid", "email": "user@example.com", ...}
```

- A preferred user name using the **preferred_username** key. This is useful when the unique, unchangeable subject is a database key or UID, and a more human-readable name exists. This is used as a hint when provisioning the OpenShift Container Platform user for the authenticated identity. For example:

```
{"sub": "014fbff9a07c", "preferred_username": "bob", ...}
```

Example 12.7. Master Configuration Using BasicAuthPasswordIdentityProvider

```

oauthConfig:
  ...
  identityProviders:
  - name: my_remote_basic_auth_provider ❶
    challenge: true ❷
    login: true ❸
    mappingMethod: claim ❹
    provider:
      apiVersion: v1
      kind: BasicAuthPasswordIdentityProvider
      url: https://www.example.com/remote-idp ❺
      ca: /path/to/ca.file ❻
      certFile: /path/to/client.crt ❼
      keyFile: /path/to/client.key ❽

```

- ❶ This provider name is prefixed to the returned user ID to form an identity name.
- ❷ When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider.
- ❸ When **true**, unauthenticated token requests from web clients (like the web console) are redirected to a login page backed by this provider.
- ❹ Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- ❺ URL accepting credentials in Basic authentication headers.
- ❻ Optional: Certificate bundle to use to validate server certificates for the configured URL.
- ❼ Optional: Client certificate to present when making requests to the configured URL.
- ❽ Key for the client certificate. Required if **certFile** is specified.

12.3.9. Request header

Set **RequestHeaderIdentityProvider** in the **identityProviders** stanza to identify users from request header values, such as **X-Remote-User**. It is typically used in combination with an authenticating proxy, which sets the request header value. This is similar to how [the remote user plug-in in OpenShift Enterprise 2](#) allowed administrators to provide Kerberos, LDAP, and many other forms of enterprise authentication.

You can also use the request header identity provider for advanced configurations such as [SAML authentication](#).

For users to authenticate using this identity provider, they must access <https://<master>/oauth/authorize> via an authenticating proxy. You can either proxy the entire master API server so that all access goes through the proxy, or you can configure the OAuth server to redirect unauthenticated requests to the proxy.

To redirect unauthenticated requests from clients expecting login flows:

1. Set the **login** parameter to **true**.
2. Set the **provider.loginURL** parameter to the proxy URL to send those clients to.

To redirect unauthenticated requests from clients expecting **WWW-Authenticate** challenges:

1. Set the **challenge** parameter to **true**.
2. Set the **provider.challengeURL** parameter to the proxy URL to send those clients to.

The **provider.challengeURL** and **provider.loginURL** parameters can include the following tokens in the query portion of the URL:

- **\${url}** is replaced with the current URL, escaped to be safe in a query parameter.
For example: **https://www.example.com/sso-login?then=\${url}**
- **\${query}** is replaced with the current query string, unescaped.
For example: **https://www.example.com/auth-proxy/oauth/authorize?\${query}**



WARNING

If you expect unauthenticated requests to reach the OAuth server, a **clientCA** parameter should be set for this identity provider, so that incoming requests are checked for a valid client certificate before the request's headers are checked for a user name. Otherwise, any direct request to the OAuth server can impersonate any identity from this provider, merely by setting a request header.

Example 12.8. Master Configuration Using RequestHeaderIdentityProvider

```
oauthConfig:
  ...
  identityProviders:
    - name: my_request_header_provider ❶
      challenge: true ❷
      login: true ❸
      mappingMethod: claim ❹
      provider:
        apiVersion: v1
        kind: RequestHeaderIdentityProvider
        challengeURL: "https://www.example.com/challenging-
proxy/oauth/authorize?${query}" ❺
        loginURL: "https://www.example.com/login-proxy/oauth/authorize?
${query}" ❻
        clientCA: /path/to/client-ca.file ❼
        clientCommonNames: ❽
        - my-auth-proxy
        headers: ❾
```

```

- X-Remote-User
- SSO-User
emailHeaders: 10
- X-Remote-User-Email
nameHeaders: 11
- X-Remote-User-Display-Name
preferredUsernameHeaders: 12
- X-Remote-User-Login

```

- 1 This provider name is prefixed to the user name in the request header to form an identity name.
- 2 **RequestHeaderIdentityProvider** can only respond to clients that request **WWW-Authenticate** challenges by redirecting to a configured **challengeURL**. The configured URL should respond with a **WWW-Authenticate** challenge.
- 3 **RequestHeaderIdentityProvider** can only respond to clients requesting a login flow by redirecting to a configured **loginURL**. The configured URL should respond with a login flow.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 Optional: URL to redirect unauthenticated `/oauth/authorize` requests to, for clients which expect interactive logins. `${url}` is replaced with the current URL, escaped to be safe in a query parameter. `${query}` is replaced with the current query string.
- 6 Optional: URL to redirect unauthenticated `/oauth/authorize` requests to, for clients which expect **WWW-Authenticate** challenges. `${url}` is replaced with the current URL, escaped to be safe in a query parameter. `${query}` is replaced with the current query string.
- 7 Optional: PEM-encoded certificate bundle. If set, a valid client certificate must be presented and validated against the certificate authorities in the specified file before the request headers are checked for user names.
- 8 Optional: list of common names (**cn**). If set, a valid client certificate with a Common Name (**cn**) in the specified list must be presented before the request headers are checked for user names. If empty, any Common Name is allowed. Can only be used in combination with **clientCA**.
- 9 Header names to check, in order, for the user identity. The first header containing a value is used as the identity. Required, case-insensitive.
- 10 Header names to check, in order, for an email address. The first header containing a value is used as the email address. Optional, case-insensitive.
- 11 Header names to check, in order, for a display name. The first header containing a value is used as the display name. Optional, case-insensitive.
- 12 Header names to check, in order, for a preferred user name, if different than the immutable identity determined from the headers specified in **headers**. The first header containing a value is used as the preferred user name when provisioning. Optional, case-insensitive.

Example 12.9. Apache Authentication Using RequestHeaderIdentityProvider

This example configures an authentication proxy on the same host as the master. Having the proxy and master on the same host is merely a convenience and may not be suitable for your environment. For example, if you were already [running a router](#) on the master, port 443 would not be available.

It is also important to note that while this reference configuration uses Apache's **mod_auth_form**, it is by no means required and other proxies can easily be used if the following requirements are met:

1. Block the **X-Remote-User** header from client requests to prevent spoofing.
2. Enforce client certificate authentication in the **RequestHeaderIdentityProvider** configuration.
3. Require the **X-Csrf-Token** header be set for all authentication request using the challenge flow.
4. Only the `/oauth/authorize` endpoint should be proxied, and redirects should not be rewritten to allow the backend server to send the client to the correct location.

Installing the Prerequisites

The **mod_auth_form** module is shipped as part of the **mod_session** package that is found in the [Optional channel](#):

```
# yum install -y httpd mod_ssl mod_session apr-util-openssl
```

Generate a CA for validating requests that submit the trusted header. This CA should be used as the file name for **clientCA** in the [master's identity provider configuration](#).

```
# oadm ca create-signer-cert \
  --cert='/etc/origin/master/proxyca.crt' \
  --key='/etc/origin/master/proxyca.key' \
  --name='openshift-proxy-signer@1432232228' \
  --serial='/etc/origin/master/proxyca.serial.txt'
```

Generate a client certificate for the proxy. This can be done using any x509 certificate tooling. For convenience, the **oadm** CLI can be used:

```
# oadm create-api-client-config \
  --certificate-authority='/etc/origin/master/proxyca.crt' \
  --client-dir='/etc/origin/master/proxy' \
  --signer-cert='/etc/origin/master/proxyca.crt' \
  --signer-key='/etc/origin/master/proxyca.key' \
  --signer-serial='/etc/origin/master/proxyca.serial.txt' \
  --user='system:proxy' 1

# pushd /etc/origin/master
# cp master.server.crt /etc/pki/tls/certs/localhost.crt 2
# cp master.server.key /etc/pki/tls/private/localhost.key
# cp ca.crt /etc/pki/CA/certs/ca.crt
# cat proxy/system\:proxy.crt \
  proxy/system\:proxy.key > \
  /etc/pki/tls/certs/authproxy.pem
# popd
```

- 1 The user name can be anything, however it is useful to give it a descriptive name as it will appear in logs.
- 2 When running the authentication proxy on a different host name than the master, it is important to generate a certificate that matches the host name instead of using the default master certificate as shown above. The value for **masterPublicURL** in the */etc/origin/master/master-config.yaml* file must be included in the **X509v3 Subject Alternative Name** in the certificate that is specified for **SSLCertificateFile**. If a new certificate needs to be created, the **oadm ca create-server-cert** command can be used.

Configuring Apache

Unlike OpenShift Enterprise 2, this proxy does not need to reside on the same host as the master. It uses a client certificate to connect to the master, which is configured to trust the **X-Remote-User** header.

Configure Apache per the following:

```
LoadModule auth_form_module modules/mod_auth_form.so
LoadModule session_module modules/mod_session.so
LoadModule request_module modules/mod_request.so

# Nothing needs to be served over HTTP. This virtual host simply
# redirects to
# HTTPS.
<VirtualHost *:80>
    DocumentRoot /var/www/html
    RewriteEngine On
    RewriteRule ^(.*)$ https://%{HTTP_HOST}$1 [R,L]
</VirtualHost>

<VirtualHost *:443>
    # This needs to match the certificates you generated. See the CN and
    X509v3
    # Subject Alternative Name in the output of:
    # openssl x509 -text -in /etc/pki/tls/certs/localhost.crt
    ServerName www.example.com

    DocumentRoot /var/www/html
    SSLEngine on
    SSLCertificateFile /etc/pki/tls/certs/localhost.crt
    SSLCertificateKeyFile /etc/pki/tls/private/localhost.key
    SSLCACertificateFile /etc/pki/CA/certs/ca.crt

    SSLProxyEngine on
    SSLProxyCACertificateFile /etc/pki/CA/certs/ca.crt
    # It's critical to enforce client certificates on the Master.
    Otherwise
    # requests could spoof the X-Remote-User header by accessing the
    Master's
    # /oauth/authorize endpoint directly.
    SSLProxyMachineCertificateFile /etc/pki/tls/certs/authproxy.pem

    # Send all requests to the console
    RewriteEngine On
```

```

RewriteRule      ^/console(.*)$      https://%
{HTTP_HOST}:8443/console$1 [R,L]

# In order to using the challenging-proxy an X-Csrftoken must be
present.
RewriteCond %{REQUEST_URI} ^/challenging-proxy
RewriteCond %{HTTP:X-Csrftoken} ^$ [NC]
RewriteRule ^.* - [F,L]

<Location /challenging-proxy/oauth/authorize>
# Insert your backend server name/ip here.
ProxyPass https://[MASTER]:8443/oauth/authorize
AuthType basic
</Location>

<Location /login-proxy/oauth/authorize>
# Insert your backend server name/ip here.
ProxyPass https://[MASTER]:8443/oauth/authorize

# mod_auth_form providers are implemented by mod_authn_dbm,
mod_authn_file,
# mod_authn_dbd, mod_authnz_ldap and mod_authn_socache.
AuthFormProvider file
AuthType form
AuthName openshift
ErrorDocument 401 /login.html
</Location>

<ProxyMatch /oauth/authorize>
AuthUserFile /etc/origin/master/htpasswd
AuthName openshift
Require valid-user
RequestHeader set X-Remote-User %{REMOTE_USER}s env=REMOTE_USER

# For ldap:
# AuthBasicProvider ldap
# AuthLDAPURL "ldap://ldap.example.com:389/ou=People,dc=my-
domain,dc=com?uid?sub?(objectClass=*)"

# It's possible to remove the mod_auth_form usage and replace it
with
# something like mod_auth_kerb, mod_auth_gssapi or even
mod_auth_mellon.
# The former would be able to support both the login and challenge
flows
# from the Master. Mellon would likely only support the login flow.

# For Kerberos
# yum install mod_auth_gssapi
# AuthType GSSAPI
# GssapiCredStore keytab:/etc/httpd.keytab
</ProxyMatch>

</VirtualHost>

RequestHeader unset X-Remote-User

```

Additional `mod_auth_form` Requirements

A sample login page is available from the [openshift_extras](#) repository. This file should be placed in the **DocumentRoot** location (`/var/www/html` by default).

Creating Users

At this point, you can create the users in the system Apache is using to store accounts information. In this example, file-backed authentication is used:

```
# yum -y install httpd-tools
# touch /etc/origin/master/htpasswd
# htpasswd /etc/origin/master/htpasswd <user_name>
```

Configuring the Master

The **identityProviders** stanza in the `/etc/origin/master/master-config.yaml` file must be updated as well:

```
identityProviders:
- name: requestheader
  challenge: true
  login: true
  provider:
    apiVersion: v1
    kind: RequestHeaderIdentityProvider
    challengeURL: "https://[MASTER]/challenging-proxy/oauth/authorize?
${query}"
    loginURL: "https://[MASTER]/login-proxy/oauth/authorize?${query}"
    clientCA: /etc/origin/master/proxyca.crt
    headers:
    - X-Remote-User
```

Restarting Services

Finally, restart the following services:

```
# systemctl restart httpd
# systemctl restart atomic-openshift-master
```

Verifying the Configuration

1. Test by bypassing the proxy. You should be able to request a token if you supply the correct client certificate and header:

```
# curl -L -k -H "X-Remote-User: joe" \
  --cert /etc/pki/tls/certs/authproxy.pem \
  https://[MASTER]:8443/oauth/token/request
```

2. If you do not supply the client certificate, the request should be denied:

```
# curl -L -k -H "X-Remote-User: joe" \
  https://[MASTER]:8443/oauth/token/request
```

3. This should show a redirect to the configured **challengeURL** (with additional query parameters):

```
# curl -k -v -H 'X-Csrf-Token: 1' \
  '<masterPublicURL>/oauth/authorize?client_id=openshift-
  challenging-client&response_type=token'
```

4. This should show a 401 response with a **WWW-Authenticate** basic challenge:


```
# curl -k -v -H 'X-Csrf-Token: 1' \
  '<redirected challengeURL from step 3 +query>'
```

5. This should show a redirect with an access token:

```
# curl -k -v -u <your_user>:<your_password> \
  -H 'X-Csrf-Token: 1' '<redirected_challengeURL_from_step_3
  +query>'
```

12.3.10. GitHub

Set **GitHubIdentityProvider** in the **identityProviders** stanza to use [GitHub](#) as an identity provider, using the [OAuth integration](#).



NOTE

Using GitHub as an identity provider requires users to get a token using **<master>/oauth/token/request** to use with command-line tools.



WARNING

Using GitHub as an identity provider allows any GitHub user to authenticate to your server. You can limit authentication to members of specific GitHub organizations with the **organizations** configuration attribute, as shown below.

Example 12.10. Master Configuration Using GitHubIdentityProvider

```
oauthConfig:
  ...
  identityProviders:
  - name: github ①
    challenge: false ②
    login: true ③
    mappingMethod: claim ④
    provider:
      apiVersion: v1
      kind: GitHubIdentityProvider
```

```

    clientID: ... 5
    clientSecret: ... 6
    organizations: 7
    - myorganization1
    - myorganization2

```

- 1 This provider name is prefixed to the GitHub numeric user ID to form an identity name. It is also used to build the callback URL.
- 2 **GitHubIdentityProvider** cannot be used to send **WWW-Authenticate** challenges.
- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to GitHub to log in.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 The client ID of a [registered GitHub OAuth application](#). The application must be configured with a callback URL of `<master>/oauth2callback/<identityProviderName>`.
- 6 The client secret issued by GitHub. This value may also be provided in an [environment variable](#), [external file](#), or [encrypted file](#).
- 7 Optional list of organizations. If specified, only GitHub users that are members of at least one of the listed organizations will be allowed to log in. If the GitHub OAuth application configured in **clientID** is not owned by the organization, an organization owner must grant third-party access in order to use this option. This can be done during the first GitHub login by the organization's administrator, or from the GitHub organization settings.

12.3.11. GitLab

Set **GitLabIdentityProvider** in the **identityProviders** stanza to use [GitLab.com](#) or any other GitLab instance as an identity provider, using the [OAuth integration](#). The OAuth provider feature requires GitLab version 7.7.0 or higher.

Example 12.11. Master Configuration Using GitLabIdentityProvider

```

oauthConfig:
  ...
  identityProviders:
  - name: gitlab 1
    challenge: true 2
    login: true 3
    mappingMethod: claim 4
    provider:
      apiVersion: v1
      kind: GitLabIdentityProvider
      url: ... 5
      clientID: ... 6
      clientSecret: ... 7
      ca: ... 8

```

- 1 This provider name is prefixed to the GitLab numeric user ID to form an identity name. It is also used to build the callback URL.
- 2 When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider. This uses the [Resource Owner Password Credentials](#) grant flow to obtain an access token from GitLab.
- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to GitLab to log in.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 The host URL of a GitLab OAuth provider. This could either be **https://gitlab.com/** or any other self hosted instance of GitLab.
- 6 The client ID of a [registered GitLab OAuth application](#). The application must be configured with a callback URL of **<master>/oauth2callback/<identityProviderName>**.
- 7 The client secret issued by GitLab. This value may also be provided in an [environment variable, external file, or encrypted file](#).
- 8 CA is an optional trusted certificate authority bundle to use when making requests to the GitLab instance. If empty, the default system roots are used.

12.3.12. Google

Set **GoogleIdentityProvider** in the **identityProviders** stanza to use Google as an identity provider, using [Google's OpenID Connect integration](#).



NOTE

Using Google as an identity provider requires users to get a token using **<master>/oauth/token/request** to use with command-line tools.



WARNING

Using Google as an identity provider allows any Google user to authenticate to your server. You can limit authentication to members of a specific hosted domain with the **hostedDomain** configuration attribute, as shown below.

Example 12.12. Master Configuration Using GoogleIdentityProvider

```
oauthConfig:
  ...
identityProviders:
  - name: google 1
```

```

challenge: false 2
login: true 3
mappingMethod: claim 4
provider:
  apiVersion: v1
  kind: GoogleIdentityProvider
  clientId: ... 5
  clientSecret: ... 6
  hostedDomain: "" 7

```

- 1 This provider name is prefixed to the Google numeric user ID to form an identity name. It is also used to build the redirect URL.
- 2 **GoogleIdentityProvider** cannot be used to send **WWW-Authenticate** challenges.
- 3 When **true**, unauthenticated token requests from web clients (like the web console) are redirected to Google to log in.
- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 The client ID of a [registered Google project](#). The project must be configured with a redirect URI of `<master>/oauth2callback/<identityProviderName>`.
- 6 The client secret issued by Google. This value may also be provided in an [environment variable](#), [external file](#), or [encrypted file](#).
- 7 Optional [hosted domain](#) to restrict sign-in accounts to. If empty, any Google account is allowed to authenticate.

12.3.13. OpenID connect

Set **OpenIDIdentityProvider** in the **identityProviders** stanza to integrate with an OpenID Connect identity provider using an [Authorization Code Flow](#).



NOTE

ID Token and **UserInfo** decryptions are not supported.

By default, the **openid** scope is requested. If required, extra scopes can be specified in the **extraScopes** field.

Claims are read from the JWT **id_token** returned from the OpenID identity provider and, if specified, from the JSON returned by the **UserInfo** URL.

At least one claim must be configured to use as the user's identity. The [standard identity claim](#) is **sub**.

You can also indicate which claims to use as the user's preferred user name, display name, and email address. If multiple claims are specified, the first one with a non-empty value is used. The [standard claims](#) are:

sub	The user identity.
preferred_username	The preferred user name when provisioning a user.
email	Email address.
name	Display name.



NOTE

Using an OpenID Connect identity provider requires users to get a token using **<master>/oauth/token/request** to use with command-line tools.

Example 12.13. Standard Master Configuration Using OpenIDIdentityProvider

```

oauthConfig:
  ...
  identityProviders:
    - name: my_openid_connect ❶
      challenge: true ❷
      login: true ❸
      mappingMethod: claim ❹
      provider:
        apiVersion: v1
        kind: OpenIDIdentityProvider
        clientID: ... ❺
        clientSecret: ... ❻
        claims:
          id:
            - sub ❼
          preferredUsername:
            - preferred_username
          name:
            - name
          email:
            - email
        urls:
          authorize: https://myidp.example.com/oauth2/authorize ❽
          token: https://myidp.example.com/oauth2/token ❾

```

- ❶ This provider name is prefixed to the value of the identity claim to form an identity name. It is also used to build the redirect URL.
- ❷ When **true**, unauthenticated token requests from non-web clients (like the CLI) are sent a **WWW-Authenticate** challenge header for this provider. This requires the OpenID provider to support the [Resource Owner Password Credentials](#) grant flow.
- ❸ When **true**, unauthenticated token requests from web clients (like the web console) are redirected to the authorize URL to log in.

- 4 Controls how mappings are established between this provider's identities and user objects, [as described above](#).
- 5 The client ID of a client registered with the OpenID provider. The client must be allowed to redirect to `<master>/oauth2callback/<identityProviderName>`.
- 6 The client secret. This value may also be provided in an [environment variable, external file, or encrypted file](#).
- 7 Use the value of the `sub` claim in the returned `id_token` as the user's identity.
- 8 [Authorization Endpoint](#) described in the OpenID spec. Must use `https`.
- 9 [Token Endpoint](#) described in the OpenID spec. Must use `https`.

A custom certificate bundle, extra scopes, extra authorization request parameters, and `userInfo` URL can also be specified:

Example 12.14. Full Master Configuration Using OpenIDIdentityProvider

```

oauthConfig:
  ...
  identityProviders:
  - name: my_openid_connect
    challenge: false
    login: true
    mappingMethod: claim
    provider:
      apiVersion: v1
      kind: OpenIDIdentityProvider
      clientID: ...
      clientSecret: ...
      ca: my-openid-ca-bundle.crt 1
      extraScopes: 2
      - email
      - profile
      extraAuthorizeParameters: 3
      include_granted_scopes: "true"
      claims:
        id: 4
        - custom_id_claim
        - sub
        preferredUsername: 5
        - preferred_username
        - email
        name: 6
        - nickname
        - given_name
        - name
        email: 7
        - custom_email_claim
        - email
      urls:

```

```

authorize: https://myidp.example.com/oauth2/authorize
token: https://myidp.example.com/oauth2/token
userInfo: https://myidp.example.com/oauth2/userinfo 8

```

- 1 Certificate bundle to use to validate server certificates for the configured URLs. If empty, system trusted roots are used.
- 2 Optional list of scopes to request, in addition to the **openid** scope, during the authorization token request.
- 3 Optional map of extra parameters to add to the authorization token request.
- 4 List of claims to use as the identity. First non-empty claim is used. At least one claim is required. If none of the listed claims have a value, authentication fails.
- 5 List of claims to use as the preferred user name when provisioning a user for this identity. First non-empty claim is used.
- 6 List of claims to use as the display name. First non-empty claim is used.
- 7 List of claims to use as the email address. First non-empty claim is used.
- 8 [UserInfo Endpoint](#) described in the OpenID spec. Must use **https**.

12.4. TOKEN OPTIONS

The OAuth server generates two kinds of tokens:

Access tokens	Longer-lived tokens that grant access to the API.
Authorize codes	Short-lived tokens whose only use is to be exchanged for an access token.

Use the **tokenConfig** stanza to set token options:

Example 12.15. Master Configuration Token Options

```

oauthConfig:
  ...
  tokenConfig:
    accessTokenMaxAgeSeconds: 86400 1
    authorizeTokenMaxAgeSeconds: 300 2

```

- 1 Set **accessTokenMaxAgeSeconds** to control the lifetime of access tokens. The default lifetime is 24 hours.
- 2 Set **authorizeTokenMaxAgeSeconds** to control the lifetime of authorize codes. The default lifetime is five minutes.

12.5. GRANT OPTIONS

When the OAuth server receives token requests for a client to which the user has not previously granted permission, the action that the OAuth server takes is dependent on the OAuth client's grant strategy.

When the OAuth client requesting token does not provide its own grant strategy, the server-wide default strategy is used. To configure the default strategy, set the **method** value in the **grantConfig** stanza. Valid values for **method** are:

auto	Auto-approve the grant and retry the request.
prompt	Prompt the user to approve or deny the grant.
deny	Auto-deny the grant and return a failure error to the client.

Example 12.16. Master Configuration Grant Options

```
oauthConfig:
  ...
  grantConfig:
    method: auto
```

12.6. SESSION OPTIONS

The OAuth server uses a signed and encrypted cookie-based session during login and redirect flows.

Use the **sessionConfig** stanza to set session options:

Example 12.17. Master Configuration Session Options

```
oauthConfig:
  ...
  sessionConfig:
    sessionMaxAgeSeconds: 300 1
    sessionName: ssn 2
    sessionSecretsFile: "... " 3
```

- 1** Controls the maximum age of a session; sessions auto-expire once a token request is complete. If [auto-grant](#) is not enabled, sessions must last as long as the user is expected to take to approve or reject a client authorization request.
- 2** Name of the cookie used to store the session.
- 3** File name containing serialized **SessionSecrets** object. If empty, a random signing and encryption secret is generated at each server start.

If no **sessionSecretsFile** is specified, a random signing and encryption secret is generated at each start of the master server. This means that any logins in progress will have their sessions invalidated if

the master is restarted. It also means that if multiple masters are configured, they will not be able to decode sessions generated by one of the other masters.

To specify the signing and encryption secret to use, specify a **sessionSecretsFile**. This allows you separate secret values from the configuration file and keep the configuration file distributable, for example for debugging purposes.

Multiple secrets can be specified in the **sessionSecretsFile** to enable rotation. New sessions are signed and encrypted using the first secret in the list. Existing sessions are decrypted and authenticated by each secret until one succeeds.

Example 12.18. Session Secret Configuration:

```
apiVersion: v1
kind: SessionSecrets
secrets: ❶
- authentication: "..." ❷
  encryption: "..." ❸
- authentication: "..."
  encryption: "..."
...
```

- ❶ List of secrets used to authenticate and encrypt cookie sessions. At least one secret must be specified. Each secret must set an authentication and encryption secret.
- ❷ Signing secret, used to authenticate sessions using HMAC. Recommended to use a secret with 32 or 64 bytes.
- ❸ Encrypting secret, used to encrypt sessions. Must be 16, 24, or 32 characters long, to select AES-128, AES-192, or AES-256.

12.7. PREVENTING CLI VERSION MISMATCH WITH USER AGENT

OpenShift Container Platform implements a user agent that can be used to prevent an application developer's CLI accessing the OpenShift Container Platform API.

User agents for the OpenShift Container Platform CLI are constructed from a set of values within OpenShift Container Platform:

```
<command>/<version> (<platform>/<architecture>) <client>/<git_commit>
```

So, for example, when:

- <command> = **oc**
- <version> = The client version. For example, **v3.3.0**. Requests made against the Kubernetes API at **/api** receive the Kubernetes version, while requests made against the OpenShift Container Platform API at **/oapi** receive the OpenShift Container Platform version (as specified by **oc version**)
- <platform> = **linux**

- <architecture> = **amd64**
- <client> = **openshift**, or **kubernetes** depending on if the request is made against the Kubernetes API at **/api**, or the OpenShift Container Platform API at **/oapi**
- <git_commit> = The Git commit of the client version (for example, **f034127**)

the user agent will be:

```
oc/v3.3.0 (linux/amd64) openshift/f034127
```

As an OpenShift Container Platform administrator, you can prevent clients from accessing the API with the **userAgentMatching** configuration setting of a master configuration. So, if a client is using a particular library or binary, they will be prevented from accessing the API.

The following user agent example denies the Kubernetes 1.2 client binary, OpenShift Origin 1.1.3 binary, and the POST and PUT **httpVerbs**:

```
policyConfig:
  userAgentMatchingConfig:
    defaultRejectionMessage: "Your client is too old. Go to
https://example.org to update it."
    deniedClients:
      - regex: '\w+/v(?:1\.1\.1|1\.0\.1) \(.+/.+\) openshift/\w{7}'
      - regex: '\w+/v(?:1\.1\.3) \(.+/.+\) openshift/\w{7}'
    httpVerbs:
      - POST
      - PUT
      - regex: '\w+/v1\.2\.0 \(.+/.+\) kubernetes/\w{7}'
    httpVerbs:
      - POST
      - PUT
    requiredClients: null
```

Administrators can also deny clients that do not exactly match the expected clients:

```
policyConfig:
  userAgentMatchingConfig:
    defaultRejectionMessage: "Your client is too old. Go to
https://example.org to update it."
    deniedClients: []
    requiredClients:
      - regex: '\w+/v1\.1\.3 \(.+/.+\) openshift/\w{7}'
      - regex: '\w+/v1\.2\.0 \(.+/.+\) kubernetes/\w{7}'
    httpVerbs:
      - POST
      - PUT
```



NOTE

When the client's user agent mismatches the configuration, errors occur. To ensure that mutating requests match, enforce a whitelist. Rules are mapped to specific verbs, so you can ban mutating requests while allowing non-mutating requests.

CHAPTER 13. SYNCING GROUPS WITH LDAP

13.1. OVERVIEW

As an OpenShift Container Platform administrator, you can use groups to manage users, change their permissions, and enhance collaboration. Your organization may have already created user groups and stored them in an LDAP server. OpenShift Container Platform can sync those LDAP records with internal OpenShift Container Platform records, enabling you to manage your groups in one place. OpenShift Container Platform currently supports group sync with LDAP servers using three common schemas for defining group membership: RFC 2307, Active Directory, and augmented Active Directory.



NOTE

You must have [cluster-admin privileges](#) to sync groups.

13.2. CONFIGURING LDAP SYNC

Before you can [run LDAP sync](#), you need a sync configuration file. This file contains LDAP client configuration details:

- Configuration for connecting to your LDAP server.
- Sync configuration options that are dependent on the schema used in your LDAP server.

A sync configuration file can also contain an administrator-defined list of name mappings that maps OpenShift Container Platform Group names to groups in your LDAP server.

13.2.1. LDAP Client Configuration

Example 13.1. LDAP Client Configuration

```
url: ldap://10.0.0.0:389 ①
bindDN: cn=admin,dc=example,dc=com ②
bindPassword: password ③
insecure: false ④
ca: my-ldap-ca-bundle.crt ⑤
```

- ① The connection protocol, IP address of the LDAP server hosting your database, and the port to connect to, formatted as **scheme://host:port**.
- ② Optional distinguished name (DN) to use as the Bind DN. OpenShift Container Platform uses this if elevated privilege is required to retrieve entries for the sync operation.
- ③ Optional password to use to bind. OpenShift Container Platform uses this if elevated privilege is necessary to retrieve entries for the sync operation. This value may also be provided in an [environment variable](#), [external file](#), or [encrypted file](#).
- ④ When **true**, no TLS connection is made to the server. When **false**, secure LDAP (**ldaps://**) URLs connect using TLS, and insecure LDAP (**ldap://**) URLs are upgraded to TLS.
- ⑤ The certificate bundle to use for validating server certificates for the configured URL. If empty, OpenShift Container Platform uses system-trusted roots. This only applies if **insecure** is set to **false**.

13.2.2. LDAP Query Definition

Sync configurations consist of LDAP query definitions for the entries that are required for synchronization. The specific definition of an LDAP query depends on the schema used to store membership information in the LDAP server.

Example 13.2. LDAP Query Definition

```
baseDN: ou=users,dc=example,dc=com ❶
scope: sub ❷
derefAliases: never ❸
timeout: 0 ❹
filter: (objectClass=inetOrgPerson) ❺
pageSize: 0 ❻
```

- ❶ The distinguished name (DN) of the branch of the directory where all searches will start from. It is required that you specify the top of your directory tree, but you can also specify a subtree in the directory.
- ❷ The scope of the search. Valid values are **base**, **one**, or **sub**. If this is left undefined, then a scope of **sub** is assumed. Descriptions of the scope options can be found in the [table below](#).
- ❸ The behavior of the search with respect to aliases in the LDAP tree. Valid values are **never**, **search**, **base**, or **always**. If this is left undefined, then the default is to **always** dereference aliases. Descriptions of the dereferencing behaviors can be found in the [table below](#).
- ❹ The time limit allowed for the search by the client, in seconds. A value of 0 imposes no client-side limit.
- ❺ A valid LDAP search filter. If this is left undefined, then the default is **(objectClass=*)**.
- ❻ The optional maximum size of response pages from the server, measured in LDAP entries. If set to 0, no size restrictions will be made on pages of responses. Setting paging sizes is necessary when queries return more entries than the client or server allow by default.

Table 13.1. LDAP Search Scope Options

LDAP Search Scope	Description
base	Only consider the object specified by the base DN given for the query.
one	Consider all of the objects on the same level in the tree as the base DN for the query.
sub	Consider the entire subtree rooted at the base DN given for the query.

Table 13.2. LDAP Dereferencing Behaviors

Dereferencing Behavior	Description
never	Never dereference any aliases found in the LDAP tree.
search	Only dereference aliases found while searching.
base	Only dereference aliases while finding the base object.
always	Always dereference all aliases found in the LDAP tree.

13.2.3. User-Defined Name Mapping

A user-defined name mapping explicitly maps the names of OpenShift Container Platform Groups to unique identifiers that find groups on your LDAP server. The mapping uses normal YAML syntax. A user-defined mapping can contain an entry for every group in your LDAP server or only a subset of those groups. If there are groups on the LDAP server that do not have a user-defined name mapping, the default behavior during sync is to use the attribute specified as the Group's name.

Example 13.3. User-Defined Name Mapping

```
groupUIDNameMapping:
  "cn=group1,ou=groups,dc=example,dc=com": firstgroup
  "cn=group2,ou=groups,dc=example,dc=com": secondgroup
  "cn=group3,ou=groups,dc=example,dc=com": thirdgroup
```

13.3. RUNNING LDAP SYNC

Once you have created a [sync configuration file](#), then sync can begin. OpenShift Container Platform allows administrators to perform a number of different sync types with the same server.



NOTE

By default, all group synchronization or pruning operations are dry-run, so you must set the **--confirm** flag on the **sync-groups** command in order to make changes to OpenShift Container Platform Group records.


To sync all groups from the LDAP server with OpenShift Container Platform:

```
$ oadm groups sync --sync-config=config.yaml --confirm
```

To sync all Groups already in OpenShift Container Platform that correspond to groups in the LDAP server specified in the configuration file:

```
$ oadm groups sync --type=openshift --sync-config=config.yaml --confirm
```

To sync a subset of LDAP groups with OpenShift Container Platform, you can use whitelist files, blacklist files, or both:

**NOTE**

Any combination of blacklist files, whitelist files, or whitelist literals will work; whitelist literals can be included directly in the command itself. This applies to groups found on LDAP servers, as well as Groups already present in OpenShift Container Platform. Your files must contain one unique group identifier per line.

```
$ oadm groups sync --whitelist=<whitelist_file> \
    --sync-config=config.yaml \
    --confirm
$ oadm groups sync --blacklist=<blacklist_file> \
    --sync-config=config.yaml \
    --confirm
$ oadm groups sync <group_unique_identifier> \
    --sync-config=config.yaml \
    --confirm
$ oadm groups sync <group_unique_identifier> \
    --whitelist=<whitelist_file> \
    --blacklist=<blacklist_file> \
    --sync-config=config.yaml \
    --confirm
$ oadm groups sync --type=openshift \
    --whitelist=<whitelist_file> \
    --sync-config=config.yaml \
    --confirm
```

13.4. RUNNING A GROUP PRUNING JOB

An administrator can also choose to remove groups from OpenShift Container Platform records if the records on the LDAP server that created them are no longer present. The prune job will accept the same sync configuration file and white- or black-lists as used for the sync job.

For example, if groups had previously been synchronized from LDAP using some *config.yaml* file, and some of those groups no longer existed on the LDAP server, the following command would determine which Groups in OpenShift Container Platform corresponded to the deleted groups in LDAP and then remove them from OpenShift Container Platform:

```
$ oadm groups prune --sync-config=config.yaml --confirm
```

13.5. SYNC EXAMPLES

This section contains examples for the [RFC 2307](#), [Active Directory](#), and [augmented Active Directory](#) schemas. All of the following examples synchronize a group named **admins** that has two members: **Jane** and **Jim**. Each example explains:

- How the group and users are added to the LDAP server.
- What the LDAP sync configuration file looks like.
- What the resulting Group record in OpenShift Container Platform will be after synchronization.

**NOTE**

These examples assume that all users are direct members of their respective groups. Specifically, no groups have other groups as members. See [Nested Membership Sync Example](#) for information on how to sync nested groups.

13.5.1. RFC 2307

In the RFC 2307 schema, both users (Jane and Jim) and groups exist on the LDAP server as first-class entries, and group membership is stored in attributes on the group. The following snippet of **ldif** defines the users and group for this schema:

Example 13.4. LDAP Entries Using RFC 2307 Schema: *rfc2307.ldif*

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com

dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups

dn: cn=admins,ou=groups,dc=example,dc=com ❶
objectClass: groupOfNames
cn: admins
owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com ❷
member: cn=Jim,ou=users,dc=example,dc=com
```

❶ The group is a first-class entry in the LDAP server.

❷ Members of a group are listed with an identifying reference as attributes on the group.

To sync this group, you must first create the configuration file. The RFC 2307 schema requires you to provide an LDAP query definition for both user and group entries, as well as the attributes with which to represent them in the internal OpenShift Container Platform records.

For clarity, the Group you create in OpenShift Container Platform should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of a Group by their e-mail, and use the name of the group as the common name. The following configuration file creates these relationships:



NOTE

If using user-defined name mappings, your [configuration file](#) will differ.

Example 13.5. LDAP Sync Configuration Using RFC 2307 Schema: *rfc2307_config.yaml*

```
kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389 ❶
insecure: false ❷
rfc2307:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn ❸
  groupNameAttributes: [ cn ] ❹
  groupMembershipAttributes: [ member ] ❺
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  userUIDAttribute: dn ❻
  userNameAttributes: [ mail ] ❼
  tolerateMemberNotFoundErrors: false
  tolerateMemberOutOfScopeErrors: false
```

- ❶ The IP address and host of the LDAP server where this group's record is stored.
- ❷ When **true**, no TLS connection is made to the server. When **false**, secure LDAP (**ldaps://**) URLs connect using TLS, and insecure LDAP (**ldap://**) URLs are upgraded to TLS.
- ❸ The attribute that uniquely identifies a group on the LDAP server. You cannot specify **groupsQuery** filters when using DN for groupUIDAttribute. For fine-grained filtering, use the [whitelist / blacklist method](#).
- ❹ The attribute to use as the name of the Group.
- ❺ The attribute on the group that stores the membership information.
- ❻ The attribute that uniquely identifies a user on the LDAP server. You cannot specify **usersQuery** filters when using DN for userUIDAttribute. For fine-grained filtering, use the [whitelist / blacklist method](#).

- 7 The attribute to use as the name of the user in the OpenShift Container Platform Group record.

To run sync with the *rfc2307_config.yaml* file:

```
$ oadm groups sync --sync-config=rfc2307_config.yaml --confirm
```

OpenShift Container Platform creates the following Group record as a result of the above sync operation:

Example 13.6. OpenShift Container Platform Group Created Using *rfc2307_config.yaml*

```
apiVersion: v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 1
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com 2
    openshift.io/ldap.url: LDAP_SERVER_IP:389 3
  creationTimestamp:
    name: admins 4
users: 5
- jane.smith@example.com
- jim.adams@example.com
```

- 1 The last time this Group was synchronized with the LDAP server, in ISO 6801 format.
- 2 The unique identifier for the group on the LDAP server.
- 3 The IP address and host of the LDAP server where this Group's record is stored.
- 4 The name of the Group as specified by the sync file.
- 5 The users that are members of the Group, named as specified by the sync file.

13.5.1.1. RFC2307 with User-Defined Name Mappings

When syncing groups with user-defined name mappings, the configuration file changes to contain these mappings as shown below.

Example 13.7. LDAP Sync Configuration Using RFC 2307 Schema With User-Defined Name Mappings: *rfc2307_config_user_defined.yaml*

```
kind: LDAPSyncConfig
apiVersion: v1
groupUIDNameMapping:
  "cn=admins,ou=groups,dc=example,dc=com": Administrators 1
rfc2307:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
```

```

    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn 2
  groupNameAttributes: [ cn ] 3
  groupMembershipAttributes: [ member ]
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  userUIDAttribute: dn 4
  userNameAttributes: [ mail ]
  tolerateMemberNotFoundErrors: false
  tolerateMemberOutOfScopeErrors: false

```

- 1** The user-defined name mapping.
- 2** The unique identifier attribute that is used for the keys in the user-defined name mapping. You cannot specify **groupsQuery** filters when using DN for groupUIDAttribute. For fine-grained filtering, use the [whitelist / blacklist method](#).
- 3** The attribute to name OpenShift Container Platform Groups with if their unique identifier is not in the user-defined name mapping.
- 4** The attribute that uniquely identifies a user on the LDAP server. You cannot specify **usersQuery** filters when using DN for userUIDAttribute. For fine-grained filtering, use the [whitelist / blacklist method](#).

To run sync with the *rfc2307_config_user_defined.yaml* file:

```
$ oadm groups sync --sync-config=rfc2307_config_user_defined.yaml --confirm
```

OpenShift Container Platform creates the following Group record as a result of the above sync operation:

Example 13.8. OpenShift Container Platform Group Created Using *rfc2307_config_user_defined.yaml*

```

apiVersion: v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com
    openshift.io/ldap.url: LDAP_SERVER_IP:389
  creationTimestamp:
  name: Administrators 1
users:
- jane.smith@example.com
- jim.adams@example.com

```

- 1 The name of the Group as specified by the user-defined name mapping.

13.5.2. RFC 2307 with User-Defined Error Tolerances

By default, if the groups being synced contain members whose entries are outside of the scope defined in the member query, the group sync fails with an error:

```
Error determining LDAP group membership for "<group>": membership lookup
for user "<user>" in group "<group>" failed because of "search for entry
with dn="<user-dn>" would search outside of the base dn specified (dn="
<base-dn>")".
```

This often indicates a mis-configured **baseDN** in the **usersQuery** field. However, in cases where the **baseDN** intentionally does not contain some of the members of the group, setting **tolerateMemberOutOfScopeErrors: true** allows the group sync to continue. Out of scope members will be ignored.

Similarly, when the group sync process fails to locate a member for a group, it fails outright with errors:

```
Error determining LDAP group membership for "<group>": membership lookup
for user "<user>" in group "<group>" failed because of "search for entry
with base dn="<user-dn>" refers to a non-existent entry".
```

```
Error determining LDAP group membership for "<group>": membership lookup
for user "<user>" in group "<group>" failed because of "search for entry
with base dn="<user-dn>" and filter "<filter>" did not return any results".
```

This often indicates a mis-configured **usersQuery** field. However, in cases where the group contains member entries that are known to be missing, setting **tolerateMemberNotFoundErrors: true** allows the group sync to continue. Problematic members will be ignored.



WARNING

Enabling error tolerances for the LDAP group sync causes the sync process to ignore problematic member entries. If the LDAP group sync is not configured correctly, this could result in synced OpenShift Container Platform groups missing members.

Example 13.9. LDAP Entries Using RFC 2307 Schema With Problematic Group Membership: *rfc2307_problematic_users.ldif*

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users
```

```
dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
```

```

objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com

dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups

dn: cn=admins,ou=groups,dc=example,dc=com
objectClass: groupOfNames
cn: admins
owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com
member: cn=Jim,ou=users,dc=example,dc=com
member: cn=INVALID,ou=users,dc=example,dc=com ❶
member: cn=Jim,ou=OUTOFSCOPE,dc=example,dc=com ❷

```

- ❶ A member that does not exist on the LDAP server.
- ❷ A member that may exist, but is not under the **baseDN** in the user query for the sync job.

In order to tolerate the errors in the above example, the following additions to your sync configuration file must be made:

Example 13.10. LDAP Sync Configuration Using RFC 2307 Schema Tolerating Errors: *rfc2307_config_tolerating.yaml*

```

kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
rfc2307:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
  groupUIDAttribute: dn
  groupNameAttributes: [ cn ]
  groupMembershipAttributes: [ member ]
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"

```



```

scope: sub
derefAliases: never
userUIDAttribute: dn ❶
userNameAttributes: [ mail ]
tolerateMemberNotFoundErrors: true ❷
tolerateMemberOutOfScopeErrors: true ❸

```

- ❷ When **true**, the sync job tolerates groups for which some members were not found, and members whose LDAP entries are not found are ignored. The default behavior for the sync job is to fail if a member of a group is not found.
- ❸ When **true**, the sync job tolerates groups for which some members are outside the user scope given in the **usersQuery** base DN, and members outside the member query scope are ignored. The default behavior for the sync job is to fail if a member of a group is out of scope.
- ❶ The attribute that uniquely identifies a user on the LDAP server. You cannot specify **usersQuery** filters when using DN for userUIDAttribute. For fine-grained filtering, use the [whitelist / blacklist method](#).

To run sync with the *rfc2307_config_tolerating.yaml* file:

```
$ oadm groups sync --sync-config=rfc2307_config_tolerating.yaml --confirm
```

OpenShift Container Platform creates the following group record as a result of the above sync operation:

Example 13.11. OpenShift Container Platform Group Created Using *rfc2307_config.yaml*

```

apiVersion: v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com
    openshift.io/ldap.url: LDAP_SERVER_IP:389
  creationTimestamp:
  name: admins
users: ❶
- jane.smith@example.com
- jim.adams@example.com

```

- ❶ The users that are members of the group, as specified by the sync file. Members for which lookup encountered tolerated errors are absent.

13.5.3. Active Directory

In the Active Directory schema, both users (Jane and Jim) exist in the LDAP server as first-class entries, and group membership is stored in attributes on the user. The following snippet of **ldif** defines the users and group for this schema:

Example 13.12. LDAP Entries Using Active Directory Schema: *active_directory.ldif*

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
memberOf: admins 1

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
memberOf: admins
```

- 1 The user's group memberships are listed as attributes on the user, and the group does not exist as an entry on the server. The **memberOf** attribute does not have to be a literal attribute on the user; in some LDAP servers, it is created during search and returned to the client, but not committed to the database.

To sync this group, you must first create the configuration file. The Active Directory schema requires you to provide an LDAP query definition for user entries, as well as the attributes to represent them with in the internal OpenShift Container Platform Group records.

For clarity, the Group you create in OpenShift Container Platform should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of a Group by their e-mail, but define the name of the Group by the name of the group on the LDAP server. The following configuration file creates these relationships:

Example 13.13. LDAP Sync Configuration Using Active Directory Schema: *active_directory_config.yaml*

```
kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
activeDirectory:
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
```

```

    filter: (objectclass=inetOrgPerson)
    pageSize: 0
    userNameAttributes: [ mail ] ❶
    groupMembershipAttributes: [ memberOf ] ❷

```

- ❶ The attribute to use as the name of the user in the OpenShift Container Platform Group record.
- ❷ The attribute on the user that stores the membership information.

To run sync with the *active_directory_config.yaml* file:

```
$ oadm groups sync --sync-config=active_directory_config.yaml --confirm
```

OpenShift Container Platform creates the following Group record as a result of the above sync operation:

Example 13.14. OpenShift Container Platform Group Created Using *active_directory_config.yaml*

```

apiVersion: v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 ❶
    openshift.io/ldap.uid: admins ❷
    openshift.io/ldap.url: LDAP_SERVER_IP:389 ❸
  creationTimestamp:
    name: admins ❹
users: ❺
- jane.smith@example.com
- jim.adams@example.com

```

- ❶ The last time this Group was synchronized with the LDAP server, in ISO 6801 format.
- ❷ The unique identifier for the group on the LDAP server.
- ❸ The IP address and host of the LDAP server where this Group's record is stored.
- ❹ The name of the group as listed in the LDAP server.
- ❺ The users that are members of the Group, named as specified by the sync file.

13.5.4. Augmented Active Directory

In the augmented Active Directory schema, both users (Jane and Jim) and groups exist in the LDAP server as first-class entries, and group membership is stored in attributes on the user. The following snippet of *ldif* defines the users and group for this schema:

Example 13.15. LDAP Entries Using Augmented Active Directory Schema:
augmented_active_directory.ldif

```

dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
memberOf: cn=admins,ou=groups,dc=example,dc=com 1

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
memberOf: cn=admins,ou=groups,dc=example,dc=com

dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups

dn: cn=admins,ou=groups,dc=example,dc=com 2
objectClass: groupOfNames
cn: admins
owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com
member: cn=Jim,ou=users,dc=example,dc=com

```

- 1 The user's group memberships are listed as attributes on the user.
- 2 The group is a first-class entry on the LDAP server.

To sync this group, you must first create the configuration file. The augmented Active Directory schema requires you to provide an LDAP query definition for both user entries and group entries, as well as the attributes with which to represent them in the internal OpenShift Container Platform Group records.

For clarity, the Group you create in OpenShift Container Platform should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of a Group by their e-mail, and use the name of the Group as the common name. The following configuration file creates these relationships.

Example 13.16. LDAP Sync Configuration Using Augmented Active Directory Schema: *augmented_active_directory_config.yaml*

```
kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
augmentedActiveDirectory:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn ❶
  groupNameAttributes: [ cn ] ❷
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  userNameAttributes: [ mail ] ❸
  groupMembershipAttributes: [ memberOf ] ❹
```

- ❶ The attribute that uniquely identifies a group on the LDAP server. You cannot specify **groupsQuery** filters when using DN for groupUIDAttribute. For fine-grained filtering, use the [whitelist / blacklist method](#).
- ❷ The attribute to use as the name of the Group.
- ❸ The attribute to use as the name of the user in the OpenShift Container Platform Group record.
- ❹ The attribute on the user that stores the membership information.

To run sync with the *augmented_active_directory_config.yaml* file:

```
$ oadm groups sync --sync-config=augmented_active_directory_config.yaml --confirm
```

OpenShift Container Platform creates the following Group record as a result of the above sync operation:

Example 13.17. OpenShift Group Created Using *augmented_active_directory_config.yaml*

```
apiVersion: v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 ❶
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com ❷
    openshift.io/ldap.url: LDAP_SERVER_IP:389 ❸
  creationTimestamp:
  name: admins ❹
```

```
users: 5
- jane.smith@example.com
- jim.adams@example.com
```

- 1 The last time this Group was synchronized with the LDAP server, in ISO 6801 format.
- 2 The unique identifier for the group on the LDAP server.
- 3 The IP address and host of the LDAP server where this Group's record is stored.
- 4 The name of the Group as specified by the sync file.
- 5 The users that are members of the Group, named as specified by the sync file.

13.6. NESTED MEMBERSHIP SYNC EXAMPLE

Groups in OpenShift Container Platform do not nest. The LDAP server must flatten group membership before the data can be consumed. Microsoft's Active Directory Server supports this feature via the [LDAP_MATCHING_RULE_IN_CHAIN](#) rule, which has the OID **1.2.840.113556.1.4.1941**. Furthermore, only explicitly [whitelisted](#) groups can be synced when using this matching rule.

This section has an example for the augmented Active Directory schema, which synchronizes a group named **admins** that has one user **Jane** and one group **otheradmins** as members. The **otheradmins** group has one user member: **Jim**. This example explains:

- How the group and users are added to the LDAP server.
- What the LDAP sync configuration file looks like.
- What the resulting Group record in OpenShift Container Platform will be after synchronization.

In the augmented Active Directory schema, both users (**Jane** and **Jim**) and groups exist in the LDAP server as first-class entries, and group membership is stored in attributes on the user or the group. The following snippet of **ldif** defines the users and groups for this schema:

LDAP Entries Using Augmented Active Directory Schema With Nested Members: *augmented_active_directory_nested.ldif*

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
memberOf: cn=admins,ou=groups,dc=example,dc=com 1
```

```

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
memberOf: cn=otheradmins,ou=groups,dc=example,dc=com 2

dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups

dn: cn=admins,ou=groups,dc=example,dc=com 3
objectClass: group
cn: admins
owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com
member: cn=otheradmins,ou=groups,dc=example,dc=com

dn: cn=otheradmins,ou=groups,dc=example,dc=com 4
objectClass: group
cn: otheradmins
owner: cn=admin,dc=example,dc=com
description: Other System Administrators
memberOf: cn=admins,ou=groups,dc=example,dc=com 5 6
member: cn=Jim,ou=users,dc=example,dc=com

```

1 2 5 The user's and group's memberships are listed as attributes on the object.

3 4 The groups are first-class entries on the LDAP server.

6 The **otheradmins** group is a member of the **admins** group.

To sync nested groups with Active Directory, you must provide an LDAP query definition for both user entries and group entries, as well as the attributes with which to represent them in the internal OpenShift Container Platform Group records. Furthermore, certain changes are required in this configuration:

- The **oadm groups sync** command must explicitly [whitelist](#) groups.
- The user's **groupMembershipAttributes** must include **"memberOf:1.2.840.113556.1.4.1941:"** to comply with the [LDAP_MATCHING_RULE_IN_CHAIN](#) rule.
- The **groupUIDAttribute** must be set to **dn**.
- The **groupsQuery**:
 - Must not set **filter**.
 - Must set a valid **derefAliases**.

- Should not set **baseDN** as that value is ignored.
- Should not set **scope** as that value is ignored.

For clarity, the Group you create in OpenShift Container Platform should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of a Group by their e-mail, and use the name of the Group as the common name. The following configuration file creates these relationships:


LDAP Sync Configuration Using Augmented Active Directory Schema With Nested Members: *augmented_active_directory_config_nested.yaml*

```
kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
augmentedActiveDirectory:
  groupsQuery: ❶
    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn ❷
  groupNameAttributes: [ cn ] ❸
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    filter: (objectclass=inetOrgPerson)
    pageSize: 0
  userNameAttributes: [ mail ] ❹
  groupMembershipAttributes: [ "memberOf:1.2.840.113556.1.4.1941:" ] ❺
```

- ❶ **groupsQuery** filters cannot be specified. The **groupsQuery** base DN and scope values are ignored. **groupsQuery** must set a valid **derefAliases**.
- ❷ The attribute that uniquely identifies a group on the LDAP server. It must be set to **dn**.
- ❸ The attribute to use as the name of the Group.
- ❹ The attribute to use as the name of the user in the OpenShift Container Platform Group record. **mail** or **sAMAccountName** are preferred choices in most installations.
- ❺ The attribute on the user that stores the membership information. Note the use of [LDAP_MATCHING_RULE_IN_CHAIN](#).

To run sync with the *augmented_active_directory_config_nested.yaml* file:

```
$ oadm groups sync \
  'cn=admins,ou=groups,dc=example,dc=com' \
  --sync-config=augmented_active_directory_config_nested.yaml \
  --confirm
```



NOTE

You must explicitly [whitelist](#) the **cn=admins,ou=groups,dc=example,dc=com** group.

OpenShift Container Platform creates the following Group record as a result of the above sync operation:

OpenShift Group Created Using *augmented_active_directory_config_nested.yaml*

```
apiVersion: v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 ❶
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com ❷
    openshift.io/ldap.url: LDAP_SERVER_IP:389 ❸
  creationTimestamp:
    name: admins ❹
users: ❺
- jane.smith@example.com
- jim.adams@example.com
```

- ❶ The last time this Group was synchronized with the LDAP server, in ISO 6801 format.
- ❷ The unique identifier for the group on the LDAP server.
- ❸ The IP address and host of the LDAP server where this Group's record is stored.
- ❹ The name of the Group as specified by the sync file.
- ❺ The users that are members of the Group, named as specified by the sync file. Note that members of nested groups are included since the group membership was flattened by the Microsoft Active Directory Server.

13.7. LDAP SYNC CONFIGURATION SPECIFICATION

The object specification for the configuration file is below. Note that the different schema objects have different fields. For example, [v1.ActiveDirectoryConfig](#) has no **groupsQuery** field whereas [v1.RFC2307Config](#) and [v1.AugmentedActiveDirectoryConfig](#) both do.

13.7.1. v1.LDAPSyncConfig

LDAPSyncConfig holds the necessary configuration options to define an LDAP group sync.

Name	Description	Schema
kind	String value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: http://releases.k8s.io/HEAD/docs/devel/api-conventions.md#types-kinds	string

Name	Description	Schema
apiVersion	Defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: http://releases.k8s.io/HEAD/docs/devel/api-conventions.md#resources	string
url	Host is the scheme, host and port of the LDAP server to connect to: scheme://host:port	string
bindDN	Optional DN to bind to the LDAP server with.	string
bindPassword	Optional password to bind with during the search phase.	v1.StringSource
insecure	If true , indicates the connection should not use TLS. Cannot be set to true with a URL scheme of ldaps:// . If false , ldaps:// URLs connect using TLS, and ldap:// URLs are upgraded to a TLS connection using StartTLS as specified in https://tools.ietf.org/html/rfc2830 .	boolean.
ca	Optional trusted certificate authority bundle to use when making requests to the server. If empty, the default system roots are used.	string
groupUIDNameMapping	Optional direct mapping of LDAP group UIDs to OpenShift Container Platform Group names.	object
rfc2307	Holds the configuration for extracting data from an LDAP server set up in a fashion similar to RFC2307: first-class group and user entries, with group membership determined by a multi-valued attribute on the group entry listing its members.	v1.RFC2307Config

Name	Description	Schema
activeDirectory	Holds the configuration for extracting data from an LDAP server set up in a fashion similar to that used in Active Directory: first-class user entries, with group membership determined by a multi-valued attribute on members listing groups they are a member of.	v1.ActiveDirectoryConfig
augmentedActiveDirectory	Holds the configuration for extracting data from an LDAP server set up in a fashion similar to that used in Active Directory as described above, with one addition: first-class group entries exist and are used to hold metadata but not group membership.	v1.AugmentedActiveDirectoryConfig

13.7.2. v1.StringSource

StringSource allows specifying a string inline, or externally via environment variable or file. When it contains only a string value, it marshals to a simple JSON string.

Name	Description	Schema
value	Specifies the cleartext value, or an encrypted value if keyFile is specified.	string
env	Specifies an environment variable containing the cleartext value, or an encrypted value if the keyFile is specified.	string
file	References a file containing the cleartext value, or an encrypted value if a keyFile is specified.	string
keyFile	References a file containing the key to use to decrypt the value.	string

13.7.3. v1.LDAPQuery

LDAPQuery holds the options necessary to build an LDAP query.

Name	Description	Schema
baseDN	DN of the branch of the directory where all searches should start from.	string
scope	The (optional) scope of the search. Can be base (only the base object), one (all objects on the base level), sub (the entire subtree). Defaults to sub if not set.	string
derefAliases	The (optional) behavior of the search with regards to aliases. Can be never (never dereference aliases), search (only dereference in searching), base (only dereference in finding the base object), always (always dereference). Defaults to always if not set.	string
timeout	Holds the limit of time in seconds that any request to the server can remain outstanding before the wait for a response is given up. If this is 0 , no client-side limit is imposed.	integer
filter	A valid LDAP search filter that retrieves all relevant entries from the LDAP server with the base DN.	string
pageSize	Maximum preferred page size, measured in LDAP entries. A page size of 0 means no paging will be done.	integer

13.7.4. v1.RFC2307Config

RFC2307Config holds the necessary configuration options to define how an LDAP group sync interacts with an LDAP server using the RFC2307 schema.

Name	Description	Schema
groupsQuery	Holds the template for an LDAP query that returns group entries.	v1.LDAPQuery

Name	Description	Schema
groupUIDAttribute	Defines which attribute on an LDAP group entry will be interpreted as its unique identifier. (ldapGroupUID)	string
groupNameAttributes	Defines which attributes on an LDAP group entry will be interpreted as its name to use for an OpenShift Container Platform group.	string array
groupMembershipAttributes	Defines which attributes on an LDAP group entry will be interpreted as its members. The values contained in those attributes must be queryable by your UserUIDAttribute .	string array
usersQuery	Holds the template for an LDAP query that returns user entries.	v1.LDAPQuery
userUIDAttribute	Defines which attribute on an LDAP user entry will be interpreted as its unique identifier. It must correspond to values that will be found from the GroupMembershipAttributes .	string
userNameAttributes	Defines which attributes on an LDAP user entry will be used, in order, as its OpenShift Container Platform user name. The first attribute with a non-empty value is used. This should match your PreferredUsername setting for your LDAPPasswordIdentityProvider .	string array

Name	Description	Schema
tolerateMemberNotFoundErrors	Determines the behavior of the LDAP sync job when missing user entries are encountered. If true , an LDAP query for users that does not find any will be tolerated and an only and error will be logged. If false , the LDAP sync job will fail if a query for users doesn't find any. The default value is 'false'. Misconfigured LDAP sync jobs with this flag set to 'true' can cause group membership to be removed, so it is recommended to use this flag with caution.	boolean
tolerateMemberOutOfScopeErrors	Determines the behavior of the LDAP sync job when out-of-scope user entries are encountered. If true , an LDAP query for a user that falls outside of the base DN given for the all user query will be tolerated and only an error will be logged. If false , the LDAP sync job will fail if a user query would search outside of the base DN specified by the all user query. Misconfigured LDAP sync jobs with this flag set to true can result in groups missing users, so it is recommended to use this flag with caution.	boolean

13.7.5. v1.ActiveDirectoryConfig

ActiveDirectoryConfig holds the necessary configuration options to define how an LDAP group sync interacts with an LDAP server using the Active Directory schema.

Name	Description	Schema
usersQuery	Holds the template for an LDAP query that returns user entries.	v1.LDAPQuery
userNameAttributes	Defines which attributes on an LDAP user entry will be interpreted as its OpenShift Container Platform user name.	string array

Name	Description	Schema
groupMembershipAttributes	Defines which attributes on an LDAP user entry will be interpreted as the groups it is a member of.	string array

13.7.6. v1.AugmentedActiveDirectoryConfig

AugmentedActiveDirectoryConfig holds the necessary configuration options to define how an LDAP group sync interacts with an LDAP server using the augmented Active Directory schema.


Name	Description	Schema
usersQuery	Holds the template for an LDAP query that returns user entries.	v1.LDAPQuery
userNameAttributes	Defines which attributes on an LDAP user entry will be interpreted as its OpenShift Container Platform user name.	string array
groupMembershipAttributes	Defines which attributes on an LDAP user entry will be interpreted as the groups it is a member of.	string array
groupsQuery	Holds the template for an LDAP query that returns group entries.	v1.LDAPQuery
groupUIDAttribute	Defines which attribute on an LDAP group entry will be interpreted as its unique identifier. (ldapGroupUID)	string
groupNameAttributes	Defines which attributes on an LDAP group entry will be interpreted as its name to use for an OpenShift Container Platform group.	string array

CHAPTER 14. CONFIGURING LDAP FAILOVER

OpenShift Container Platform provides an [authentication provider](#) for use with Lightweight Directory Access Protocol (LDAP) setups, but it can connect to only a single LDAP server. During OpenShift Container Platform installation, you can configure the System Security Services Daemon (SSSD) for LDAP failover to ensure access to your cluster if one LDAP server fails.

The setup for this configuration is advanced and requires a separate authentication server, also called an **remote basic authentication server**, for OpenShift Container Platform to communicate with. You configure this server to pass extra attributes, such as email addresses, to OpenShift Container Platform so it can display them in the web console.

This topic describes how to complete this set up on a dedicated physical or virtual machine (VM), but you can also configure SSSD in containers.



IMPORTANT

You must complete all sections of this topic.

14.1. PREREQUISITES FOR CONFIGURING BASIC REMOTE AUTHENTICATION

- Before starting setup, you need to know the following information about your LDAP server:
 - Whether the directory server is powered by [FreeIPA](#), Active Directory, or another LDAP solution.
 - The Uniform Resource Identifier (URI) for the LDAP server, for example, **ldap.example.com**.
 - The location of the CA certificate for the LDAP server.
 - Whether the LDAP server corresponds to RFC 2307 or RFC2307bis for user groups.
- Prepare the servers:
 - **remote-basic.example.com**: A VM to use as the remote basic authentication server.
 - Select an operating system that includes SSSD version 1.12.0 for this server such as Red Hat Enterprise Linux 7.0 or later.
 - **openshift.example.com**: A new installation of OpenShift Container Platform.
 - You must not have an authentication method configured for this cluster.
 - Do not start OpenShift Container Platform on this cluster.

14.2. GENERATING AND SHARING CERTIFICATES WITH THE REMOTE BASIC AUTHENTICATION SERVER

Complete the following steps on the first master host listed in the Ansible host inventory file, by default **/etc/ansible/hosts**.

1. To ensure that communication between the remote basic authentication server and OpenShift Container Platform is trustworthy, create a set of Transport Layer Security (TLS) certificates to use during the other phases of this set up. Run the following command:

```
# openshift start \
  --public-master=https://openshift.example.com:8443 \
  --write-config=/etc/origin/
```

The output includes the **/etc/origin/master/ca.crt** and **/etc/origin/master/ca.key** signing certificates.

2. Use the signing certificate to generate keys to use on the remote basic authentication server:

```
# mkdir -p /etc/origin/remote-basic/
# oc adm ca create-server-cert \
  --cert='/etc/origin/remote-basic/remote-basic.example.com.crt' \
  --key='/etc/origin/remote-basic/remote-basic.example.com.key' \
  --hostnames=remote-basic.example.com \ 1
  --signer-cert='/etc/origin/master/ca.crt' \
  --signer-key='/etc/origin/master/ca.key' \
  --signer-serial='/etc/origin/master/ca.serial.txt'
```


1

A comma-separated list of all the host names and interface IP addresses that need to access the remote basic authentication server.



NOTE

The certificate files that you generate are valid for two years. You can alter this period by changing the **--expire-days** and **--signer-expire-days** values, but for security reasons, do not make them greater than 730.



IMPORTANT

If you do not list all host names and interface IP addresses that need to access the remote basic authentication server, the HTTPS connection will fail.

3. Copy the necessary certificates and key to the remote basic authentication server:

```
# scp /etc/origin/master/ca.crt \
  root@remote-basic.example.com:/etc/pki/CA/certs/

# scp /etc/origin/remote-basic/remote-basic.example.com.crt \
  root@remote-basic.example.com:/etc/pki/tls/certs/

# scp /etc/origin/remote-basic/remote-basic.example.com.key \
  root@remote-basic.example.com:/etc/pki/tls/private/
```

14.3. CONFIGURING SSSD FOR LDAP FAILOVER

Complete these steps on the remote basic authentication server.

You can configure the SSSD to retrieve attributes, such as email addresses and display names, and pass them to OpenShift Container Platform to display in the web interface. In the following steps, you configure the SSSD to provide email addresses to OpenShift Container Platform:

1. Install the required SSSD and the web server components:

```
# yum install -y sssd \
    sssd-dbus \
    realmd \
    httpd \
    mod_session \
    mod_ssl \
    mod_lookup_identity \
    mod_authnz_pam \
    php \
    mod_php
```

2. Set up SSSD to authenticate this VM against the LDAP server. If the LDAP server is a FreeIPA or Active Directory environment, then use **realmd** to join this machine to the domain.

```
# realm join ldap.example.com
```

For more advanced cases, see the [System-Level Authentication Guide](#)

3. To use SSSD to manage failover situations for LDAP, add more entries to the **/etc/sss/sss.conf** file on the **ldap_uri** line. Systems that are enrolled with FreeIPA can automatically handle failover by using DNS SRV records.
4. Modify the **[domain/DOMAINNAME]** section of the **/etc/sss/sss.conf** file and add this attribute:

```
[domain/example.com]
...
ldap_user_extra_attrs = mail ❶
```

- ❶ Specify the correct attribute to retrieve email addresses for your LDAP solution. For IPA, specify **mail**. Other LDAP solutions might use another attribute, such as **email**.

5. Confirm that the **domain** parameter in the **/etc/sss/sss.conf** file contains only the domain name listed in the **[domain/DOMAINNAME]** section.

```
domains = example.com
```

6. Grant Apache permission to retrieve the email attribute. Add the following lines to the **[ifp]** section of the **/etc/sss/sss.conf** file:

```
[ifp]
user_attributes = +mail
allowed_uids = apache, root
```

7. To ensure that all of the changes are applied properly, restart SSSD:

```
$ systemctl restart sssd.service
```

- Test that the user information can be retrieved properly:

```
$ getent passwd <username>
username:*:12345:12345:Example User:/home/username:/usr/bin/bash
```

- Confirm that the mail attribute you specified returns an email address from your domain:

```
# dbus-send --print-reply --system --
dest=org.freedesktop.sssd.infopipe \
  /org/freedesktop/sss/infopipe
org.freedesktop.sssd.infopipe.GetUserAttr \
  string:username \ ❶
  array:string:mail ❷

method return time=1528091855.672691 sender=:1.2787 ->
destination=:1.2795 serial=13 reply_serial=2
array [
  dict entry(
    string "mail"
    variant
      array [
        string "username@example.com"
      ]
  )
]
```

❶ Provide a user name in your LDAP solution.

❷ Specify the attribute that you configured.

- Attempt to log into the VM as an LDAP user and confirm that you can log in using LDAP credentials. You can use either the local console or a remote service like SSH to log in.



IMPORTANT

By default, all users can log into the remote basic authentication server by using their LDAP credentials. You can change this behavior:

- If you use IPA joined systems, [configure host-based access control](#).
- If you use Active Directory joined systems, use a [group policy object](#).
- For other cases, see the [SSSD configuration](#) documentation.

14.4. CONFIGURING APACHE TO USE SSSD

- Create a `/etc/pam.d/openshift` file that contains the following contents:


```
auth required pam_sss.so
account required pam_sss.so
```

This configuration enables PAM, the pluggable authentication module, to use **pam_sss.so** to determine authentication and access control when an authentication request is issued for the **openshift** stack.

2. Edit the `/etc/httpd/conf.modules.d/55-authnz_pam.conf` file and uncomment the following line:

```
LoadModule authnz_pam_module modules/mod_authnz_pam.so
```

3. To configure the Apache `httpd.conf` file for remote basic authentication, create the ***openshift-remote-basic-auth.conf*** file in the `/etc/httpd/conf.d` directory. Use the following template to provide your required settings and values:



IMPORTANT

Carefully review the template and customize its contents to fit your environment.

```
LoadModule request_module modules/mod_request.so
LoadModule php7_module modules/libphp7.so

# Nothing needs to be served over HTTP. This virtual host simply
# redirects to
# HTTPS.
<VirtualHost *:80>
    DocumentRoot /var/www/html
    RewriteEngine On
    RewriteRule ^(.*)$ https://%{HTTP_HOST}$1 [R,L]
</VirtualHost>

<VirtualHost *:443>
    # This needs to match the certificates you generated. See the CN
    # and X509v3
    # Subject Alternative Name in the output of:
    # openssl x509 -text -in /etc/pki/tls/certs/remote-
    # basic.example.com.crt
    ServerName remote-basic.example.com

    DocumentRoot /var/www/html

    # Secure all connections with TLS
    SSLEngine on
    SSLCertificateFile /etc/pki/tls/certs/remote-basic.example.com.crt
    SSLCertificateKeyFile /etc/pki/tls/private/remote-
    basic.example.com.key
    SSLCACertificateFile /etc/pki/CA/certs/ca.crt

    # Require that TLS clients provide a valid certificate
    SSLVerifyClient require
    SSLVerifyDepth 10

    # Other SSL options that may be useful
    # SSLCertificateChainFile ...
    # SSLCARevocationFile ...

    # Send logs to a specific location to make them easier to find
    ErrorLog logs/remote_basic_error_log
    TransferLog logs/remote_basic_access_log
    LogLevel warn

    # PHP script that turns the Apache REMOTE_USER env var
```

```

# into a JSON formatted response that OpenShift understands
<Location /check_user.php>
# all requests not using SSL are denied
SSLRequireSSL
# denies access when SSLRequireSSL is applied
SSLOptions +StrictRequire
# Require both a valid basic auth user (so REMOTE_USER is always
set)
# and that the CN of the TLS client matches that of the
OpenShift master
<RequireAll>
    Require valid-user
    Require expr %{SSL_CLIENT_S_DN_CN} == 'system:openshift-
master'
</RequireAll>
# Use basic auth since OpenShift will call this endpoint with a
basic challenge
AuthType Basic
AuthName openshift
AuthBasicProvider PAM
AuthPAMService openshift

# Store attributes in environment variables. Specify the email
attribute that
# you confirmed.
LookupOutput Env
LookupUserAttr mail REMOTE_USER_MAIL
LookupUserGECOS REMOTE_USER_DISPLAY_NAME

# Other options that might be useful

# While REMOTE_USER is used as the sub field and serves as the
immutable ID,
# REMOTE_USER_PREFERRED_USERNAME could be used to have a
different username
# LookupUserAttr <attr_name> REMOTE_USER_PREFERRED_USERNAME

# Group support may be added in a future release
# LookupUserGroupsIter REMOTE_USER_GROUP
</Location>

# Deny everything else
<Location ~ "^(?!\/check_user\.php)\.*$">
    Deny from all
</Location>
</VirtualHost>

```

4. Create the **check_user.php** script in the **/var/www/html** directory. Include the following code:

```

<?php
// Get the user based on the Apache var, this should always be
// set because we 'Require valid-user' in the configuration
$user = apache_getenv('REMOTE_USER');

// However, we assume it may not be set and
// build an error response by default

```

```

$data = array(
    'error' => 'remote PAM authentication failed'
);

// Build a success response if we have a user
if (!empty($user)) {
    $data = array(
        'sub' => $user
    );
    // Map of optional environment variables to optional JSON fields
    $env_map = array(
        'REMOTE_USER_MAIL' => 'email',
        'REMOTE_USER_DISPLAY_NAME' => 'name',
        'REMOTE_USER_PREFERRED_USERNAME' => 'preferred_username'
    );

    // Add all non-empty environment variables to JSON data
    foreach ($env_map as $env_name => $json_name) {
        $env_data = apache_getenv($env_name);
        if (!empty($env_data)) {
            $data[$json_name] = $env_data;
        }
    }
}

// We always output JSON from this script
header('Content-Type: application/json', true);

// Write the response as JSON
echo json_encode($data);
?>

```

5. Enable Apache to load the module. Modify the `/etc/httpd/conf.modules.d/55-lookup_identity.conf` file and uncomment the following line:

```
LoadModule lookup_identity_module modules/mod_lookup_identity.so
```

6. Set an SELinux boolean so that SELinux allows Apache to connect to SSSD over D-BUS:

```
# setsebool -P httpd_dbus_sssd on
```

7. Set a boolean to tell SELinux that it is acceptable for Apache to contact the PAM subsystem:

```
# setsebool -P allow_httpd_mod_auth_pam on
```

8. Start Apache:

```
# systemctl start httpd.service
```

14.5. CONFIGURING OPENSIFT CONTAINER PLATFORM TO USE SSSD AS THE BASIC REMOTE AUTHENTICATION SERVER

Modify the default configuration of your cluster to use the new identity provider that you created. Complete the following steps on the first master host listed in the Ansible host inventory file.

1. Open the `/etc/origin/master/master-config.yaml` file.
2. Locate the **identityProviders** section and replace it with the following code:

```
identityProviders:
- name: sssd
  challenge: true
  login: true
  mappingMethod: claim
  provider:
    apiVersion: v1
    kind: BasicAuthPasswordIdentityProvider
    url: https://remote-basic.example.com/check_user.php
    ca: /etc/origin/master/ca.crt
    certFile: /etc/origin/master/openshift-master.crt
    keyFile: /etc/origin/master/openshift-master.key
```

3. Start OpenShift Container Platform with the updated configuration:

```
# openshift start \
--public-master=https://openshift.example.com:8443 \
--master-config=/etc/origin/master/master-config.yaml \
--node-config=/etc/origin/node-node1.example.com/node-
config.yaml
```

4. Test a login by using the **oc** CLI:

```
oc login https://openshift.example.com:8443
```

You can log in only with valid LDAP credentials.

5. List the identities and confirm that an email address is displayed for each user name. Run the following command:

```
$ oc get identity -o yaml
```

CHAPTER 15. CONFIGURING THE SDN

15.1. OVERVIEW

The [OpenShift Container Platform SDN](#) enables communication between pods across the OpenShift Container Platform cluster, establishing a *pod network*. Two [SDN plug-ins](#) are currently available (**ovs-subnet** and **ovs-multitenant**), which provide different methods for configuring the pod network.

15.2. CONFIGURING THE POD NETWORK WITH ANSIBLE

For initial [advanced installations](#), the **ovs-subnet** plug-in is installed and configured by default, though it can be [overridden during installation](#) using the [os_sdn_network_plugin_name](#) parameter, which is configurable in the Ansible inventory file.

Example 15.1. Example SDN Configuration with Ansible

```
# Configure the multi-tenant SDN plugin (default is 'redhat/openshift-ovs-subnet')
# os_sdn_network_plugin_name='redhat/openshift-ovs-multitenant'

# Disable the OpenShift SDN plugin
# openshift_use_openshift_sdn=False

# Configure SDN cluster network CIDR block. This network block should
# be a private block and should not conflict with existing network
# blocks in your infrastructure that pods may require access to.
# Can not be changed after deployment.
#osm_cluster_network_cidr=10.1.0.0/16

# default subdomain to use for exposed routes
#openshift_master_default_subdomain=apps.test.example.com

# Configure SDN cluster network and kubernetes service CIDR blocks.
These
# network blocks should be private and should not conflict with network
blocks
# in your infrastructure that pods may require access to. Can not be
changed
# after deployment.
#osm_cluster_network_cidr=10.1.0.0/16
#openshift_portal_net=172.30.0.0/16

# Configure number of bits to allocate to each host's subnet e.g. 8
# would mean a /24 network on the host.
#osm_host_subnet_length=8

# This variable specifies the service proxy implementation to use:
# either iptables for the pure-iptables version (the default),
# or userspace for the userspace proxy.
#openshift_node_proxy_mode=iptables
```



For initial [quick installations](#), the **ovs-subnet** plug-in is installed and configured by default as well, and can be [reconfigured post-installation](#) using the **networkConfig** stanza of the **master-config.yaml** file.

15.3. CONFIGURING THE POD NETWORK ON MASTERS

Cluster administrators can control pod network settings on masters by modifying parameters in the **networkConfig** section of the [master configuration file](#) (located at `/etc/origin/master/master-config.yaml` by default):

```
networkConfig:
  clusterNetworkCIDR: 10.128.0.0/14 ❶
  hostSubnetLength: 9 ❷
  networkPluginName: "redhat/openshift-ovs-subnet" ❸
  serviceNetworkCIDR: 172.30.0.0/16 ❹
```

- ❶ Cluster network for node IP allocation
- ❷ Number of bits for pod IP allocation within a node
- ❸ Set to **redhat/openshift-ovs-subnet** for the **ovs-subnet** plug-in or **redhat/openshift-ovs-multitenant** for the **ovs-multitenant** plug-in
- ❹ Service IP allocation for the cluster



IMPORTANT

The **serviceNetworkCIDR** and **hostSubnetLength** values cannot be changed after the cluster is first created, and **clusterNetworkCIDR** can only be changed to be a larger network that still contains the original network. For example, given the default value of **10.128.0.0/14**, you could change **clusterNetworkCIDR** to **10.128.0.0/9** (i.e., the entire upper half of net 10) but not to **10.64.0.0/16**, because that does not overlap the original value.

15.4. CONFIGURING THE POD NETWORK ON NODES

Cluster administrators can control pod network settings on nodes by modifying parameters in the **networkConfig** section of the [node configuration file](#) (located at `/etc/origin/node/node-config.yaml` by default):

```
networkConfig:
  mtu: 1450 ❶
  networkPluginName: "redhat/openshift-ovs-subnet" ❷
```

- ❶ Maximum transmission unit (MTU) for the pod overlay network
- ❷ Set to **redhat/openshift-ovs-subnet** for the **ovs-subnet** plug-in or **redhat/openshift-ovs-multitenant** for the **ovs-multitenant** plug-in

15.5. MIGRATING BETWEEN SDN PLUG-INS

If you are already using one SDN plug-in and want to switch to another:

1. Change the **networkPluginName** parameter on all [masters](#) and [nodes](#) in their configuration files.
2. Restart the **atomic-openshift-master** service on masters and the **atomic-openshift-node** service on nodes.
3. If you are switching from an OpenShift Container Platform SDN plug-in to a third-party plug-in, then clean up OpenShift Container Platform SDN-specific artifacts:

```
$ oc delete clusternetwork --all
$ oc delete hostsubnets --all
$ oc delete netnamespaces --all
```

When switching from the **ovs-subnet** to the **ovs-multitenant** OpenShift Container Platform SDN plug-in, all the existing projects in the cluster will be fully isolated (assigned unique VNIDs). Cluster administrators can choose to [modify the project networks](#) using the administrator CLI.

Check VNIDs by running:

```
$ oc get netnamespace
```

15.6. EXTERNAL ACCESS TO THE CLUSTER NETWORK


If a host that is external to OpenShift Container Platform requires access to the cluster network, you have two options:

1. Configure the host as an OpenShift Container Platform node but mark it [unschedulable](#) so that the master does not schedule containers on it.
2. Create a tunnel between your host and a host that is on the cluster network.

Both options are presented as part of a practical use-case in the documentation for configuring [routing from an edge load-balancer to containers within OpenShift Container Platform SDN](#).


15.7. USING FLANNEL

As an alternative to the default SDN, OpenShift Container Platform also provides Ansible playbooks for installing **flannel**-based networking. This is useful if running OpenShift Container Platform within a cloud provider platform that also relies on SDN, such as Red Hat OpenStack Platform, and you want to avoid encapsulating packets twice through both platforms.



IMPORTANT

This is only supported for OpenShift Container Platform on Red Hat OpenStack Platform.



IMPORTANT

Neutron port security must be configured, even when security groups are not being used.

To enable **flannel** within your OpenShift Container Platform cluster:

1. Neutron port security controls must be configured to be compatible with Flannel. The default configuration of Red Hat OpenStack Platform disables user control of **port_security**. Configure Neutron to allow users to control the **port_security** setting on individual ports.

- a. On the Neutron servers, add the following to the `/etc/neutron/plugins/ml2/ml2_conf.ini` file:

```
[ml2]
...
extension_drivers = port_security
```

- b. Then, restart the Neutron services:

```
service neutron-dhcp-agent restart
service neutron-ovs-cleanup restart
service neutron-metadata-agent restart
service neutron-l3-agent restart
service neutron-plugin-openvswitch-agent restart
service neutron-vpn-agent restart
service neutron-server restart
```

2. Set the following variables in your Ansible inventory file before running the installation:

```
openshift_use_openshift_sdn=false ❶
openshift_use_flannel=true ❷
flannel_interface=eth0
```

- ❶ Set **openshift_use_openshift_sdn** to **false** to disable the default SDN.
- ❷ Set **openshift_use_flannel** to **true** to enable **flannel** in place.

3. Optionally, you can specify the interface to use for inter-host communication using the **flannel_interface** variable. Without this variable, the OpenShift Container Platform installation uses the default interface.

CHAPTER 16. CONFIGURING NUAGE SDN

16.1. NUAGE SDN AND OPENSIFT CONTAINER PLATFORM

Nuage Networks Virtualized Services Platform (VSP) provides virtual networking and software-defined networking (SDN) infrastructure to Docker container environments that simplifies IT operations and expands OpenShift Container Platform's native networking capabilities.

Nuage Networks VSP supports Docker-based applications running on OpenShift Container Platform to accelerate the provisioning of virtual networks between pods and traditional workloads, and to enable security policies across the entire cloud infrastructure. VSP allows for the automation of security appliances to include granular security and microsegmentation policies for container applications.

Integrating VSP with the OpenShift Container Platform application workflow allows business applications to be quickly turned up and updated by removing the network lag faced by DevOps teams. VSP supports different workflows with OpenShift Container Platform in order to accommodate scenarios where users can choose ease-of-use or complete control using policy-based automation.

See [Networking](#) for more information on how VSP is integrated with OpenShift Container Platform.

16.2. DEVELOPER WORKFLOW

This workflow is used in developer environments and requires little input from the developer in setting up the networking. In this workflow, **nuage-openshift-monitor** is responsible for creating the VSP constructs (Zone, Subnets, etc.) needed to provide appropriate policies and networking for pods created in an OpenShift Container Platform project. When a project is created, a default zone and default subnet for that project are created by **nuage-openshift-monitor**. When the default subnet created for a given project gets depleted, **nuage-openshift-monitor** dynamically creates additional subnets.



NOTE

A separate VSP Zone is created for each OpenShift Container Platform project ensuring isolation amongst the projects.

16.3. OPERATIONS WORKFLOW

This workflow is used by operations teams rolling out applications. In this workflow, the network and security policies are first configured on the VSD in accordance with the rules set by the organization to deploy applications. Administrative users can potentially create multiple zones and subnets and map them to the same project using labels. While spinning up the pods, the user can use the Nuage Labels to specify what network a pod needs to attach to and what network policies need to be applied to it. This allows for deployments where inter- and intra-project traffic can be controlled in a fine-grained manner. For example, inter-project communication is enabled on a project by project basis. This may be used to connect projects to common services that are deployed in a shared project.

16.4. INSTALLATION

The VSP integration with OpenShift Container Platform works for both virtual machines (VMs) and bare metal OpenShift Container Platform installations.

16.4.1. Installation for a Single Master

In the Ansible nodes file, specify the following parameters in order to set up Nuage VSP as the network plug-in:

```
# Nuage specific parameters
openshift_use_openshift_sdn=False
openshift_use_nuage=True
os_sdn_network_plugin_name='nuage/vsp-openshift'
openshift_node_proxy_mode='userspace'
nuage_openshift_monitor_rest_server_port=9443

# VSP related parameters
vsd_api_url=https://192.168.103.200:8443
vsp_version=v4_0
enterprise=nuage
domain=openshift
vsc_active_ip=192.168.103.201
vsc_standby_ip=192.168.103.202
uplink_interface=eth0

# rpm locations
nuage_openshift_rpm=http://location_of_rpm_server/openshift/RPMS/x86_64/nuage-openshift-monitor-4.0.X.1830.el7.centos.x86_64.rpm
vrs_rpm=http://location_of_rpm_server/openshift/RPMS/x86_64/nuage-openvswitch-4.0.X.225.el7.x86_64.rpm
plugin_rpm=http://location_of_rpm_server/openshift/RPMS/x86_64/vsp-openshift-4.0.X1830.el7.centos.x86_64.rpm

# Optional parameters
nuage_interface_mtu=1460
nuage_master_adminusername=admin
nuage_master_adminuserpasswd=admin
nuage_master_cspadminpasswd=admin
nuage_openshift_monitor_log_dir=/var/log/nuage-openshift-monitor

# Required for brownfield install (where a {product-title} cluster exists
without Nuage as the networking plugin)
nuage_dockker_bridge=lbr0
```

16.4.2. Installation for Multiple Masters (HA)

An environment with High Availability (HA) can be configured with multiple masters and multiple nodes.

Nuage VSP integration in multi-master mode only supports the native HA configuration method described in this section. This can be combined with any load balancing solution, the default being HAProxy. The inventory file contains three master hosts, the nodes, an etcd server, and a host that functions as the HAProxy to balance the master API on all master hosts. The HAProxy host is defined in the [lb] section of the inventory file enabling Ansible to automatically install and configure HAProxy as the load balancing solution.

In the Ansible nodes file, the following parameters need to be specified in order to setup Nuage VSP as the network plug-in:

```
# Create and OSEv3 group that contains masters, nodes, load-balancers,
and etcd hosts
masters
```

```

nodes
etcd
lb

# Nuage specific parameters
openshift_use_openshift_sdn=False
openshift_use_nuage=True
os_sdn_network_plugin_name='nuage/vsp-openshift'
openshift_node_proxy_mode='userspace'

# VSP related parameters
vsd_api_url=https://192.168.103.200:8443
vsp_version=v4_0
enterprise=nuage
domain=openshift
vsc_active_ip=192.168.103.201
vsc_standby_ip=192.168.103.202
uplink_interface=eth0

# rpm locations
nuage_openshift_rpm=http://location_of_rpm_server/openshift/RPMS/x86_64/nuage-openshift-monitor-4.0.X.1830.el7.centos.x86_64.rpm
vrs_rpm=http://location_of_rpm_server/openshift/RPMS/x86_64/nuage-openvswitch-4.0.X.225.el7.x86_64.rpm
plugin_rpm=http://location_of_rpm_server/openshift/RPMS/x86_64/vsp-openshift-4.0.X.1830.el7.centos.x86_64.rpm

# Required for Nuage Monitor REST server and HA
openshift_master_cluster_method=native
openshift_master_cluster_hostname=lb.nuageopenshift.com
openshift_master_cluster_public_hostname=lb.nuageopenshift.com
nuage_openshift_monitor_rest_server_port=9443

# Optional parameters
nuage_interface_mtu=1460
nuage_master_adminusername='admin's user-name'
nuage_master_adminuserpasswd='admin's password'
nuage_master_cspadminpasswd='csp admin password'
nuage_openshift_monitor_log_dir=/var/log/nuage-openshift-monitor

# Required for brownfield install (where a {product-title} cluster exists
without Nuage as the networking plugin)
nuage_dockker_bridge=lbr0

# Specify master hosts
[masters]
fqdn_of_master_1
fqdn_of_master_2
fqdn_of_master_3

# Specify load balancer host
[lb]
fqdn_of_load_balancer

```

CHAPTER 17. CONFIGURING FOR AWS

17.1. OVERVIEW

OpenShift Container Platform can be configured to access an [AWS EC2 infrastructure](#), including [using AWS volumes as persistent storage](#) for application data. After AWS is configured properly, some additional configurations will need to be completed on the OpenShift Container Platform hosts.

17.2. CONFIGURING AWS VARIABLES

To set the required AWS variables, create a `/etc/aws/aws.conf` file with the following contents on all of your OpenShift Container Platform hosts, both masters and nodes:

```
[Global]
Zone = us-east-1c 1
```

- 1 This is the Availability Zone of your AWS Instance and where your EBS Volume resides; this information is obtained from the AWS Management Console.

17.3. CONFIGURING OPENSIFT CONTAINER PLATFORM MASTERS FOR AWS

You can set the AWS configuration on your OpenShift Container Platform master hosts in two ways:

- [using Ansible and the advanced installation tool](#)
- [manually, by modifying the `master-config.yaml` file](#)

17.3.1. Configuring OpenShift Container Platform for AWS with Ansible

During [advanced installations](#), AWS can be configured using [the `openshift_cloudprovider_aws_access_key`, `openshift_cloudprovider_aws_secret_key`, and `openshift_cloudprovider_kind` parameters](#), which are configurable in the inventory file.

Example 17.1. Example AWS Configuration with Ansible

```
# Cloud Provider Configuration
#
# Note: You may make use of environment variables rather than store
# sensitive configuration within the ansible inventory.
# For example:
#openshift_cloudprovider_aws_access_key="{{
lookup('env', 'AWS_ACCESS_KEY_ID') }}"
#openshift_cloudprovider_aws_secret_key="{{
lookup('env', 'AWS_SECRET_ACCESS_KEY') }}"
#
# AWS
#openshift_cloudprovider_kind=aws
# Note: IAM profiles may be used instead of storing API credentials on
```

```
disk.
#openshift_cloudprovider_aws_access_key=aws_access_key_id
#openshift_cloudprovider_aws_secret_key=aws_secret_access_key
```

NOTE

When Ansible configures AWS, the following files are created for you:

- */etc/aws/aws.conf*
- */etc/origin/master/master-config.yaml*
- */etc/origin/node/node-config.yaml*
- */etc/sysconfig/atomic-openshift-master*
- */etc/sysconfig/atomic-openshift-node*

17.3.2. Manually Configuring OpenShift Container Platform Masters for AWS

Edit or [create](#) the master configuration file on all masters (*/etc/origin/master/master-config.yaml* by default) and update the contents of the **apiServerArguments** and **controllerArguments** sections:

```
kubernetesMasterConfig:
  ...
  apiServerArguments:
    cloud-provider:
      - "aws"
    cloud-config:
      - "/etc/aws/aws.conf"
  controllerArguments:
    cloud-provider:
      - "aws"
    cloud-config:
      - "/etc/aws/aws.conf"
```

IMPORTANT

When triggering a containerized installation, only the directories of */etc/origin* and */var/lib/origin* are mounted to the master and node container. Therefore, **aws.conf** should be in */etc/origin/* instead of */etc/*.

17.3.3. Manually Configuring OpenShift Container Platform Nodes for AWS

Edit or [create](#) the node configuration file on all nodes (*/etc/origin/node/node-config.yaml* by default) and update the contents of the **kubeletArguments** section:

```
kubeletArguments:
  cloud-provider:
    - "aws"
  cloud-config:
    - "/etc/aws/aws.conf"
```


**IMPORTANT**

When triggering a containerized installation, only the directories of */etc/origin* and */var/lib/origin* are mounted to the master and node container. Therefore, *aws.conf* should be in */etc/origin/* instead of */etc/*.

17.4. SETTING KEY VALUE ACCESS PAIRS

Make sure the following environment variables are set in the */etc/sysconfig/atomic-openshift-master* file on masters and the */etc/sysconfig/atomic-openshift-node* file on nodes:

```
AWS_ACCESS_KEY_ID=<key_ID>
AWS_SECRET_ACCESS_KEY=<secret_key>
```

**NOTE**

Access keys are obtained when setting up your AWS IAM user.

17.5. APPLYING CONFIGURATION CHANGES

Start or restart OpenShift Container Platform services on all master and node hosts to apply your configuration changes:

```
# systemctl restart atomic-openshift-master
# systemctl restart atomic-openshift-node
```

Switching from not using a cloud provider to using a cloud provider produces an error message. Adding the cloud provider tries to delete the node because the node switches from using the **hostname** as the **externalID** (which would have been the case when no cloud provider was being used) to using the AWS **instance-id** (which is what the AWS cloud provider specifies). To resolve this issue:

1. Log in to the CLI as a cluster administrator.
2. Delete the nodes:

```
$ oc delete node <node_name>
```

3. On each node host, restart the **atomic-openshift-node** service.
4. Add back any [labels on each node](#) that you previously had.

CHAPTER 18. CONFIGURING FOR OPENSTACK

18.1. OVERVIEW

When deployed on [OpenStack](#), OpenShift Container Platform can be configured to access OpenStack infrastructure, including [using OpenStack Cinder volumes as persistent storage](#) for application data.

18.2. CONFIGURING OPENSTACK VARIABLES

To set the required OpenStack variables, create a `/etc/cloud.conf` file with the following contents on all of your OpenShift Container Platform hosts, both masters and nodes:

```
[Global]
auth-url = <OS_AUTH_URL>
username = <OS_USERNAME>
password = <password>
domain-id = <OS_USER_DOMAIN_ID>
tenant-id = <OS_TENANT_ID>
region = <OS_REGION_NAME>

[LoadBalancer]
subnet-id = <UUID of the load balancer subnet>
```

Consult your OpenStack administrators for values of the `OS_` variables, which are commonly used in OpenStack configuration.

18.3. CONFIGURING OPENSIFT CONTAINER PLATFORM MASTERS FOR OPENSTACK

You can set an OpenStack configuration on your OpenShift Container Platform master and node hosts in two different ways:

- [Using Ansible and the advanced installation tool](#)
- Manually, by [modifying the `master-config.yaml`](#) and [node-config.yaml](#) files.

18.3.1. Configuring OpenShift Container Platform for OpenStack with Ansible

During [advanced installations](#), OpenStack can be configured using [the following parameters](#), which are configurable in the inventory file:

- `openshift_cloudprovider_kind`
- `openshift_cloudprovider_openstack_auth_url`
- `openshift_cloudprovider_openstack_username`
- `openshift_cloudprovider_openstack_password`
- `openshift_cloudprovider_openstack_domain_id`
- `openshift_cloudprovider_openstack_domain_name`

- `openshift_cloudprovider_openstack_tenant_id`
- `openshift_cloudprovider_openstack_tenant_name`
- `openshift_cloudprovider_openstack_region`
- `openshift_cloudprovider_openstack_lb_subnet_id`


Example 18.1. Example OpenStack Configuration with Ansible

```
# Cloud Provider Configuration
#
# Note: You may make use of environment variables rather than store
# sensitive configuration within the ansible inventory.
# For example:
#openshift_cloudprovider_openstack_username="{{ lookup('env', 'USERNAME')
}}"
#openshift_cloudprovider_openstack_password="{{ lookup('env', 'PASSWORD')
}}"
#
# Openstack
#openshift_cloudprovider_kind=openstack
#openshift_cloudprovider_openstack_auth_url=http://openstack.example.com
:35357/v2.0/
#openshift_cloudprovider_openstack_username=username
#openshift_cloudprovider_openstack_password=password
#openshift_cloudprovider_openstack_domain_id=domain_id
#openshift_cloudprovider_openstack_domain_name=domain_name
#openshift_cloudprovider_openstack_tenant_id=tenant_id
#openshift_cloudprovider_openstack_tenant_name=tenant_name
#openshift_cloudprovider_openstack_region=region
#openshift_cloudprovider_openstack_lb_subnet_id=subnet_id
```

18.3.2. Manually Configuring OpenShift Container Platform Masters for OpenStack

Edit or [create](#) the master configuration file on all masters (`/etc/origin/master/master-config.yaml` by default) and update the contents of the **apiServerArguments** and **controllerArguments** sections:

```
kubernetesMasterConfig:
  ...
  apiServerArguments:
    cloud-provider:
      - "openstack"
    cloud-config:
      - "/etc/cloud.conf"
  controllerArguments:
    cloud-provider:
      - "openstack"
    cloud-config:
      - "/etc/cloud.conf"
```

**IMPORTANT**

When triggering a containerized installation, only the directories of */etc/origin* and */var/lib/origin* are mounted to the master and node container. Therefore, **cloud.conf** should be in */etc/origin/* instead of */etc/*.

18.3.3. Manually Configuring OpenShift Container Platform Nodes for OpenStack

Edit or [create](#) the node configuration file on all nodes (*/etc/origin/node/node-config.yaml* by default) and update the contents of the **kubeletArguments** and **nodeName** sections:

```
nodeName:
  <instance_name> ❶

kubeletArguments:
  cloud-provider:
    - "openstack"
  cloud-config:
    - "/etc/cloud.conf"
```

❶ Name of the OpenStack instance where the node runs (i.e., name of the virtual machine)

**IMPORTANT**

When triggering a containerized installation, only the directories of */etc/origin* and */var/lib/origin* are mounted to the master and node container. Therefore, **cloud.conf** should be in */etc/origin/* instead of */etc/*.

CHAPTER 19. CONFIGURING FOR GCE

19.1. OVERVIEW

OpenShift Container Platform can be configured to access a [Google Compute Engine \(GCE\) infrastructure](#), including [using GCE volumes as persistent storage](#) for application data. After GCE is configured properly, some additional configurations will need to be completed on the OpenShift Container Platform hosts.

19.2. CONFIGURING MASTERS

You can set the GCE configuration on your OpenShift Container Platform master hosts in two ways:

- [Using Ansible and the advanced installation tool.](#)
- [Manually by modifying the *master-config.yaml* file.](#)

19.2.1. Configuring OpenShift Container Platform Masters for GCE with Ansible

During [advanced installations](#), GCE can be configured using the `openshift_cloudprovider_kind` parameter, which is configurable in the inventory file.

Example GCE Configuration with Ansible

```
# Cloud Provider Configuration
#
openshift_cloudprovider_kind=gce
```

NOTE

When Ansible configures GCE, the following files are created for you:

- `/etc/origin/cloudprovider/gce.conf`
- `/etc/origin/master/master-config.yaml`
- `/etc/origin/node/node-config.yaml`

The advanced installation configures multizone support by default. If you want single-zone support, edit the `/etc/origin/cloudprovider/gce.conf` as shown in [Configuring Multizone Support in a GCE Deployment](#).


19.2.2. Manually Configuring OpenShift Container Platform Masters for GCE

To configure the OpenShift Container Platform masters for GCE:

1. Edit or [create](#) the master configuration file (`/etc/origin/master/master-config.yaml` by default) on all masters and update the contents of the `apiServerArguments` and `controllerArguments` sections:

```
kubernetesMasterConfig:
  ...
  apiServerArguments:
```

```
cloud-provider:
  - "gce"
cloud-config:
  - "/etc/origin/cloudprovider/gce.conf"
controllerArguments:
  cloud-provider:
    - "gce"
  cloud-config:
    - "/etc/origin/cloudprovider/gce.conf"
```

**IMPORTANT**

When triggering a containerized installation, only the directories of **/etc/origin** and **/var/lib/origin** are mounted to the master and node container. Therefore, **master-config.yaml** should be in **/etc/origin/master** instead of **/etc/**.

2. Start or restart the OpenShift Container Platform services:


```
# systemctl restart atomic-openshift-master
```

19.3. CONFIGURING NODES

To configure the OpenShift Container Platform nodes for GCE:

1. Edit or [create](#) the node configuration file (**/etc/origin/node/node-config.yaml** by default) on all nodes and update the contents of the **kubeletArguments** section:

```
kubeletArguments:
  cloud-provider:
    - "gce"
  cloud-config:
    - "/etc/origin/cloudprovider/gce.conf"
```

**IMPORTANT**

When triggering a containerized installation, only the directories of **/etc/origin** and **/var/lib/origin** are mounted to the master and node container. Therefore, **node-config.yaml** should be in **/etc/origin/node** instead of **/etc/**.

2. Start or restart the OpenShift Container Platform services all nodes.

```
# systemctl restart atomic-openshift-node
```

19.4. CONFIGURING MULTIZONE SUPPORT IN A GCE DEPLOYMENT

If manually configuring GCE, multizone support is not configured by default.

**NOTE**

The advanced installation configures multizone support by default.

If you want multizone support:

1. Edit or create a **`/etc/origin/cloudprovider/gce.conf`** file on all of your OpenShift Container Platform hosts, both masters and nodes.
2. Add the following contents:

```
[Global]
multizone = true
```

3. Start or restart the OpenShift Container Platform services on the master and all nodes.

```
# systemctl restart atomic-openshift-master
# systemctl restart atomic-openshift-node
```

To return to single-zone support, set the **`multizone`** value to **`false`**.

CHAPTER 20. CONFIGURING PERSISTENT STORAGE

20.1. OVERVIEW

The Kubernetes [persistent volume](#) framework allows you to provision an OpenShift Container Platform cluster with persistent storage using networked storage available in your environment. This can be done after completing the initial OpenShift Container Platform installation depending on your application needs, giving users a way to request those resources without having any knowledge of the underlying infrastructure.

These topics show how to configure persistent volumes in OpenShift Container Platform using the following supported volume plug-ins:

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

20.2. PERSISTENT STORAGE USING NFS

20.2.1. Overview

OpenShift Container Platform clusters can be provisioned with [persistent storage](#) using NFS. Persistent volumes (PVs) and persistent volume claims (PVCs) provide a convenient method for sharing a volume across a project. While the NFS-specific information contained in a PV definition could also be defined directly in a pod definition, doing so does not create the volume as a distinct cluster resource, making the volume more susceptible to conflicts.

This topic covers the specifics of using the NFS persistent storage type. Some familiarity with OpenShift Container Platform and [NFS](#) is beneficial. See the [Persistent Storage](#) concept topic for details on the OpenShift Container Platform persistent volume (PV) framework in general.

20.2.2. Provisioning

Storage must exist in the underlying infrastructure before it can be mounted as a volume in OpenShift Container Platform. To provision NFS volumes, a list of NFS servers and export paths are all that is required.

You must first create an object definition for the PV:

Example 20.1. PV Object Definition Using NFS


```

apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv0001 ❶
spec:
  capacity:
    storage: 5Gi ❷
  accessModes:
    - ReadWriteOnce ❸
  nfs: ❹
    path: /tmp ❺
    server: 172.17.0.2 ❻
  persistentVolumeReclaimPolicy: Recycle ❼

```

- ❶ The name of the volume. This is the PV identity in various **oc <command> pod** commands.
- ❷ The amount of storage allocated to this volume.
- ❸ Though this appears to be related to controlling access to the volume, it is actually used similarly to labels and used to match a PVC to a PV. Currently, no access rules are enforced based on the **accessModes**.
- ❹ The volume type being used, in this case the **nfs** plug-in.
- ❺ The path that is exported by the NFS server.
- ❻ The host name or IP address of the NFS server.
- ❼ The reclaim policy for the PV. This defines what happens to a volume when released from its claim. Valid options are **Retain** (default) and **Recycle**. See [Reclaiming Resources](#).



NOTE

Each NFS volume must be mountable by all schedulable nodes in the cluster.

Save the definition to a file, for example **nfs-pv.yaml**, and create the PV:

```

$ oc create -f nfs-pv.yaml
persistentvolume "pv0001" created

```

Verify that the PV was created:

```

# oc get pv
NAME          CLAIM          REASON          AGE          LABELS          CAPACITY          ACCESSMODES          STATUS
pv0001        31s            <none>          5368709120   RWO             Available

```

The next step can be to create a PVC, which binds to the new PV:

Example 20.2. PVC Object Definition

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: nfs-claim1
spec:
  accessModes:
    - ReadWriteOnce ❶
  resources:
    requests:
      storage: 1Gi ❷

```

❶ As mentioned above for PVs, the **accessModes** do not enforce security, but rather act as labels to match a PV to a PVC.

❷ This claim looks for PVs offering **1Gi** or greater capacity.

Save the definition to a file, for example ***nfs-claim.yaml***, and create the PVC:

```
# oc create -f nfs-claim.yaml
```


20.2.3. Enforcing Disk Quotas

You can use disk partitions to enforce disk quotas and size constraints. Each partition can be its own export. Each export is one PV. OpenShift Container Platform enforces unique names for PVs, but the uniqueness of the NFS volume's server and path is up to the administrator.

Enforcing quotas in this way allows the developer to request persistent storage by a specific amount (for example, 10Gi) and be matched with a corresponding volume of equal or greater capacity.

20.2.4. NFS Volume Security

This section covers NFS volume security, including matching permissions and SELinux considerations. The user is expected to understand the basics of POSIX permissions, process UIDs, supplemental groups, and SELinux.



NOTE

See the full [Volume Security](#) topic before implementing NFS volumes.

Developers request NFS storage by referencing, in the **volumes** section of their pod definition, either a PVC by name or the NFS volume plug-in directly.

The **/etc/exports** file on the NFS server contains the accessible NFS directories. The target NFS directory has POSIX owner and group IDs. The OpenShift Container Platform NFS plug-in mounts the container's NFS directory with the same POSIX ownership and permissions found on the exported NFS directory. However, the container is not run with its effective UID equal to the owner of the NFS mount, which is the desired behavior.


As an example, if the target NFS directory appears on the NFS server as:

```
# ls -lZ /opt/nfs -d
```

```
drwxrws---. nfsnobody 5555 unconfined_u:object_r:usr_t:s0 /opt/nfs

# id nfsnobody
uid=65534(nfsnobody) gid=65534(nfsnobody) groups=65534(nfsnobody)
```

Then the container must match SELinux labels, and either run with a UID of **65534** (**nfsnobody** owner) or with **5555** in its supplemental groups in order to access the directory.



NOTE

The owner ID of 65534 is used as an example. Even though NFS's **root_squash** maps **root** (0) to **nfsnobody** (65534), NFS exports can have arbitrary owner IDs. Owner 65534 is not required for NFS exports.

20.2.4.1. Group IDs

The recommended way to handle NFS access (assuming it is not an option to change permissions on the NFS export) is to use supplemental groups. Supplemental groups in OpenShift Container Platform are used for shared storage, of which NFS is an example. In contrast, block storage, such as Ceph RBD or iSCSI, use the **fsGroup** SCC strategy and the **fsGroup** value in the pod's **securityContext**.



NOTE

It is generally preferable to use supplemental group IDs to gain access to persistent storage versus using [user IDs](#). Supplemental groups are covered further in the full [Volume Security](#) topic.

Because the group ID on the [example target NFS directory](#) shown above is 5555, the pod can define that group ID using **supplementalGroups** under the pod-level **securityContext** definition. For example:

```
spec:
  containers:
    - name:
      ...
  securityContext: ❶
    supplementalGroups: [5555] ❷
```

- ❶ **securityContext** must be defined at the pod level, not under a specific container.
- ❷ An array of GIDs defined for the pod. In this case, there is one element in the array; additional GIDs would be comma-separated.

Assuming there are no custom SCCs that might satisfy the pod's requirements, the pod likely matches the **restricted** SCC. This SCC has the **supplementalGroups** strategy set to **RunAsAny**, meaning that any supplied group ID is accepted without range checking.

As a result, the above pod passes admissions and is launched. However, if group ID range checking is desired, a custom SCC, as described in [pod security and custom SCCs](#), is the preferred solution. A custom SCC can be created such that minimum and maximum group IDs are defined, group ID range checking is enforced, and a group ID of 5555 is allowed.

**NOTE**

To use a custom SCC, you must first add it to the appropriate service account. For example, use the **default** service account in the given project unless another has been specified on the pod specification. See [Add an SCC to a User, Group, or Project](#) for details.

20.2.4.2. User IDs

User IDs can be defined in the container image or in the pod definition. The full [Volume Security](#) topic covers controlling storage access based on user IDs, and should be read prior to setting up NFS persistent storage.

**NOTE**

It is generally preferable to use [supplemental group IDs](#) to gain access to persistent storage versus using user IDs.

In the [example target NFS directory](#) shown above, the container needs its UID set to 65534 (ignoring group IDs for the moment), so the following can be added to the pod definition:

```
spec:
  containers: ❶
  - name:
    ...
    securityContext:
      runAsUser: 65534 ❷
```


❶ Pods contain a **securityContext** specific to each container (shown here) and a pod-level **securityContext** which applies to all containers defined in the pod.

❷ 65534 is the **nfsnobody** user.

Assuming the **default** project and the **restricted** SCC, the pod's requested user ID of 65534 is not allowed, and therefore the pod fails. The pod fails for the following reasons:

- It requests 65534 as its user ID.
- All SCCs available to the pod are examined to see which SCC allows a user ID of 65534 (actually, all policies of the SCCs are checked but the focus here is on user ID).
- Because all available SCCs use **MustRunAsRange** for their **runAsUser** strategy, UID range checking is required.
- 65534 is not included in the SCC or project's user ID range.

It is generally considered a good practice not to modify the predefined SCCs. The preferred way to fix this situation is to create a custom SCC, as described in the full [Volume Security](#) topic. A custom SCC can be created such that minimum and maximum user IDs are defined, UID range checking is still enforced, and the UID of 65534 is allowed.

**NOTE**

To use a custom SCC, you must first add it to the appropriate service account. For example, use the **default** service account in the given project unless another has been specified on the pod specification. See [Add an SCC to a User, Group, or Project](#) for details.

20.2.4.3. SELinux**NOTE**

See the full [Volume Security](#) topic for information on controlling storage access in conjunction with using SELinux.

By default, SELinux does not allow writing from a pod to a remote NFS server. The NFS volume mounts correctly, but is read-only.

To enable writing to NFS volumes with SELinux enforcing on each node, run:

```
# setsebool -P virt_use_nfs 1
```

The **-P** option above makes the bool persistent between reboots.

The **virt_use_nfs** boolean is defined by the **docker-selinux** package. If an error is seen indicating that this bool is not defined, ensure this package has been installed.

20.2.4.4. Export Settings

In order to enable arbitrary container users to read and write the volume, each exported volume on the NFS server should conform to the following conditions:

- Each export must be:

```
/<example_fs> *(rw,root_squash)
```

- The firewall must be configured to allow traffic to the mount point.
 - For NFSv4, configure the default port **2049** (**nfs**) and port **111** (**portmapper**).

NFSv4

```
# iptables -I INPUT 1 -p tcp --dport 2049 -j ACCEPT
# iptables -I INPUT 1 -p tcp --dport 111 -j ACCEPT
```

- For NFSv3, there are three ports to configure: **2049** (**nfs**), **20048** (**mountd**), and **111** (**portmapper**).

NFSv3

```
# iptables -I INPUT 1 -p tcp --dport 2049 -j ACCEPT
# iptables -I INPUT 1 -p tcp --dport 20048 -j ACCEPT
# iptables -I INPUT 1 -p tcp --dport 111 -j ACCEPT
```

- The NFS export and directory must be set up so that it is accessible by the target pods. Either set the export to be owned by the container's primary UID, or supply the pod group access using **supplementalGroups**, as shown in [Group IDs](#) above. See the full [Volume Security](#) topic for additional pod security information as well.

20.2.5. Reclaiming Resources

NFS implements the OpenShift Container Platform **Recyclable** plug-in interface. Automatic processes handle reclamation tasks based on policies set on each persistent volume.

By default, PVs are set to **Retain**. NFS volumes which are set to **Recycle** are scrubbed (i.e., `rm -rf` is run on the volume) after being released from their claim (i.e., after the user's **PersistentVolumeClaim** bound to the volume is deleted). Once recycled, the NFS volume can be bound to a new claim.

Once claim to a PV is released (that is, the PVC is deleted), the PV object should not be re-used. Instead, a new PV should be created with the same basic volume details as the original.

For example, the administrator creates a PV named **nfs1**:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: nfs1
spec:
  capacity:
    storage: 1Mi
  accessModes:
    - ReadWriteMany
  nfs:
    server: 192.168.1.1
    path: "/"
```

The user creates **PVC1**, which binds to **nfs1**. The user then deletes **PVC1**, releasing claim to **nfs1**, which causes **nfs1** to be **Released**. If the administrator wishes to make the same NFS share available, they should create a new PV with the same NFS server details, but a different PV name:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: nfs2
spec:
  capacity:
    storage: 1Mi
  accessModes:
    - ReadWriteMany
  nfs:
    server: 192.168.1.1
    path: "/"
```

Deleting the original PV and re-creating it with the same name is discouraged. Attempting to manually change the status of a PV from **Released** to **Available** causes errors and potential data loss.



NOTE

A PV with retention policy of **Recycle** scrubs (**rm -rf**) the data and marks it as **Available** for claim. The **Recycle** retention policy is deprecated starting in OpenShift Container Platform 3.6 and should be avoided. Anyone using recycler should use dynamic provision and volume deletion instead.

20.2.6. Automation

Clusters can be provisioned with persistent storage using NFS in the following ways:

- [Enforce storage quotas](#) using disk partitions.
- Enforce security by [restricting volumes](#) to the project that has a claim to them.
- Configure [reclamation of discarded resources](#) for each PV.

There are many ways that you can use scripts to automate the above tasks. You can use an [example Ansible playbook](#) to help you get started.

20.2.7. Additional Configuration and Troubleshooting

Depending on what version of NFS is being used and how it is configured, there may be additional configuration steps needed for proper export and security mapping. The following are some that may apply:

NFSv4 mount incorrectly shows all files with ownership of nobody:nobody	<ul style="list-style-type: none"> • Could be attributed to the ID mapping settings (/etc/idmapd.conf) on your NFS • See this Red Hat Solution.
Disabling ID mapping on NFSv4	<ul style="list-style-type: none"> • On both the NFS client and server, run: <pre># echo 'Y' > /sys/module/nfsd/parameters/nfs4_disable _idmapping</pre>

20.3. PERSISTENT STORAGE USING GLUSTERFS

20.3.1. Overview

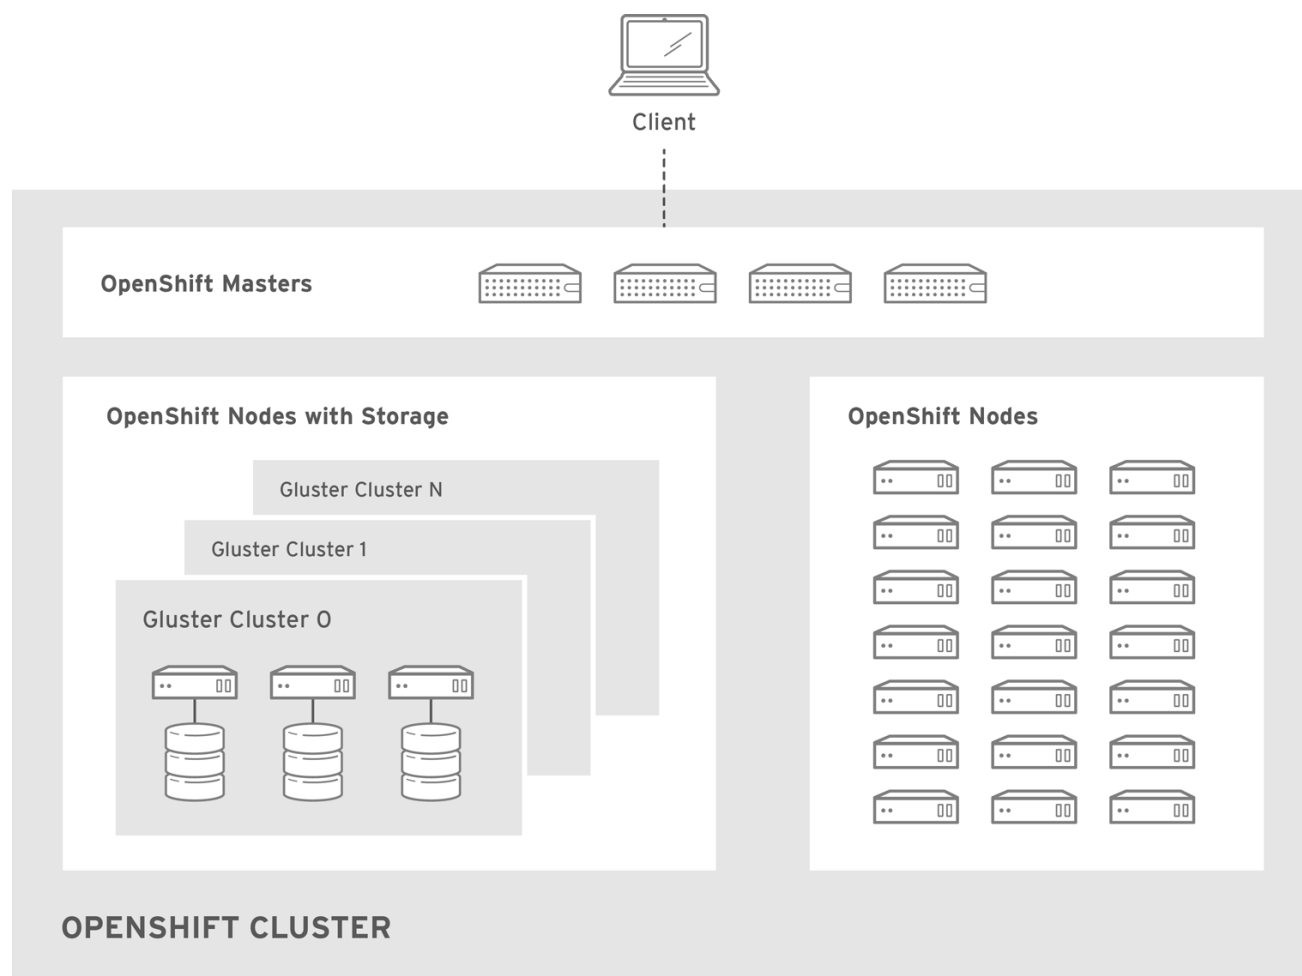
You can configure your OpenShift Container Platform cluster to use Red Hat Gluster Storage as persistent storage for containerized applications. There are two deployment solutions available when using Red Hat Gluster Storage, using either a containerized or dedicated storage cluster. This topic focuses mainly on the persistent volume plug-in solution using a dedicated Red Hat Gluster Storage cluster.

20.3.1.1. Containerized Red Hat Gluster Storage

Starting with the Red Hat Gluster Storage 3.1 update 3 release, you can deploy containerized Red Hat

Gluster Storage directly on OpenShift Container Platform. Containerized Red Hat Gluster Storage converged with OpenShift Container Platform addresses the use case where containerized applications require both shared file storage and the flexibility of a converged infrastructure with compute and storage instances being scheduled and run from the same set of hardware.

Figure 20.1. Architecture - Red Hat Gluster Storage Container Converged with OpenShift



OPENSIFT_412816_0716

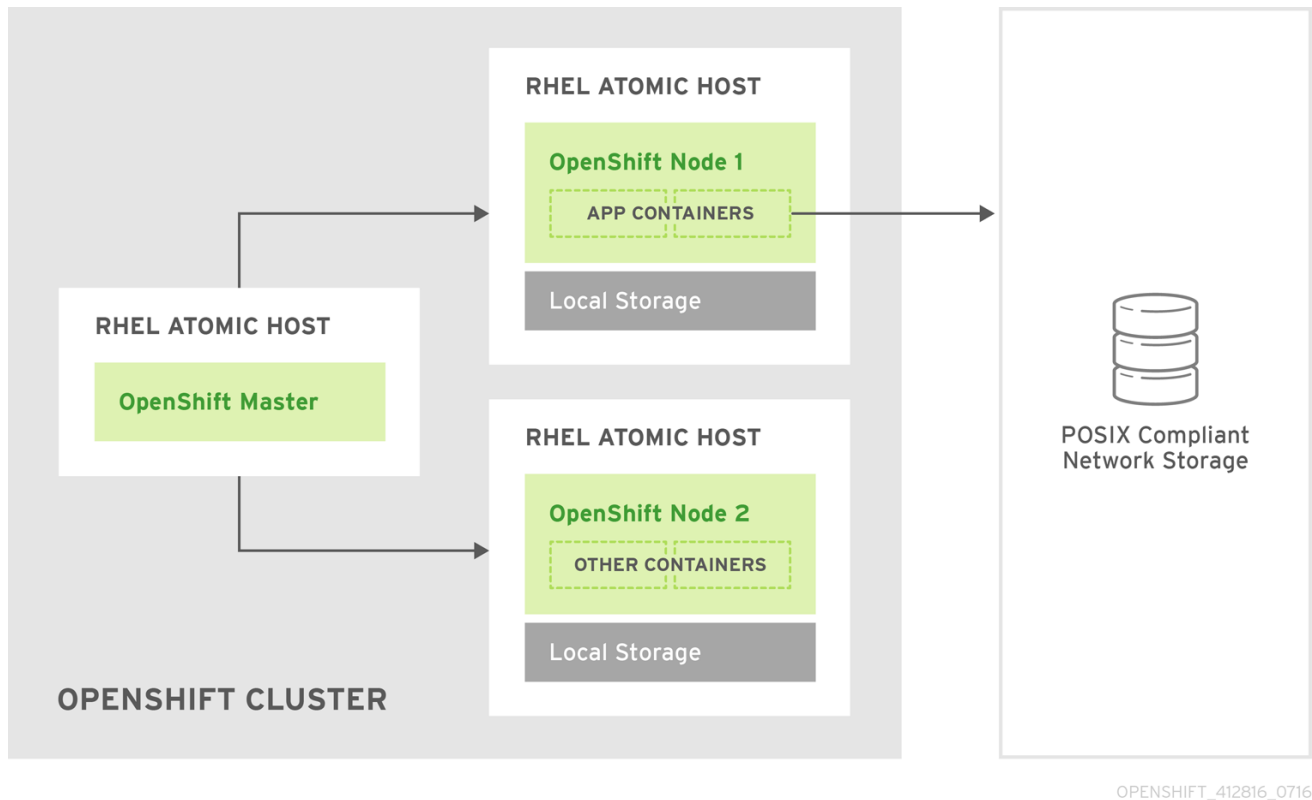
Step-by-step instructions for this containerized solution are provided separately in the following Red Hat Gluster Storage documentation:

[Container-Native Storage for OpenShift Container Platform](#)

20.3.1.2. Dedicated Storage Cluster


If you have a dedicated Red Hat Gluster Storage cluster available in your environment, you can configure OpenShift Container Platform's Gluster volume plug-in. The dedicated storage cluster delivers persistent Red Hat Gluster Storage file storage for containerized applications over the network. The applications access storage served out from the storage clusters through common storage protocols.

Figure 20.2. Architecture - Dedicated Red Hat Gluster Storage Cluster Using the OpenShift Container Platform Volume Plug-in



This solution is a conventional deployment where containerized compute applications run on an OpenShift Container Platform cluster. The remaining sections in this topic provide the step-by-step instructions for the dedicated Red Hat Gluster Storage solution.

This topic presumes some familiarity with OpenShift Container Platform and GlusterFS; see the [Red Hat Gluster Storage 3 Administration Guide](#) for more on GlusterFS. See the [Persistent Storage](#) topic for details on the OpenShift Container Platform PV framework in general.



IMPORTANT

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

20.3.2. Support Requirements

The following requirements must be met to create a supported integration of Red Hat Gluster Storage and OpenShift Container Platform.

20.3.2.1. Supported Operating Systems

The following table lists the supported versions of OpenShift Container Platform with Red Hat Gluster Storage Server.

Red Hat Gluster Storage	OpenShift Container Platform
3.1.3	3.1 or later

20.3.2.2. Environment Requirements

The environment requirements for OpenShift Container Platform and Red Hat Gluster Storage are described in this section.

Red Hat Gluster Storage

- All installations of Red Hat Gluster Storage must have valid subscriptions to Red Hat Network channels and Subscription Management repositories.
- Red Hat Gluster Storage installations must adhere to the requirements laid out in the [Red Hat Gluster Storage Installation Guide](#).
- Red Hat Gluster Storage installations must be completely up to date with the latest patches and upgrades. Refer to the [Red Hat Gluster Storage 3.1 Installation Guide](#) to upgrade to the latest version.
- The versions of OpenShift Container Platform and Red Hat Gluster Storage integrated must be compatible, according to the information in [Supported Operating Systems](#).
- A fully-qualified domain name (FQDN) must be set for each hypervisor and Red Hat Gluster Storage server node. Ensure that correct DNS records exist, and that the FQDN is resolvable via both forward and reverse DNS lookup.

Red Hat OpenShift Enterprise

- All installations of OpenShift Container Platform must have valid subscriptions to Red Hat Network channels and Subscription Management repositories.
- OpenShift Container Platform installations must adhere to the requirements laid out in the [Installation and Configuration](#) documentation.
- The OpenShift Container Platform cluster must be up and running.
- A user with **cluster-admin** permissions must be created.
- All OpenShift Container Platform nodes on RHEL systems must have the **glusterfs-fuse** RPM installed, which should match the version of Red Hat Gluster Storage server running in the containers. For more information on installing **glusterfs-fuse**, see [Native Client](#) in the Red Hat Gluster Storage Administration Guide.

20.3.3. Provisioning

To provision GlusterFS volumes the following are required:

- An existing storage device in your underlying infrastructure.
- A distinct list of servers (IP addresses) in the Gluster cluster, to be defined as endpoints.
- A service, to persist the endpoints (optional).
- An existing Gluster volume to be referenced in the persistent volume object.
- **glusterfs-fuse** installed on each schedulable OpenShift Container Platform node in your cluster:

```
# yum install glusterfs-fuse
```

**NOTE**

Persistent volumes (PVs) and persistent volume claims (PVCs) can share volumes across a single project. While the GlusterFS-specific information contained in a PV definition could also be defined directly in a pod definition, doing so does not create the volume as a distinct cluster resource, making the volume more susceptible to conflicts.

20.3.3.1. Creating Gluster Endpoints

An endpoints definition defines the GlusterFS cluster as **EndPoints** and includes the IP addresses of your Gluster servers. The port value can be any numeric value within the accepted range of ports. Optionally, you can create a [service](#) that persists the endpoints.

1. Define the following service:

Example 20.3. Gluster Service Definition

```
apiVersion: v1
kind: Service
metadata:
  name: glusterfs-cluster ❶
spec:
  ports:
  - port: 1
```

- ❶ This name must be defined in the endpoints definition. If using a service, then the endpoints name must match the service name.

2. Save the service definition to a file, for example *gluster-service.yaml*, then create the service:

```
$ oc create -f gluster-service.yaml
```

3. Verify that the service was created:

```
# oc get services
NAME                                CLUSTER_IP          EXTERNAL_IP  PORT(S)
SELECTOR      AGE
glusterfs-cluster  172.30.205.34      <none>       1/TCP
<none>         44s
```

4. Define the Gluster endpoints:

Example 20.4. Gluster Endpoints Definition

```
apiVersion: v1
kind: Endpoints
metadata:
  name: glusterfs-cluster ❶
subsets:
  - addresses:
    - ip: 192.168.122.221 ❷
  ports:
  - port: 1
```

```
- addresses:
  - ip: 192.168.122.222 ③
  ports:
    - port: 1 ④
```

① This name must match the service name from step 1.

② ③ The **ip** values must be the actual IP addresses of a Gluster server, not fully-qualified host names.

④ The port number is ignored.


5. Save the endpoints definition to a file, for example ***gluster-endpoints.yaml***, then create the endpoints:

```
$ oc create -f gluster-endpoints.yaml
endpoints "glusterfs-cluster" created
```

6. Verify that the endpoints were created:

```
$ oc get endpoints
NAME                ENDPOINTS                                     AGE
docker-registry     10.1.0.3:5000                                4h
glusterfs-cluster   192.168.122.221:1,192.168.122.222:1         11s
kubernetes           172.16.35.3:8443                             4d
```

20.3.3.2. Creating the Persistent Volume



NOTE

GlusterFS does not support the 'Recycle' reclaim policy.

1. Next, define the PV in an object definition before creating it in OpenShift Container Platform:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: gluster-default-volume ①
spec:
  capacity:
    storage: 2Gi ②
  accessModes: ③
    - ReadWriteMany
  glusterfs: ④
    endpoints: glusterfs-cluster ⑤
    path: myVol1 ⑥
    readOnly: false
  persistentVolumeReclaimPolicy: Retain ⑦
```

- 1 The name of the volume. This is how it is identified via [persistent volume claims](#) or from pods.
- 2 The amount of storage allocated to this volume.
- 3 **accessModes** are used as labels to match a PV and a PVC. They currently do not define any form of access control.
- 4 The volume type being used, in this case the **glusterfs** plug-in.
- 5 The endpoints name that defines the Gluster cluster created in [Creating Gluster Endpoints](#).
- 6 The Gluster volume that will be accessed, as shown in the **gluster volume status** command.
- 7 The volume reclaim policy **Retain** indicates that the volume will be preserved after the pods accessing it terminates. For GlusterFS, the accepted values include **Retain**, and **Delete**.



NOTE

Endpoints are name-spaced. Each project accessing the Gluster volume needs its own endpoints.

1. Save the definition to a file, for example **gluster-pv.yaml**, and create the persistent volume:

```
# oc create -f gluster-pv.yaml
```

2. Verify that the persistent volume was created:

```
# oc get pv
NAME                                LABELS            CAPACITY    ACCESSMODES
STATUS    CLAIM    REASON    AGE
gluster-default-volume    <none>    2147483648    RWX
Available                                2s
```

20.3.3.3. Creating the Persistent Volume Claim

Developers request GlusterFS storage by referencing either a PVC or the Gluster volume plug-in directly in the **volumes** section of a pod spec. A PVC exists only in the user's project and can only be referenced by pods within that project. Any attempt to access a PV across a project causes the pod to fail.

1. Create a PVC that will bind to the new PV:

Example 20.5. PVC Object Definition

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gluster-claim
spec:
  accessModes:
    - ReadWriteMany 1
```

```
resources:
  requests:
    storage: 1Gi 2
```

- 1 **accessModes** do not enforce security, but rather act as labels to match a PV to a PVC.
- 2 This claim will look for PVs offering **1Gi** or greater capacity.

2. Save the definition to a file, for example ***gluster-claim.yaml***, and create the PVC:

```
# oc create -f gluster-claim.yaml
```



NOTE

PVs and PVCs make sharing a volume across a project simpler. The gluster-specific information contained in the PV definition can also be defined directly in a pod specification.

20.3.4. Gluster Volume Security

This section covers Gluster volume security, including matching permissions and SELinux considerations. Understanding the basics of POSIX permissions, process UIDs, supplemental groups, and SELinux is presumed.



NOTE

See the full [Volume Security](#) topic before implementing Gluster volumes.

As an example, assume that the target Gluster volume, **HadoopVol** is mounted under ***/mnt/glusterfs/***, with the following POSIX permissions and SELinux labels:

```
# ls -lZ /mnt/glusterfs/
drwxrwx---. yarn hadoop system_u:object_r:fusefs_t:s0 HadoopVol

# id yarn
uid=592(yarn) gid=590(hadoop) groups=590(hadoop)
```

In order to access the **HadoopVol** volume, containers must match the SELinux label, and run with a UID of 592 or 590 in their supplemental groups. The OpenShift Container Platform GlusterFS plug-in mounts the volume in the container with the same POSIX ownership and permissions found on the target gluster mount, namely the owner will be **592** and group ID will be **590**. However, the container is not run with its effective UID equal to **592**, nor with its GID equal to **590**, which is the desired behavior. Instead, a container's UID and supplemental groups are determined by Security Context Constraints (SCCs) and the project defaults.

20.3.4.1. Group IDs

Configure Gluster volume access by using supplemental groups, assuming it is not an option to change permissions on the Gluster mount. Supplemental groups in OpenShift Container Platform are used for shared storage, such as GlusterFS. In contrast, block storage, such as Ceph RBD or iSCSI, use the **fsGroup** SCC strategy and the **fsGroup** value in the pod's **securityContext**.

**NOTE**

Use supplemental group IDs instead of [user IDs](#) to gain access to persistent storage. Supplemental groups are covered further in the full [Volume Security](#) topic.

The group ID on the [target Gluster mount example above](#) is 590. Therefore, a pod can define that group ID using **supplementalGroups** under the pod-level **securityContext** definition. For example:

```
spec:
  containers:
    - name:
      ...
      securityContext: ❶
        supplementalGroups: [590] ❷
```

❶ **securityContext** must be defined at the pod level, not under a specific container.

❷ An array of GIDs defined at the pod level.

Assuming there are no custom SCCs that satisfy the pod's requirements, the pod matches the **restricted** SCC. This SCC has the **supplementalGroups** strategy set to **RunAsAny**, meaning that any supplied group IDs are accepted without range checking.

As a result, the above pod will pass admissions and can be launched. However, if group ID range checking is desired, use a custom SCC, as described in [pod security and custom SCCs](#). A custom SCC can be created to define minimum and maximum group IDs, enforce group ID range checking, and allow a group ID of **590**.

20.3.4.2. User IDs

User IDs can be defined in the container image or in the pod definition. The full [Volume Security](#) topic covers controlling storage access based on user IDs, and should be read prior to setting up NFS persistent storage.

**NOTE**

Use [supplemental group IDs](#) instead of user IDs to gain access to persistent storage.

In the [target Gluster mount example above](#), the container needs a UID set to **592**, so the following can be added to the pod definition:

```
spec:
  containers: ❶
    - name:
      ...
      securityContext:
        runAsUser: 592 ❷
```

❶ Pods contain a **securityContext** specific to each container and a pod-level **securityContext**, which applies to all containers defined in the pod.

❷ The UID defined on the Gluster mount.

With the **default** project and the **restricted** SCC, a pod's requested user ID of **592** will not be allowed, and the pod will fail. This is because:

- The pod requests **592** as its user ID.
- All SCCs available to the pod are examined to see which SCC will allow a user ID of **592**.
- Because all available SCCs use **MustRunAsRange** for their **runAsUser** strategy, UID range checking is required.
- **592** is not included in the SCC or project's user ID range.

Do not modify the predefined SCCs. Instead, [create a custom SCC](#) so that minimum and maximum user IDs are defined, UID range checking is still enforced, and the UID of **592** will be allowed.

20.3.4.3. SELinux




NOTE

See the full [Volume Security](#) topic for information on controlling storage access in conjunction with using SELinux.

By default, SELinux does not allow writing from a pod to a remote Gluster server.

To enable writing to GlusterFS volumes with SELinux enforcing on each node, run:

```
$ sudo setsebool -P virt_sandbox_use_fusefs on
```



NOTE

The **virt_sandbox_use_fusefs** boolean is defined by the **docker-selinux** package. If you get an error saying it is not defined, please ensure that this package is installed.

The **-P** option makes the bool persistent between reboots.

20.4. PERSISTENT STORAGE USING OPENSTACK CINDER

20.4.1. Overview

You can provision your OpenShift Container Platform cluster with [persistent storage](#) using [OpenStack Cinder](#). Some familiarity with Kubernetes and OpenStack is assumed.



IMPORTANT

Before creating persistent volumes using Cinder, OpenShift Container Platform must first be properly [configured for OpenStack](#).

The Kubernetes [persistent volume](#) framework allows administrators to provision a cluster with persistent storage and gives users a way to request those resources without having any knowledge of the underlying infrastructure. OpenStack Cinder volumes can be [provisioned dynamically](#). Persistent

volumes are not bound to a single project or namespace; they can be shared across the OpenShift Container Platform cluster. [Persistent volume claims](#), however, are specific to a project or namespace and can be requested by users.

For a detailed example, see the guide for [WordPress and MySQL using persistent volumes](#).



IMPORTANT

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

20.4.2. Provisioning

Storage must exist in the underlying infrastructure before it can be mounted as a volume in OpenShift Container Platform. After ensuring OpenShift Container Platform is [configured for OpenStack](#), all that is required for Cinder is a Cinder volume ID and the **PersistentVolume** API.

20.4.2.1. Creating the Persistent Volume



NOTE

Cinder does not support the 'Recycle' reclaim policy.

You must define your persistent volume in an object definition before creating it in OpenShift Container Platform:

Example 20.6. Persistent Volume Object Definition Using Cinder

```
apiVersion: "v1"
kind: "PersistentVolume"
metadata:
  name: "pv0001" ❶
spec:
  capacity:
    storage: "5Gi" ❷
  accessModes:
    - "ReadWriteOnce"
  cinder: ❸
    fsType: "ext3" ❹
    volumeID: "f37a03aa-6212-4c62-a805-9ce139fab180" ❺
```

- ❶ The name of the volume. This will be how it is identified via [persistent volume claims](#) or from pods.
- ❷ The amount of storage allocated to this volume.
- ❸ This defines the volume type being used, in this case the **cinder** plug-in.
- ❹ File system type to mount.
- ❺ This is the Cinder volume that will be used.

**IMPORTANT**

Changing the value of the **fstype** parameter after the volume has been formatted and provisioned can result in data loss and pod failure.

Save your definition to a file, for example ***cinder-pv.yaml***, and create the persistent volume:

```
# oc create -f cinder-pv.yaml
persistentvolume "pv0001" created
```

Verify that the persistent volume was created:

```
# oc get pv
NAME          LABELS          CAPACITY  ACCESSMODES  STATUS      CLAIM          REASON
AGE
pv0001        <none>          5Gi       RWO           Available
2s
```

Users can then [request storage using persistent volume claims](#), which can now utilize your new persistent volume.

**IMPORTANT**

Persistent volume claims only exist in the user's namespace and can only be referenced by a pod within that same namespace. Any attempt to access a persistent volume from a different namespace causes the pod to fail.

20.4.2.2. Volume Format

Before OpenShift Container Platform mounts the volume and passes it to a container, it checks that it contains a file system as specified by the **fsType** parameter in the persistent volume definition. If the device is not formatted with the file system, all data from the device is erased and the device is automatically formatted with the given file system.

This allows using unformatted Cinder volumes as persistent volumes, because OpenShift Container Platform formats them before the first use.

20.5. PERSISTENT STORAGE USING CEPH RADOS BLOCK DEVICE (RBD)**20.5.1. Overview**

OpenShift Container Platform clusters can be provisioned with [persistent storage](#) using Ceph RBD.

Persistent volumes (PVs) and [persistent volume claims \(PVCs\)](#) can share volumes across a single project. While the Ceph RBD-specific information contained in a PV definition could also be defined directly in a pod definition, doing so does not create the volume as a distinct cluster resource, making the volume more susceptible to conflicts.

This topic presumes some familiarity with OpenShift Container Platform and [Ceph RBD](#). See the [Persistent Storage](#) concept topic for details on the OpenShift Container Platform persistent volume (PV) framework in general.

**NOTE**

Project and *namespace* are used interchangeably throughout this document. See [Projects and Users](#) for details on the relationship.

**IMPORTANT**

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

20.5.2. Provisioning

To provision Ceph volumes, the following are required:

- An existing storage device in your underlying infrastructure.
- The Ceph key to be used in an OpenShift Container Platform secret object.
- The Ceph image name.
- The file system type on top of the block storage (e.g., ext4).
- **ceph-common** installed on each schedulable OpenShift Container Platform node in your cluster:

```
# yum install ceph-common
```

20.5.2.1. Creating the Ceph Secret

Define the authorization key in a secret configuration, which is then converted to base64 for use by OpenShift Container Platform.

**NOTE**

In order to use Ceph storage to back a persistent volume, the secret must be created in the same project as the PVC and pod. The secret cannot simply be in the default project.

1. Run **ceph auth get-key** on a Ceph MON node to display the key value for the **client.admin** user:

```
apiVersion: v1
kind: Secret
metadata:
  name: ceph-secret
data:
  key: QVFB0FF2S1ZheUJQRVJBQWgvs2cwT1laQUhPQno3akZwekxxdGc9PQ==
```

2. Save the secret definition to a file, for example **ceph-secret.yaml**, then create the secret:

```
$ oc create -f ceph-secret.yaml
```

3. Verify that the secret was created:

```
# oc get secret ceph-secret
NAME                TYPE                DATA      AGE
```

ceph-secret	Opaque	1	23d
-------------	--------	---	-----

20.5.2.2. Creating the Persistent Volume



NOTE

Ceph RBD does not support the 'Recycle' reclaim policy.

Developers request Ceph RBD storage by referencing either a PVC, or the Gluster volume plug-in directly in the **volumes** section of a pod specification. A PVC exists only in the user's namespace and can be referenced only by pods within that same namespace. Any attempt to access a PV from a different namespace causes the pod to fail.


1. Define the PV in an object definition before creating it in OpenShift Container Platform:

Example 20.7. Persistent Volume Object Definition Using Ceph RBD

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: ceph-pv ❶
spec:
  capacity:
    storage: 2Gi ❷
  accessModes:
    - ReadWriteOnce ❸
  rbd: ❹
    monitors: ❺
      - 192.168.122.133:6789
    pool: rbd
    image: ceph-image
    user: admin
    secretRef:
      name: ceph-secret ❻
    fsType: ext4 ❼
    readOnly: false
  persistentVolumeReclaimPolicy: Retain
```

- ❶ The name of the PV that is referenced in pod definitions or displayed in various **oc** volume commands.
- ❷ The amount of storage allocated to this volume.
- ❸ **accessModes** are used as labels to match a PV and a PVC. They currently do not define any form of access control. All block storage is defined to be single user (non-shared storage).
- ❹ The volume type being used, in this case the **rbd** plug-in.
- ❺ An array of Ceph monitor IP addresses and ports.
- ❻ The Ceph secret used to create a secure connection from OpenShift Container Platform to the Ceph server.

- 7 The file system type mounted on the Ceph RBD block device.



IMPORTANT

Changing the value of the **fstype** parameter after the volume has been formatted and provisioned can result in data loss and pod failure.

2. Save your definition to a file, for example **ceph-pv.yaml**, and create the PV:

```
# oc create -f ceph-pv.yaml
```

3. Verify that the persistent volume was created:

```
# oc get pv
NAME                                LABELS                                CAPACITY    ACCESSMODES
STATUS      CLAIM      REASON    AGE
ceph-pv                                <none>      2147483648  RW0
Available                                     2s
```

4. Create a PVC that will bind to the new PV:

Example 20.8. PVC Object Definition

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: ceph-claim
spec:
  accessModes: 1
  - ReadWriteOnce
  resources:
    requests:
      storage: 2Gi 2
```

- 1 The **accessModes** do not enforce access right, but instead act as labels to match a PV to a PVC.
- 2 This claim looks for PVs offering **2Gi** or greater capacity.

5. Save the definition to a file, for example **ceph-claim.yaml**, and create the PVC:

```
# oc create -f ceph-claim.yaml
```

20.5.3. Ceph Volume Security



NOTE

See the full [Volume Security](#) topic before implementing Ceph RBD volumes.

A significant difference between shared volumes (NFS and GlusterFS) and block volumes (Ceph RBD, iSCSI, and most cloud storage), is that the user and group IDs defined in the pod definition or container image are applied to the target physical storage. This is referred to as managing ownership of the block device. For example, if the Ceph RBD mount has its owner set to **123** and its group ID set to **567**, and if the pod defines its **runAsUser** set to **222** and its **fsGroup** to be **7777**, then the Ceph RBD physical mount's ownership will be changed to **222:7777**.



NOTE

Even if the user and group IDs are not defined in the pod specification, the resulting pod may have defaults defined for these IDs based on its matching SCC, or its project. See the full [Volume Security](#) topic which covers storage aspects of SCCs and defaults in greater detail.

A pod defines the group ownership of a Ceph RBD volume using the **fsGroup** stanza under the pod's **securityContext** definition:


```
spec:
  containers:
    - name:
      ...
  securityContext: 1
    fsGroup: 7777 2
```

- 1 The **securityContext** must be defined at the pod level, not under a specific container.
- 2 All containers in the pod will have the same fsGroup ID.

20.6. PERSISTENT STORAGE USING AWS ELASTIC BLOCK STORE

20.6.1. Overview

OpenShift Container Platform supports AWS Elastic Block Store volumes (EBS). You can provision your OpenShift Container Platform cluster with [persistent storage](#) using [AWS EC2](#). Some familiarity with Kubernetes and AWS is assumed.



IMPORTANT

Before creating persistent volumes using AWS, OpenShift Container Platform must first be properly [configured for AWS ElasticBlockStore](#).

The Kubernetes [persistent volume](#) framework allows administrators to provision a cluster with persistent storage and gives users a way to request those resources without having any knowledge of the underlying infrastructure. AWS Elastic Block Store volumes can be [provisioned dynamically](#). Persistent volumes are not bound to a single project or namespace; they can be shared across the OpenShift Container Platform cluster. [Persistent volume claims](#), however, are specific to a project or namespace and can be requested by users.



IMPORTANT

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

20.6.2. Provisioning

Storage must exist in the underlying infrastructure before it can be mounted as a volume in OpenShift Container Platform. After ensuring OpenShift is [configured for AWS Elastic Block Store](#), all that is required for OpenShift and AWS is an AWS EBS volume ID and the **PersistentVolume** API.

20.6.2.1. Creating the Persistent Volume



NOTE


AWS does not support the 'Recycle' reclaim policy.

You must define your persistent volume in an object definition before creating it in OpenShift Container Platform:

Example 20.9. Persistent Volume Object Definition Using AWS

```
apiVersion: "v1"
kind: "PersistentVolume"
metadata:
  name: "pv0001" 1
spec:
  capacity:
    storage: "5Gi" 2
  accessModes:
    - "ReadWriteOnce"
  awsElasticBlockStore: 3
    fsType: "ext4" 4
    volumeID: "vol-f37a03aa" 5
```

- 1 The name of the volume. This will be how it is identified via [persistent volume claims](#) or from pods.
- 2 The amount of storage allocated to this volume.
- 3 This defines the volume type being used, in this case the **awsElasticBlockStore** plug-in.
- 4 File system type to mount.
- 5 This is the AWS volume that will be used.



IMPORTANT

Changing the value of the **fsType** parameter after the volume has been formatted and provisioned can result in data loss and pod failure.

Save your definition to a file, for example **aws-pv.yaml**, and create the persistent volume:

```
# oc create -f aws-pv.yaml
persistentvolume "pv0001" created
```

Verify that the persistent volume was created:

```
# oc get pv
NAME          LABELS          CAPACITY  ACCESSMODES  STATUS   CLAIM          REASON
AGE
pv0001        <none>          5Gi       RWO           Available
2s
```

Users can then [request storage using persistent volume claims](#), which can now utilize your new persistent volume.



IMPORTANT

Persistent volume claims only exist in the user's namespace and can only be referenced by a pod within that same namespace. Any attempt to access a persistent volume from a different namespace causes the pod to fail.

20.6.2.2. Volume Format

Before OpenShift Container Platform mounts the volume and passes it to a container, it checks that it contains a file system as specified by the **fsType** parameter in the persistent volume definition. If the device is not formatted with the file system, all data from the device is erased and the device is automatically formatted with the given file system.

This allows using unformatted AWS volumes as persistent volumes, because OpenShift Container Platform formats them before the first use.

20.7. PERSISTENT STORAGE USING GCE PERSISTENT DISK

20.7.1. Overview

OpenShift Container Platform supports GCE Persistent Disk volumes (gcePD). You can provision your OpenShift Container Platform cluster with [persistent storage](#) using [GCE](#). Some familiarity with Kubernetes and GCE is assumed.



IMPORTANT

Before creating persistent volumes using GCE, OpenShift Container Platform must first be properly [configured for GCE Persistent Disk](#).

The Kubernetes [persistent volume](#) framework allows administrators to provision a cluster with persistent storage and gives users a way to request those resources without having any knowledge of the underlying infrastructure. GCE Persistent Disk volumes can be [provisioned dynamically](#). Persistent volumes are not bound to a single project or namespace; they can be shared across the OpenShift Container Platform cluster. [Persistent volume claims](#), however, are specific to a project or namespace and can be requested by users.




IMPORTANT

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

20.7.2. Provisioning

Storage must exist in the underlying infrastructure before it can be mounted as a volume in OpenShift Container Platform. After ensuring OpenShift Container Platform is [configured for GCE PersistentDisk](#), all that is required for OpenShift Container Platform and GCE is an GCE Persistent Disk volume ID and the **PersistentVolume** API.

20.7.2.1. Creating the Persistent Volume



NOTE

GCE does not support the 'Recycle' reclaim policy.

You must define your persistent volume in an object definition before creating it in OpenShift Container Platform:

Example 20.10. Persistent Volume Object Definition Using GCE

```
apiVersion: "v1"
kind: "PersistentVolume"
metadata:
  name: "pv0001" 1
spec:
  capacity:
    storage: "5Gi" 2
  accessModes:
    - "ReadWriteOnce"
  gcePersistentDisk: 3
    fsType: "ext4" 4
    pdName: "pd-disk-1" 5
```

- 1 The name of the volume. This will be how it is identified via [persistent volume claims](#) or from pods.
- 2 The amount of storage allocated to this volume.
- 3 This defines the volume type being used, in this case the **gcePersistentDisk** plug-in.
- 4 File system type to mount.
- 5 This is the GCE Persistent Disk volume that will be used.



IMPORTANT

Changing the value of the **fsType** parameter after the volume has been formatted and provisioned can result in data loss and pod failure.

Save your definition to a file, for example **gce-pv.yaml**, and create the persistent volume:

```
# oc create -f gce-pv.yaml
persistentvolume "pv0001" created
```

Verify that the persistent volume was created:

```
# oc get pv
NAME          LABELS          CAPACITY  ACCESSMODES  STATUS   CLAIM   REASON
AGE
pv0001        <none>          5Gi       RWO           Available
2s
```

Users can then [request storage using persistent volume claims](#), which can now utilize your new persistent volume.



IMPORTANT

Persistent volume claims only exist in the user's namespace and can only be referenced by a pod within that same namespace. Any attempt to access a persistent volume from a different namespace causes the pod to fail.

20.7.2.2. Volume Format

Before OpenShift Container Platform mounts the volume and passes it to a container, it checks that it contains a file system as specified by the **fsType** parameter in the persistent volume definition. If the device is not formatted with the file system, all data from the device is erased and the device is automatically formatted with the given file system.

This allows using unformatted GCE volumes as persistent volumes, because OpenShift Container Platform formats them before the first use.

20.7.2.3. Multi-zone Configuration

In multi-zone configurations, if the volume being created does not exist in the master node's zone, you must specify **failure-domain.beta.kubernetes.io/region** and **failure-domain.beta.kubernetes.io/zone** PV labels to match the zone where the GCE volume exists.

Example 20.11. Persistent Volume Object With Failure Domain

```
apiVersion: "v1"
kind: "PersistentVolume"
metadata:
  name: "pv0001"
  labels:
    failure-domain.beta.kubernetes.io/region: "us-central1" 1
    failure-domain.beta.kubernetes.io/zone: "us-central1-a" 2
spec:
  capacity:
    storage: "5Gi"
  accessModes:
    - "ReadWriteOnce"
  gcePersistentDisk:
    fsType: "ext4"
    pdName: "pd-disk-1"
```

1 The region in which the volume exists.


2 The zone in which the volume exists.

20.8. PERSISTENT STORAGE USING ISCSI

20.8.1. Overview

You can provision your OpenShift Container Platform cluster with [persistent storage](#) using [iSCSI](#). Some familiarity with Kubernetes and iSCSI is assumed.

The Kubernetes [persistent volume](#) framework allows administrators to provision a cluster with persistent storage and gives users a way to request those resources without having any knowledge of the underlying infrastructure.



IMPORTANT

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

20.8.2. Provisioning

Verify that the storage exists in the underlying infrastructure before mounting it as a volume in OpenShift Container Platform. All that is required for the iSCSI is the iSCSI target portal, a valid iSCSI Qualified Name (IQN), a valid LUN number, the filesystem type, and the **PersistentVolume** API.

Optionally, multipath portals and Challenge Handshake Authentication Protocol (CHAP) configuration can be provided.



NOTE

iSCSI does not support the 'Recycle' reclaim policy.

Example 20.12. Persistent Volume Object Definition

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: iscsi-pv
spec:
  capacity:
    storage: 1Gi
  accessModes:
    - ReadWriteOnce
  iscsi:
    targetPortal: 10.16.154.81
    iqn: iqn.2014-12.example.server:storage.target00
    lun: 0
    fsType: 'ext4'
    readOnly: false
    chapAuthDiscovery: true
    chapAuthSession: true
    secretRef:
      name: chap-secret
```

20.8.2.1. Enforcing Disk Quotas

Use LUN partitions to enforce disk quotas and size constraints. Each LUN is one persistent volume. Kubernetes enforces unique names for persistent volumes.

Enforcing quotas in this way allows the end user to request persistent storage by a specific amount (e.g, 10Gi) and be matched with a corresponding volume of equal or greater capacity.

20.8.2.2. iSCSI Volume Security

Users request storage with a **PersistentVolumeClaim**. This claim only lives in the user's namespace and can only be referenced by a pod within that same namespace. Any attempt to access a persistent volume across a namespace causes the pod to fail.

Each iSCSI LUN must be accessible by all nodes in the cluster.

20.9. PERSISTENT STORAGE USING FIBRE CHANNEL

20.9.1. Overview

You can provision your OpenShift Container Platform cluster with [persistent storage](#) using [Fibre Channel](#). Some familiarity with Kubernetes and Fibre Channel is assumed.

The Kubernetes [persistent volume](#) framework allows administrators to provision a cluster with persistent storage and gives users a way to request those resources without having any knowledge of the underlying infrastructure.



IMPORTANT

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

20.9.2. Provisioning

Storage must exist in the underlying infrastructure before it can be mounted as a volume in OpenShift Container Platform. All that is required for Fibre Channel persistent storage is the targetWWNs (array of Fibre Channel target's World Wide Names), a valid LUN number, and filesystem type, and the **PersistentVolume** API. Note, the number of LUNs must correspond to the number of Persistent Volumes that are created. In the example below, we have LUN as 2, therefore we have created two Persistent Volume definitions.



NOTE

Fiber Channel does not support the 'Recycle' reclaim policy.

Example 20.13. Persistent Volumes Object Definition

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv0001
spec:
  capacity:
    storage: 1Gi
```

```

accessModes:
  - ReadWriteOnce
fc:
  targetWWNs: ['500a0981891b8dc5', '500a0981991b8dc5']
  lun: 2
  fsType: ext4

apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv0002
spec:
  capacity:
    storage: 1Gi
  accessModes:
    - ReadOnlyMany
  fc:
    targetWWNs: ['500a0981891b8dc5', '500a0981991b8dc5']
    lun: 2
    fsType: ext4

```



IMPORTANT

Changing the value of the **fstype** parameter after the volume has been formatted and provisioned can result in data loss and pod failure.

20.9.2.1. Enforcing Disk Quotas

Use LUN partitions to enforce disk quotas and size constraints. Each LUN is one persistent volume. Kubernetes enforces unique names for persistent volumes.

Enforcing quotas in this way allows the end user to request persistent storage by a specific amount (e.g, 10Gi) and be matched with a corresponding volume of equal or greater capacity.

20.9.2.2. Fibre Channel Volume Security

Users request storage with a **PersistentVolumeClaim**. This claim only lives in the user's namespace and can only be referenced by a pod within that same namespace. Any attempt to access a persistent volume across a namespace causes the pod to fail.

Each Fibre Channel LUN must be accessible by all nodes in the cluster.

20.10. DYNAMICALLY PROVISIONING PERSISTENT VOLUMES

20.10.1. Overview

You can provision your OpenShift Container Platform cluster with storage dynamically when running in a cloud environment. The Kubernetes [persistent volume](#) framework allows administrators to provision a cluster with persistent storage and gives users a way to request those resources without having any knowledge of the underlying infrastructure.

Many storage types are available for use as persistent volumes in OpenShift Container Platform. While

all of them can be statically provisioned by an administrator, some types of storage can be created dynamically using an API. These types of storage can be provisioned in an OpenShift Container Platform cluster using the new and experimental dynamic storage feature.



IMPORTANT

Dynamic provisioning of persistent volumes is currently a Technology Preview feature, introduced in OpenShift Container Platform 3.1.1. This feature is experimental and expected to change in the future as it matures and feedback is received from users. New ways to provision the cluster are planned and the means by which one accesses this feature is going to change. Backwards compatibility is not guaranteed.

20.10.2. Enabling Provisioner Plug-ins

OpenShift Container Platform provides the following *provisioner plug-ins*, which have generic implementations for dynamic provisioning that use the cluster's configured cloud provider's API to create new storage resources:

Storage Type	Provisioner Plug-in Name	Required Cloud Configuration	Notes
OpenStack Cinder	kubernetes.io/cinder	Configuring for OpenStack	
AWS Elastic Block Store (EBS)	kubernetes.io/aws-ebs	Configuring for AWS	For dynamic provisioning when using multiple clusters in different zones, each node must be tagged with Key=KubernetesCluster, Value=clusterid .
GCE Persistent Disk (gcePD)	kubernetes.io/gce-pd	Configuring for GCE	In multi-zone configurations, it is recommended that you run one OpenShift Container Platform cluster per GCE project to avoid the creation of PVs in zones where no node from the current cluster exists.



IMPORTANT

For any chosen provisioner plug-ins, the relevant cloud configuration must also be set up, per **Required Cloud Configuration** in the above table.

When your OpenShift Container Platform cluster is configured for EBS, GCE, or Cinder, the associated provisioner plug-in is implied and automatically enabled. No additional OpenShift Container Platform configuration by the cluster administration is required for dynamic provisioning.

For example, if your OpenShift Container Platform cluster is configured to run in AWS, the EBS provisioner plug-in is automatically available for creating [dynamically provisioned storage requested by a user](#).

Future provisioner plug-ins will include the many types of storage a single provider offers. AWS, for example, has several types of EBS volumes to offer, each with its own performance characteristics; there is also an NFS-like storage option. More provisioner plug-ins will be implemented for the supported storage types available in OpenShift Container Platform.

20.10.3. Requesting Dynamically Provisioned Storage

Users can request dynamically provisioned storage by including a storage class annotation in their [persistent volume claim](#):

Example 20.14. Persistent Volume Claim Requesting Dynamic Storage

```
kind: "PersistentVolumeClaim"
apiVersion: "v1"
metadata:
  name: "claim1"
  annotations:
    volume.alpha.kubernetes.io/storage-class: "foo" 1
spec:
  accessModes:
    - "ReadWriteOnce"
  resources:
    requests:
      storage: "3Gi"
```

- 1 The value of the **volume.alpha.kubernetes.io/storage-class** annotation is not meaningful at this time. The presence of the annotation, with any arbitrary value, triggers provisioning using the single implied [provisioner plug-in per cloud](#).

20.10.3.1. Volume Owner Information

For dynamically provisioned storage, OpenShift Container Platform defines three key/value pairs, collectively known as the *volume owner information*, and arranges for the storage to associate this triplet with the provisioned volume. The keys are normally not visible to OpenShift Container Platform users, while the values are taken from user-visible PV and PVC objects.

Keys

kubernetes.io/created-for/pv/name

Name of the **PersistentVolume**.

**NOTE**

There is no key for the PV namespace because that has value **default** and cannot be changed.

kubernetes.io/created-for/pvc/namespace, kubernetes.io/created-for/pvc/name

Namespace and name, respectively, of the **PersistentVolumeClaim**.

Other Terms for Volume Owner Information

Each storage type saves the volume owner information in its own way. When communicating with the storage administrator, use these specific terms to avoid confusion:

Term for Key/Value Pairs	Storage Type
tags	AWS EBS
metadata	OpenStack Cinder
JSON-in-description	GCE PD

Using Volume Owner Information

The main benefit of saving the volume owner information is to enable storage administrators to recognize volumes dynamically created by OpenShift Container Platform.

Example scenarios:

- OpenShift Container Platform terminates unexpectedly and the dynamically provisioned AWS EBS contains useful data that must be recovered. The OpenShift Container Platform users provide the storage administrators with a list of affected projects and their PVCs:

Project Name	PVC Name
app-server	a-pv-01
	a-pv-02
notifications	n-pv-01

The storage administrators search for the orphaned volumes, matching project names and PVC names to the **kubernetes.io/created-for/pvc/namespace** and **kubernetes.io/created-for/pvc/name** tags, respectively. They find them and arrange to make them available again for data-recovery efforts.

- The users do not explicitly delete the dynamically provisioned storage volumes when they are finished with a project. The storage administrators find the defunct volumes and delete them. Unlike the preceding scenario, they need match only the project names to **kubernetes.io/created-for/pvc/namespace**.

20.10.4. Volume Recycling

Volumes created dynamically by a provisioner have their **`persistentVolumeReclaimPolicy`** set to **`Delete`**. When a persistent volume claim is deleted, its backing persistent volume is considered released of its claim, and that resource can be reclaimed by the cluster. Dynamic provisioning utilizes the provider's API to delete the volume from the provider and then removes the persistent volume from the cluster.

20.11. VOLUME SECURITY

20.11.1. Overview

This topic provides a general guide on pod security as it relates to volume security. For information on pod-level security in general, see [Managing Security Context Constraints \(SCC\)](#) and the [Security Context Constraint](#) concept topic. For information on the OpenShift Container Platform persistent volume (PV) framework in general, see the [Persistent Storage](#) concept topic.

Accessing persistent storage requires coordination between the cluster and/or storage administrator and the end developer. The cluster administrator creates PVs, which abstract the underlying physical storage. The developer creates pods and, optionally, PVCs, which bind to PVs, based on matching criteria, such as capacity.

Multiple persistent volume claims (PVCs) within the same project can bind to the same PV. However, once a PVC binds to a PV, that PV cannot be bound by a claim outside of the first claim's project. If the underlying storage needs to be accessed by multiple projects, then each project needs its own PV, which can point to the same physical storage. In this sense, a bound PV is tied to a project. For a detailed PV and PVC example, see the guide for [WordPress and MySQL using NFS](#).

For the cluster administrator, granting pods access to PVs involves:

- knowing the group ID and/or user ID assigned to the actual storage,
- understanding SELinux considerations, and
- ensuring that these IDs are allowed in the range of legal IDs defined for the project and/or the SCC that matches the requirements of the pod.

Group IDs, the user ID, and SELinux values are defined in the **`SecurityContext`** section in a pod definition. Group IDs are global to the pod and apply to all containers defined in the pod. User IDs can also be global, or specific to each container. Four sections control access to volumes:

- [**`supplementalGroups`**](#)
- [**`fsGroup`**](#)
- [**`runAsUser`**](#)
- [**`seLinuxOptions`**](#)

20.11.2. SCCs, Defaults, and Allowed Ranges

SCCs influence whether or not a pod is given a default user ID, **`fsGroup`** ID, supplemental group ID, and SELinux label. They also influence whether or not IDs supplied in the pod definition (or in the image) will be validated against a range of allowable IDs. If validation is required and fails, then the pod will also fail.

SCCs define strategies, such as **runAsUser**, **supplementalGroups**, and **fsGroup**. These strategies help decide whether the pod is authorized. Strategy values set to **RunAsAny** are essentially stating that the pod can do what it wants regarding that strategy. Authorization is skipped for that strategy and no OpenShift Container Platform default is produced based on that strategy. Therefore, IDs and SELinux labels in the resulting container are based on container defaults instead of OpenShift Container Platform policies.

For a quick summary of **RunAsAny**:

- Any ID defined in the pod definition (or image) is allowed.
- Absence of an ID in the pod definition (and in the image) results in the container assigning an ID, which is **root** (0) for Docker.
- No SELinux labels are defined, so Docker will assign a unique label.

For these reasons, SCCs with **RunAsAny** for ID-related strategies should be protected so that ordinary developers do not have access to the SCC. On the other hand, SCC strategies set to **MustRunAs** or **MustRunAsRange** trigger ID validation (for ID-related strategies), and cause default values to be supplied by OpenShift Container Platform to the container when those values are not supplied directly in the pod definition or image.

CAUTION

Allowing access to SCCs with a **RunAsAny fsGroup** strategy can also prevent users from accessing their block devices. Pods need to specify an **fsGroup** in order to take over their block devices. Normally, this is done when the SCC **fsGroup** strategy is set to **MustRunAs**. If a user's pod is assigned an SCC with a **RunAsAny fsGroup** strategy, then the user may face **permission denied** errors until they discover that they need to specify an **fsGroup** themselves.

SCCs may define the range of allowed IDs (user or groups). If range checking is required (for example, using **MustRunAs**) and the allowable range is not defined in the SCC, then the project determines the ID range. Therefore, projects support ranges of allowable ID. However, unlike SCCs, projects do not define strategies, such as **runAsUser**.

Allowable ranges are helpful not only because they define the boundaries for container IDs, but also because the minimum value in the range becomes the default value for the ID in question. For example, if the SCC ID strategy value is **MustRunAs**, the minimum value of an ID range is **100**, and the ID is absent from the pod definition, then 100 is provided as the default for this ID.

As part of pod admission, the SCCs available to a pod are examined (roughly, in priority order followed by most restrictive) to best match the requests of the pod. Setting a SCC's strategy type to **RunAsAny** is less restrictive, whereas a type of **MustRunAs** is more restrictive. All of these strategies are evaluated. To see which SCC was assigned to a pod, use the **oc get pod** command:

```
# oc get pod <pod_name> -o yaml
...
metadata:
  annotations:
    openshift.io/scc: nfs-scc ❶
    name: nfs-pod1 ❷
    namespace: default ❸
...
```

- 1 Name of the SCC that the pod used (in this case, a custom SCC).
- 2 Name of the pod.
- 3 Name of the project. "Namespace" is interchangeable with "project" in OpenShift Container Platform. See [Projects and Users](#) for details.

It may not be immediately obvious which SCC was matched by a pod, so the command above can be very useful in understanding the UID, supplemental groups, and SELinux relabeling in a live container.

Any SCC with a strategy set to **RunAsAny** allows specific values for that strategy to be defined in the pod definition (and/or image). When this applies to the user ID (**runAsUser**) it is prudent to restrict access to the SCC to prevent a container from being able to run as root.

Because pods often match the **restricted** SCC, it is worth knowing the security this entails. The **restricted** SCC has the following characteristics:

- User IDs are constrained due to the **runAsUser** strategy being set to **MustRunAsRange**. This forces user ID validation.
- Because a range of allowable user IDs is not defined in the SCC (see **oc export scc restricted** for more details), the project's **openshift.io/sa.scc.uid-range** range will be used for range checking and for a default ID, if needed.
- A default user ID is produced when a user ID is not specified in the pod definition due to **runAsUser** being set to **MustRunAsRange**.
- An SELinux label is required (**seLinuxContext** set to *MustRunAs*), which uses the project's default MCS label.
- **fsGroup** IDs are constrained to a single value due to the **FSGroup** strategy being set to **MustRunAs**, which dictates that the value to use is the minimum value of the first range specified.
- Because a range of allowable **fsGroup** IDs is not defined in the SCC, the minimum value of the project's **openshift.io/sa.scc.supplemental-groups** range (or the same range used for user IDs) will be used for validation and for a default ID, if needed.
- A default **fsGroup** ID is produced when a **fsGroup** ID is not specified in the pod and the matching SCC's **FSGroup** is set to **MustRunAs**.
- Arbitrary supplemental group IDs are allowed because no range checking is required. This is a result of the **supplementalGroups** strategy being set to **RunAsAny**.
- Default supplemental groups are not produced for the running pod due to **RunAsAny** for the two group strategies above. Therefore, if no groups are defined in the pod definition (or in the image), the container(s) will have no supplemental groups predefined.

The following shows the **default** project and a custom SCC (**my-custom-scc**), which summarizes the interactions of the SCC and the project:

```
$ oc get project default -o yaml 1
...
metadata:
  annotations: 2
```

```

    openshift.io/sa.scc.mcs: s0:c1,c0 3
    openshift.io/sa.scc.supplemental-groups: 1000000000/10000 4
    openshift.io/sa.scc.uid-range: 1000000000/10000 5

$ oc get scc my-custom-scc -o yaml
...
fsGroup:
  type: MustRunAs 6
  ranges:
    - min: 5000
      max: 6000
runAsUser:
  type: MustRunAsRange 7
  uidRangeMin: 65534
  uidRangeMax: 65634
seLinuxContext: 8
  type: MustRunAs
  SELinuxOptions: 9
    user: <selinux-user-name>
    role: ...
    type: ...
    level: ...
supplementalGroups:
  type: MustRunAs 10
  ranges:
    - min: 5000
      max: 6000

```

- 1 **default** is the name of the project.
- 2 Default values are only produced when the corresponding SCC strategy is not **RunAsAny**.
- 3 SELinux default when not defined in the pod definition or in the SCC.
- 4 Range of allowable group IDs. ID validation only occurs when the SCC strategy is **RunAsAny**. There can be more than one range specified, separated by commas. See below for [supported formats](#).
- 5 Same as <4> but for user IDs. Also, only a single range of user IDs is supported.
- 6 10 **MustRunAs** enforces group ID range checking and provides the container's groups default. Based on this SCC definition, the default is 5000 (the minimum ID value). If the range was omitted from the SCC, then the default would be 1000000000 (derived from the project). The other supported type, **RunAsAny**, does not perform range checking, thus allowing any group ID, and produces no default groups.
- 7 **MustRunAsRange** enforces user ID range checking and provides a UID default. **Based on this SCC, the default UID is 65534 (the minimum value). If the minimum *and maximum range were omitted from the SCC, the default user ID would be *1000000000 (derived from the project). *MustRunAsNonRoot and RunAsAny are *the other supported types.** The range of allowed IDs can be defined to include *any user IDs required for the target storage.
- 8 When set to **MustRunAs**, the container is created with the SCC's SELinux options, or the MCS default defined in the project. A type of **RunAsAny** indicates that SELinux context is not required, and, if not defined in the pod, is not set in the container.

- 9** The SELinux user name, role name, type, and labels can be defined here.

Two formats are supported for allowed ranges:

1. **M/N**, where **M** is the starting ID and **N** is the count, so the range becomes **M** through (and including) **M+N-1**.
2. **M-N**, where **M** is again the starting ID and **N** is the ending ID. The default group ID is the starting ID in the first range, which is **1000000000** in this project. If the SCC did not define a minimum group ID, then the project's default ID is applied.

20.11.3. Supplemental Groups



NOTE

Read [SCCs, Defaults, and Allowed Ranges](#) before working with supplemental groups.

TIP

It is generally preferable to use group IDs (supplemental or [fsGroup](#)) to gain access to persistent storage versus using [user IDs](#).

Supplemental groups are regular Linux groups. When a process runs in Linux, it has a UID, a GID, and one or more supplemental groups. These attributes can be set for a container's main process. The **supplementalGroups** IDs are typically used for controlling access to shared storage, such as NFS and GlusterFS, whereas [fsGroup](#) is used for controlling access to block storage, such as Ceph RBD and iSCSI.

The OpenShift Container Platform shared storage plug-ins mount volumes such that the POSIX permissions on the mount match the permissions on the target storage. For example, if the target storage's owner ID is **1234** and its group ID is **5678**, then the mount on the host node and in the container will have those same IDs. Therefore, the container's main process must match one or both of those IDs in order to access the volume.

For example, consider the following NFS export.

On an OpenShift Container Platform node:



NOTE

showmount requires access to the ports used by **rpcbind** and **rpc.mount** on the NFS server


```
# showmount -e <nfs-server-ip-or-hostname>
Export list for f21-nfs.vm:
/opt/nfs *
```

On the NFS server:

```
# cat /etc/exports
/opt/nfs *(rw,sync,root_squash)
...
```

```
# ls -lZ /opt/nfs -d
drwxrws---. nfsnobody 5555 unconfined_u:object_r:usr_t:s0 /opt/nfs

# id nfsnobody
uid=65534(nfsnobody) gid=65534(nfsnobody) groups=65534(nfsnobody)
```

**NOTE**

In the above, the owner is 65534 (**nfsnobody**), but the suggestions and examples in this topic apply to any non-root owner.

The **/opt/nfs/** export is accessible by UID **65534** and the group **5555**. In general, containers should not run as root, so in this NFS example, containers which are not run as UID **65534** or are not members the group **5555** will not be able to access the NFS export.

Often, the SCC matching the pod does not allow a specific user ID to be specified, thus using supplemental groups is a more flexible way to grant storage access to a pod. For example, to grant NFS access to the export above, the group **5555** can be defined in the pod definition:

```
apiVersion: v1
kind: Pod
...
spec:
  containers:
    - name: ...
      volumeMounts:
        - name: nfs 1
          mountPath: /usr/share/... 2
      securityContext: 3
        supplementalGroups: [5555] 4
  volumes:
    - name: nfs 5
      nfs:
        server: <nfs_server_ip_or_host>
        path: /opt/nfs 6
```

- 1** Name of the volume mount. Must match the name in the **volumes** section.
- 2** NFS export path as seen in the container.
- 3** Pod global security context. Applies to all containers in the pod. Each container can also define its **securityContext**, however group IDs are global to the pod and cannot be defined for individual containers.
- 4** Supplemental groups, which is an array of IDs, is set to 5555. This grants group access to the export.
- 5** Name of the volume. Must match the name in the **volumeMounts** section.
- 6** Actual NFS export path on the NFS server.

All containers in the above pod (assuming the matching SCC or project allows the group **5555**) will be

members of the group **5555** and have access to the volume, regardless of the container's user ID. However, the assumption above is critical. Sometimes, the SCC does not define a range of allowable group IDs but requires group ID validation (due to **supplementalGroups** set to **MustRunAs**; note this is not the case for the **restricted** SCC). The project will not likely allow a group ID of **5555**, unless the project has been customized for access to this NFS export. So in this scenario, the above pod will fail because its group ID of **5555** is not within the SCC's or the project's range of allowed group IDs.

Supplemental Groups and Custom SCCs

To remedy the situation in [the previous example](#), a custom SCC can be created such that:

- a minimum and max group ID are defined,
- ID range checking is enforced, and
- the group ID of **5555** is allowed.

It is better to create new SCCs versus modifying a predefined SCC, or changing the range of allowed IDs in the predefined projects.

The easiest way to create a new SCC is to export an existing SCC and customize the YAML file to meet the requirements of the new SCC. For example:

1. Use the **restricted** SCC as a template for the new SCC:

```
$ oc export scc restricted > new-scc.yaml
```

2. Edit the **new-scc.yaml** file to your desired specifications.
3. Create the new SCC:

```
$ oc create -f new-scc.yaml
```



NOTE

The **oc edit scc** command can be used to modify an instantiated SCC.

Here is a fragment of a new SCC named **nfs-scc**:

```
$ oc export scc nfs-scc


allowHostDirVolumePlugin: false ❶
...
kind: SecurityContextConstraints
metadata:
  ...
  name: nfs-scc ❷
priority: 9 ❸
...
supplementalGroups:
  type: MustRunAs ❹
  ranges:
```

```
- min: 5000 5
  max: 6000
...
```

- 1 The **allow*** booleans are the same as for the **restricted** SCC.
- 2 Name of the new SCC.
- 3 Numerically larger numbers have greater priority. Nil or omitted is the lowest priority. Higher priority SCCs sort before lower priority SCCs and thus have a better chance of matching a new pod.
- 4 **supplementalGroups** is a strategy and it is set to **MustRunAs**, which means group ID checking is required.
- 5 Multiple ranges are supported. The allowed group ID range here is 5000 through 5999, with the default supplemental group being 5000.

When the same pod shown earlier runs against this new SCC (assuming, of course, the pod has access to the new SCC), it will start because the group **5555**, supplied in the pod definition, is now allowed by the custom SCC.

20.11.4. fsGroup



NOTE

Read [SCCs, Defaults, and Allowed Ranges](#) before working with supplemental groups.

TIP

It is generally preferable to use group IDs ([supplemental](#) or **fsGroup**) to gain access to persistent storage versus using [user IDs](#).

fsGroup defines a pod's "file system group" ID, which is added to the container's supplemental groups. The **supplementalGroups** ID applies to shared storage, whereas the **fsGroup** ID is used for block storage.

Block storage, such as Ceph RBD, iSCSI, and various cloud storage, is typically dedicated to a single pod which has requested the block storage volume, either directly or using a PVC. Unlike shared storage, block storage is taken over by a pod, meaning that user and group IDs supplied in the pod definition (or image) are applied to the actual, physical block device. Typically, block storage is not shared.

A **fsGroup** definition is shown below in the following pod definition fragment:

```
kind: Pod
...
spec:
  containers:
    - name: ...
      securityContext: 1
        fsGroup: 5555 2
  ...
```


- 1 As with **supplementalGroups**, **fsGroup** must be defined globally to the pod, not per container.
- 2 5555 will become the group ID for the volume's group permissions and for all new files created in the volume.

As with **supplementalGroups**, all containers in the above pod (assuming the matching SCC or project allows the group 5555) will be members of the group 5555, and will have access to the block volume, regardless of the container's user ID. If the pod matches the **restricted** SCC, whose **fsGroup** strategy is **MustRunAs**, then the pod will fail to run. However, if the SCC has its **fsGroup** strategy set to **RunAsAny**, then any **fsGroup** ID (including 5555) will be accepted. Note that if the SCC has its **fsGroup** strategy set to **RunAsAny** and no **fsGroup** ID is specified, the "taking over" of the block storage does not occur and permissions may be denied to the pod.

fsGroups and Custom SCCs

To remedy the situation in the previous example, a custom SCC can be created such that:

- a minimum and maximum group ID are defined,
- ID range checking is enforced, and
- the group ID of 5555 is allowed.

It is better to create new SCCs versus modifying a predefined SCC, or changing the range of allowed IDs in the predefined projects.

Consider the following fragment of a new SCC definition:

```
# oc export scc new-scc
...
kind: SecurityContextConstraints
...
fsGroup:
  type: MustRunAs 1
  ranges: 2
  - max: 6000
    min: 5000 3
...
```

- 1 **MustRunAs** triggers group ID range checking, whereas **RunAsAny** does not require range checking.
- 2 The range of allowed group IDs is 5000 through, and including, 5999. Multiple ranges are supported. The allowed group ID range here is 5000 through 5999, with the default **fsGroup** being 5000.
- 3 The minimum value (or the entire range) can be omitted from the SCC, and thus range checking and generating a default value will defer to the project's **openshift.io/sa.scc.supplemental-groups** range. **fsGroup** and **supplementalGroups** use the same group field in the project; there is not a separate range for **fsGroup**.


When the pod shown above runs against this new SCC (assuming, of course, the pod has access to the new SCC), it will start because the group 5555, supplied in the pod definition, is allowed by the custom SCC. Additionally, the pod will "take over" the block device, so when the block storage is viewed by a

process outside of the pod, it will actually have **5555** as its group ID.

Currently the list of volumes which support block ownership (block) management include:

- AWS Elastic Block Store
- OpenStack Cinder
- Ceph RBD
- GCE Persistent Disk
- iSCSI
- emptyDir
- gitRepo

20.11.5. User IDs



NOTE

Read [SCCs, Defaults, and Allowed Ranges](#) before working with supplemental groups.

TIP

It is generally preferable to use group IDs ([supplemental](#) or [fsGroup](#)) to gain access to persistent storage versus using user IDs.

User IDs can be defined in the container image or in the pod definition. In the pod definition, a single user ID can be defined globally to all containers, or specific to individual containers (or both). A user ID is supplied as shown in the pod definition fragment below:

```
spec:
  containers:
  - name: ...
    securityContext:
      runAsUser: 65534
```

ID 65534 in the above is container-specific and matches the owner ID on the export. If the NFS export's owner ID was **54321**, then that number would be used in the pod definition. Specifying **securityContext** outside of the container definition makes the ID global to all containers in the pod.

Similar to group IDs, user IDs may be validated according to policies set in the SCC and/or project. If the SCC's **runAsUser** strategy is set to **RunAsAny**, then any user ID defined in the pod definition or in the image is allowed.



WARNING

This means even a UID of **0** (root) is allowed.

If, instead, the **runAsUser** strategy is set to **MustRunAsRange**, then a supplied user ID will be validated against a range of allowed IDs. If the pod supplies no user ID, then the default ID is the minimum value of the range of allowable user IDs.

Returning to the earlier [NFS example](#), the container needs its UID set to **65534**, which is shown in the pod fragment above. Assuming the **default** project and the **restricted** SCC, the pod's requested user ID of **65534** will **not** be allowed, and therefore the pod will fail. The pod fails because:

- it requests **65534** as its user ID,
- all available SCCs use **MustRunAsRange** for their **runAsUser** strategy, so UID range checking is required, and
- **65534** is not included in the SCC or project's user ID range.

To address this situation, the recommended path would be to create a new SCC with the appropriate user ID range. A new project could also be created with the appropriate user ID range defined. There are other, less-preferred options:

- The **restricted** SCC could be modified to include **65534** within its minimum and maximum user ID range. This is not recommended as you should avoid modifying the predefined SCCs if possible.
- The **restricted** SCC could be modified to use **RunAsAny** for the **runAsUser** value, thus eliminating ID range checking. This is strongly not recommended, as containers could run as root.
- The **default** project's UID range could be changed to allow a user ID of **65534**. This is not generally advisable because only a single range of user IDs can be specified.

User IDs and Custom SCCs

It is good practice to avoid modifying the predefined SCCs if possible. The preferred approach is to create a custom SCC that better fits an organization's security needs, or [create a new project](#) that supports the desired user IDs.

To remedy the situation in the previous example, a custom SCC can be created such that:

- a minimum and maximum user ID is defined,
- UID range checking is still enforced, and
- the UID of **65534** will be allowed.

For example:

```
$ oc export scc nfs-scc

allowHostDirVolumePlugin: false ❶
...
kind: SecurityContextConstraints
metadata:
  ...
  name: nfs-scc ❷
priority: 9 ❸
requiredDropCapabilities: null
```

```
runAsUser:
  type: MustRunAsRange 4
  uidRangeMax: 65534 5
  uidRangeMin: 65534
...
```

- 1 The **allow*** bools are the same as for the **restricted** SCC.
- 2 The name of this new SCC is **nfs-scc**.
- 3 Numerically larger numbers have greater priority. Nil or omitted is the lowest priority. Higher priority SCCs sort before lower priority SCCs, and thus have a better chance of matching a new pod.
- 4 The **runAsUser** strategy is set to **MustRunAsRange**, which means UID range checking is enforced.
- 5 The UID range is 65534 through 65534 (a range of one value).

Now, with **runAsUser: 65534** shown in the previous pod definition fragment, the pod matches the new **nfs-scc** and is able to run with a UID of 65534.

20.11.6. SELinux Options

All predefined SCCs, except for the **privileged** SCC, set the **seLinuxContext** to **MustRunAs**. So the SCCs most likely to match a pod's requirements will force the pod to use an SELinux policy. The SELinux policy used by the pod can be defined in the pod itself, in the image, in the SCC, or in the project (which provides the default).

SELinux labels can be defined in a pod's **securityContext.seLinuxOptions** section, and supports **user**, **role**, **type**, and **level**:



NOTE

Level and MCS label are used interchangeably in this topic.

```
...
securityContext: 1
  seLinuxOptions:
    level: "s0:c123,c456" 2
...
```

- 1 **level** can be defined globally for the entire pod, or individually for each container.
- 2 SELinux level label.

Here are fragments from an SCC and from the **default** project:

```
$ oc export scc scc-name
...
seLinuxContext:
  type: MustRunAs 1
```

```
# oc export project default
...
metadata:
  annotations:
    openshift.io/sa.scc.mcs: s0:c1,c0 2
...
```

- 1 **MustRunAs** causes volume relabeling.
- 2 If the label is not provided in the pod or in the SCC, then the default comes from the project.

All predefined SCCs, except for the **privileged** SCC, set the **seLinuxContext** to **MustRunAs**. This forces pods to use MCS labels, which can be defined in the pod definition, the image, or provided as a default.

The SCC determines whether or not to require an SELinux label and can provide a default label. If the **seLinuxContext** strategy is set to **MustRunAs** and the pod (or image) does not define a label, OpenShift Container Platform defaults to a label chosen from the SCC itself or from the project.

If **seLinuxContext** is set to **RunAsAny**, then no default labels are provided, and the container determines the final label. In the case of Docker, the container will use a unique MCS label, which will not likely match the labeling on existing storage mounts. Volumes which support SELinux management will be relabeled so that they are accessible by the specified label and, depending on how exclusionary the label is, only that label.

This means two things for unprivileged containers:

- The volume will be given a **type** which is accessible by unprivileged containers. This **type** is usually **svirt_sandbox_file_t**.
- If a **level** is specified, the volume will be labeled with the given MCS label.

For a volume to be accessible by a pod, the pod must have both categories of the volume. So a pod with **s0:c1,c2** will be able to access a volume with **s0:c1,c2**. A volume with **s0** will be accessible by all pods.

If pods fail authorization, or if the storage mount is failing due to permissions errors, then there is a possibility that SELinux enforcement is interfering. One way to check for this is to run:

```
# ausearch -m avc --start recent
```

This examines the log file for AVC (Access Vector Cache) errors.

20.12. SELECTOR-LABEL VOLUME BINDING

20.12.1. Overview

This guide provides the steps necessary to enable binding of persistent volume claims (PVCs) to persistent volumes (PVs) via **selector** and **label** attributes. By implementing selectors and labels, regular users are able to target [provisioned storage](#) by identifiers defined by a cluster administrator.

20.12.2. Motivation

In cases of statically provisioned storage, developers seeking persistent storage are required to know a

handful identifying attributes of a PV in order to deploy and bind a PVC. This creates several problematic situations. Regular users might have to contact a cluster administrator to either deploy the PVC or provide the PV values. PV attributes alone do not convey the intended use of the storage volumes, nor do they provide methods by which volumes can be grouped.

Selector and label attributes can be used to abstract away PV details from the user while providing cluster administrators a way of identifying volumes by a descriptive and customizable tag. Through the selector-label method of binding, users are only required to know which labels are defined by the administrator.



NOTE

The selector-label feature is currently only available for *statically* provisioned storage and is currently not implemented for storage provisioned dynamically.

20.12.3. Deployment

This section reviews how to define and deploy PVCs.

20.12.3.1. Prerequisites

1. A running OpenShift Container Platform 3.3+ cluster
2. A volume provided by a supported [storage provider](#)
3. A user with a cluster-admin role binding

20.12.3.2. Define the Persistent Volume and Claim

1. As the **cluster-admin** user, define the PV. For this example, we will be using a [GlusterFS](#) volume. See the appropriate [storage provider](#) for your provider's configuration.

Example 20.15. Persistent Volume with Labels

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: gluster-volume
  labels: ❶
    volume-type: ssd
    aws-availability-zone: us-east-1
spec:
  capacity:
    storage: 2Gi
  accessModes:
    - ReadWriteMany
  glusterfs:
    endpoints: glusterfs-cluster
    path: myVol1
    readOnly: false
  persistentVolumeReclaimPolicy: Recycle
```

- ❶ A PVC whose selectors match *all* of a PV's labels will be bound, assuming a PV is available.

2. Define the PVC:

Example 20.16. Persistent Volume Claim with Selectors

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gluster-claim
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 1Gi
  selector: ❶
    matchLabels: ❷
      volume-type: ssd
      aws-availability-zone: us-east-1

```

❶ Begin **selectors** section.

❷ List all labels by which the user is requesting storage. Must match *all* labels of targeted PV.

20.12.3.3. Deploy the Persistent Volume and Claim

As the **cluster-admin** user, create the persistent volume:

Example 20.17. Create the Persistent Volume

```

# oc create -f gluster-pv.yaml
persistentVolume "gluster-volume" created

# oc get pv
NAME                                LABELS                                CAPACITY    ACCESSMODES    STATUS
CLAIM      REASON    AGE
gluster-volume    map[]    2147483648    RWX
Available                                2s

```

Once the PV is created, any user whose selectors match *all* its labels can create their PVC.

Example 20.18. Create the Persistent Volume Claim

```

# oc create -f gluster-pvc.yaml
persistentVolumeClaim "gluster-claim" created
# oc get pvc
NAME            LABELS    STATUS    VOLUME
gluster-claim    Bound    gluster-volume

```

CHAPTER 21. PERSISTENT STORAGE EXAMPLES

21.1. OVERVIEW

The following sections provide detailed, comprehensive instructions on setting up and configuring common storage use cases. These examples cover both the administration of persistent volumes and their security, and how to claim against the volumes as a user of the system.

- [Sharing an NFS PV Across Two Pods](#)
- [Ceph-RBD Block Storage Volume](#)
- [Shared Storage Using a GlusterFS Volume](#)
- [Backing Docker Registry with GlusterFS Storage](#)
- [Binding Persistent Volumes by Labels](#)

21.2. SHARING AN NFS MOUNT ACROSS TWO PERSISTENT VOLUME CLAIMS

21.2.1. Overview

The following use case describes how a cluster administrator wanting to leverage shared storage for use by two separate containers would configure the solution. This example highlights the use of NFS, but can easily be adapted to other shared storage types, such as GlusterFS. In addition, this example will show configuration of pod security as it relates to shared storage.

[Persistent Storage Using NFS](#) provides an explanation of persistent volumes (PVs), persistent volume claims (PVCs), and using NFS as persistent storage. This topic shows an end-to-end example of using an existing NFS cluster and OpenShift Container Platform persistent store, and assumes an existing NFS server and exports exist in your OpenShift Container Platform infrastructure.



NOTE

All **oc** commands are executed on the OpenShift Container Platform master host.

21.2.2. Creating the Persistent Volume

Before creating the PV object in OpenShift Container Platform, the persistent volume (PV) file is defined:

Example 21.1. Persistent Volume Object Definition Using NFS

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: nfs-pv 1
spec:
  capacity:
    storage: 1Gi 2
  accessModes:
    - ReadWriteMany 3
```



```

persistentVolumeReclaimPolicy: Retain ❷
nfs: ❸
  path: /opt/nfs ❹
  server: nfs.f22 ❺
  readOnly: false

```

- ❶ The name of the PV, which is referenced in pod definitions or displayed in various **oc** volume commands.
- ❷ The amount of storage allocated to this volume.
- ❸ **accessModes** are used as labels to match a PV and a PVC. They currently do not define any form of access control.
- ❹ The volume reclaim policy **Retain** indicates that the volume will be preserved after the pods accessing it terminates.
- ❺ This defines the volume type being used, in this case the **NFS** plug-in.
- ❻ This is the NFS mount path.
- ❼ This is the NFS server. This can also be specified by IP address.

Save the PV definition to a file, for example *nfs-pv.yaml*, and create the persistent volume:

```

# oc create -f nfs-pv.yaml
persistentvolume "nfs-pv" created

```

Verify that the persistent volume was created:

```

# oc get pv
NAME          LABELS          CAPACITY  ACCESSMODES  STATUS    CLAIM
REASON      AGE
nfs-pv      <none>         1Gi       RWX           Available
37s

```

21.2.3. Creating the Persistent Volume Claim

A persistent volume claim (PVC) specifies the desired access mode and storage capacity. Currently, based on only these two attributes, a PVC is bound to a single PV. Once a PV is bound to a PVC, that PV is essentially tied to the PVC's project and cannot be bound to by another PVC. There is a one-to-one mapping of PVs and PVCs. However, multiple pods in the same project can use the same PVC. This is the use case we are highlighting in this example.

Example 21.2. PVC Object Definition

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: nfs-pvc ❶
spec:
  accessModes:

```

```
- ReadWriteMany      2
resources:
  requests:
    storage: 1Gi      3
```

- 1 The claim name is referenced by the pod under its **volumes** section.
- 2 As mentioned above for PVs, the **accessModes** do not enforce access right, but rather act as labels to match a PV to a PVC.
- 3 This claim will look for PVs offering **1Gi** or greater capacity.

Save the PVC definition to a file, for example **nfs-pvc.yaml**, and create the PVC:

```
# oc create -f nfs-pvc.yaml
persistentvolumeclaim "nfs-pvc" created
```

Verify that the PVC was created and bound to the expected PV:

```
# oc get pvc
NAME          LABELS      STATUS      VOLUME      CAPACITY  ACCESSMODES
AGE
nfs-pvc       <none>      Bound       nfs-pv       1Gi       RWX
24s
```

1

- 1 The claim, **nfs-pvc**, was bound to the **nfs-pv** PV.

21.2.4. Ensuring NFS Volume Access

Access is necessary to a node in the NFS server. On this node, examine the NFS export mount:

```
[root@nfs nfs]# ls -lZ /opt/nfs/
total 8
-rw-r--r--. 1 root 100003 system_u:object_r:usr_t:s0 10 Oct 12 23:27
test2b
```

- 1 the owner has ID 0.
- 2 the group has ID 100003.

In order to access the NFS mount, the container must match the SELinux label, and either run with a UID of 0, or with 100003 in its supplemental groups range. Gain access to the volume by matching the NFS mount's groups, which will be defined in the pod definition below.

By default, SELinux does not allow writing from a pod to a remote NFS server. To enable writing to NFS volumes with SELinux enforcing on each node, run:

■

```
# setsebool -P virt_use_nfs on
```

21.2.5. Creating the Pod

A pod definition file or a template file can be used to define a pod. Below is a pod specification that creates a single container and mounts the NFS volume for read-write access:

Example 21.3. Pod Object Definition

```
apiVersion: v1
kind: Pod
metadata:
  name: hello-openshift-nfs-pod ❶
  labels:
    name: hello-openshift-nfs-pod
spec:
  containers:
    - name: hello-openshift-nfs-pod
      image: openshift/hello-openshift ❷
      ports:
        - name: web
          containerPort: 80
      volumeMounts:
        - name: nfsvol ❸
          mountPath: /usr/share/nginx/html ❹
  securityContext:
    supplementalGroups: [100003] ❺
    privileged: false
  volumes:
    - name: nfsvol
      persistentVolumeClaim:
        claimName: nfs-pvc ❻
```

❶ The name of this pod as displayed by `oc get pod`.

❷ The image run by this pod.

❸ The name of the volume. This name must be the same in both the **containers** and **volumes** sections.

❹ The mount path as seen in the container.

❺ The group ID to be assigned to the container.

❻ The PVC that was created in the previous step.

Save the pod definition to a file, for example *nfs.yaml*, and create the pod:

```
# oc create -f nfs.yaml
pod "hello-openshift-nfs-pod" created
```

Verify that the pod was created:

```
# oc get pods
NAME                                READY    STATUS    RESTARTS   AGE
hello-openshift-nfs-pod            1/1      Running   0           4s
```

More details are shown in the **oc describe pod** command:

```
[root@ose70 nfs]# oc describe pod hello-openshift-nfs-pod
Name:      hello-openshift-nfs-pod
Namespace: default 1
Image(s):  fedora/S3
Node:      ose70.rh7/192.168.234.148 2
Start Time: Mon, 21 Mar 2016 09:59:47 -0400
Labels:    name=hello-openshift-nfs-pod
Status:     Running
Reason:
Message:
IP:        10.1.0.4
Replication Controllers: <none>
Containers:
  hello-openshift-nfs-pod:
    Container ID:
docker://a3292104d6c28d9cf49f440b2967a0fc5583540fc3b062db598557b93893bc6f
    Image:      fedora/S3
    Image ID:
docker://403d268c640894cbd76d84a1de3995d2549a93af51c8e16e89842e4c3ed6a00a
    QoS Tier:
      cpu:      BestEffort
      memory:   BestEffort
    State:      Running
      Started:  Mon, 21 Mar 2016 09:59:49 -0400
    Ready:      True
    Restart Count: 0
    Environment Variables:
Conditions:
  Type      Status
  Ready     True
Volumes:
  nfsvol:
    Type: PersistentVolumeClaim (a reference to a PersistentVolumeClaim in
the same namespace)
    ClaimName: nfs-pvc 3
    ReadOnly: false
  default-token-a06zb:
    Type: Secret (a secret that should populate this volume)
    SecretName: default-token-a06zb
Events: 4
  FirstSeen LastSeen Count From      SubobjectPath
Reason      Message
-----
4m 4m 1 {scheduler }
Scheduled Successfully assigned hello-openshift-nfs-pod to ose70.rh7
4m 4m 1 {kubelet ose70.rh7} implicitly required container POD
```

```

Pulled Container image "openshift3/ose-pod:v3.1.0.4" already present on
machine
 4m 4m 1 {kubelet ose70.rh7} implicitly required container POD
Created Created with docker id 866a37108041
 4m 4m 1 {kubelet ose70.rh7} implicitly required container POD
Started Started with docker id 866a37108041
 4m 4m 1 {kubelet ose70.rh7} spec.containers{hello-openshift-nfs-pod}
Pulled Container image "fedora/S3" already present on machine
 4m 4m 1 {kubelet ose70.rh7} spec.containers{hello-openshift-nfs-pod}
Created Created with docker id a3292104d6c2
 4m 4m 1 {kubelet ose70.rh7} spec.containers{hello-openshift-nfs-pod}
Started Started with docker id a3292104d6c2

```

- ❶ The project (namespace) name.
- ❷ The IP address of the OpenShift Container Platform node running the pod.
- ❸ The PVC name used by the pod.
- ❹ The list of events resulting in the pod being launched and the NFS volume being mounted. The container will not start correctly if the volume cannot mount.

There is more internal information, including the SCC used to authorize the pod, the pod's user and group IDs, the SELinux label, and more, shown in the **oc get pod <name> -o yaml** command:

```

[root@ose70 nfs]# oc get pod hello-openshift-nfs-pod -o yaml
apiVersion: v1
kind: Pod
metadata:
  annotations:
    openshift.io/scc: restricted ❶
  creationTimestamp: 2016-03-21T13:59:47Z
  labels:
    name: hello-openshift-nfs-pod
    name: hello-openshift-nfs-pod
  namespace: default ❷
  resourceVersion: "2814411"
  selflink: /api/v1/namespaces/default/pods/hello-openshift-nfs-pod
  uid: 2c22d2ea-ef6d-11e5-adc7-000c2900f1e3
spec:
  containers:
  - image: fedora/S3
    imagePullPolicy: IfNotPresent
    name: hello-openshift-nfs-pod
    ports:
    - containerPort: 80
      name: web
      protocol: TCP
    resources: {}
    securityContext:
      privileged: false
    terminationMessagePath: /dev/termination-log
    volumeMounts:
    - mountPath: /usr/share/S3/html
      name: nfsvol

```

```

    - mountPath: /var/run/secrets/kubernetes.io/serviceaccount
      name: default-token-a06zb
      readOnly: true
  dnsPolicy: ClusterFirst
  host: ose70.rh7
  imagePullSecrets:
    - name: default-dockercfg-xvdew
  nodeName: ose70.rh7
  restartPolicy: Always
  securityContext:
    supplementalGroups:
      - 100003 ❸
  serviceAccount: default
  serviceAccountName: default
  terminationGracePeriodSeconds: 30
  volumes:
    - name: nfsvol
      persistentVolumeClaim:
        claimName: nfs-pvc ❹
    - name: default-token-a06zb
      secret:
        secretName: default-token-a06zb
  status:
    conditions:
      - lastProbeTime: null
        lastTransitionTime: 2016-03-21T13:59:49Z
        status: "True"
        type: Ready
    containerStatuses:
      - containerID:
  docker://a3292104d6c28d9cf49f440b2967a0fc5583540fc3b062db598557b93893bc6f
        image: fedora/S3
        imageID:
  docker://403d268c640894cbd76d84a1de3995d2549a93af51c8e16e89842e4c3ed6a00a
        lastState: {}
        name: hello-openshift-nfs-pod
        ready: true
        restartCount: 0
        state:
          running:
            startedAt: 2016-03-21T13:59:49Z
  hostIP: 192.168.234.148
  phase: Running
  podIP: 10.1.0.4
  startTime: 2016-03-21T13:59:47Z

```

- ❶ The SCC used by the pod.
- ❷ The project (namespace) name.
- ❸ The supplemental group ID for the pod (all containers).
- ❹ The PVC name used by the pod.

21.2.6. Creating an Additional Pod to Reference the Same PVC

This pod definition, created in the same namespace, uses a different container. However, we can use the same backing storage by specifying the claim name in the volumes section below:

Example 21.4. Pod Object Definition

```
apiVersion: v1
kind: Pod
metadata:
  name: busybox-nfs-pod ❶
  labels:
    name: busybox-nfs-pod
spec:
  containers:
    - name: busybox-nfs-pod
      image: busybox ❷
      command: ["sleep", "60000"]
      volumeMounts:
        - name: nfsvol-2 ❸
          mountPath: /usr/share/busybox ❹
          readOnly: false
      securityContext:
        supplementalGroups: [100003] ❺
        privileged: false
  volumes:
    - name: nfsvol-2
      persistentVolumeClaim:
        claimName: nfs-pvc ❻
```

- ❶ The name of this pod as displayed by `oc get pod`.
- ❷ The image run by this pod.
- ❸ The name of the volume. This name must be the same in both the **containers** and **volumes** sections.
- ❹ The mount path as seen in the container.
- ❺ The group ID to be assigned to the container.
- ❻ The PVC that was created earlier and is also being used by a different container.

Save the pod definition to a file, for example *nfs-2.yaml*, and create the pod:

```
# oc create -f nfs-2.yaml
pod "busybox-nfs-pod" created
```

Verify that the pod was created:

```
# oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
busybox-nfs-pod	1/1	Running	0	3s

More details are shown in the **oc describe pod** command:

```
[root@ose70 nfs]# oc describe pod busybox-nfs-pod
Name:      busybox-nfs-pod
Namespace:  default
Image(s):  busybox
Node:      ose70.rh7/192.168.234.148
Start Time:   Mon, 21 Mar 2016 10:19:46 -0400
Labels:      name=busybox-nfs-pod
Status:      Running
Reason:
Message:
IP:          10.1.0.5
Replication Controllers: <none>
Containers:
  busybox-nfs-pod:
    Container ID:
docker://346d432e5a4824ebf5a47fceb4247e0568ecc64eadcc160e9bab481aecfb0594
    Image:      busybox
    Image ID:
docker://17583c7dd0dae6244203b8029733bdb7d17fccbb2b5d93e2b24cf48b8bfd06e2
    QoS Tier:
      cpu:      BestEffort
      memory:   BestEffort
    State:      Running
      Started:   Mon, 21 Mar 2016 10:19:48 -0400
    Ready:      True
    Restart Count: 0
    Environment Variables:
Conditions:
  Type      Status
  Ready     True
Volumes:
  nfsvol-2:
    Type: PersistentVolumeClaim (a reference to a PersistentVolumeClaim in
the same namespace)
    ClaimName: nfs-pvc
    ReadOnly: false
  default-token-32d2z:
    Type: Secret (a secret that should populate this volume)
    SecretName: default-token-32d2z
Events:
  FirstSeen  LastSeen  Count  From              SubobjectPath      Reason      Message
  -----
  4m         4m         1  {scheduler }      Scheduled           Successfully assigned busybox-
nfs-pod to ose70.rh7
  4m         4m         1  {kubelet ose70.rh7} implicitly required container POD Pulled
Container image "openshift3/ose-pod:v3.1.0.4" already present on machine
  4m         4m         1  {kubelet ose70.rh7} implicitly required container POD Created
Created with docker id 249b7d7519b1
  4m         4m         1  {kubelet ose70.rh7} implicitly required container POD Started
Started with docker id 249b7d7519b1
  4m         4m         1  {kubelet ose70.rh7} spec.containers{busybox-nfs-pod} Pulled
Container image "busybox" already present on machine
  4m         4m         1  {kubelet ose70.rh7} spec.containers{busybox-nfs-pod} Created
```



```
Created with docker id 346d432e5a48
4m 4m 1 {kubelet ose70.rh7} spec.containers{busybox-nfs-pod} Started
Started with docker id 346d432e5a48
```

As you can see, both containers are using the same storage claim that is attached to the same NFS mount on the back end.

21.3. COMPLETE EXAMPLE USING CEPH RBD

21.3.1. Overview

This topic provides an end-to-end example of using an existing Ceph cluster as an OpenShift Container Platform persistent store. It is assumed that a working Ceph cluster is already set up. If not, consult the [Overview of Red Hat Ceph Storage](#).

[Persistent Storage Using Ceph Rados Block Device](#) provides an explanation of persistent volumes (PVs), persistent volume claims (PVCs), and using Ceph RBD as persistent storage.



NOTE

All **oc** ... commands are executed on the OpenShift Container Platform master host.

21.3.2. Installing the ceph-common Package

The **ceph-common** library must be installed on **all schedulable** OpenShift Container Platform nodes:



NOTE

The OpenShift Container Platform all-in-one host is not often used to run pod workloads and, thus, is not included as a schedulable node.

```
# yum install -y ceph-common
```

21.3.3. Creating the Ceph Secret

The **ceph auth get-key** command is run on a Ceph **MON** node to display the key value for the **client.admin** user:

Example 21.5. Ceph Secret Definition

```
apiVersion: v1
kind: Secret
metadata:
  name: ceph-secret
data:
  key: QVFB0FF2S1ZheUJQRVJBQWgvS2cwT1laQUhPQno3akZwekxxdGc9PQ== 1
```

- 1 This base64 key is generated on one of the Ceph MON nodes using the **ceph auth get-key client.admin | base64** command, then copying the output and pasting it as the secret key's value.

Save the secret definition to a file, for example ***ceph-secret.yaml***, then create the secret:

```
$ oc create -f ceph-secret.yaml
secret "ceph-secret" created
```

Verify that the secret was created:

```
# oc get secret ceph-secret
NAME          TYPE      DATA   AGE
ceph-secret   Opaque    1       23d
```

21.3.4. Creating the Persistent Volume

Next, before creating the PV object in OpenShift Container Platform, define the persistent volume file:

Example 21.6. Persistent Volume Object Definition Using Ceph RBD

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: ceph-pv
spec:
  capacity:
    storage: 2Gi
  accessModes:
    - ReadWriteOnce
  rbd:
    monitors:
      - 192.168.122.133:6789
    pool: rbd
    image: ceph-image
    user: admin
    secretRef:
      name: ceph-secret
    fsType: ext4
    readOnly: false
  persistentVolumeReclaimPolicy: Recycle
```

- 1 The name of the PV, which is referenced in pod definitions or displayed in various **oc** volume commands.
- 2 The amount of storage allocated to this volume.
- 3 **accessModes** are used as labels to match a PV and a PVC. They currently do not define any form of access control. All block storage is defined to be single user (non-shared storage).
- 4 This defines the volume type being used. In this case, the **rbd** plug-in is defined.
- 5 This is an array of Ceph monitor IP addresses and ports.
- 6 This is the Ceph secret, defined above. It is used to create a secure connection from OpenShift Container Platform to the Ceph server.

- 7 This is the file system type mounted on the Ceph RBD block device.

Save the PV definition to a file, for example ***ceph-pv.yaml***, and create the persistent volume:

```
# oc create -f ceph-pv.yaml
persistentvolume "ceph-pv" created
```

Verify that the persistent volume was created:

```
# oc get pv
NAME                                LABELS            CAPACITY          ACCESSMODES        STATUS
CLAIM          REASON          AGE
ceph-pv        <none>          2147483648        RWO                Available
2s
```

21.3.5. Creating the Persistent Volume Claim

A persistent volume claim (PVC) specifies the desired access mode and storage capacity. Currently, based on only these two attributes, a PVC is bound to a single PV. Once a PV is bound to a PVC, that PV is essentially tied to the PVC's project and cannot be bound to by another PVC. There is a one-to-one mapping of PVs and PVCs. However, multiple pods in the same project can use the same PVC.

Example 21.7. PVC Object Definition

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: ceph-claim
spec:
  accessModes: 1
    - ReadWriteOnce
  resources:
    requests:
      storage: 2Gi 2
```

- 1 As mentioned above for PVs, the **accessModes** do not enforce access right, but rather act as labels to match a PV to a PVC.
- 2 This claim will look for PVs offering **2Gi** or greater capacity.

Save the PVC definition to a file, for example ***ceph-claim.yaml***, and create the PVC:

```
# oc create -f ceph-claim.yaml
persistentvolumeclaim "ceph-claim" created

#and verify the PVC was created and bound to the expected PV:
# oc get pvc
```

NAME	LABELS	STATUS	VOLUME	CAPACITY	ACCESSMODES	AGE
ceph-claim	<none>	Bound	ceph-pv	1Gi	RWX	21s

- 1 the claim was bound to the **ceph-pv** PV.

21.3.6. Creating the Pod

A pod definition file or a template file can be used to define a pod. Below is a pod specification that creates a single container and mounts the Ceph RBD volume for read-write access:

Example 21.8. Pod Object Definition

```

apiVersion: v1
kind: Pod
metadata:
  name: ceph-pod1
spec:
  containers:
  - name: ceph-busybox
    image: busybox
    command: ["sleep", "60000"]
    volumeMounts:
    - name: ceph-vol1
      mountPath: /usr/share/busybox
      readOnly: false
  volumes:
  - name: ceph-vol1
    persistentVolumeClaim:
      claimName: ceph-claim

```

- 1 The name of this pod as displayed by **oc get pod**.
- 2 The image run by this pod. In this case, we are telling **busybox** to sleep.
- 3 5 The name of the volume. This name must be the same in both the **containers** and **volumes** sections.
- 4 The mount path as seen in the container.
- 6 The PVC that is bound to the Ceph RBD cluster.

Save the pod definition to a file, for example **ceph-pod1.yaml**, and create the pod:

```

# oc create -f ceph-pod1.yaml
pod "ceph-pod1" created

#verify pod was created
# oc get pod

```

NAME	READY	STATUS	RESTARTS	AGE
ceph-pod1	1/1	Running	0	2m

1

- 1 After a minute or so, the pod will be in the **Running** state.

21.3.7. Defining Group and Owner IDs (Optional)

When using block storage, such as Ceph RBD, the physical block storage is **managed** by the pod. The group ID defined in the pod becomes the group ID of **both** the Ceph RBD mount inside the container, and the group ID of the actual storage itself. Thus, it is usually unnecessary to define a group ID in the pod specification. However, if a group ID is desired, it can be defined using **fsGroup**, as shown in the following pod definition fragment:

Example 21.9. Group ID Pod Definition

```
...
spec:
  containers:
    - name:
      ...
  securityContext: 1
    fsGroup: 7777 2
  ...
```


- 1 **securityContext** must be defined at the pod level, not under a specific container.
- 2 All containers in the pod will have the same **fsGroup** ID.

21.4. COMPLETE EXAMPLE USING GLUSTERFS

21.4.1. Overview

This topic provides an end-to-end example of how to use an existing Gluster cluster as an OpenShift Container Platform persistent store. It is assumed that a working Gluster cluster is already set up. If not, consult the [Red Hat Gluster Storage Administration Guide](#).

[Persistent Storage Using GlusterFS](#) provides an explanation of persistent volumes (PVs), persistent volume claims (PVCs), and using GlusterFS as persistent storage.




NOTE

All **oc** ... commands are executed on the OpenShift Container Platform master host.

21.4.2. Installing the glusterfs-fuse Package

The **glusterfs-fuse** library must be installed on all **schedulable** OpenShift Container Platform nodes:

```
# yum install -y glusterfs-fuse
```

**NOTE**

The OpenShift Container Platform all-in-one host is often not used to run pod workloads and, thus, is not included as a schedulable node.

21.4.3. Creating the Gluster Endpoints and Gluster Service for Persistence

The named endpoints define each node in the Gluster-trusted storage pool:

Example 21.10. GlusterFS Endpoint Definition

```
apiVersion: v1
kind: Endpoints
metadata:
  name: gluster-cluster 1
subsets:
- addresses: 2
  - ip: 192.168.122.21
  ports: 3
    - port: 1
      protocol: TCP
- addresses:
  - ip: 192.168.122.22
  ports:
    - port: 1
      protocol: TCP
```

- 1** The endpoints name. If using a service, then the endpoints name must match the service name.
- 2** An array of IP addresses for each node in the Gluster pool. Currently, host names are not supported.
- 3** The port numbers are ignored, but must be legal port numbers. The value 1 is commonly used.

Save the endpoints definition to a file, for example ***gluster-endpoints.yaml***, then create the endpoints object:

```
# oc create -f gluster-endpoints.yaml
endpoints "gluster-cluster" created
```

Verify that the endpoints were created:

```
# oc get endpoints gluster-cluster
NAME                ENDPOINTS                                AGE
gluster-cluster     192.168.122.21:1,192.168.122.22:1     1m
```

**NOTE**

To persist the Gluster endpoints, you also need to create a service.

**NOTE**

Endpoints are name-spaced. Each project accessing the Gluster volume needs its own endpoints.

Example 21.11. GlusterFS Service Definition

```
apiVersion: v1
kind: Service
metadata:
  name: gluster-cluster ❶
spec:
  ports:
    - port: 1 ❷
```

❶ The name of the service. If using a service, then the endpoints name must match the service name.

❷ The port should match the same port used in the endpoints.

Save the service definition to a file, for example ***gluster-service.yaml***, then create the endpoints object:

```
# oc create -f gluster-service.yaml
endpoints "gluster-cluster" created
```

Verify that the service was created:

```
# oc get service gluster-cluster
```

NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
gluster-cluster	10.0.0.130	<none>	1/TCP	9s

21.4.4. Creating the Persistent Volume

Next, before creating the PV object, define the persistent volume in OpenShift Container Platform:

Persistent Volume Object Definition Using GlusterFS

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: gluster-pv ❶
spec:
  capacity:
    storage: 1Gi ❷
  accessModes:
    - ReadWriteMany ❸
  glusterfs: ❹
    endpoints: gluster-cluster ❺
```

```

    path: /HadoopVol 6
    readOnly: false
    persistentVolumeReclaimPolicy: Retain 7

```

- 1** The name of the PV, which is referenced in pod definitions or displayed in various **oc** volume commands.
- 2** The amount of storage allocated to this volume.
- 3** **accessModes** are used as labels to match a PV and a PVC. They currently do not define any form of access control.
- 4** This defines the volume type being used. In this case, the **glusterfs** plug-in is defined.
- 5** This references the endpoints named above.
- 6** This is the Gluster volume name, preceded by **/**.
- 7** The volume reclaim policy **Retain** indicates that the volume will be preserved after the pods accessing it terminates. For GlusterFS, the accepted values include **Retain**, and **Delete**.

Save the PV definition to a file, for example **gluster-pv.yaml**, and create the persistent volume:

```

# oc create -f gluster-pv.yaml
persistentvolume "gluster-pv" created

```

Verify that the persistent volume was created:

```

# oc get pv
NAME          LABELS          CAPACITY  ACCESSMODES  STATUS    CLAIM
REASON    AGE
gluster-pv  <none>         1Gi       RWX           Available
37s

```

21.4.5. Creating the Persistent Volume Claim

A persistent volume claim (PVC) specifies the desired access mode and storage capacity. Currently, based on only these two attributes, a PVC is bound to a single PV. Once a PV is bound to a PVC, that PV is essentially tied to the PVC's project and cannot be bound to by another PVC. There is a one-to-one mapping of PVs and PVCs. However, multiple pods in the same project can use the same PVC.

Example 21.12. PVC Object Definition

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gluster-claim 1
spec:
  accessModes:
    - ReadWriteMany 2
  resources:
    requests:
      storage: 1Gi 3

```


**NOTE**

The `virt_sandbox_use_fusefs` boolean is defined by the `docker-selinux` package. If you get an error saying it is not defined, ensure that this package is installed.

21.4.7. Creating the Pod using NGINX Web Server image

A pod definition file or a template file can be used to define a pod. Below is a pod specification that creates a single container and mounts the Gluster volume for read-write access:

**NOTE**

The NGINX image may require to run in privileged mode to create the mount and run properly. An easy way to accomplish this is to simply add your user to the **privileged** Security Context Constraint (SCC):

```
$ oadm policy add-scc-to-user privileged myuser
```

Then, add the **privileged: true** to the containers `securityContext:` section of the YAML file (as seen in the example below).

[Managing Security Context Constraints](#) provides additional information regarding SCCs.

Example 21.13. Pod Object Definition using NGINX image

```
apiVersion: v1
kind: Pod
metadata:
  name: gluster-pod1
  labels:
    name: gluster-pod1 1
spec:
  containers:
    - name: gluster-pod1
      image: nginx 2
      ports:
        - name: web
          containerPort: 80
      securityContext:
        privileged: true
      volumeMounts:
        - name: gluster-vol1 3
          mountPath: /usr/share/nginx/html 4
          readOnly: false
      securityContext:
        supplementalGroups: [590] 5
  volumes:
    - name: gluster-vol1 6
      persistentVolumeClaim:
        claimName: gluster-claim 7
```

1 The name of this pod as displayed by `oc get pod`.

- 2 The image run by this pod. In this case, we are using a standard NGINX image.
- 3 6 The name of the volume. This name must be the same in both the **containers** and **volumes** sections.
- 4 The mount path as seen in the container.
- 5 The **SupplementalGroup** ID (Linux Groups) to be assigned at the pod level and as discussed this should match the POSIX permissions on the Gluster volume.
- 7 The PVC that is bound to the Gluster cluster.

Save the pod definition to a file, for example *gluster-pod1.yaml*, and create the pod:

```
# oc create -f gluster-pod1.yaml
pod "gluster-pod1" created
```

Verify the pod was created:

```
# oc get pod
NAME          READY   STATUS    RESTARTS   AGE
gluster-pod1  1/1     Running   0           31s
```

1

- 1 After a minute or so, the pod will be in the **Running** state.

More details are shown in the **oc describe pod** command:

```
# oc describe pod gluster-pod1
Name:      gluster-pod1
Namespace: default 1
Security Policy: privileged
Node:      ose1.rhs/192.168.122.251
Start Time: Wed, 24 Aug 2016 12:37:45 -0400
Labels:    name=gluster-pod1
Status:    Running
IP:        172.17.0.2 2
Controllers: <none>
Containers:
  gluster-pod1:
    Container ID:
docker://e67ed01729e1dc7369c5112d07531a27a7a02a7eb942f17d1c5fce32d8c31a2d
    Image:      nginx
    Image ID:
docker://sha256:4efb2fcdb1ab05fb03c9435234343c1cc65289eeb016be86193e88d3a5d84f6b
    Port:      80/TCP
    State:     Running
      Started: Wed, 24 Aug 2016 12:37:52 -0400
    Ready:     True
    Restart Count: 0
```

```

    Volume Mounts:
      /usr/share/nginx/html/test from glustervol (rw)
      /var/run/secrets/kubernetes.io/serviceaccount from default-token-
1n70u (ro)
    Environment Variables: <none>
Conditions:
  Type      Status
  Initialized True
  Ready      True
  PodSchedul True
Volumes:
  glustervol:
    Type: PersistentVolumeClaim (a reference to a PersistentVolumeClaim in
the same namespace)
    ClaimName: gluster-claim ❸
    ReadOnly: false
  default-token-1n70u:
    Type: Secret (a volume populated by a Secret)
    SecretName: default-token-1n70u
QoS Tier: BestEffort
Events: ❹
  FirstSeen LastSeen Count From           SubobjectPath      Type    Reason  Message
  -----
  ---
  10s 10s 1 {default-scheduler }      Normal    Scheduled Successfully
assigned gluster-pod1 to ose1.rhs
  9s 9s 1 {kubelet ose1.rhs} spec.containers{gluster-pod1} Normal
Pulling pulling image "nginx"
  4s 4s 1 {kubelet ose1.rhs} spec.containers{gluster-pod1} Normal
Pulled Successfully pulled image "nginx"
  3s 3s 1 {kubelet ose1.rhs} spec.containers{gluster-pod1} Normal
Created Created container with docker id e67ed01729e1
  3s 3s 1 {kubelet ose1.rhs} spec.containers{gluster-pod1} Normal
Started Started container with docker id e67ed01729e1

```

- ❶ The project (namespace) name.
- ❷ The IP address of the OpenShift Container Platform node running the pod.
- ❸ The PVC name used by the pod.
- ❹ The list of events resulting in the pod being launched and the Gluster volume being mounted.

There is more internal information, including the SCC used to authorize the pod, the pod's user and group IDs, the SELinux label, and more shown in the **oc get pod <name> -o yaml** command:

```

# oc get pod gluster-pod1 -o yaml
apiVersion: v1
kind: Pod
metadata:
  annotations:
    openshift.io/scc: privileged ❶
  creationTimestamp: 2016-08-24T16:37:45Z
  labels:
    name: gluster-pod1

```

```

name: gluster-pod1
namespace: default ❷
resourceVersion: "482"
selfLink: /api/v1/namespaces/default/pods/gluster-pod1
uid: 15afda77-6a19-11e6-aadb-525400f7256d
spec:
  containers:
  - image: nginx
    imagePullPolicy: Always
    name: gluster-pod1
    ports:
    - containerPort: 80
      name: web
      protocol: TCP
    resources: {}
    securityContext:
      privileged: true ❸
    terminationMessagePath: /dev/termination-log
    volumeMounts:
    - mountPath: /usr/share/nginx/html
      name: glustervol
    - mountPath: /var/run/secrets/kubernetes.io/serviceaccount
      name: default-token-1n70u
      readOnly: true
  dnsPolicy: ClusterFirst
  host: ose1.rhs
  imagePullSecrets:
  - name: default-dockercfg-20xg9
  nodeName: ose1.rhs
  restartPolicy: Always
  securityContext:
    supplementalGroups:
    - 590 ❹
  serviceAccount: default
  serviceAccountName: default
  terminationGracePeriodSeconds: 30
  volumes:
  - name: glustervol
    persistentVolumeClaim:
      claimName: gluster-claim ❺
  - name: default-token-1n70u
    secret:
      secretName: default-token-1n70u
status:
  conditions:
  - lastProbeTime: null
    lastTransitionTime: 2016-08-24T16:37:45Z
    status: "True"
    type: Initialized
  - lastProbeTime: null
    lastTransitionTime: 2016-08-24T16:37:53Z
    status: "True"
    type: Ready
  - lastProbeTime: null
    lastTransitionTime: 2016-08-24T16:37:45Z
    status: "True"

```

```
    type: PodScheduled
  containerStatuses:
  - containerID:
docker://e67ed01729e1dc7369c5112d07531a27a7a02a7eb942f17d1c5fce32d8c31a2d
    image: nginx
    imageID:
docker://sha256:4efb2fcdb1ab05fb03c9435234343c1cc65289eeb016be86193e88d3a5
d84f6b
    lastState: {}
    name: gluster-pod1
    ready: true
    restartCount: 0
    state:
      running:
        startedAt: 2016-08-24T16:37:52Z
  hostIP: 192.168.122.251
  phase: Running
  podIP: 172.17.0.2
  startTime: 2016-08-24T16:37:45Z
```

- ❶ The SCC used by the pod.
- ❷ The project (namespace) name.
- ❸ The security context level requested, in this case privileged
- ❹ The supplemental group ID for the pod (all containers).
- ❺ The PVC name used by the pod.

21.5. BACKING DOCKER REGISTRY WITH GLUSTERFS STORAGE

21.5.1. Overview

This topic reviews how to attach a GlusterFS persistent volume to the Docker Registry.

It is assumed that the Docker registry service has already been started and the Gluster volume has been created.

21.5.2. Prerequisites

- The [docker-registry](#) was deployed **without** configuring storage.
- A Gluster volume exists and **glusterfs-fuse** is installed on schedulable nodes.
- Definitions written for GlusterFS [endpoints and service](#), [persistent volume \(PV\)](#), and [persistent volume claim \(PVC\)](#).
 - For this guide, these will be:
 - ***gluster-endpoints-service.yaml***
 - ***gluster-endpoints.yaml***
 - ***gluster-pv.yaml***

■ *gluster-pvc.yaml*

- A user with the [cluster-admin](#) role binding.
 - For this guide, that user is **admin**.



NOTE

All **oc** commands are executed on the master node as the **admin** user.

21.5.3. Create the Gluster Persistent Volume

First, make the Gluster volume available to the registry.

```
$ oc create -f gluster-endpoints-service.yaml
$ oc create -f gluster-endpoints.yaml
$ oc create -f gluster-pv.yaml
$ oc create -f gluster-pvc.yaml
```

Check to make sure the PV and PVC were created and bound successfully. The expected output should resemble the following. Note that the PVC status is **Bound**, indicating that it has bound to the PV.

```
$ oc get pv
NAME                LABELS              CAPACITY  ACCESSMODES  STATUS      CLAIM
REASON              AGE
gluster-pv          <none>             1Gi       RWX           Available
37s
$ oc get pvc
NAME                LABELS              STATUS      VOLUME        CAPACITY  ACCESSMODES
AGE
gluster-claim       <none>             Bound       gluster-pv    1Gi       RWX
24s
```



NOTE

If either the PVC or PV failed to create or the PVC failed to bind, refer back to the [GlusterFS Persistent Storage](#) guide. **Do not** proceed until they initialize and the PVC status is **Bound**.

21.5.4. Attach the PVC to the Docker Registry

Before moving forward, ensure that the **docker-registry** service is running.

```
$ oc get svc
NAME                CLUSTER_IP          EXTERNAL_IP  PORT(S)
SELECTOR            AGE
docker-registry     172.30.167.194      <none>       5000/TCP
docker-registry=default 18m
```



NOTE

If either the **docker-registry** service or its associated pod is not running, refer back to the [docker-registry](#) setup instructions for troubleshooting before continuing.

Then, attach the PVC:

```
$ oc volume deploymentconfigs/docker-registry --add --name=registry-storage -t pvc \
    --claim-name=gluster-claim --overwrite
```

[Deploying a Docker Registry](#) provides more information on using the Docker registry.

21.5.5. Known Issues

21.5.5.1. Pod Cannot Resolve the Volume Host

In non-production cases where the **dnsmasq** server is located on the same node as the OpenShift Container Platform master service, pods might not resolve to the host machines when mounting the volume, causing errors in the **docker-registry-1-deploy** pod. This can happen when **dnsmasq.service** fails to start because of a collision with OpenShift Container Platform DNS on port 53. To run the DNS server on the master host, some configurations needs to be changed.

In **/etc/dnsmasq.conf**, add:

```
# Reverse DNS record for master
host-record=master.example.com,<master-IP>
# Wildcard DNS for OpenShift Applications - Points to Router
address=/apps.example.com/<master-IP>
# Forward .local queries to SkyDNS
server=/local/127.0.0.1#8053
# Forward reverse queries for service network to SkyDNS.
# This is for default OpenShift SDN - change as needed.
server=/17.30.172.in-addr.arpa/127.0.0.1#8053
```

With these settings, **dnsmasq** will pull from the **/etc/hosts** file on the master node.

Add the appropriate host names and IPs for all necessary hosts.

In **master-config.yaml**, change **bindAddress** to:

```
dnsConfig:
  bindAddress: 127.0.0.1:8053
```

When pods are created, they receive a copy of **/etc/resolv.conf**, which typically contains only the master DNS server so they can resolve external DNS requests. To enable internal DNS resolution, insert the **dnsmasq** server at the top of the server list. This way, **dnsmasq** will attempt to resolve requests internally first.

In **/etc/resolv.conf** all scheduled nodes:

```
nameserver 192.168.1.100 ❶
nameserver 192.168.1.1    ❷
```

❶ Add the internal DNS server.

❷ Pre-existing external DNS server.

Once the configurations are changed, restart the OpenShift Container Platform master and **dnsmasq** services.

```
$ systemctl restart atomic-openshift-master
$ systemctl restart dnsmasq
```

21.6. BINDING PERSISTENT VOLUMES BY LABELS

21.6.1. Overview

This topic provides an end-to-end example for binding [persistent volume claims \(PVCs\)](#) to [persistent volumes \(PVs\)](#), by defining labels in the PV and matching selectors in the PVC. This feature is available for all [storage options](#). It is assumed that a OpenShift Container Platform cluster contains persistent storage resources which are available for binding by PVCs.

A Note on Labels and Selectors

Labels are an OpenShift Container Platform feature that support user-defined tags (key-value pairs) as part of an object's specification. Their primary purpose is to enable the arbitrary grouping of objects by defining identical labels among them. These labels can then be targeted by selectors to match all objects with specified label values. It is this functionality we will take advantage of to enable our PVC to bind to our PV. For a more in-depth look at labels, see [Pods and Services](#).



NOTE

For this example, we will be using modified [GlusterFS](#) PV and PVC specifications. However, implementation of selectors and labels is generic across for all storage options. See the [relevant storage option](#) for your volume provider to learn more about its unique configuration.

21.6.1.1. Assumptions

It is assumed that you have:

- An existing OpenShift Container Platform cluster with at least one **master** and one **node**
- At least one supported [storage volume](#)
- A user with **cluster-admin** privileges

21.6.2. Defining Specifications



NOTE

These specifications are tailored to **GlusterFS**. Consult the [relevant storage option](#) for your volume provider to learn more about its unique configuration.

21.6.2.1. Persistent Volume with Labels

Example 21.14. glusterfs-pv.yaml

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

```

metadata:
  name: gluster-volume
  labels: ❶
    storage-tier: gold
    aws-availability-zone: us-east-1
spec:
  capacity:
    storage: 2Gi
  accessModes:
    - ReadWriteMany
  glusterfs:
    endpoints: glusterfs-cluster ❷
    path: myVol1
    readOnly: false
  persistentVolumeReclaimPolicy: Retain

```

❶ Use labels to identify common attributes or characteristics shared among volumes. In this case, we defined the Gluster volume to have a custom attribute (key) named **storage-tier** with a value of **gold** assigned. A claim will be able to select a PV with **storage-tier=gold** to match this PV.

❷ Endpoints define the Gluster trusted pool and are discussed below.

21.6.2.2. Persistent Volume Claim with Selectors

A claim with a **selector** stanza (see example below) attempts to match existing, unclaimed, and non-prebound PVs. The existence of a PVC selector ignores a PV's capacity. However, **accessModes** are still considered in the matching criteria.

It is important to note that a claim must match **all** of the key-value pairs included in its **selector** stanza. If no PV matches the claim, then the PVC will remain unbound (Pending). A PV can subsequently be created and the claim will automatically check for a label match.

Example 21.15. glusterfs-pvc.yaml

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gluster-claim
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 1Gi
  selector: ❶
    matchLabels:
      storage-tier: gold
      aws-availability-zone: us-east-1

```

❶ The **selector** stanza defines all labels necessary in a PV in order to match this claim.

21.6.2.3. Volume Endpoints

To attach the PV to the Gluster volume, endpoints should be configured before creating our objects.

Example 21.16. glusterfs-ep.yaml

```
apiVersion: v1
kind: Endpoints
metadata:
  name: glusterfs-cluster
subsets:
  - addresses:
    - ip: 192.168.122.221
    ports:
    - port: 1
  - addresses:
    - ip: 192.168.122.222
    ports:
    - port: 1
```

21.6.2.4. Deploy the PV, PVC, and Endpoints

For this example, run the **oc** commands as a **cluster-admin** privileged user. In a production environment, cluster clients might be expected to define and create the PVC.

```
# oc create -f glusterfs-ep.yaml
endpoints "glusterfs-cluster" created
# oc create -f glusterfs-pv.yaml
persistentvolume "gluster-volume" created
# oc create -f glusterfs-pvc.yaml
persistentvolumeclaim "gluster-claim" created
```

Lastly, confirm that the PV and PVC bound successfully.

```
# oc get pv,pvc
NAME                                CAPACITY  ACCESSMODES  STATUS  CLAIM
REASON    AGE
gluster-volume  2Gi      RWX          Bound   gfs-
trial/gluster-claim              7s
NAME                                STATUS  VOLUME          CAPACITY  ACCESSMODES
AGE
gluster-claim  Bound   gluster-volume  2Gi      RWX
7s
```



NOTE

PVCs are local to a project, whereas PVs are a cluster-wide, global resource. Developers and non-administrator users may not have access to see all (or any) of the available PVs.

CHAPTER 22. WORKING WITH HTTP PROXIES

22.1. OVERVIEW

Production environments can deny direct access to the Internet and instead have an HTTP or HTTPS proxy available. Configuring OpenShift Container Platform to use these proxies can be as simple as setting standard environment variables in configuration or JSON files. This can be done during an [advanced installation](#) or configured after installation.

The proxy configuration must be the same on each host in the cluster. Therefore, when setting up the proxy or modifying it, you must update the files on each OpenShift Container Platform host to the same values. Then, you must restart OpenShift Container Platform services on each host in the cluster.

The **NO_PROXY**, **HTTP_PROXY**, and **HTTPS_PROXY** environment variables are found in each host's */etc/sysconfig/atomic-openshift-master* file (for single master configuration), */etc/sysconfig/atomic-openshift-master-api*, or */etc/sysconfig/atomic-openshift-master-controllers* files (for multi-master configuration) and */etc/sysconfig/atomic-openshift-node*.

22.2. CONFIGURING NO_PROXY

The **NO_PROXY** environment variable lists all of the OpenShift Container Platform components and all IP addresses that are managed by OpenShift Container Platform.

NO_PROXY accepts a comma-separated list of hosts, IP addresses, or IP ranges in CIDR format:

For master hosts

- Node host name
- Master IP or host name

For node hosts

- Master IP or host name

For the Docker service

- Registry service IP and host name

NO_PROXY also includes the SDN network and service IP addresses as found in the *master-config.yaml* file.

/etc/origin/master/master-config.yaml

```
networkConfig:
  clusterNetworkCIDR: 10.1.0.0/16
  serviceNetworkCIDR: 172.30.0.0/16
```

OpenShift Container Platform does not accept ***** as a wildcard attached to a domain suffix. For example, this works:

```
NO_PROXY=.example.com
```

However, this does not:

```
NO_PROXY=*.example.com
```

The only wildcard **NO_PROXY** accepts is a single `*` character, which matches all hosts, and effectively disables the proxy.

Each name in this list is matched as either a domain which contains the host name as a suffix, or the host name itself.

For instance, **example.com** would match **example.com**, **example.com:80**, and **www.example.com**.

22.3. CONFIGURING HOSTS FOR PROXIES

1. Edit the proxy environment variables in the OpenShift Container Platform control files. Ensure all of the files in the cluster are correct.

```
HTTP_PROXY=http://<user>:<password>@<ip_addr>:<port>/
HTTPS_PROXY=https://<user>:<password>@<ip_addr>:<port>/
NO_PROXY=master.hostname.example.com,10.1.0.0/16,172.30.0.0/16 1
```

- 1 Supports host names and CIDRs. Must include the SDN network and service IP ranges **10.1.0.0/16**, **172.30.0.0/16** by default.

2. Restart the master or node host as appropriate:

```
# systemctl restart atomic-openshift-master
# systemctl restart atomic-openshift-node
```

For multi-master installations:

```
# systemctl restart atomic-openshift-master-controllers
# systemctl restart atomic-openshift-master-api
```

22.4. CONFIGURING HOSTS FOR PROXIES USING ANSIBLE

During [advanced installations](#), the **NO_PROXY**, **HTTP_PROXY**, and **HTTPS_PROXY** environment variables can be configured using the [openshift_no_proxy](#), [openshift_http_proxy](#), and [openshift_https_proxy parameters](#), which are configurable in the inventory file.

Example 22.1. Example Proxy Configuration with Ansible

```
# Global Proxy Configuration
# These options configure HTTP_PROXY, HTTPS_PROXY, and NO_PROXY
environment
# variables for docker and master services.
openshift_http_proxy=http://<user>:<password>@<ip_addr>:<port>
openshift_https_proxy=https://<user>:<password>@<ip_addr>:<port>
openshift_no_proxy='.hosts.example.com,some-host.com'
#
# Most environments do not require a proxy between OpenShift masters,
nodes, and
```

```
# etcd hosts. So automatically add those host names to the
openshift_no_proxy list.
# If all of your hosts share a common domain you may wish to disable
this and
# specify that domain above.
# openshift_generate_no_proxy_hosts=True
```

NOTE

There are [additional proxy settings](#) that can be [configured for builds](#) using Ansible parameters. For example:

The `openshift_builddefaults_git_http_proxy` and `openshift_builddefaults_git_https_proxy` parameters allow you to [use a proxy for Git cloning](#)

The `openshift_builddefaults_http_proxy` and `openshift_builddefaults_https_proxy` parameters can make environment variables available to the [Docker build strategy](#) and [Custom build strategy](#) processes.

22.5. PROXYING DOCKER PULL

OpenShift Container Platform node hosts need to perform push and pull operations to Docker registries. If you have a registry that does not need a proxy for nodes to access, include the **NO_PROXY** parameter with:

- the registry's host name,
- the registry service's IP address, and
- the service name.

This blacklists that registry, leaving the external HTTP proxy as the only option.

1. Retrieve the registry service's IP address `docker_registry_ip` by running:

```
$ oc describe svc/docker-registry -n default

Name:      docker-registry
Namespace: default
Labels:    docker-registry=default
Selector:  docker-registry=default
Type:      ClusterIP
IP:        172.30.163.183 1
Port:      5000-tcp 5000/TCP
Endpoints: 10.1.0.40:5000
Session Affinity: ClientIP
No events.
```

- 1** Registry service IP.

2. Edit the `/etc/sysconfig/docker` file and add the **NO_PROXY** variables in shell format, replacing `<docker_registry_ip>` with the IP address from the previous step.

```
HTTP_PROXY=http://<user>:<password>@<ip_addr>:<port>/
HTTPS_PROXY=https://<user>:<password>@<ip_addr>:<port>/
NO_PROXY=master.hostname.example.com,<docker_registry_ip>,docker-
registry.default.svc.cluster.local
```

3. Restart the Docker service:

```
# systemctl restart docker
```

22.6. USING MAVEN BEHIND A PROXY

Using Maven with proxies requires using the **HTTP_PROXY_NONPROXYHOSTS** variable.

See the [Red Hat JBoss Enterprise Application Platform for OpenShift documentation](#) for information about configuring your OpenShift Container Platform environment for Red Hat JBoss Enterprise Application Platform, including the step for setting up Maven behind a proxy.

22.7. CONFIGURING S2I BUILDS FOR PROXIES

S2I builds fetch dependencies from various locations. You can [use a `.s2i/environment` file](#) to specify simple shell variables and OpenShift Container Platform will react accordingly when seeing build images.

The following are the supported proxy environment variables with example values:

```
HTTP_PROXY=http://USERNAME:PASSWORD@10.0.1.1:8080/
HTTPS_PROXY=https://USERNAME:PASSWORD@10.0.0.1:8080/
NO_PROXY=master.hostname.example.com
```

22.8. CONFIGURING DEFAULT TEMPLATES FOR PROXIES

The [example templates](#) available in OpenShift Container Platform by default do not include settings for HTTP proxies. For existing applications based on these templates, modify the **source** section of the application's build configuration and add proxy settings:

```
...
source:
  type: Git
  git:
    uri: https://github.com/openshift/ruby-hello-world
    httpProxy: http://proxy.example.com
    httpsProxy: https://proxy.example.com
...
```

This is similar to the process for [using proxies for Git cloning](#).

22.9. SETTING PROXY ENVIRONMENT VARIABLES IN PODS

You can set the **NO_PROXY**, **HTTP_PROXY**, and **HTTPS_PROXY** environment variables in the **templates.spec.containers** stanza in a deployment configuration to pass proxy connection

information. The same can be done for configuring a Pod's proxy at runtime:

```
...
containers:
- env:
  - name: "HTTP_PROXY"
    value: "http://<user>:<password>@<ip_addr>:<port>"
...
```


You can also use the **oc set env** command to update an existing deployment configuration with a new environment variable:

```
$ oc set env dc/frontend HTTP_PROXY=http://<user>:<password>@<ip_addr>:
<port>
```

If you have a [ConfigChange trigger](#) set up in your OpenShift Container Platform instance, the changes happen automatically. Otherwise, manually redeploy your application for the changes to take effect.

22.10. GIT REPOSITORY ACCESS

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:
  type: Git
  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](#).

CHAPTER 23. CONFIGURING GLOBAL BUILD DEFAULTS AND OVERRIDES

23.1. OVERVIEW

Developers can define settings in specific build configurations within their projects, such as [configuring a proxy for Git cloning](#). Rather than requiring developers to define certain settings in each of their build configurations, cluster administrators can use admission control plug-ins to configure global build defaults and overrides that automatically use these settings in any build.

The settings from these plug-ins are not set in the build configurations or builds themselves, but rather are only used during the build process. This allows administrators to change the global configuration at any time, and any builds that are re-run from existing build configurations or builds will get the new settings.

The **BuildDefaults** admission control plug-in allows administrators to set global defaults for settings such as the Git HTTP and HTTPS proxy, as well as default environment variables. These defaults do not overwrite values that have been configured for a specific build. However, if those values are not present on the build definition, they are set to the default value.

The **BuildOverrides** admission control plug-in allows administrators to override a setting in a build, regardless of the value stored in the build. It currently supports overriding the **forcePull** flag on a build strategy to enforce always refreshing the local image during a build by pulling the image from the registry. This ensures that a user can only build with an image that they are allowed to pull.

23.2. SETTING GLOBAL BUILD DEFAULTS

You can set global build defaults two ways:

- [using Ansible and the advanced installation tool](#)
- [manually by modifying the *master-config.yaml* file](#)

23.2.1. Configuring Global Build Defaults with Ansible

During [advanced installations](#), the **BuildDefaults** plug-in can be configured using [the following parameters](#), which are configurable in the inventory file:

- **openshift_builddefaults_http_proxy**
- **openshift_builddefaults_https_proxy**
- **openshift_builddefaults_no_proxy**
- **openshift_builddefaults_git_http_proxy**
- **openshift_builddefaults_git_https_proxy**

Example 23.1. Example Build Defaults Configuration with Ansible

```
# These options configure the BuildDefaults admission controller which
injects
# environment variables into Builds. These values will default to the
```

```

global proxy
# config values. You only need to set these if they differ from the
# global settings
# above. See BuildDefaults
# documentation at
https://docs.openshift.org/latest/admin_guide/build_defaults_overrides.html
#openshift_builddefaults_http_proxy=http://USER:PASSWORD@HOST:PORT
openshift_builddefaults_https_proxy=https://USER:PASSWORD@HOST:PORT
openshift_builddefaults_no_proxy=build_defaults
openshift_builddefaults_git_http_proxy=http://USER:PASSWORD@HOST:PORT
openshift_builddefaults_git_https_proxy=https://USER:PASSWORD@HOST:PORT
# Or you may optionally define your own serialized as json
#openshift_builddefaults_json='{ "BuildDefaults": { "configuration":
{ "apiVersion": "v1", "env":
[ { "name": "HTTP_PROXY", "value": "http://proxy.example.com.redhat.com:3128"
}, { "name": "NO_PROXY", "value": "ose3-
master.example.com" } ], "gitHTTPProxy": "http://proxy.example.com:3128", "ki
nd": "BuildDefaultsConfig" } } }'

```

NOTE

There are [additional proxy settings](#) that can be [configured for builds](#) using Ansible parameters. For example: - The **openshift_builddefaults_git_http_proxy** and **openshift_builddefaults_git_https_proxy** parameters allow you to [use a proxy for git cloning](#) - The **openshift_builddefaults_http_proxy** and **openshift_builddefaults_https_proxy** parameters can make environment variables available to the [Docker build strategy](#) and [Custom build strategy](#) processes.

23.2.2. Manually Setting Global Build Defaults

To configure the **BuildDefaults** plug-in, add a configuration for it in the **/etc/origin/master/master-config.yaml** file on masters:

```

admissionConfig:
  pluginConfig:
    BuildDefaults:
      configuration:
        apiVersion: v1
        kind: BuildDefaultsConfig
        gitHTTPProxy: http://my.proxy:8080 ①
        gitHTTPSProxy: https://my.proxy:8443 ②
        gitNoProxy: somedomain.com, otherdomain.com ③
        env:
          - name: HTTP_PROXY ④
            value: http://my.proxy:8080
          - name: HTTPS_PROXY ⑤
            value: https://my.proxy:8443
          - name: BUILD_LOGLEVEL ⑥
            value: 4
          - name: CUSTOM_VAR ⑦
            value: custom_value

```

```

imageLabels:
- name: url 8
  value: https://containers.example.org
- name: vendor
  value: ExampleCorp Ltd.
nodeSelector: 9
  key1: value1
  key2: value2
annotations: 10
  key1: value1
  key2: value2
resources: 11
  requests:
    cpu: "100m"
    memory: "256Mi"
  limits:
    cpu: "100m"
    memory: "256Mi"

```

- 1** Sets the HTTP proxy to use when cloning source code from a Git repository.
- 2** Sets the HTTPS proxy to use when cloning source code from a Git repository.
- 3** Sets the list of domains for which proxying should not be performed.
- 4** Default environment variable that sets the HTTP proxy to use during the build. This can be used for downloading dependencies during the assemble and build phases.
- 5** Default environment variable that sets the HTTPS proxy to use during the build. This can be used for downloading dependencies during the assemble and build phases.
- 6** Default environment variable that sets the build log level during the build.
- 7** Additional default environment variable that will be added to every build.
- 8** Labels to be applied to every image built. These can be overridden in **BuildConfig**.
- 9** Build pods will only run on nodes with the **key1=value1** and **key2=value2** labels. Users can define a different set of **nodeSelectors** for their builds, causing these values to be ignored.
- 10** Build pods will have these annotations added to them.
- 11** Sets the default resources to the build pod if the **BuildConfig** does not have related resource defined.

Restart the master service for the changes to take effect:

```
# systemctl restart atomic-openshift-master
```

23.3. SETTING GLOBAL BUILD OVERRIDES

To configure the **BuildOverrides** plug-in, add a configuration for it in the **/etc/origin/master/master-config.yaml** file on masters:

■

```
kubernetesMasterConfig:
  admissionConfig:
    pluginConfig:
      BuildOverrides:
        configuration:
          apiVersion: v1
          kind: BuildOverridesConfig
          forcePull: true ❶
```

- ❶ Force all builds to pull their builder image and any source images before starting the build.

Restart the master service for the changes to take effect:

```
# systemctl restart atomic-openshift-master
```

CHAPTER 24. CONFIGURING PIPELINE EXECUTION

24.1. OVERVIEW

The first time a user creates a build configuration using the [Pipeline](#) build strategy, OpenShift Container Platform looks for the a template named **jenkins-ephemeral** in the **openshift** namespace and instantiates it within the user's project. The **jenkins-ephemeral** template that ships with OpenShift Container Platform creates, upon instantiation:

- a deployment configuration for Jenkins using the official OpenShift Container Platform Jenkins image
- a service and route for accessing the Jenkins deployment
- a new Jenkins service account
- rolebindings to grant the service account edit access to the project

Cluster administrators can control what is created by either modifying the content of the built-in template, or by editing the cluster configuration to direct the cluster to a different template location.

To modify the content of the default template:

```
$ oc edit template jenkins-ephemeral -n openshift
```

To use a different template, such as the **jenkins-persistent** template which uses persistent storage for Jenkins, add the following to your master configuration file:

```
jenkinsPipelineConfig:
  autoProvisionEnabled: true 1
  templateNamespace: openshift 2
  templateName: jenkins-pipeline 3
  serviceName: jenkins-pipeline-svc 4
  parameters: 5
    key1: value1
    key2: value2
```

- 1 Defaults to **false**. If **false**, then no template will be instantiated.
- 2 Namespace containing the template to be instantiated.
- 3 Name of the template to be instantiated.
- 4 Name of the service to be created by the template upon instantiation.
- 5 Optional values to pass to the template during instantiation.



IMPORTANT

When a Pipeline build configuration is created, OpenShift Container Platform instantiates the **jenkinsPipelineConfig** template **only if** no existing service name in the project matches the **serviceName** field. This means **serviceName** must be chosen such that it is unique in the project.

CHAPTER 25. CONFIGURING ROUTING

25.1. OVERVIEW

After installing OpenShift Container Platform and [deploying a router](#), you can modify the router to suit your needs using the examples in this topic.

25.2. CONFIGURING ROUTE TIMEOUTS

You can configure the default timeouts for an existing route when you have services in need of a low timeout, as required for Service Level Availability (SLA) purposes, or a high timeout, for cases with a slow back end.

Using the **oc annotate** command, add the timeout to the route:

```
# oc annotate route <route_name> \
  --overwrite haproxy.router.openshift.io/timeout=<timeout><time_unit>
```

For example, to set a route named **myroute** to a timeout of two seconds:

```
# oc annotate route myroute --overwrite
haproxy.router.openshift.io/timeout=2s
```

Supported time units are microseconds (us), milliseconds (ms), seconds (s), minutes (m), hours (h), or days (d).

25.3. CONFIGURING NATIVE CONTAINER ROUTING

This section describes how to set up container networking using existing switches and routers and the kernel networking stack in Linux. The setup requires that the network administrator or a script modifies the router or routers when new nodes are added to the cluster.



NOTE

The procedure outlined in this topic can be adapted to any type of router.


25.3.1. Network Overview

The following describes a general network setup:

- 11.11.0.0/16 is the container network.
- The 11.11.x.0/24 subnet is reserved for each node and assigned to the Docker Linux bridge.
- Each node has a route to the router for reaching anything in the 11.11.0.0/16 range, except the local subnet.
- The router has routes for each node, so it can be directed to the right node.
- Existing nodes do not need any changes when new nodes are added, unless the network topology is modified.

- IP forwarding is enabled on each node.

The following diagram shows the container networking setup described in this topic. It uses one Linux node with two network interface cards serving as a router, two switches, and three nodes connected to these switches.



Node setup

1. Assign an unused 11.11.x.0/24 subnet IP address to the Linux bridge on the node:

```
# brctl addbr lbr0
# ip addr add 11.11.1.1/24 dev lbr0
# ip link set dev lbr0 up
```

1. Modify the Docker startup script to use the new bridge. By default, the startup script is the **/etc/sysconfig/docker** file:

```
# docker -d -b lbr0 --other-options
```

2. Add a route to the router for the 11.11.0.0/16 network:

```
# ip route add 11.11.0.0/16 via 192.168.2.2 dev p3p1
```

3. Enable IP forwarding on the node:

```
# sysctl -w net.ipv4.ip_forward=1
```

Router setup

The following procedure assumes a Linux box with multiple NICs is used as a router. Modify the steps as required to use the syntax for a particular router:

1. Enable IP forwarding on the router:


```
# sysctl -w net.ipv4.ip_forward=1
```

2. Add a route for each node added to the cluster:

```
# ip route add <node_subnet> via <node_ip_address> dev <interface  
through which node is L2 accessible>  
# ip route add 11.11.1.0/24 via 192.168.2.1 dev p3p1  
# ip route add 11.11.2.0/24 via 192.168.3.3 dev p3p2  
# ip route add 11.11.3.0/24 via 192.168.3.4 dev p3p2
```

CHAPTER 26. ROUTING FROM EDGE LOAD BALANCERS

26.1. OVERVIEW

[Pods](#) inside of an OpenShift Container Platform cluster are only reachable via their IP addresses on the cluster network. An edge load balancer can be used to accept traffic from outside networks and proxy the traffic to pods inside the OpenShift Container Platform cluster. In cases where the load balancer is not part of the cluster network, routing becomes a hurdle as the internal cluster network is not accessible to the edge load balancer.

To solve this problem where the OpenShift Container Platform cluster is using [OpenShift SDN](#) as the cluster networking solution, there are two ways to achieve network access to the pods.

26.2. INCLUDING THE LOAD BALANCER IN THE SDN

If possible, run an OpenShift Container Platform node instance on the load balancer itself that uses OpenShift Container Platform SDN as the networking plug-in. This way, the edge machine gets its own Open vSwitch bridge that the SDN automatically configures to provide access to the pods and nodes that reside in the cluster. The *routing table* is dynamically configured by the SDN as pods are created and deleted, and thus the routing software is able to reach the pods.

Mark the load balancer machine as an [unschedulable node](#) so that no pods end up on the load balancer itself:

```
$ oadm manage-node <load_balancer_hostname> --schedulable=false
```

If the load balancer comes packaged as a container, then it is even easier to integrate with OpenShift Container Platform: Simply run the load balancer as a pod with the [host port exposed](#). The pre-packaged [HAProxy router](#) in OpenShift Container Platform runs in precisely this fashion.

26.3. ESTABLISHING A TUNNEL USING A RAMP NODE

In some cases, the previous solution is not possible. For example, an **F5 BIG-IP®** host cannot run an OpenShift Container Platform node instance or the OpenShift Container Platform SDN because **F5®** uses a custom, incompatible Linux kernel and distribution.

Instead, to enable **F5 BIG-IP®** to reach pods, you can choose an existing node within the cluster network as a *ramp node* and establish a tunnel between the **F5 BIG-IP®** host and the designated ramp node. Because it is otherwise an ordinary OpenShift Container Platform node, the ramp node has the necessary configuration to route traffic to any pod on any node in the cluster network. The ramp node thus assumes the role of a gateway through which the **F5 BIG-IP®** host has access to the entire cluster network.

Following is an example of establishing an **ipip** tunnel between an **F5 BIG-IP®** host and a designated ramp node.

On the F5 BIG-IP® host:

1. Set the following variables:

```
# F5_IP=10.3.89.66 1
# RAMP_IP=10.3.89.89 2
# TUNNEL_IP1=10.3.91.216 3
```

```
# CLUSTER_NETWORK=10.128.0.0/14 4
```

- 1 2 The **F5_IP** and **RAMP_IP** variables refer to the **F5 BIG-IP®** host's and the ramp node's IP addresses, respectively, on a shared, internal network.
- 3 An arbitrary, non-conflicting IP address for the **F5®** host's end of the **ipip** tunnel.
- 4 The overlay network CIDR that the OpenShift Container Platform SDN uses to assign addresses to pods.

2. Delete any old route, self, tunnel and SNAT pool:

```
# tmsh delete net route $CLUSTER_NETWORK || true
# tmsh delete net self SDN || true
# tmsh delete net tunnels tunnel SDN || true
# tmsh delete ltm snatpool SDN_snatpool || true
```

3. Create the new tunnel, self, route and SNAT pool and use the SNAT pool in the virtual servers:

```
# tmsh create net tunnels tunnel SDN \
  \{ description "OpenShift SDN" local-address \
    $F5_IP profile ipip remote-address $RAMP_IP \}
# tmsh create net self SDN \{ address \
  ${TUNNEL_IP1}/24 allow-service all vlan SDN \}
# tmsh create net route $CLUSTER_NETWORK interface SDN
# tmsh create ltm snatpool SDN_snatpool members add { $TUNNEL_IP1 }
# tmsh modify ltm virtual ose-vserver source-address-translation {
  type snat pool SDN_snatpool }
# tmsh modify ltm virtual https-ose-vserver source-address-
translation { type snat pool SDN_snatpool }
```

On the ramp node:

1. Set the following variables:

```
# F5_IP=10.3.89.66
# TUNNEL_IP1=10.3.91.216
# TUNNEL_IP2=10.3.91.217 1
# CLUSTER_NETWORK=10.128.0.0/1 2
```

- 1 A second, arbitrary IP address for the ramp node's end of the **ipip** tunnel.
- 2 The overlay network CIDR that the OpenShift Container Platform SDN uses to assign addresses to pods.

2. Delete any old tunnel:

```
# ip tunnel del tun1 || true
```

3. Create the **ipip** tunnel on the ramp node, using a suitable L2-connected interface (e.g., **eth0**):

```
# ip tunnel add tun1 mode ipip \
  remote $F5_IP dev eth0
```

```
# ip addr add $TUNNEL_IP2 dev tun1
# ip link set tun1 up
# ip route add $TUNNEL_IP1 dev tun1
# ping -c 5 $TUNNEL_IP1
```

4. SNAT the tunnel IP with an unused IP from the SDN subnet:

```
# source /run/openshift-sdn/config.env
# tap1=$(ip -o -4 addr list tun0 | awk '{print $4}' | cut -d/ -f1 |
head -n 1)
# subaddr=$(echo ${OPENSIFT_SDN_TAP1_ADDR:-"$tap1"} | cut -d "." -f
1,2,3)
# export RAMP_SDN_IP=${subaddr}.254
```

5. Assign this **RAMP_SDN_IP** as an additional address to **tun0** (the local SDN's gateway):

```
# ip addr add ${RAMP_SDN_IP} dev tun0
```

6. Modify the OVS rules for SNAT:


```
# ipflowopts="cookie=0x999,ip"
# arpflowopts="cookie=0x999, table=0, arp"
#
# ovs-ofctl -O OpenFlow13 add-flow br0 \

"${ipflowopts},nw_src=${TUNNEL_IP1},actions=mod_nw_src:${RAMP_SDN_IP
},resubmit(,0)"
# ovs-ofctl -O OpenFlow13 add-flow br0 \

"${ipflowopts},nw_dst=${RAMP_SDN_IP},actions=mod_nw_dst:${TUNNEL_IP1
},resubmit(,0)"
# ovs-ofctl -O OpenFlow13 add-flow br0 \
    "${arpflowopts}, arp_tpa=${RAMP_SDN_IP}, actions=output:2"
# ovs-ofctl -O OpenFlow13 add-flow br0 \
    "${arpflowopts}, priority=200, in_port=2,
arp_spa=${RAMP_SDN_IP}, arp_tpa=${CLUSTER_NETWORK},
actions=goto_table:5"
# ovs-ofctl -O OpenFlow13 add-flow br0 \
    "arp, table=5, priority=300, arp_tpa=${RAMP_SDN_IP},
actions=output:2"
# ovs-ofctl -O OpenFlow13 add-flow br0 \
    "ip,table=5,priority=300,nw_dst=${RAMP_SDN_IP},actions=output:2"
# ovs-ofctl -O OpenFlow13 add-flow br0
"${ipflowopts},nw_dst=${TUNNEL_IP1},actions=output:2"
```

7. Optionally, if you do not plan on configuring the ramp node to be highly available, mark the ramp node as unschedulable. Skip this step if you do plan to follow the next section and plan on creating a highly available ramp node.

```
$ oadm manage-node <ramp_node_hostname> --schedulable=false
```



NOTE

The [F5 router plug-in](#) integrates with F5 BIG-IP®.

26.3.1. Configuring a Highly-Available Ramp Node

You can use OpenShift Container Platform's **ipfailover** feature, which uses **keepalived** internally, to make the ramp node highly available from **F5 BIG-IP®**'s point of view. To do so, first bring up two nodes, for example called **ramp-node-1** and **ramp-node-2**, on the same L2 subnet.

Then, choose some unassigned IP address from within the same subnet to use for your virtual IP, or *VIP*. This will be set as the **RAMP_IP** variable with which you will configure your tunnel on **F5 BIG-IP®**.

For example, suppose you are using the **10.20.30.0/24** subnet for your ramp nodes, and you have assigned **10.20.30.2** to **ramp-node-1** and **10.20.30.3** to **ramp-node-2**. For your VIP, choose some unassigned address from the same **10.20.30.0/24** subnet, for example **10.20.30.4**. Then, to configure **ipfailover**, mark both nodes with a label, such as **f5rampnode**:

```
$ oc label node ramp-node-1 f5rampnode=true
$ oc label node ramp-node-2 f5rampnode=true
```

Similar to instructions from the [ipfailover documentation](#), you must now create a service account and add it to the **privileged** SCC. First, create the **f5ipfailover** service account:

```
$ oc create serviceaccount f5ipfailover -n default
```

Next, you can add the **f5ipfailover** service to the **privileged** SCC. To add the **f5ipfailover** in the **default** namespace to the **privileged** SCC, run:

```
$ oadm policy add-scc-to-user privileged
system:serviceaccount:default:f5ipfailover
```

Finally, configure **ipfailover** using your chosen VIP (the **RAMP_IP** variable) and the **f5ipfailover** service account, assigning the VIP to your two nodes using the **f5rampnode** label you set earlier:

```
# RAMP_IP=10.20.30.4
# IFNAME=eth0 ❶
# oadm ipfailover <name-tag> \
  --virtual-ips=$RAMP_IP \
  --interface=$IFNAME \
  --watch-port=0 \
  --replicas=2 \
  --service-account=f5ipfailover \
  --selector='f5rampnode=true'
```

❶ The interface where **RAMP_IP** should be configured.

With the above setup, the VIP (the **RAMP_IP** variable) is automatically re-assigned when the ramp node host that currently has it assigned fails.

CHAPTER 27. AGGREGATING CONTAINER LOGS

27.1. OVERVIEW

As an OpenShift Container Platform cluster administrator, you can deploy the EFK stack to aggregate logs for a range of OpenShift Container Platform services. Application developers can view the logs of the projects for which they have view access. The EFK stack aggregates logs from hosts and applications, whether coming from multiple containers or even deleted pods.

The EFK stack is a modified version of the [ELK stack](#) and is comprised of:

- [Elasticsearch](#): An object store where all logs are stored.
- [Fluentd](#): Gathers logs from nodes and feeds them to Elasticsearch.
- [Kibana](#): A web UI for Elasticsearch.

Once deployed in a cluster, the stack aggregates logs from all nodes and projects into Elasticsearch, and provides a Kibana UI to view any logs. Cluster administrators can view all logs, but application developers can only view logs for projects they have permission to view. The stack components communicate securely.



NOTE

[Managing Docker Container Logs](#) discusses the use of **json-file** logging driver options to manage container logs and prevent filling node disks.


27.2. PRE-DEPLOYMENT CONFIGURATION

1. Ensure that you have deployed a router for the cluster.
2. Ensure that you have [the necessary storage](#) for Elasticsearch. Note that each Elasticsearch replica requires its own storage volume. See [Elasticsearch](#) for more information.
3. Ansible-based installs should create the **logging-deployer-template** template in the **openshift** project. Otherwise you can create it with the following command:

```
$ oc apply -n openshift -f \
  /usr/share/ansible/openshift-
  ansible/roles/openshift_hosted_templates/files/v1.3/enterprise/loggi
  ng-deployer.yaml
```

4. Create a new project. Once implemented in a single project, the EFK stack collects logs for every project within your OpenShift Container Platform cluster. The examples in this topic use **logging** as an example project:


```
$ oadm new-project logging --node-selector=""
$ oc project logging
```

**NOTE**

Specifying an empty [node selector](#) on the project is recommended, as Fluentd should be deployed throughout the cluster and any selector would restrict where it is deployed. To control component placement, specify node selectors per component to be applied to their deployment configurations.

5. Create the logging [service accounts](#) and custom roles:

```
$ oc new-app logging-deployer-account-template
```

**NOTE**

If you deployed logging previously, for example in a different project, then it is normal for the cluster roles to fail to be created because they already exist.

6. Enable the deployer service account to create an OAuthClient (normally a cluster administrator privilege) for Kibana to use later when authenticating against the master.

```
$ oadm policy add-cluster-role-to-user oauth-editor \
    system:serviceaccount:logging:logging-deployer ❶
```

- ❶ Use the project you created earlier (for example, **logging**) when specifying this service account.

7. Enable the Fluentd service account to mount and read system logs by adding it to the **privileged** security context, and also enable it to read pod metadata by giving it the **cluster-reader** role:

```
$ oadm policy add-scc-to-user privileged \
    system:serviceaccount:logging:aggregated-logging-fluentd ❶
$ oadm policy add-cluster-role-to-user cluster-reader \
    system:serviceaccount:logging:aggregated-logging-fluentd ❷
```

- ❶ ❷ Use the project you created earlier (for example, **logging**) when specifying this service account.

8. Ensure that port 9300 is open. By default the Elasticsearch service uses port 9300 for TCP communication between nodes in a cluster.

27.3. SPECIFYING DEPLOYER PARAMETERS

Parameters for the EFK deployment may be specified in the form of a [ConfigMap](#), a [secret](#), or template parameters (which are passed to the deployer in environment variables). The deployer looks for each value first in a **logging-deployer** ConfigMap, then a **logging-deployer** secret, then as an environment variable. Any or all may be omitted if not needed. If you are specifying values within a ConfigMap, the values will not be reflected in the output from **oc new-app**, but will still precede the corresponding template values within the deployer pod.

The available parameters are outlined below. Typically, you should at least specify the host name at which Kibana should be exposed to client browsers, and also the master URL where client browsers will be directed to for authenticating to OpenShift Container Platform.

1. Create a [ConfigMap](#) to provide most deployer parameters. An invocation supplying the most important parameters might be:

```
$ oc create configmap logging-deployer \
  --from-literal kibana-hostname=kibana.example.com \
  --from-literal public-master-url=https://master.example.com:8443 \
  --from-literal es-cluster-size=3 \
  --from-literal es-instance-ram=8G
```

2. Edit the ConfigMap YAML file after creating it:

```
$ oc edit configmap logging-deployer
```

Other parameters are available. Read the [ElasticSearch section](#) before choosing ElasticSearch parameters for the deployer, and the [Fluentd section](#) for some possible parameters:

Parameter	Description
<i>kibana-hostname</i>	The external host name for web clients to reach Kibana.
<i>public-master-url</i>	The external URL for the master; used for OAuth purposes.
<i>es-cluster-size</i> (default: 1)	The number of instances of Elasticsearch to deploy. Redundancy requires at least three, and more can be used for scaling.
<i>es-instance-ram</i> (default: 8G)	Amount of RAM to reserve per Elasticsearch instance. The default is 8G (for 8GB), and it must be at least 512M. Possible suffixes are G,g,M,m.
<i>es-pvc-prefix</i> (default: logging-es-)	Prefix for the names of persistent volume claims to be used as storage for Elasticsearch instances; a number will be appended per instance (for example, logging-es-1). If they do not already exist, they will be created with size <i>es-pvc-size</i> .
<i>es-pvc-size</i>	Size of the persistent volume claim to create per Elasticsearch instance, 100G, for example. If omitted, no PVCs are created and ephemeral volumes are used instead.
<i>es-pvc-dynamic</i>	Set to true to have created persistent volume claims annotated so that their backing storage can be dynamically provisioned (if that is available for your cluster).
<i>storage-group</i>	Number of a supplemental group ID for access to Elasticsearch storage volumes; backing volumes should allow access by this group ID (defaults to 65534).

Parameter	Description
<i>fluentd-nodeselector</i> (default: logging-infra-fluentd=true)	A node selector that specifies which nodes are eligible targets for deploying Fluentd instances. All nodes where Fluentd should run (typically, all) must have this label before Fluentd will be able to run and collect logs.
<i>es-nodeselector</i>	A node selector that specifies which nodes are eligible targets for deploying Elasticsearch instances. This can be used to place these instances on nodes reserved and/or optimized for running them. For example, the selector could be node-type=infrastructure . At least one active node must have this label before Elasticsearch will deploy.
<i>kibana-nodeselector</i>	A node selector that specifies which nodes are eligible targets for deploying Kibana instances.
<i>curator-nodeselector</i>	A node selector that specifies which nodes are eligible targets for deploying Curator instances.
<i>enable-ops-cluster</i>	If set to true , configures a second Elasticsearch cluster and Kibana for operations logs. Fluentd splits logs between the main cluster and a cluster reserved for operations logs (which consists of /var/log/messages on nodes and the logs from the projects default , openshift , and openshift-infra). This means a second Elasticsearch and Kibana are deployed. The deployments are distinguishable by the -ops included in their names and have parallel deployment options listed below.
<i>kibana-ops-hostname, es-ops-instance-ram, es-ops-pvc-size, es-ops-pvc-prefix, es-ops-cluster-size, es-ops-nodeselector, kibana-ops-nodeselector, curator-ops-nodeselector</i>	Parallel parameters for the ops log cluster.
<i>image-pull-secret</i>	Specify the name of an existing pull secret to be used for pulling component images from an authenticated registry.

3. Create a [secret](#) to provide security-related files to the deployer. Providing the secret is optional, and the objects will be randomly generated if not supplied.

You can supply the following files when creating a new secret, for example:

```
$ oc create secret generic logging-deployer \
  --from-file kibana.crt=/path/to/cert \
  --from-file kibana.key=/path/to/key
```

File Name	Description
<i>kibana.crt</i>	A browser-facing certificate for the Kibana server.
<i>kibana.key</i>	A key to be used with the Kibana certificate.
<i>kibana-ops.crt</i>	A browser-facing certificate for the Ops Kibana server.
<i>kibana-ops.key</i>	A key to be used with the Ops Kibana certificate.
<i>server-tls.json</i>	JSON TLS options to override the Kibana server defaults. Refer to Node.JS docs for available options.
<i>ca.crt</i>	A certificate for a CA that will be used to sign all certificates generated by the deployer.
<i>ca.key</i>	A matching CA key.

27.4. DEPLOYING THE EFK STACK

The EFK stack is deployed using a [template](#) to create a deployer pod that reads the deployment parameters and manages the deployment.

Run the deployer, optionally specifying parameters (described in the table below), for example:

Without template parameters:

```
$ oc new-app logging-deployer-template
```

With parameters:

```
$ oc new-app logging-deployer-template \
  --param IMAGE_VERSION=<tag> \
  --param MODE=install
```

Replace **<tag>** with **v3.3** for the latest version.

Parameter Name	Description
<i>IMAGE_PREFIX</i>	The prefix for logging component images. For example, setting the prefix to registry.access.redhat.com/openshift3/ creates registry.access.redhat.com/openshift3/logging-deployer:latest .
<i>IMAGE_VERSION</i>	The version for logging component images. For example, setting the version to v3.3 creates registry.access.redhat.com/openshift3/logging-deployer:v3.3 .

Parameter Name	Description
MODE (default: install)	Mode to run the deployer in; one of install , uninstall , reinstall , upgrade , migrate , start , stop .

Running the deployer creates a deployer pod and prints its name. Wait until the pod is running. This can take up to a few minutes for OpenShift Container Platform to retrieve the deployer image from the registry. Watch its process with:

```
$ oc get pod/<pod_name> -w
```

It will eventually enter **Running** status and end in **Complete** status. If it takes too long to start, retrieve more details about the pod and any associated events with:

```
$ oc describe pod/<pod_name>
```

Check the logs if the deployment does not complete successfully:

```
$ oc logs -f <pod_name>
```

Once deployment completes successfully, you may need to [label the nodes for Fluentd to deploy on](#), and may have other adjustments to make to the deployed components. These tasks are described in the next section.

27.5. UNDERSTANDING AND ADJUSTING THE DEPLOYMENT

This section describes adjustments that you can make to deployed components.

27.5.1. Ops Cluster



NOTE

The logs for the **default**, **openshift**, and **openshift-infra** projects are automatically aggregated and grouped into the **.operations** item in the Kibana interface.

The project where you have deployed the EFK stack (**logging**, as documented here) is *not* aggregated into **.operations** and is found under its ID.

If you set **enable-ops-cluster** to **true** for the deployer, Fluentd is configured to split logs between the main Elasticsearch cluster and another cluster reserved for operations logs (which are defined as node system logs and the projects **default**, **openshift**, and **openshift-infra**). Therefore, a separate Elasticsearch cluster, a separate Kibana, and a separate Curator are deployed to index, access, and manage operations logs. These deployments are set apart with names that include **-ops**. Keep these separate deployments in mind if you enabled this option. Most of the following discussion also applies to the operations cluster if present, just with the names changed to include **-ops**.

27.5.2. Elasticsearch

A highly-available environment requires at least three replicas of Elasticsearch; each on a different host. Elasticsearch replicas require their own storage, but an OpenShift Container Platform deployment

configuration shares storage volumes between all its pods. So, when scaled up, the EFK deployer ensures each replica of Elasticsearch has its own deployment configuration.

It is possible to scale your cluster up after creation by adding more deployments from a template; however, scaling up (or down) requires the correct procedure and an awareness of clustering parameters (to be described in a separate section). It is best to indicate the desired scale at first deployment.

Refer to [Elastic's documentation](#) for considerations involved in choosing storage and network location as directed below.

Viewing all Elasticsearch Deployments

To view all current Elasticsearch deployments:

```
$ oc get dc --selector logging-infra=elasticsearch
```

Node Selector

Because Elasticsearch can use a lot of resources, all members of a cluster should have low latency network connections to each other and to any remote storage. Ensure this by directing the instances to dedicated nodes, or a dedicated region within your cluster, using a [node selector](#).

To configure a node selector, specify the **es-nodeselector** configuration option at deployment. This applies to all Elasticsearch deployments; if you need to individualize the node selectors, you must manually edit each deployment configuration after deployment.

Persistent Elasticsearch Storage

By default, the deployer creates an ephemeral deployment in which all of a pod's data is lost upon restart. For production usage, specify a persistent storage volume for each Elasticsearch deployment configuration. You can create the necessary [persistent volume claims](#) before deploying or have them created for you. The PVCs must be named based on the **es-pvc-prefix** setting, which defaults to **logging-es-**; each PVC name will have a sequence number added to it, so **logging-es-1**, **logging-es-2**, and so on. If a PVC needed for the deployment exists already, it is used; if not, and **es-pvc-size** has been specified, it is created with a request for that size.



WARNING

Using NFS storage as a volume or a persistent volume (or via NAS such as Gluster) is not supported for Elasticsearch storage, as Lucene relies on file system behavior that NFS does not supply. Data corruption and other problems can occur. If NFS storage is a requirement, you can allocate a large file on a volume to serve as a storage device and mount it locally on one host. For example, if your NFS storage volume is mounted at **/nfs/storage**:

```
$ truncate -s 1T /nfs/storage/elasticsearch-1
$ mkfs.xfs /nfs/storage/elasticsearch-1
$ mount -o loop /nfs/storage/elasticsearch-1 /usr/local/es-
storage
$ chown 1000:1000 /usr/local/es-storage
```

Then, use **/usr/local/es-storage** as a host-mount as described below. Use a different backing file as storage for each Elasticsearch replica.

This loopback must be maintained manually outside of OpenShift Container Platform, on the node. You must not maintain it from inside a container.

It is possible to use a local disk volume (if available) on each node host as storage for an Elasticsearch replica. Doing so requires some preparation as follows.

1. The relevant service account must be given the privilege to mount and edit a local volume:

```
$ oadm policy add-scc-to-user privileged \
    system:serviceaccount:logging:aggregated-logging-
    elasticsearch 1
```

- 1** Use the project you created earlier (for example, **logging**) when specifying this service account.

2. Each Elasticsearch replica definition must be patched to claim that privilege, for example:

```
$ for dc in $(oc get deploymentconfig --selector logging-
infra=elasticsearch -o name); do
    oc scale $dc --replicas=0
    oc patch $dc \
        -p '{"spec":{"template":{"spec":{"containers":
[{"name":"elasticsearch","securityContext":{"privileged":
true}}}}}}'
done
```

3. The Elasticsearch replicas must be located on the correct nodes to use the local storage, and should not move around even if those nodes are taken down for a period of time. This requires giving each Elasticsearch replica a node selector that is unique to a node where an administrator

has allocated storage for it. To configure a node selector, edit each Elasticsearch deployment configuration and add or edit the **nodeSelector** section to specify a unique label that you have applied for each desired node:

```
apiVersion: v1
kind: DeploymentConfig
spec:
  template:
    spec:
      nodeSelector:
        logging-es-node: "1" 1
```

- 1 This label should uniquely identify a replica with a single node that bears that label, in this case **logging-es-node=1**. Use the **oc label** command to apply labels to nodes as needed.

To automate applying the node selector you can instead use the **oc patch** command:

```
$ oc patch dc/logging-es-<suffix> \
  -p '{"spec":{"template":{"spec":{"nodeSelector":{"logging-es-
node":"1"}}}}}'
```

4. Once these steps are taken, a local host mount can be applied to each replica as in this example (where we assume storage is mounted at the same path on each node):

```
$ for dc in $(oc get deploymentconfig --selector logging-
infra=elasticsearch -o name); do
  oc set volume $dc \
    --add --overwrite --name=elasticsearch-storage \
    --type=hostPath --path=/usr/local/es-storage
  oc deploy --latest $dc
  oc scale $dc --replicas=1
done
```

Changing the Scale of Elasticsearch

If you need to scale up the number of Elasticsearch instances your cluster uses, it is not as simple as scaling up an Elasticsearch deployment configuration. This is due to the nature of persistent volumes and how Elasticsearch is configured to store its data and recover the cluster. Instead, scaling up requires creating a deployment configuration for each Elasticsearch cluster node.

By far the simplest way to change the scale of Elasticsearch is to reinstall the whole deployment. Assuming you have supplied persistent storage for the deployment, this should not be very disruptive. Simply re-run the deployer with the updated **es-cluster-size** configuration value and the **MODE=reinstall** template parameter. For example:

```
$ oc edit configmap logging-deployer
[change es-cluster-size value to 5]
$ oc new-app logging-deployer-template --param MODE=reinstall
```

If you previously deployed using template parameters rather than a ConfigMap, this would be a good time to create a ConfigMap instead for future deployer execution.

If you do not wish to reinstall, for instance because you have made customizations that you would like to preserve, then it is possible to add new Elasticsearch deployment configurations to the cluster using a template supplied by the deployer. This requires a more complicated procedure however.

During installation, the deployer [creates templates](#) with the Elasticsearch configurations provided to it: **logging-es-template** (and **logging-es-ops-template** if the deployer was run with **ENABLE_OPS_CLUSTER=true**). You can use these for scaling, but you need to adjust the size-related parameters in the templates:

Parameter	Description
NODE_QUORUM	The quorum required to elect a new master. Should be more than half the intended cluster size.
RECOVER_AFTER_NODES	When restarting the cluster, require this many nodes to be present before starting recovery. Defaults to one less than the cluster size to allow for one missing node.
RECOVER_EXPECTED_NODES	When restarting the cluster, wait for this number of nodes to be present before starting recovery. By default, the same as the cluster size.

The node quorum and recovery settings in the template were set based on the **es-[ops-]cluster-size** value initially provided to the deployer. Since the cluster size is changing, those values need to be overridden.

1. The existing deployment configurations for that cluster also need to have the three environment variable values above updated. To edit each of the configurations for the cluster in series, you may use the following command:

```
$ oc edit $(oc get dc -l component=es[-ops] -o name)
```

Edit the environment variables supplied so that the next time they restart, they will begin with the correct values. For example, for a cluster of size 5, you would set **NODE_QUORUM** to **3**, **RECOVER_AFTER_NODES** to **4**, and **RECOVER_EXPECTED_NODES** to **5**.

2. Create additional deployment configurations by running the following command against the Elasticsearch cluster you want to scale up for (**logging-es-template** or **logging-es-ops-template**), overriding the parameters as above.

```
$ oc new-app logging-es[-ops]-template \
  --param NODE_QUORUM=3 \
  --param RECOVER_AFTER_NODES=4 \
  --param RECOVER_EXPECTED_NODES=5
```

These deployments will be named differently, but all will have the **logging-es** prefix.

3. Each new deployment configuration is created without a persistent volume. If you want to attach a persistent volume to it, after creation you can use the **oc set volume** command to do so, for example:

```
$ oc volume dc/logging-es-<suffix> \
  --add --overwrite --name=elasticsearch-storage \
  --type=persistentVolumeClaim --claim-name=<your_pvc>
```

4. After the intended number of deployment configurations are created, scale up each new one to deploy it:

```
$ oc scale --replicas=1 dc/logging-es-<suffix>
```

27.5.3. Fluentd

Fluentd is deployed as a DaemonSet that deploys replicas according to a node label selector (which you can specify with the deployer parameter **fluentd-nodeselector**; the default is **logging-infra-fluentd**).

Once you have Elasticsearch running as desired, label the nodes intended for Fluentd deployment to feed their logs into ES. The example below would label a node named **node.example.com** using the default Fluentd node selector:

```
$ oc label node/node.example.com logging-infra-fluentd=true
```

Alternatively, you can label all nodes with:

```
$ oc label node --all logging-infra-fluentd=true
```



NOTE

Labeling nodes requires cluster administrator capability.

Having Fluentd Use the Systemd Journal as the Log Source

By default, Fluentd reads from **/var/log/messages** and **/var/log/containers/<container>.log** for system logs and container logs, respectively. You can instead use the systemd journal as the log source. There are three deployer configuration parameters available in the deployer ConfigMap:

Parameter	Description
use-journal	The default is empty, which tells the deployer to have Fluentd check which log driver Docker is using. If Docker is using --log-driver=journald , Fluentd reads from the systemd journal, otherwise, it assumes docker is using the json-file log driver and reads from the /var/log file sources. You can specify the use-journal option as true or false to be explicit about which log source to use. Using the systemd journal requires docker-1.10 or later, and Docker must be configured to use --log-driver=journald .
journal-source	The default is empty, so that when using the systemd journal, Fluentd first looks for /var/log/journal , and if that is not available, uses /run/log/journal as the journal source. You can specify journal-source with an explicit journal path. For example, if you want Fluentd to always read logs from the transient in-memory journal, set journal-source=/run/log/journal .

Parameter	Description
journal-read-from-head	If this setting is false , Fluentd starts reading from the end of the journal, ignoring historical logs. If this setting is true , Fluentd starts reading logs from the beginning of the journal.

**NOTE**

As of OpenShift Container Platform 3.3, Fluentd no longer reads historical log files when using the JSON file log driver. In situations where clusters have a large number of log files and are older than the EFK deployment, this avoids delays when pushing the most recent logs into Elasticsearch. Curator deleting logs are migrated soon after they are added to Elasticsearch.

**NOTE**

It may require several minutes, or hours, depending on the size of your journal, before any new log entries are available in Elasticsearch, when using **journal-read-from-head=true**.

Having Fluentd Send Logs to Another Elasticsearch

**NOTE**

The use of **ES_COPY** is being deprecated. To configure FluentD to send a copy of its logs to an external aggregator, use [Fluentd Secure Forward](#) instead.

You can configure Fluentd to send a copy of each log message to both the Elasticsearch instance included with OpenShift Container Platform aggregated logging, *and* to an external Elasticsearch instance. For example, if you already have an Elasticsearch instance set up for auditing purposes, or data warehousing, you can send a copy of each log message to that Elasticsearch.

This feature is controlled via environment variables on Fluentd, which can be modified as described below.

If its environment variable **ES_COPY** is **true**, Fluentd sends a copy of the logs to another Elasticsearch. The names for the copy variables are just like the current **ES_HOST**, **OPS_HOST**, and other variables, except that they add **_COPY**: **ES_COPY_HOST**, **OPS_COPY_HOST**, and so on. There are some additional parameters added:

- **ES_COPY_SCHEME**, **OPS_COPY_SCHEME** - can use either **http** or **https** - defaults to **https**
- **ES_COPY_USERNAME**, **OPS_COPY_USERNAME** - user name to use to authenticate to Elasticsearch using username/password auth
- **ES_COPY_PASSWORD**, **OPS_COPY_PASSWORD** - password to use to authenticate to Elasticsearch using username/password auth

**NOTE**

Sending logs directly to an AWS Elasticsearch instance is not supported. Use [Fluentd Secure Forward](#) to direct logs to an instance of Fluentd that you control and that is configured with the **fluent-plugin-aws-elasticsearch-service** plug-in.

To set the parameters:

1. Edit the template for the Fluentd daemonset:

```
$ oc edit -n logging template logging-fluentd-template
```

Add or edit the environment variable **ES_COPY** to have the value **"true"** (with the quotes), and add or edit the **COPY** variables listed above.

2. Recreate the Fluentd daemonset from the template:

```
$ oc delete daemonset logging-fluentd
$ oc new-app logging-fluentd-template
```

Configuring Fluentd to Send Logs to an External Log Aggregator

You can configure Fluentd to send a copy of its logs to an external log aggregator, and not the default Elasticsearch, using the **secure-forward** plug-in. From there, you can further process log records after the locally hosted Fluentd has processed them.

The deployer provides a **secure-forward.conf** section in the Fluentd configmap for configuring the external aggregator:

```
<store>
@type secure_forward
self_hostname pod-${HOSTNAME}
shared_key thisisasharedkey
secure yes
enable_strict_verification yes
ca_cert_path /etc/fluent/keys/your_ca_cert
ca_private_key_path /etc/fluent/keys/your_private_key
ca_private_key_passphrase passphrase
<server>
  host ose1.example.com
  port 24284
</server>
<server>
  host ose2.example.com
  port 24284
  standby
</server>
<server>
  host ose3.example.com
  port 24284
  standby
</server>
</store>
```

This can be updated using the **oc edit** command:

```
$ oc edit configmap/logging-fluentd
```

Certificates to be used in **secure-forward.conf** can be added to the existing secret that is mounted on the Fluentd pods. The **your_ca_cert** and **your_private_key** values must match what is specified in **secure-forward.conf** in **configmap/logging-fluentd**:

```
$ oc patch secrets/logging-fluentd --type=json \
  --patch "[{'op':'add','path':'/data/your_ca_cert','value':'$(base64
/path/to/your_ca_cert.pem)'}]"
$ oc patch secrets/logging-fluentd --type=json \
  --patch "[{'op':'add','path':'/data/your_private_key','value':'$(base64
/path/to/your_private_key.pem)'}]"
```



NOTE

Avoid using secret names such as 'cert', 'key', and 'ca' so that the values do not conflict with the keys generated by the Deployer pod for Fluentd to talk to the OpenShift Container Platform hosted Elasticsearch.

When configuring the external aggregator, it must be able to accept messages securely from Fluentd.

If the external aggregator is another Fluentd process, it must have the **fluent-plugin-secure-forward** plug-in installed and make use of the input plug-in it provides:

```
<source>
  @type secure_forward

  self_hostname ${HOSTNAME}
  bind 0.0.0.0
  port 24284

  shared_key thisisasharedkey

  secure yes
  cert_path      /path/for/certificate/cert.pem
  private_key_path /path/for/certificate/key.pem
  private_key_passphrase secret_foo_bar_baz
</source>
```

Further explanation of how to set up the **fluent-plugin-secure-forward** plug-in can be [found here](#).

Throttling logs in Fluentd

For projects that are especially verbose, an administrator can throttle down the rate at which the logs are read in by Fluentd before being processed.

**WARNING**

Throttling can contribute to log aggregation falling behind for the configured projects; log entries can be lost if a pod is deleted before Fluentd catches up.

**NOTE**

Throttling does not work when using the systemd journal as the log source. The throttling implementation depends on being able to throttle the reading of the individual log files for each project. When reading from the journal, there is only a single log source, no log files, so no file-based throttling is available. There is not a method of restricting the log entries that are read into the Fluentd process.

To tell Fluentd which projects it should be restricting, edit the throttle configuration in its ConfigMap after deployment:

```
$ oc edit configmap/logging-fluentd
```

The format of the ***throttle-config.yaml*** key is a YAML file that contains project names and the desired rate at which logs are read in on each node. The default is 1000 lines at a time per node. For example:

```
logging:
  read_lines_limit: 500

test-project:
  read_lines_limit: 10

.operations:
  read_lines_limit: 100
```

When you make changes to any part of the EFK stack, specifically Elasticsearch or Fluentd, you should first scale Elasticsearch down to zero and scale Fluentd so it does not match any other nodes. Then, make the changes and scale Elasticsearch and Fluentd back.

To scale Elasticsearch to zero:

```
$ oc scale --replicas=0 dc/<ELASTICSEARCH_DC>
```

Change nodeSelector in the daemonset configuration to match zero:

Get the fluentd node selector:

```
$ oc get ds logging-fluentd -o yaml |grep -A 1 Selector
nodeSelector:
  logging-infra-fluentd: "true"
```

Use the oc patch command to modify the daemonset nodeSelector:

```
$ oc patch ds logging-fluentd -p '{"spec":{"template":{"spec":
```

```
{"nodeSelector":{"nonexistlabel":"true"}}}}}'
```

Get the fluentd node selector:

```
$ oc get ds logging-fluentd -o yaml |grep -A 1 Selector
nodeSelector:
  "nonexistlabel: "true"
```

Scale Elasticsearch back up from zero:

```
$ oc scale --replicas=# dc/<ELASTICSEARCH_DC>
```

Change nodeSelector in the daemonset configuration back to logging-infra-fluentd: "true".

Use the **oc patch** command to modify the daemonset nodeSelector:

```
oc patch ds logging-fluentd -p '{"spec":{"template":{"spec":
{"nodeSelector":{"logging-infra-fluentd":"true"}}}}}'
```

27.5.4. Kibana

To access the Kibana console from the OpenShift Container Platform web console, add the **loggingPublicURL** parameter in the */etc/origin/master/master-config.yaml* file, with the URL of the Kibana console (the **kibana-hostname** parameter). The value must be an HTTPS URL:

```
...
assetConfig:
  ...
  loggingPublicURL: "https://kibana.example.com"
...
```

Setting the **loggingPublicURL** parameter creates a **View Archive** button on the OpenShift Container Platform web console under the **Browse** → **Pods** → **<pod_name>** → **Logs** tab. This links to the Kibana console.

You can scale the Kibana deployment as usual for redundancy:

```
$ oc scale dc/logging-kibana --replicas=2
```

You can see the user interface by visiting the site specified at the **KIBANA_HOSTNAME** variable.

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

27.5.5. Curator

Curator allows administrators to configure scheduled Elasticsearch maintenance operations to be performed automatically on a per-project basis. It is scheduled to perform actions daily based on its configuration. Only one Curator pod is recommended per Elasticsearch cluster. Curator is configured via a YAML configuration file with the following structure:

```
$PROJECT_NAME:
  $ACTION:
```

```

    $UNIT: $VALUE
$PROJECT_NAME:
    $ACTION:
        $UNIT: $VALUE
    ...

```

The available parameters are:

Variable Name	Description
\$PROJECT_NAME	The actual name of a project, such as myapp-devel . For OpenShift Container Platform operations logs, use the name .operations as the project name.
\$ACTION	The action to take, currently only delete is allowed.
\$UNIT	One of days , weeks , or months .
\$VALUE	An integer for the number of units.
.defaults	Use .defaults as the \$PROJECT_NAME to set the defaults for projects that are not specified.
runhour	(Number) the hour of the day in 24-hour format at which to run the Curator jobs. For use with .defaults .
runminute	(Number) the minute of the hour at which to run the Curator jobs. For use with .defaults .

For example, to configure Curator to:

- delete indices in the **myapp-dev** project older than **1 day**
- delete indices in the **myapp-qe** project older than **1 week**
- delete **operations** logs older than **8 weeks**
- delete all other projects indices after they are **30 days** old
- run the Curator jobs at midnight every day

Use:

```

myapp-dev:
  delete:
    days: 1

myapp-qe:
  delete:
    weeks: 1

```

```
.operations:
  delete:
    weeks: 8

.defaults:
  delete:
    days: 30
  runhour: 0
  runminute: 0
```



IMPORTANT

When you use **month** as the **\$UNIT** for an operation, Curator starts counting at the first day of the current month, not the current day of the current month. For example, if today is April 15, and you want to delete indices that are 2 months older than today (`delete: months: 2`), Curator does not delete indices that are dated older than February 15; it deletes indices older than February 1. That is, it goes back to the first day of the current month, then goes back two whole months from that date. If you want to be exact with Curator, it is best to use days (for example, **`delete: days: 30`**).

27.5.5.1. Creating the Curator Configuration

The deployer provides a ConfigMap from which Curator reads its configuration. You may edit or replace this ConfigMap to reconfigure Curator. Currently the **logging-curator** ConfigMap is used to configure both your ops and non-ops Curator instances. Any **.operations** configurations will be in the same location as your application logs configurations.

1. To edit the provided ConfigMap to configure your Curator instances:

```
$ oc edit configmap/logging-curator
```

2. To replace the provided ConfigMap instead:

```
$ create /path/to/mycuratorconfig.yaml
$ oc create configmap logging-curator -o yaml \
  --from-file=config.yaml=/path/to/mycuratorconfig.yaml | \
  oc replace -f -
```

3. After you make your changes, redeploy Curator:

```
$ oc deploy --latest dc/logging-curator
$ oc deploy --latest dc/logging-curator-ops
```

27.6. CLEANUP

Remove everything generated during the deployment while leaving other project contents intact:

```
$ oc new-app logging-deployer-template --param MODE=uninstall
```

27.7. UPGRADING

To upgrade the EFK logging stack, see [Manual Upgrades](#).

27.8. TROUBLESHOOTING KIBANA

Using the Kibana console with OpenShift Container Platform can cause problems that are easily solved, but are not accompanied with useful error messages. Check the following troubleshooting sections if you are experiencing any problems when deploying Kibana on OpenShift Container Platform:

Login Loop

The OAuth2 proxy on the Kibana console must share a secret with the master host's OAuth2 server. If the secret is not identical on both servers, it can cause a login loop where you are continuously redirected back to the Kibana login page.

To fix this issue, delete the current OAuthClient, and create a new one, using the same template as before:

```
$ oc delete oauthclient/kibana-proxy
$ oc new-app logging-support-template
```

Cryptic Error When Viewing the Console

When attempting to visit the Kibana console, you may receive a browser error instead:

```
{"error":"invalid_request","error_description":"The request is missing a
required parameter,
includes an invalid parameter value, includes a parameter more than once,
or is otherwise malformed."}
```

This can be caused by a mismatch between the OAuth2 client and server. The return address for the client must be in a whitelist so the server can securely redirect back after logging in.

Fix this issue by replacing the OAuthClient entry:

```
$ oc delete oauthclient/kibana-proxy
$ oc new-app logging-support-template
```

If the problem persists, check that you are accessing Kibana at a URL listed in the OAuth client. This issue can be caused by accessing the URL at a forwarded port, such as 1443 instead of the standard 443 HTTPS port. You can adjust the server whitelist by editing the OAuth client:

```
$ oc edit oauthclient/kibana-proxy
```

503 Error When Viewing the Console

If you receive a proxy error when viewing the Kibana console, it could be caused by one of two issues.

First, Kibana may not be recognizing pods. If Elasticsearch is slow in starting up, Kibana may timeout trying to reach it. Check whether the relevant service has any endpoints:

```
$ oc describe service logging-kibana
Name:                logging-kibana
[...]
Endpoints:           <none>
```


If any Kibana pods are live, endpoints will be listed. If they are not, check the state of the Kibana pods and deployment. You may need to scale the deployment down and back up again.

The second possible issue may be caused if the route for accessing the Kibana service is masked. This can happen if you perform a test deployment in one project, then deploy in a different project without completely removing the first deployment. When multiple routes are sent to the same destination, the default router will only route to the first created. Check the problematic route to see if it is defined in multiple places:

```
$ oc get route --all-namespaces --selector logging-infra=support
```

F-5 Load Balancer and X-Forwarded-For Enabled

If you are attempting to use a F-5 load balancer in front of Kibana with **X-Forwarded-For** enabled, this can cause an issue in which the Elasticsearch **Searchguard** plug-in is unable to correctly accept connections from Kibana.

Example Kibana Error Message

```
Kibana: Unknown error while connecting to Elasticsearch
Error: Unknown error while connecting to Elasticsearch
Error: UnknownHostException[No trusted proxies]
```

To configure Searchguard to ignore the extra header:

1. Scale down all Fluentd pods.
2. Scale down Elasticsearch after the Fluentd pods have terminated.
3. Add **searchguard.http.xforwardedfor.header: DUMMY** to the Elasticsearch configuration section.

```
$ oc edit configmap/logging-elasticsearch 1
```

1 This approach requires that Elasticsearch's configurations are within a ConfigMap.

4. Scale Elasticsearch back up.
5. Scale up all Fluentd pods.

27.9. SENDING LOGS TO AN EXTERNAL ELASTICSEARCH INSTANCE

Fluentd sends logs to the value of the **ES_HOST**, **ES_PORT**, **OPS_HOST**, and **OPS_PORT** environment variables of the Elasticsearch deployment configuration. The application logs are directed to the **ES_HOST** destination, and operations logs to **OPS_HOST**.



NOTE

Sending logs directly to an AWS Elasticsearch instance is not supported. Use [Fluentd Secure Forward](#) to direct logs to an instance of Fluentd that you control and that is configured with the **fluent-plugin-aws-elasticsearch-service** plug-in.

To direct logs to a specific Elasticsearch instance, edit the deployment configuration and replace the value of the above variables with the desired instance:

```
$ oc edit dc/<deployment_configuration>
```

For an external Elasticsearch instance to contain both application and operations logs, you can set **ES_HOST** and **OPS_HOST** to the same destination, while ensuring that **ES_PORT** and **OPS_PORT** also have the same value.

If your externally hosted Elasticsearch instance does not use TLS, update the **_CLIENT_CERT**, **_CLIENT_KEY**, and **_CA** variables to be empty. If it does use TLS, but not mutual TLS, update the **_CLIENT_CERT** and **_CLIENT_KEY** variables to be empty and patch or recreate the **logging-fluentd** secret with the appropriate **_CA** value for communicating with your Elasticsearch instance. If it uses Mutual TLS as the provided Elasticsearch instance does, patch or recreate the **logging-fluentd** secret with your client key, client cert, and CA.

Since Fluentd is deployed by a DaemonSet, update the **logging-fluentd-template** template, delete your current DaemonSet, and recreate it with **oc new-app logging-fluentd-template** after seeing all previous Fluentd pods have terminated.



NOTE

If you are not using the provided Kibana and Elasticsearch images, you will not have the same multi-tenant capabilities and your data will not be restricted by user access to a particular project.

27.10. PERFORMING ADMINISTRATIVE ELASTICSEARCH OPERATIONS

As of the Deployer version 3.2.0, an administrator certificate, key, and CA that can be used to communicate with and perform administrative operations on Elasticsearch are provided within the **logging-elasticsearch** secret.



NOTE

To confirm whether or not your EFK installation provides these, run:

```
$ oc describe secret logging-elasticsearch
```

If they are not available, refer to [Manual Upgrades](#) to ensure you are on the latest version first.

1. Connect to an Elasticsearch pod that is in the cluster on which you are attempting to perform maintenance.
2. To find a pod in a cluster use either:

```
$ oc get pods -l component=es -o name | head -1
$ oc get pods -l component=es-ops -o name | head -1
```

3. Connect to a pod:

```
$ oc rsh <your_Elasticsearch_pod>
```

4. Once connected to an Elasticsearch container, you can use the certificates mounted from the secret to communicate with Elasticsearch per its 1.5 [Document APIs](#).
Fluentd sends its logs to Elasticsearch using the index format **{project_name}.
{project_uuid}.YYYY.MM.DD** where YYYY.MM.DD is the date of the log record.

For example, to delete all logs for the **logging** project with uuid **3b3594fa-2ccd-11e6-acb7-0eb6b35eae3** from June 15, 2016, we can run:

```
$ curl --key /etc/elasticsearch/secret/admin-key \  
  --cert /etc/elasticsearch/secret/admin-cert \  
  --cacert /etc/elasticsearch/secret/admin-ca -XDELETE \  
  "https://localhost:9200/logging.3b3594fa-2ccd-11e6-acb7-  
  0eb6b35eae3.2016.06.15"
```

CHAPTER 28. AGGREGATE LOGGING SIZING GUIDELINES

28.1. OVERVIEW

The [Elasticsearch, Fluentd, and Kibana](#) (EFK) stack aggregates logs from nodes and applications running inside your OpenShift Container Platform installation. Once deployed it uses [Fluentd](#) to aggregate event logs from all nodes, projects, and pods into [Elasticsearch \(ES\)](#). It also provides a centralized [Kibana](#) web UI where users and administrators can create rich visualizations and dashboards with the aggregated data.

Fluentd [bulk uploads](#) logs to an index, in JSON format, then Elasticsearch routes your search requests to the appropriate shards.

28.2. INSTALLATION

The general procedure for installing an aggregate logging stack in OpenShift Container Platform is described in [Aggregating Container Logs](#). There are some important things to keep in mind while going through the installation guide:

In order for the logging pods to spread evenly across your cluster, an empty [node selector](#) should be used.

```
$ oadm new-project logging --node-selector=""
```

In conjunction with node labeling, which is done later, this controls pod placement across the logging project. You can now create the logging project.

```
$ oc project logging
```

If you are installing in a PoC or testing environment, a local openshift-ansible template install is recommended.

```
$ oc create -f
${OPENSIFT_ANSIBLE_REPO}/roles/openshift_examples/files/examples/${VERSION}/
infrastructure-templates/origin/logging-deployer.yaml
```

Elasticsearch (ES) should be deployed with a cluster size of at least three for resiliency to node failures. This is specified by passing the **ES_CLUSTER_SIZE** parameter to the installer.

```
$ oc new-app logging-deployer-template \
    --param ES_CLUSTER_SIZE=3 \
    --param PUBLIC_MASTER_URL=$PUBLIC_MASTER_URL \
    --param KIBANA_HOSTNAME=$KIBANA_URL
```

Refer to [Deploying the EFK Stack](#) for a full list of parameters.

If you do not have an existing Kibana installation, you can use **kibana.example.com** as a value to **KIBANA_HOSTNAME**.

Now, ensure Fluentd pod spreading through labeling.

```
$ oc label nodes --all logging-infra-fluentd=true
```

This operation requires the **cluster-admin** default role.

Installation can take some time depending on whether the images were already retrieved from the registry or not, and on the size of your cluster.

Inside the **logging** namespace, you can check your deployment with **oc get all**.

```
$ oc get all
```

NAME	REVISION	REPLICAS
TRIGGERED BY		
logging-curator	1	1
logging-es-6cvk237t	1	1
logging-es-e5x4t4ai	1	1
logging-es-xmwvnorv	1	1
logging-kibana	1	1
NAME	DESIRED	CURRENT
AGE		
logging-curator-1	1	1 3d
logging-es-6cvk237t-1	1	1 3d
logging-es-e5x4t4ai-1	1	1 3d
logging-es-xmwvnorv-1	1	1 3d
logging-kibana-1	1	1 3d
NAME	HOST/PORT	PATH
SERVICE	TERMINATION	LABELS
logging-kibana		kibana.example.com
logging-kibana	reencrypt	component=support, logging-
		infra=support, provider=openshift
logging-kibana-ops		kibana-ops.example.com
logging-kibana-ops	reencrypt	component=support, logging-
		infra=support, provider=openshift
NAME	CLUSTER-IP	EXTERNAL-IP
PORT(S)	AGE	
logging-es	172.24.155.177	<none>
9200/TCP	3d	
logging-es-cluster	None	<none>
9300/TCP	3d	
logging-es-ops	172.27.197.57	<none>
9200/TCP	3d	
logging-es-ops-cluster	None	<none>
9300/TCP	3d	
logging-kibana	172.27.224.55	<none>
443/TCP	3d	
logging-kibana-ops	172.25.117.77	<none>
443/TCP	3d	
NAME	READY	STATUS
RESTARTS	AGE	
logging-curator-1-6s7wy	1/1	Running 0
3d		
logging-deployer-un6ut	0/1	Completed 0
3d		
logging-es-6cvk237t-1-cn3pw	1/1	Running 0
3d		
logging-es-e5x4t4ai-1-v933h	1/1	Running 0
3d		
logging-es-xmwvnorv-1-adr5x	1/1	Running 0

```

3d
logging-fluentd-156xn          1/1          Running      0
3d
logging-fluentd-40biz          1/1          Running      0
3d
logging-fluentd-7dtom          0/1          Pending      0
2d
logging-fluentd-8k847          1/1          Running      0
3d

```

You should end up with a similar setup to the below.

```

$ oc get pods -o wide

NAME                                READY    STATUS    RESTARTS   AGE
NODE
logging-curator-1-6s7wy             1/1     Running   0           3d
ip-172-31-24-239.us-west-2.compute.internal
logging-deployer-un6ut              0/1     Completed 0           3d
ip-172-31-6-152.us-west-2.compute.internal
logging-es-6cvk237t-1-cnpw3         1/1     Running   0           3d
ip-172-31-24-238.us-west-2.compute.internal
logging-es-e5x4t4ai-1-v933h         1/1     Running   0           3d
ip-172-31-24-235.us-west-2.compute.internal
logging-es-xmwvnr-1-adr5x           1/1     Running   0           3d
ip-172-31-24-233.us-west-2.compute.internal
logging-fluentd-156xn               1/1     Running   0           3d
ip-172-31-24-241.us-west-2.compute.internal
logging-fluentd-40biz               1/1     Running   0           3d
ip-172-31-24-236.us-west-2.compute.internal
logging-fluentd-7dtom               0/1     Pending   0           2d
ip-172-31-24-243.us-west-2.compute.internal
logging-fluentd-8k847               1/1     Running   0           3d
ip-172-31-24-237.us-west-2.compute.internal
logging-fluentd-9a3qx               1/1     Running   0           3d
ip-172-31-24-231.us-west-2.compute.internal
logging-fluentd-abvgj               1/1     Running   0           3d
ip-172-31-24-228.us-west-2.compute.internal
logging-fluentd-bh74n               1/1     Running   0           3d
ip-172-31-24-238.us-west-2.compute.internal
...
...

```

Notice how the pods are placed in different cluster nodes.

By default the amount of RAM allocated to each ES instance is 8GB. **ES_INSTANCE_RAM** is the parameter used in the [openshift-ansible](#) template. Keep in mind that **half** of this value will be passed to the individual elasticsearch pods java processes [heap size](#).

[Learn more about installing EFK.](#)

28.3. SYSTEMD-JOURNALD AND RSYSLOG

Rate-limiting

In Red Hat Enterprise Linux (RHEL) 7 the **systemd-journald.socket** unit creates **/dev/log** during the boot process, and then passes input to **systemd-journald.service**. Every **syslog()** call goes to the journal.

Rsyslog uses the **imjournal** module as a default input mode for journal files. Refer to [Interaction of rsyslog and journal](#) for detailed information about this topic.

A simple test harness was developed, which uses [logger](#) across the cluster nodes to make entries of different sizes at different rates in the system log. During testing simulations under a default Red Hat Enterprise Linux (RHEL) 7 installation with **systemd-219-19.el7.x86_64** at certain logging rates (approximately 40 log lines per second), we encountered the default rate limit of **rsyslogd**. After adjusting these limits, entries stopped being written to journald due to local journal file corruption. [This issue is resolved in later versions of systemd](#).

Scaling up

As you scale up your project, the default logging environment might need some adjustments. After updating to **systemd-219-22.el7.x86_64**, we added:

```
$IMUXSockRateLimitInterval 0
$IMJournalRateLimitInterval 0
```

to **/etc/rsyslog.conf** and:

```
# Disable rate limiting
RateLimitInterval=1s
RateLimitBurst=10000
Storage=volatile
Compress=no
MaxRetentionSec=5s
```

to **/etc/systemd/journald.conf**.

Now, restart the services.

```
$ systemctl restart systemd-journald.service
$ systemctl restart rsyslog.service
```

These settings account for the bursty nature of uploading in bulk.

After removing the rate limit, you may see increased CPU utilization on the system logging daemons as it processes any messages that would have previously been throttled.

Rsyslog is configured (see **ratelimit.interval**, **ratelimit.burst**) to rate-limit entries read from the journal at 10,000 messages in 300 seconds. A good rule of thumb is to ensure that the rsyslog rate-limits account for the systemd-journald rate-limits.

28.4. SCALING UP EFK LOGGING

If you do not indicate the desired scale at first deployment, the least disruptive way of adjusting your cluster is by re-running the deployer with the updated **ES_CLUSTER_SIZE** value and using the **MODE=reinstall** template parameter. Refer to the [Performing Administrative Elasticsearch Operations](#) section for more in-depth information.

```
$ oc edit configmap logging-deployer
[change es-cluster-size value to 5]

$ oc new-app logging-deployer-template --param MODE=reinstall
```

28.5. STORAGE CONSIDERATIONS

An Elasticsearch index is a collection of shards and its corresponding replica shards. This is how ES implements high availability internally, therefore there is little need to use hardware based mirroring RAID variants. RAID 0 can still be used to increase overall disk performance.

Every search request needs to hit a copy of every shard in the index. Each ES instance requires its own individual storage, but an OpenShift Container Platform deployment can only provide volumes shared by all of its pods, which again means that Elasticsearch shouldn't be implemented with a single node.

A [persistent volume](#) should be added to each Elasticsearch deployment configuration so that we have one volume per [replica shard](#). On OpenShift Container Platform this is often achieved through [Persistent Volume Claims](#)

- 1 volume per shard
- 1 volume per replica shard

The PVCs must be named based on the **es-pvc-prefix** setting. Refer to [Persistent Elasticsearch Storage](#) for more details.

Below are capacity planning guidelines for OpenShift Container Platform aggregate logging. **Example scenario**

Assumptions:

1. Which application: Apache
2. Bytes per line: 256
3. Lines per second load on application: 1
4. Raw text data → JSON

Baseline (256 characters per minute → 15KB/min)

Logging Infra Pods	Storage Throughput
3 es 1 kibana 1 curator 1 fluentd	6 pods total: 90000 x 86400 = 7,7 GB/day
3 es 1 kibana 1 curator 11 fluentd	16 pods total: 225000 x 86400 = 24,0 GB/day
3 es 1 kibana 1 curator 20 fluentd	25 pods total: 225000 x 86400 = 32,4 GB/day

Calculating total logging throughput and disk space required for your logging environment requires knowledge of your application. For example, if one of your applications on average logs 10 lines-per-second, each 256 bytes-per-line, calculate per-application throughput and disk space as follows:


```
(bytes-per-line * (lines-per-second) = 2560 bytes per app per second
(2560) * (number-of-pods-per-node,100) = 256,000 bytes per second per node
256k * (number-of-nodes) = total logging throughput per cluster
```

Fluentd ships any logs from **`/var/log/messages`** and **`/var/lib/docker/containers/`** to Elasticsearch. [Learn more](#).

Local SSD drives are recommended in order to achieve the best performance. In Red Hat Enterprise Linux (RHEL) 7, the [deadline](#) IO scheduler is the default for all block devices except SATA disks. For SATA disks, the default IO scheduler is **`cfq`**.

Sizing storage for ES is greatly dependent on how you optimize your indices. Therefore, consider how much data you need in advance and that you are aggregating application log data.

CHAPTER 29. ENABLING CLUSTER METRICS

29.1. OVERVIEW

The [kubelet](#) exposes metrics that can be collected and stored in back-ends by [Heapster](#).

As an OpenShift Container Platform administrator, you can view a cluster's metrics from all containers and components in one user interface. These metrics are also used by [horizontal pod autoscalers](#) in order to determine when and how to scale.

This topic describes using [Hawkular Metrics](#) as a metrics engine which stores the data persistently in a [Cassandra](#) database. When this is configured, CPU, memory and network-based metrics are viewable from the OpenShift Container Platform web console and are available for use by [horizontal pod autoscalers](#).

Heapster retrieves a list of all nodes from the master server, then contacts each node individually through the `/stats` endpoint. From there, Heapster scrapes the metrics for CPU, memory and network usage, then exports them into Hawkular Metrics.

Browsing individual pods in the web console displays separate sparkline charts for memory and CPU. The time range displayed is selectable, and these charts automatically update every 30 seconds. If there are multiple containers on the pod, then you can select a specific container to display its metrics.

If [resource limits](#) are defined for your project, then you can also see a donut chart for each pod. The donut chart displays usage against the resource limit. For example: **145 Available of 200 MiB**, with the donut chart showing **55 MiB Used**.

29.2. BEFORE YOU BEGIN

The components for cluster metrics must be deployed to the **openshift-infra** project. This allows [horizontal pod autoscalers](#) to discover the Heapster service and use it to retrieve metrics that can be used for autoscaling.

All of the following commands in this topic must be executed under the **openshift-infra** project. To switch to the **openshift-infra** project:

```
$ oc project openshift-infra
```

To enable cluster metrics, you must next configure the following:

- [Service Accounts](#)
- [Metrics Data Storage](#)
- [Metrics Deployer](#)

29.3. SERVICE ACCOUNTS

You must configure [service accounts](#) for:

- [Metrics Deployer](#)
- [Heapster](#)

29.3.1. Metrics Deployer Service Account

The [Metrics Deployer](#) will be discussed in a later step, but you must first set up a service account for it:

1. Create a **metrics-deployer** service account:

```
$ oc create -f - <<API
apiVersion: v1
kind: ServiceAccount
metadata:
  name: metrics-deployer
secrets:
- name: metrics-deployer
API
```


2. Before it can deploy components, the **metrics-deployer** service account must also be granted the **edit** permission for the **openshift-infra** project:

```
$ oadm policy add-role-to-user \
  edit system:serviceaccount:openshift-infra:metrics-deployer
```

29.3.2. Heapster Service Account

The Heapster component requires access to the master server to list all available nodes and access the `/stats` endpoint for each node. Before it can do this, the Heapster service account requires the **cluster-reader** permission:

```
$ oadm policy add-cluster-role-to-user \
  cluster-reader system:serviceaccount:openshift-infra:heapster
```



NOTE

The Heapster service account is created automatically during the [Deploying the Metrics Components](#) step.

29.4. METRICS DATA STORAGE

You can store the metrics data to either [persistent storage](#) or to a temporary [pod volume](#).

29.4.1. Persistent Storage

Running OpenShift Container Platform cluster metrics with persistent storage means that your metrics will be stored to a [persistent volume](#) and be able to survive a pod being restarted or recreated. This is ideal if you require your metrics data to be guarded from data loss. For production environments it is highly recommended to configure persistent storage for your metrics pods.

The size of the persisted volume can be specified with the **CASSANDRA_PV_SIZE** [template parameter](#). By default it is set to 10 GB, which may or may not be sufficient for the size of the cluster you are using. If you require more space, for instance 100 GB, you could specify it with something like this:

```
$ oc process -f metrics-deployer.yaml -v \
  HAWKULAR_METRICS_HOSTNAME=hawkular-
  metrics.example.com,CASSANDRA_PV_SIZE=100Gi \
```

```
| oc create -f -
```

The size requirement of the Cassandra storage is dependent on the number of pods. It is the administrator's responsibility to ensure that the size requirements are sufficient for their setup and to monitor usage to ensure that the disk does not become full.

If you would like to use [dynamically provisioned persistent volumes](#) you must set the **DYNAMICALLY_PROVISION_STORAGE** template option to **true** for the Metrics Deployer.

29.4.2. Capacity Planning for Cluster Metrics

After the metrics deployer runs, the output of **oc get pods** should resemble the following:

```
# oc get pods -n openshift-infra
NAME                                READY   STATUS
RESTARTS      AGE
hawkular-cassandra-1-l5y4g          1/1     Running    0
17h
hawkular-metrics-1t9so              1/1     Running    0
17h
heapster-febru                      1/1     Running    0
17h
```

OpenShift Container Platform metrics are stored using the Cassandra database, which is deployed with settings of **MAX_HEAP_SIZE=512M** and **NEW_HEAP_SIZE=100M**. These values should cover most OpenShift Container Platform metrics installations, but you can modify them in the Cassandra Dockerfile prior to deploying cluster metrics.

By default, metrics data is stored for 7 days. You can configure this with the **METRIC_DURATION** parameter in the **metrics.yaml** configuration file. After 7 days, Cassandra begins to purge the oldest metrics data. Metrics data for deleted pods and projects is not automatically purged; it is only removed once the data is 7 days old.

Example 29.1. Data Accumulated by 10 Nodes and 1000 Pods

In a test scenario including 10 nodes and 1000 pods, a 24 hour period accumulated 2.5GB of metrics data. Therefore, the capacity planning formula for metrics data in this scenario is:

$$(((2.5 \times 10^9) \div 1000) \div 24) \div 10^6 = \sim 0.125 \text{ MB/hour per pod.}$$

Example 29.2. Data Accumulated by 120 Nodes and 10000 Pods

In a test scenario including 120 nodes and 10000 pods, a 24 hour period accumulated 25GB of metrics data. Therefore, the capacity planning formula for metrics data in this scenario is:

$$(((11.410 \times 10^9) \div 1000) \div 24) \div 10^6 = 0.475 \text{ MB/hour}$$

	1000 pods	10000 pods
Cassandra storage data accumulated over 24 hours (default metrics parameters)	2.5GB	11.4GB

If the default value of 7 days for **METRIC_DURATION** and 10 seconds for **METRIC_RESOLUTION** are preserved, then weekly storage requirements for the Cassandra pod would be:

	1000 pods	10000 pods
Cassandra storage data accumulated over 7 days (default metrics parameters)	20GB	90GB

In the previous table, an additional 10% was added to the expected storage space as a buffer for unexpected monitored pod usage.



WARNING

If the Cassandra persisted volume runs out of sufficient space, then data loss will occur.

For cluster metrics to work with persistent storage, ensure that the persistent volume has the **ReadWriteOnce** access mode. If this mode is not active, then the persistent volume claim cannot locate the persistent volume, and Cassandra fails to start.


To use persistent storage with the metric components, ensure that a [persistent volume](#) of sufficient size is available. The creation of [persistent volume claims](#) is handled by the [Metrics Deployer](#).

OpenShift Container Platform metrics also supports dynamically-provisioned persistent volumes. To use this feature with OpenShift Container Platform metrics, it is necessary to add an additional flag to the metrics-deployer: **DYNAMICALLY_PROVISION_STORAGE=true**. You can use EBS, GCE, and Cinder storage back-ends to [dynamically provision persistent volumes](#).

29.4.3. Non-Persistent Storage

Running OpenShift Container Platform cluster metrics with non-persistent storage means that any stored metrics will be deleted when the pod is deleted. While it is much easier to run cluster metrics with non-persistent data, running with non-persistent data does come with the risk of permanent data loss. However, metrics can still survive a container being restarted.

In order to use non-persistent storage, you must set the **USE_PERSISTENT_STORAGE** [template option](#) to **false** for the Metrics Deployer.



NOTE

When using non-persistent storage, metrics data will be written to **`/var/lib/origin/openshift.local.volumes/pods`** on the node where the Cassandra pod is running. Ensure **`/var`** has enough free space to accommodate metrics storage.

29.5. METRICS DEPLOYER

The Metrics Deployer deploys and configures all of the metrics components. You can configure it by passing in information from [secrets](#) and by passing parameters to the Metrics Deployer's [template](#).

29.5.1. Using Secrets

The Metrics Deployer will auto-generate self-signed certificates for use between its components and will generate a [re-encrypting route](#) to expose the Hawkular Metrics service. This route is what allows the web console to access the Hawkular Metrics service.

In order for the browser running the web console to trust the connection through this route, it must trust the route's certificate. This can be accomplished by [providing your own certificates](#) signed by a trusted Certificate Authority. The Metric Deployer's secret allows you to pass your own certificates which it will then use when creating the route.

If you do not wish to provide your own certificates, the router's default certificate will be used instead.

29.5.1.1. Providing Your Own Certificates

To provide your own certificate which will be used by the [re-encrypting route](#), you can pass these values as [secrets](#) to the Metrics Deployer.

The **`hawkular-metrics.pem`** value needs to contain the certificate in its **`.pem`** format. You may also need to provide the certificate for the Certificate Authority which signed this **`.pem`** file via the **`hawkular-metrics-ca.cert`** secret.

```
$ oc secrets new metrics-deployer \
    hawkular-metrics.pem=/home/openshift/metrics/hm.pem \
    hawkular-metrics-ca.cert=/home/openshift/metrics/hm-ca.cert
```

When these secrets are provided, the deployer uses these values to specify the **key**, **certificate** and **caCertificate** values for the re-encrypting route it generated.

For more information, please see the [re-encryption route documentation](#).

29.5.1.2. Using the Router's Default Certificate

If the **`hawkular-metrics.pem`** value is not specified, the re-encrypting route will use the router's default certificate, which may not be trusted by browsers.

A secret named **`metrics-deployer`** will still be required in this situation. This can be considered a "dummy" secret since the secret it specifies is not actually used by the component.

To create a "dummy" secret that does not specify a certificate value:

```
$ oc secrets new metrics-deployer nothing=/dev/null
```

29.5.1.3. Deployer Secret Options

The following table contains more advanced configuration options, detailing all the secrets which can be used by the deployer:

Secret Name	Description
<i>hawkular-metrics.pem</i>	The <i>pem</i> file to use for the Hawkular Metrics certificate used for the re-encrypting route. This certificate must contain the host name used by the route (e.g., HAWKULAR_METRICS_HOSTNAME). The default router's certificate is used for the route if this option is unspecified.
<i>hawkular-metrics-ca.cert</i>	The certificate for the CA used to sign the <i>hawkular-metrics.pem</i> . This is optional if the <i>hawkular-metrics.pem</i> does not contain the CA certificate directly.
<i>heapster.cert</i>	The certificate for Heapster to use. This is auto-generated if unspecified.
<i>heapster.key</i>	The key to use with the Heapster certificate. This is ignored if <i>heapster.cert</i> is not specified
<i>heapster_client_ca.cert</i>	The certificate that generates <i>heapster.cert</i> . This is required if <i>heapster.cert</i> is specified. Otherwise, the main CA for the OpenShift Container Platform installation is used. In order for horizontal pod autoscaling to function properly, this should not be overridden.
<i>heapster_allowed_users</i>	A file containing a comma-separated list of CN to accept from certificates signed with the specified CA. By default, this is set to allow the OpenShift Container Platform service proxy to connect. If you override this, make sure to add system:master-proxy to the list in order to allow horizontal pod autoscaling to function properly.

29.5.2. Modifying the Deployer Template

The OpenShift Container Platform installer uses a [template](#) to deploy the metrics components. The default template can be found at the following path:

```
/usr/share/ansible/openshift-
ansible/roles/openshift_hosted_templates/files/v1.3/enterprise/metrics-
deployer.yaml
```

In case you need to make any changes to this file, copy it to another directory with the file name *metrics-deployer.yaml* and refer to the new location when using it in the following sections.

29.5.2.1. Deployer Template Parameters

The deployer template parameter options and their defaults are listed in the default *metrics-deployer.yaml* file. If required, you can override these values when creating the Metrics Deployer.

Table 29.1. Template Parameters

Parameter	Description
METRIC_DURATION	The number of days metrics should be stored.
CASSANDRA_PV_SIZE	The persistent volume size for each of the Cassandra nodes.
CASSANDRA_NODES	The number of initial Cassandra nodes to deploy.
USE_PERSISTENT_STORAGE	Set to true for persistent storage; set to false to use non-persistent storage.
DYNAMICALLY_PROVISION_STORAGE	Set to true to allow for dynamically provisioned storage .
REDEPLOY	If set to true , the deployer will try to delete all the existing components before trying to redeploy.
HAWKULAR_METRICS_HOSTNAME	External host name where clients can reach Hawkular Metrics. This is the FQDN of the machine running the router pod.
MASTER_URL	Internal URL for the master, for authentication retrieval.
IMAGE_VERSION	Specify version for metrics components. For example, for openshift/origin-metrics-deployer:latest , set version to latest .
IMAGE_PREFIX	Specify prefix for metrics components. For example, for openshift/origin-metrics-deployer:latest , set prefix to openshift/origin- .
MODE	<p>Can be set to:</p> <ul style="list-style-type: none"> • preflight to perform validation before a deployment. • deploy to perform an initial deployment. • refresh to delete and redeploy all components but to keep persisted data and routes. • redeploy to delete and redeploy everything (losing all data in the process). • validate to re-run validations after a deployment.

Parameter	Description
IGNORE_PREFLIGHT	Can be set to true to disable the preflight checks. This allows the deployer to continue even if the preflight check has failed.
USER_WRITE_ACCESS	Sets whether user accounts should be able to write metrics. Defaults to false so that only Heapster can write metrics and not individual users. It is recommended to disable user write access; if enabled, any user will be able to write metrics to the system which can affect performance and can cause Cassandra disk usage to unpredictably increase.

The only required parameter is **HAWKULAR_METRICS_HOSTNAME**. This value is required when creating the deployer, because it specifies the host name for the Hawkular Metrics [route](#). This value should correspond to a fully qualified domain name. You must know the value of **HAWKULAR_METRICS_HOSTNAME** when [configuring the console](#) for metrics access.

If you are using [persistent storage](#) with Cassandra, it is the administrator's responsibility to set a sufficient disk size for the cluster using the **CASSANDRA_PV_SIZE** parameter. It is also the administrator's responsibility to monitor disk usage to make sure that it does not become full.

**WARNING**

Data loss will result if the Cassandra persisted volume runs out of sufficient space.

All of the other parameters are optional and allow for greater customization. For instance, if you have a custom install in which the Kubernetes master is not available under <https://kubernetes.default.svc:443> you can specify the value to use instead with the **MASTER_URL** parameter. To deploy a specific version of the metrics components, use the **IMAGE_VERSION** parameter.


**WARNING**

It is highly recommended to not use **latest** for the **IMAGE_VERSION**. The **latest** version corresponds to the very latest version available and can cause issues if it brings in a newer version not meant to function on the version of OpenShift Container Platform you are currently running.

29.6. DEPLOYING THE METRIC COMPONENTS

Because deploying and configuring all the metric components is handled by the Metrics Deployer, you can simply deploy everything in one step.

The following examples show you how to deploy metrics with and without persistent storage using the default template parameters. Optionally, you can specify any of the [template parameters](#) when calling these commands.



IMPORTANT

In accordance with upstream Kubernetes rules, metrics can be collected only on the default interface of **eth0**.

Example 29.3. Deploying with Persistent Storage

The following command sets the Hawkular Metrics route to use **hawkular-metrics.example.com** and is deployed using persistent storage.

You must have a persistent volume of sufficient size available.

```
$ oc new-app --as=system:serviceaccount:openshift-infra:metrics-deployer \
  -f metrics-deployer.yaml \
  -p HAWKULAR_METRICS_HOSTNAME=hawkular-metrics.example.com
```

Example 29.4. Deploying without Persistent Storage

The following command sets the Hawkular Metrics route to use **hawkular-metrics.example.com** and deploy without persistent storage.

```
$ oc new-app --as=system:serviceaccount:openshift-infra:metrics-deployer \
  -f metrics-deployer.yaml \
  -p HAWKULAR_METRICS_HOSTNAME=hawkular-metrics.example.com \
  -p USE_PERSISTENT_STORAGE=false
```



WARNING

Because this is being deployed without persistent storage, metric data loss can occur.

29.6.1. Metrics Deployer Validations

The metrics deployer does some validation both before and after deployment. If the pre-flight validation fails, the environment for deployment is considered unsuitable and the deployment is aborted. However, you can add **IGNORE_PREFLIGHT=true** to the deployer parameters if you believe the validation has erred.

If post-deployment validation fails, the deployer finishes in an **Error** state, which indicates that you should check the deployer logs for issues that may require addressing. For example, the validation may detect that the external **hawkular-metrics** route is not actually in use, because the route was already created somewhere else. The validation output at the end of a deployment should explain as clearly as possible any issues it finds and what you can do to address them.

Once you have addressed deployment validation issues, you can re-run just the validation by running the deployer again with the **MODE=validate** parameter added, for example:

```
$ oc new-app --as=system:serviceaccount:openshift-infra:metrics-deployer \
  -f metrics-deployer.yaml \
  -p HAWKULAR_METRICS_HOSTNAME=hawkular-metrics.example.com \
  -p MODE=validate
```

There is also a diagnostic for metrics:

```
$ oadm diagnostics MetricsApiProxy
```

29.7. SETTING THE METRICS PUBLIC URL

The OpenShift Container Platform web console uses the data coming from the Hawkular Metrics service to display its graphs. The URL for accessing the Hawkular Metrics service must be configured via the **metricsPublicURL** option in the [master configuration file](#) (*/etc/origin/master/master-config.yaml*). This URL corresponds to the route created with the **HAWKULAR_METRICS_HOSTNAME** template parameter during the [deployment](#) of the metrics components.



NOTE

You must be able to resolve the **HAWKULAR_METRICS_HOSTNAME** from the browser accessing the console.

For example, if your **HAWKULAR_METRICS_HOSTNAME** corresponds to **hawkular-metrics.example.com**, then you must make the following change in the *master-config.yaml* file:

```
assetConfig:
  ...
  metricsPublicURL: "https://hawkular-
  metrics.example.com/hawkular/metrics"
```

Once you have updated and saved the *master-config.yaml* file, you must restart your OpenShift Container Platform instance.

When your OpenShift Container Platform server is back up and running, metrics will be displayed on the pod overview pages.

CAUTION

If you are using self-signed certificates, remember that the Hawkular Metrics service is hosted under a different host name and uses different certificates than the console. You may need to explicitly open a browser tab to the value specified in **metricsPublicURL** and accept that certificate.

To avoid this issue, use certificates which are configured to be acceptable by your browser.

29.8. ACCESSING HAWKULAR METRICS DIRECTLY

To access and manage metrics more directly, use the Hawkular Metrics API.



NOTE

When accessing Hawkular Metrics via the API, you will only be able to perform reads. Writing metrics has been disabled by default. If you want for individual users to also be able to write metrics, you must set the **USER_WRITE_ACCESS** [deployer template parameter](#) to **true**.

However, it is recommended to use the default configuration and only have metrics enter the system via Heapster. If write access is enabled, any user will be able to write metrics to the system, which can affect performance and cause Cassandra disk usage to unpredictably increase.

The [Hawkular Metrics documentation](#) covers how to use the API, but there are a few differences when dealing with the version of Hawkular Metrics configured for use on OpenShift Container Platform:

29.8.1. OpenShift Container Platform Projects and Hawkular Tenants

Hawkular Metrics is a multi-tenanted application. It is configured so that a project in OpenShift Container Platform corresponds to a tenant in Hawkular Metrics.

As such, when accessing metrics for a project named **MyProject** you must set the [Hawkular-Tenant](#) header to **MyProject**.

There is also a special tenant named **_system** which contains system level metrics. This requires either a **cluster-reader** or **cluster-admin** level privileges to access.

29.8.2. Authorization

The Hawkular Metrics service will authenticate the user against OpenShift Container Platform to determine if the user has access to the project it is trying to access.

Hawkular Metrics accepts a bearer token from the client and verifies that token with the OpenShift Container Platform server using a **SubjectAccessReview**. If the user has proper read privileges for the project, they are allowed to read the metrics for that project. For the **_system** tenant, the user requesting to read from this tenant must have **cluster-reader** permission.

When accessing the Hawkular Metrics API, you must pass a bearer token in the **Authorization** header.

29.9. SCALING OPENSIFT CONTAINER PLATFORM METRICS PODS

One set of metrics pods (Cassandra/Hawkular/Heapster) is able to monitor at least 10,000 pods.



NOTE

Autoscaling the metrics components, such as Hawkular and Heapster, is not supported by OpenShift Container Platform.

CAUTION

Pay attention to system load on nodes where OpenShift Container Platform metrics pods run. Use that information to determine if it is necessary to scale out a number of OpenShift Container Platform metrics pods and spread the load across multiple OpenShift Container Platform nodes. Scaling OpenShift Container Platform metrics heapster pods is not recommended.

29.9.1. Prerequisites

If persistent storage was used to deploy OpenShift Container Platform metrics, then you must [create a persistent volume \(PV\)](#) for the new Cassandra pod to use before you can scale out the number of OpenShift Container Platform metrics Cassandra pods. However, if Cassandra was deployed with dynamically provisioned PVs, then this step is not necessary.

29.9.2. Scaling the Cassandra Components

The Cassandra nodes use persistent storage, therefore scaling up or down is not possible with replication controllers.

Scaling a Cassandra cluster requires you to use the **hawkular-cassandra-node** template. By default, the Cassandra cluster is a single-node cluster.

To scale up the number of OpenShift Container Platform metrics hawkular pods to two replicas, run:

```
# oc scale -n openshift-infra --replicas=2 rc hawkular-metrics
```

Alternatively, update your inventory file and re-run the [deployment](#).



NOTE

If you add a new node to or remove an existing node from a Cassandra cluster, the data stored in the cluster rebalances across the cluster.

To scale down:

1. If remotely accessing the container, run the following for the Cassandra node you want to remove:

```
$ oc exec -it <hawkular-cassandra-pod> nodetool decommission
```

If locally accessing the container, run the following instead:

```
$ oc rsh <hawkular-cassandra-pod> nodetool decommission
```

This command can take a while to run since it copies data across the cluster. You can monitor the decommission progress with **nodetool netstats -H**.

2. Once the previous command succeeds, scale down the **rc** for the Cassandra instance to **0**.

```
# oc scale -n openshift-infra --replicas=0 rc <hawkular-cassandra-rc>
```

This will remove the Cassandra pod.



IMPORTANT

If the scale down process completed and the existing Cassandra nodes are functioning as expected, you can also delete the **rc** for this Cassandra instance and its corresponding persistent volume claim (PVC). Deleting the PVC can permanently delete any data associated with this Cassandra instance, so if the scale down did not fully and successfully complete, you will not be able to recover the lost data.

29.10. CLEANUP

You can remove everything deployed by the metrics deployer by performing the following steps:

```
$ oc delete all,sa,templates,secrets,pvc --selector="metrics-infra"
```

To remove the deployer components, perform the following steps:

```
$ oc delete sa,secret metrics-deployer
```

CHAPTER 30. CUSTOMIZING THE WEB CONSOLE

30.1. OVERVIEW

Administrators can customize the [web console](#) using extensions, which let you run scripts and load custom stylesheets when the web console loads. Extension scripts allow you to override the default behavior of the web console and customize it for your needs.

For example, extension scripts can be used to add your own company's branding or to add company-specific capabilities. A common use case for this is rebranding or white-labelling for different environments. You can use the same extension code, but provide settings that change the web console. You can change the look and feel of nearly any aspect of the user interface in this way.

30.2. LOADING EXTENSION SCRIPTS AND STYLESHEETS

To add scripts and stylesheets, edit the [master configuration file](#). The scripts and stylesheet files must exist on the Asset Server and are added with the following options:

```
assetConfig:
  ...
  extensionScripts:
    - /path/to/script1.js
    - /path/to/script2.js
    - ...
  extensionStylesheets:
    - /path/to/stylesheet1.css
    - /path/to/stylesheet2.css
    - ...
```

NOTE

Wrap extension scripts in an Immediately Invoked Function Expression (IIFE). This ensures that you do not create global variables that conflict with the names used by the web console or by other extensions. For example:

```
(function() {
  // Put your extension code here...
})();
```

Relative paths are resolved relative to the master configuration file. To pick up configuration changes, restart the server.


Custom scripts and stylesheets are read once at server start time. To make developing extensions easier, you can reload scripts and stylesheets on every request by enabling development mode with the following setting:

```
assetConfig:
  ...
  extensionDevelopment: true
```

When set, the web console reloads any changes to existing extension script or stylesheet files when you refresh the page in your browser. You still must restart the server when adding new extension

stylesheets or scripts, however. This setting is only recommended for testing changes and not for production.

The examples in the following sections show common ways you can customize the web console.



NOTE

Additional extension examples are available in the [OpenShift Origin](#) repository on GitHub.

30.2.1. Setting Extension Properties

If you have a specific extension, but want to use different text in it for each of the environments, you can define the environment in the ***master-config.yaml*** file, and use the same extension script across environments. Pass settings from the ***master-config.yaml*** file to be used by the extension using the [extensionProperties](#) mechanism:

```
assetConfig:
  extensionDevelopment: true
  extensionProperties:
    doc_url: https://docs.openshift.com
    key1: value1
    key2: value2
  extensionScripts:
```

This results in a global variable that can be accessed by the extension, as if the following code was executed:

```
window.OPENSIFT_EXTENSION_PROPERTIES = {
  doc_url: "https://docs.openshift.com",
  key1: "value1",
  key2: "value2",
}
```

30.2.2. Customizing the Logo

The following style changes the logo in the web console header:

```
#header-logo {
  background-image: url("https://www.example.com/images/logo.png");
  width: 190px;
  height: 20px;
}
```

Replace the **example.com** URL with a URL to an actual image, and adjust the width and height. The ideal height is **20px**.

Save the style to a file (for example, **logo.css**) and add it to the master configuration file:

```
assetConfig:
  ...
  extensionStylesheets:
    - /path/to/logo.css
```


30.2.3. Changing Links to Documentation

Links to external documentation are shown in various sections of the web console. The following example changes the URL for two given links to the documentation:

```
window.OPENSIFT_CONSTANTS.HELP['get_started_cli'] =
  "https://example.com/doc1.html";
window.OPENSIFT_CONSTANTS.HELP['basic_cli_operations'] =
  "https://example.com/doc2.html";
```

Save this script to a file (for example, **help-links.js**) and add it to the master configuration file:

```
assetConfig:
  ...
  extensionScripts:
    - /path/to/help-links.js
```

30.2.4. Adding or Changing Links to Download the CLI

The **About** page in the web console provides download links for the [command line interface \(CLI\)](#) tools. These links can be configured by providing both the link text and URL, so that you can choose to point them directly to file packages, or to an external page that points to the actual packages.

For example, to point directly to packages that can be downloaded, where the link text is the package platform:

```
window.OPENSIFT_CONSTANTS.CLI = {
  "Linux (32 bits)": "https://<cdn>/openshift-client-tools-linux-32bit.tar.gz",
  "Linux (64 bits)": "https://<cdn>/openshift-client-tools-linux-64bit.tar.gz",
  "Windows":        "https://<cdn>/openshift-client-tools-windows.zip",
  "Mac OS X":       "https://<cdn>/openshift-client-tools-mac.zip"
};
```

Alternatively, to point to a page that links the actual download packages, with the **Latest Release** link text:

```
window.OPENSIFT_CONSTANTS.CLI = {
  "Latest Release": "https://<cdn>/openshift-client-tools/latest.html"
};
```

Save this script to a file (for example, **cli-links.js**) and add it to the master configuration file:

```
assetConfig:
  ...
  extensionScripts:
    - /path/to/cli-links.js
```

30.2.5. Customizing the About Page

To provide a custom **About** page for the web console:

1. Write an extension that looks like:

```
angular
  .module('aboutPageExtension', ['openshiftConsole'])
  .config(function($routeProvider) {
    $routeProvider
      .when('/about', {
        templateUrl: 'extensions/about/about.html',
        controller: 'AboutController'
      });
  })
  );
```

2. Save the script to a file (for example, **about/about.js**).
3. Write a customized template.
 - a. Start from the version of **about.html** from the OpenShift Container Platform [release](#) you are using. Within the template, there are two angular scope variables available: **version.master.openshift** and **version.master.kubernetes**.
 - b. Save the custom template to a file (for example, **about/about.html**).
 - c. Modify the master configuration file:

```
assetConfig:
  ...
  extensionScripts:
    - about/about.js
  ...
  extensions:
    - name: about
      sourceDirectory: /path/to/about
```

30.2.6. Configuring Navigation Menus


30.2.6.1. Top Navigation Dropdown Menus

The top navigation bar of the web console contains the help icon and the user dropdown menus. You can add additional menu items to these using the [angular-extension-registry](#).

The available extension points are:

- **nav-help-dropdown** - the help icon dropdown menu, visible at desktop screen widths
- **nav-user-dropdown** - the user dropdown menu, visible at desktop screen widths
- **nav-dropdown-mobile** - the single menu for top navigation items at mobile screen widths

The following example extends the **nav-help-dropdown** menu, with a name of **<myExtensionModule>**:

**NOTE**


`<myExtensionModule>` is a placeholder name. Each dropdown menu extension must be unique enough so that it does not clash with any future angular modules.

```
angular
  .module('<myExtensionModule>', ['openshiftConsole'])
  .run([
    'extensionRegistry',
    function(extensionRegistry) {
      extensionRegistry
        .add('nav-help-dropdown', function() {
          return [
            {
              type: 'dom',
              node: '<li><a href="http://www.example.com/report"
target="_blank">Report a Bug</a></li>'
            }, {
              type: 'dom',
              node: '<li class="divider"></li>' // If you want a
horizontal divider to appear in the menu
            }, {
              type: 'dom',
              node: '<li><a href="http://www.example.com/status"
target="_blank">System Status</a></li>'
            }
          ];
        });
    }
  ]);

hawtioPluginLoader.addModule('<myExtensionModule>');
```

30.2.6.2. Project Left Navigation

When navigating within a project, a menu appears on the left with primary and secondary navigation. This menu structure is defined as a constant and can be overridden or modified.

**NOTE**

Significant customizations to the project navigation may affect the user experience and should be done with careful consideration. You may need to update this customization in future upgrades if you modify existing navigation items.

1. Create the configuration scripts within a file (for example, ***navigation.js***):

```
// Append a new primary nav item. This is a simple direct
navigation item
// with no secondary menu.
window.OPENSIFT_CONSTANTS.PROJECT_NAVIGATION.push({
  label: "Dashboard", // The text label
  iconClass: "fa fa-dashboard", // The icon you want to appear
  href: "/dashboard" // Where to go when this nav item
is clicked.
```

```

// Relative URLs are pre-pended
with the path
// '/project/<project-name>'
});

// Splice a primary nav item to a specific spot in the list. This
// primary item has
// a secondary menu.
window.OPENSIFT_CONSTANTS.PROJECT_NAVIGATION.splice(2, 0, { //
  Insert at the third spot
    label: "Git",
    iconClass: "fa fa-code",
    secondaryNavSections: [ // Instead of an href, a sub-menu
can be defined
      {
        items: [
          {
            label: "Branches",
            href: "/git/branches",
            prefixes: [
              "/git/branches/" // Defines prefix URL patterns
that will cause
// this nav item to show the
active state, so
// tertiary or lower pages show
the right context
            ]
          }
        ]
      },
      {
        header: "Collaboration", // Sections within a sub-menu can
have an optional header
        items: [
          {
            label: "Pull Requests",
            href: "/git/pull-requests",
            prefixes: [
              "/git/pull-requests/"
            ]
          }
        ]
      }
    ]
  });

// Add a primary item to the top of the list. This primary item is
// shown conditionally.
window.OPENSIFT_CONSTANTS.PROJECT_NAVIGATION.unshift({
  label: "Getting Started",
  iconClass: "pficon pficon-screen",
  href: "/getting-started",
  prefixes: [ // Primary nav items can also
specify prefixes to trigger
    "/getting-started/" // active state
  ],

```

```

    isValid: function() {           // Primary or secondary items can
define an isValid                  // function. If present it will be
    return isNewUser;              // the item should be shown, it
called to test whether
    }                               // the item should be shown, it
    }                               // the item should be shown, it
  });

// Modify an existing menu item
var applicationsMenu =
_.find(window.OPENSIFT_CONSTANTS.PROJECT_NAVIGATION, { label:
'Applications' });
applicationsMenu.secondaryNavSections.push({ // Add a new secondary
nav section to the Applications menu
  // my secondary nav section
});

```

2. Save the file and add it to the master configuration at `/etc/origin/master/master-config.yml`:

```

assetConfig:
  ...
  extensionScripts:
    - /path/to/navigation.js

```

3. Restart the master host:

```
# systemctl restart atomic-openshift-master
```

30.2.7. Enabling Features in Technology Preview

Sometimes features are available in Technology Preview. By default, these features are disabled in the web console and hidden from end users.

Web console features currently in Technology Preview include:

Feature Name	Description	Script Value
Pipelines	Enabling this feature will add the Pipelines navigation item underneath the Builds menu.	pipelines

To enable a Technology Preview feature:

1. Save this script to a file (for example, `tech-preview.js`):

```

window.OPENSIFT_CONSTANTS.ENABLE_TECH_PREVIEW_FEATURE.
<feature_name> = true;

```

2. Add it to the master configuration file:

```
assetConfig:
```

```
...
extensionScripts:
  - /path/to/tech-preview.js
```

30.3. SERVING STATIC FILES

You can serve other files from the Asset Server as well. For example, you might want to make the CLI executable available for download from the web console or add images to use in a custom stylesheet.

Add the directory with the files you want using the following configuration option:

```
assetConfig:
  ...
  extensions:
    - name: images
      sourceDirectory: /path/to/my_images
```

The files under the **/path/to/my_images** directory will be available under the URL **/<context>/extensions/images** in the web console.

To reference these files from a stylesheet, you should generally use a relative path. For example:

```
#header-logo {
  background-image: url("../extensions/images/my-logo.png");
}
```

30.3.1. Enabling HTML5 Mode

The web console has a special mode for supporting certain static web applications that use the HTML5 history API:

```
assetConfig:
  ...
  extensions:
    - name: my_extension
      sourceDirectory: /path/to/myExtension
      html5Mode: true
```

Setting **html5Mode** to **true** enables two behaviors:

1. Any request for a non-existent file under **/<context>/extensions/my_extension/** instead serves **/path/to/myExtension/index.html** rather than a "404 Not Found" page.
2. The element **<base href="/">** will be rewritten in **/path/to/myExtension/index.html** to use the actual base depending on the asset configuration; only this exact string is rewritten.

This is needed for JavaScript frameworks such as AngularJS that require **base** to be set in **index.html**.

30.4. CUSTOMIZING THE LOGIN PAGE

You can also change the login page, and the login provider selection page for the web console. Run the following commands to create templates you can modify:

```
$ oadm create-login-template > login-template.html
$ oadm create-provider-selection-template > provider-selection-
template.html
```

Edit the file to change the styles or add content, but be careful not to remove any required parameters inside the curly brackets.

To use your custom login page or provider selection page, set the following options in the master configuration file:

```
oauthConfig:
  ...
  templates:
    login: /path/to/login-template.html
    providerSelection: /path/to/provider-selection-template.html
```

Relative paths are resolved relative to the master configuration file. You must restart the server after changing this configuration.

When there are multiple login providers configured or when the [alwaysShowProviderSelection](#) option in the *master-config.yaml* file is set to **true**, each time a user's token to OpenShift Container Platform expires, the user is presented with this custom page before they can proceed with other tasks.

30.4.1. Example Usage

Custom login pages can be used to create Terms of Service information. They can also be helpful if you use a third-party login provider, like GitHub or Google, to show users a branded page that they trust and expect before being redirected to the authentication provider.

30.5. CUSTOMIZING THE OAUTH ERROR PAGE

When errors occur during authentication, you can change the page shown.

1. Run the following command to create a template you can modify:

```
$ oadm create-error-template > error-template.html
```

2. Edit the file to change the styles or add content.
You can use the **Error** and **ErrorCode** variables in the template. To use your custom error page, set the following option in the master configuration file:

```
oauthConfig:
  ...
  templates:
    error: /path/to/error-template.html
```

Relative paths are resolved relative to the master configuration file.

3. You must restart the server after changing this configuration.

30.6. CHANGING THE LOGOUT URL

You can change the location a console user is sent to when logging out of the console by modifying the **logoutURL** parameter in the `/etc/origin/master/master-config.yaml` file:

```
...
assetConfig:
  logoutURL: "http://www.example.com"
...
```

This can be useful when authenticating with [Request Header](#) and OAuth or [OpenID](#) identity providers, which require visiting an external URL to destroy single sign-on sessions.

30.7. CONFIGURING WEB CONSOLE CUSTOMIZATIONS WITH ANSIBLE

During [advanced installations](#), many modifications to the web console can be configured using [the following parameters](#), which are configurable in the inventory file:

- [openshift_master_logout_url](#)
- [openshift_master_extension_scripts](#)
- [openshift_master_extension_stylesheets](#)
- [openshift_master_extensions](#)
- [openshift_master_oauth_template](#)
- [openshift_master_metrics_public_url](#)
- [openshift_master_logging_public_url](#)

Example Web Console Customization with Ansible

```
# Configure logoutURL in the master config for console customization
# See:
https://docs.openshift.com/enterprise/latest/install_config/web_console_customization.html#changing-the-logout-url
#openshift_master_logout_url=http://example.com

# Configure extensionScripts in the master config for console customization
# See:
https://docs.openshift.com/enterprise/latest/install_config/web_console_customization.html#loading-custom-scripts-and-stylesheets
#openshift_master_extension_scripts=
['/path/on/host/to/script1.js', '/path/on/host/to/script2.js']

# Configure extensionStylesheets in the master config for console customization
# See:
https://docs.openshift.com/enterprise/latest/install_config/web_console_customization.html#loading-custom-scripts-and-stylesheets
#openshift_master_extension_stylesheets=
['/path/on/host/to/styleSheet1.css', '/path/on/host/to/styleSheet2.css']
```



```
# Configure extensions in the master config for console customization
# See:
https://docs.openshift.com/enterprise/latest/install\_config/web\_console\_customization.html#serving-static-files
#openshift_master_extensions=[{'name': 'images', 'sourceDirectory':
'/path/to/my_images'}]

# Configure extensions in the master config for console customization
# See:
https://docs.openshift.com/enterprise/latest/install\_config/web\_console\_customization.html#serving-static-files
#openshift_master_oauth_template=/path/on/host/to/login-template.html

# Configure metricsPublicURL in the master config for cluster metrics.
Ansible is also able to configure metrics for you.
# See:
https://docs.openshift.com/enterprise/latest/install\_config/cluster\_metrics.html
#openshift_master_metrics_public_url=https://hawkular-
metrics.example.com/hawkular/metrics

# Configure loggingPublicURL in the master config for aggregate logging.
Ansible is also able to install logging for you.
# See:
https://docs.openshift.com/enterprise/latest/install\_config/aggregate\_logging.html
#openshift_master_logging_public_url=https://kibana.example.com
```

CHAPTER 31. REVISION HISTORY: INSTALLATION AND CONFIGURATION

31.1. WED MAR 07 2018

Affected Topic	Description of Change
Aggregating Container Logs	Added to instructions to scale EFK pods when changes are made in the Understanding and Adjusting the Deployment section.

31.2. FRI AUG 25 2017

Affected Topic	Description of Change
Setting up the Registry → Deploying a Registry on Existing Clusters	Removed Technology Preview notice from the Deploying the Registry as a DaemonSet section.

31.3. TUE AUG 08 2017

Affected Topic	Description of Change
Router → Using the Default HAProxy Router	Added new section on how to manually redeploy the router certificates .
Configuring the SDN	Added information about configuring Neutron to the Using Flannel section.

31.4. FRI JUL 28 2017

Affected Topic	Description of Change
Redeploying Certificates	Added the Redeploying a New etcd CA section.

31.5. THU JUL 27 2017

Affected Topic	Description of Change
Installing a Cluster → Disconnected Installation	Added the openshift/hello-openshift dependency to the Syncing Repositories section.
Configuring Authentication and User Agent	Added a note about whitelisting users.

Affected Topic	Description of Change
Installing a Cluster → Advanced Installation	Updated Section 2.6.11, “Known Issues” Known Issues] to remove yum remove steps and instead link to Uninstalling OpenShift Container Platform .
Setting up the Registry → Securing and Exposing the Registry	Added steps to the Securing the Registry section on configuring your cluster to trust the certificate at the OS level.

31.6. WED JUL 12 2017

Affected Topic	Description of Change
Installing a Cluster → Host Preparation	Replaced the deprecated openshift_master_portal_net variable with openshift_portal_net .
Installing a Cluster → Advanced Installation	Replaced the deprecated openshift_master_portal_net variable with openshift_portal_net .

31.7. TUE JUN 13 2017

Affected Topic	Description of Change
Installing a Cluster → Advanced Installation	Added a step to verify that the etcd package is installed, if you installed multiple etcd hosts.

31.8. WED MAY 31 2017

Affected Topic	Description of Change
Syncing Groups With LDAP	Added Nested Membership Sync Example .
Installing a Cluster → Installing a Stand-alone Deployment of OpenShift Container Registry	Updated to use OpenShift Container Registry name and add clarification on the distinction between Atomic Registry.

31.9. THU MAY 25 2017

Affected Topic	Description of Change
Configuring Authentication and User Agent	Noted that after making changes to an identity provider, you must restart the master service for the changes to take effect.

31.10. MON MAY 15 2017

Affected Topic	Description of Change
Upgrading a Cluster → Overview	Added clarification that nodes and masters are forward and backward compatible across one minor version.
Master and Node Configuration	Added information about <code>openshift_master_audit_config</code> to the Audit Configuration section.

31.11. TUE MAY 09 2017

Affected Topic	Description of Change
Installing a Cluster → Advanced Installation	Updated the Configuring Schedulability on Masters section to note that masters are automatically marked unschedulable by default by the installer.
	Updated the Configuring Node Host Labels section to better describe the special region=infra label and to suggest configuring dedicated infrastructure nodes.
	Added a Configuring Deployment Type section.

31.12. TUE APR 25 2017

Affected Topic	Description of Change
Redeploying Certificates	Updated for new set of playbooks and options.

31.13. WED APR 12 2017

Affected Topic	Description of Change
Setting up a Router → Using the Default HAProxy Router	Added graphics to the Using Router Shards section.

Affected Topic	Description of Change
Setting up a Router → Deploying a Customized HAProxy Router	Expanded details in the Rebuilding Your Router section.
Installing a Cluster → Prerequisites	Specified the UDP for port 4789.
Installing → Advanced Installation	In the Known Issues multiple masters discussion, included the docker-common package in the removal process, following a failed setup play.
Configuring for OpenStack	Added openshift_cloudprovider_openstack_domain_id and openshift_cloudprovider_openstack_domain_name to the list of configurable parameters.

31.14. MON APR 03 2017

Affected Topic	Description of Change
Advanced Installation	Added the debug_level parameter to the Configuring Cluster Variables section for setting the default verbosity of journal log messages.
Master and Node Configuration	Added a new section on configuring the verbosity of journal log messages .
Setting Up a Registry → Extended Registry Configuration	Updated the list of supported registry storage drivers in the Storage section.
Setting Up a Registry → Deploying a Registry on Existing Clusters	Added list of supported registry storage drivers to the Storage for the Registry section.
Redeploying Certificates	Added Registry and Router Certificates section with instructions on redeploying these certificates manually.

31.15. MON MAR 27 2017

Affected Topic	Description of Change
Downgrading OpenShift	Moved the procedure for restoring etcd to the Backup and Restore topic.

Affected Topic	Description of Change
Setting up a Router → Configuring the HAProxy Router to Use the PROXY Protocol	New topic on configuring the HAProxy router to use the PROXY protocol.

31.16. MON MAR 20 2017

Affected Topic	Description of Change
Installing a Cluster → Prerequisites	Added Network Time Protocol (NTP) as a prerequisite.
Configuring Authentication and User Agent	Added information on using the htpasswd command to generate hashed passwords for HTPasswd authentication to the HTPasswd section.
Installing a Stand-alone Registry	Added openshift_master_default_subdomain to the openshift_master_default_subdomain file in the Advanced Installation for Stand-alone Registries section.

31.17. TUE MAR 14 2017

Affected Topic	Description of Change
Installing a Cluster → Prerequisites	Renamed instances of openshift_node_set_node_ip to openshift_set_node_ip , the correct openshift-ansible variable name.
Installing a Cluster → Disconnected Installation	Added information about disconnected installation of the stand-alone registry.
Setting up the Registry → Deploying a Registry on Existing Clusters	Added footnotes to the example file in the Use Amazon S3 as a Storage Back-end section.
Setting up the Registry → Known Issues	Added note recommending moving off NFS for production in the Image Push Errors with Scaled Registry Using Shared NFS Volume section.
Upgrading a Cluster → Performing Manual In-place Cluster Upgrades	Removed a repetitive step within the Updating the Default Image Streams and Templates section.
Upgrading a Cluster → Operating System Updates and Upgrades	Added the procedure for upgrading the operating system.

31.18. TUE MAR 07 2017

Affected Topic	Description of Change
Installing a Cluster → Advanced Installation	Updated Before You Begin section to raise minimal Ansible version to 2.2.0.
	Provided guidance for preconfigured load balancers for OpenShift Container Platform with high availability.
Redeploying Certificates	Added the Checking Certificate Expirations section.

31.19. THU FEB 16 2017

Affected Topic	Description of Change
Installing a Cluster → Disconnected Installation	Added the cluster image to the Syncing Images section.
Setting up the Registry → Deploying a Registry on Existing Clusters	Added additional URL instructions to the Deploying the Registry Console section.
Setting up the Registry → Securing and Exposing the Registry	Added a step for adding the public route host name in the -hostnames flag.
Setting up a Router → Using the Default HAProxy Router	Added the Filtering Routes to Specific Routers section.
Master and Node Configuration	Fixed the options for creating a configuration file in the Creating New Configuration Files section.
Configuring Persistent Storage → Volume Security	Added details about RunAsAny FSGroup and block device permissions.
Aggregating Container Logs	Fixed example in the Configuring Fluentd to Send Logs to an External Log Aggregator section.
	Added a version variable and <tag> to code block example in Deploying the EFK Stack section to display the correct current version to use.

31.20. MON FEB 06 2017

Affected Topic	Description of Change
Setting up a Router → Using the Default HAProxy Router	Arranged the topic to create the Creating a Router section, and added a paragraph on router options on creation.
Persistent Storage Examples → Complete Example Using GlusterFS	Clarified that, if using a service, the endpoints name must match the service name.
Configuring Persistent Storage → Persistent Storage Using GlusterFS	Clarified that, if using a service, the endpoints name must match the service name.
Setting up the Registry → Deploying a Registry on Existing Clusters	Added Important box about shutting down Cockpit to the Non-Production Use section.
Installing a Cluster → Host Preparation	Added steps on using yum-config-manager to the host registration steps.
Setting up the Registry → Deploying a Registry on Existing Clusters	Arranged the Securing the Registry Console section to include information on the certificate.
Installing a Cluster → Advanced Install	Added the Configuring a Registry Location section.
Configuring Nuage SDN	Added the Configuring Nuage SDN file.

31.21. TUE JAN 31 2017

Affected Topic	Description of Change
Installing a Cluster → Host Preparation	Added instructions for installing and using the atomic-openshift-excluder and atomic-openshift-docker-excluder scripts during cluster installations and upgrades.
Installing a Cluster → Quick Installation	
Installing a Cluster → Advanced Installation	
Upgrading a Cluster → Manual In-place Upgrades	

Affected Topic	Description of Change
Upgrading a Cluster → Automated In-place Upgrades	

31.22. MON JAN 30 2017

Affected Topic	Description of Change
Setting up the Registry → Securing and Exposing the Registry	Removed references to the deprecated --api-version flag.
Configuring Custom Certificates	Clarified custom certificate configuration locations in the Configuring Custom Certificates section.

31.23. WED JAN 25 2017

Affected Topic	Description of Change
Working with HTTP Proxies	Added step to Proxying Docker Pull for finding the registry service IP.
Setting up a Router → Using the F5 Router Plug-in	Removed references to the deprecated --credentials option.
Installing a Cluster → Prerequisites	Added information about required ports for Aggregated Logging.
Configuring Global Build Defaults and Overrides	Added notes to explain additional values in the /etc/origin/master/master-config.yaml file in the Manually Setting Global Build Defaults section.
Customizing the Web Console	Added information about setting extension properties .

31.24. WED JAN 18 2017

Affected Topic	Description of Change
Setting up the Registry → Securing and Exposing the Registry	Added note box about mounting secrets to service accounts.

31.25. MON JAN 16 2017

Affected Topic	Description of Change
Configuring Authentication and User Agent	Clarified the difference between <code>/api</code> and <code>/oapi</code> in the User Agent section.
Aggregating Container Logs	Added clarification regarding ConfigMaps and output of <code>oc new-app</code> .

31.26. MON JAN 09 2017

Affected Topic	Description of Change
Working with HTTP Proxies	Added clarifying details about HTTP proxies.

31.27. TUE DEC 20 2016

Affected Topic	Description of Change
Working with HTTP Proxies	Removed section on configuring Maven with http proxies.
Installing a Cluster → Host Preparation	Updated the path to the latest epel-release package.

31.28. TUE DEC 13 2016

Affected Topic	Description of Change
Configuring Persistent Storage → Persistent Storage Using GCE Persistent Disk	Added a new Multi-zone Configuration section.
Configuring Persistent Storage → Dynamically Provisioning Persistent Volumes	Added clarifying details about multi-zone persistent volume (PV) configuration.
Aggregating Container Logs	Added Configuring Fluentd to Send Logs to an External Log Aggregator section.
	Added Note boxes explicitly stating that sending logs directly to an AWS Elasticsearch instance is not supported.

Affected Topic	Description of Change
	Added F-5 Load Balancer and X-Forwarded-For Enabled to the Troubleshooting Kibana section.
Customizing the Web Console	Added a new Customizing the About Page section.

31.29. MON DEC 05 2016

Affected Topic	Description of Change
Installing a Cluster → Prerequisites	Removed Git access as a prerequisite because it is a requirement for development, but not for installing a cluster.
Aggregating Container Logs	Added a NOTE indicating that, as of OpenShift Container Platform 3.3, Fluentd no longer reads historical log files when using the JSON file log driver.
Installing a Cluster → Host Preparation	Added the matches option to the subscription-manager list --available command.
Installing a Cluster → Disconnected Installation	Added the matches option to the subscription-manager list --available command.
Syncing Groups With LDAP	Removed references to filtering UIDAttribute values in configurations.

31.30. MON NOV 21 2016

Affected Topic	Description of Change
Aggregating Container Logs	Clarified points around NFS storage.
Setting up the Registry → Deploying a Registry on Existing Clusters	Removed inaccurate example showing how to attach an existing NFS volume to the registry.
Setting up the Registry → Extended Registry Configuration	Described the CloudFront middleware extension .

31.31. MON NOV 14 2016

Affected Topic	Description of Change
Aggregating Container Logs	Updated the admin-cert location in an example within the Performing Administrative Elasticsearch Operations section.
Installing a Cluster → Advanced Installation	Added steps to verify the web console.

31.32. MON NOV 07 2016

Affected Topic	Description of Change
Upgrading a Cluster → Operating System Updates and Upgrades	New topic on the impacts of operating system updates and upgrades and possible solutions.
Setting up the Registry → Securing and Exposing the Registry	Added clarification to the Exposing a Non-Secure Registry section.

31.33. TUE NOV 01 2016

Affected Topic	Description of Change
Installing → Planning	Updated the Sizing Considerations section for clarity.

31.34. THU OCT 27 2016

OpenShift Container Platform 3.3.1 release.

Affected Topic	Description of Change
Upgrading a Cluster → Performing Automated In-place Cluster Upgrades	Minor updates for clarity, including converting some procedures into numbered steps for easier readability.

31.35. MON OCT 17 2016

Affected Topic	Description of Change
Configuring Pipeline Execution	Clarified Jenkins template names.

Affected Topic	Description of Change
Loading the Default Image Streams and Templates	Updated information in the Offerings by Subscription Type section on which images are provided by which subscription s.
Installing a Cluster → Advanced Installation	Added more information to the openshift_master_portal_net parameter description in the Configuring Cluster Variables section.

31.36. TUE OCT 11 2016

Affected Topic	Description of Change
Aggregating Container Logs	Fixed error in Deploying the EFK Stack section.
Configuring the SDN	Added clarifying details to the Migrating Between SDN Plug-ins section about when to clean up SDN-specific artifacts.
Configuring Persistent Storage → Persistent Storage Using Ceph Rados Block Device (RBD)	Updated the persistentVolumeReclaimPolicy setting to retain in the Persistent Volume Object Definition Using Ceph RBD example.
Installing → Advanced Installation	Replaced ansible_sudo with ansible_become .

31.37. TUE OCT 04 2016

Affected Topic	Description of Change
Configuring the SDN	Added clarifying details to the Migrating Between SDN Plug-ins section about when to clean up SDN-specific artifacts.
Configuring the SDN	Added that oc get netnamespace can be run to check VNIDs.
Setting up the Registry → Known Issues	Added troubleshooting guidance on Image Pruning Failures .
Installing → Prerequisites	Added information about disabling dnsmasq .
Installing → Advanced Installation	Added example for a multi-master install with etcd on the same hosts.

Affected Topic	Description of Change
Configuring Persistent Storage → Persistent Storage Using Ceph Rados Block Device (RBD)	Updated the persistentVolumeReclaimPolicy setting to retain in the Persistent Volume Object Definition Using Ceph RBD example .
Persistent Storage Examples → Binding Persistent Volumes by Labels	Updated the persistentVolumeReclaimPolicy setting to retain in the glusterfs-pv.yaml example , since recycle is not supported in this case.
Persistent Storage Examples → Complete Example Using GlusterFS	Updated the GlusterFS persistent storage example to use NGNIX instead of busybox.
Configuring Persistent Storage → Volume Security	Fixed formatting of the oc get project default -o yaml example output within the SCCs, Defaults, and Allowed Ranges section.
Configuring Persistent Storage → Volume Security	Removed no_root_squash from the NFS example, as it is not a recommended option.

31.38. TUE SEP 27 2016

OpenShift Container Platform 3.3 initial release.

Affected Topic	Description of Change
Configuring the SDN	Added that oc get netnamespace can be run to check VNIDs.
Setting up the Registry → Securing and Exposing the Registry	Added two new sections on Exposing a Secure Registry and Exposing a Non-Secure Registry.
Customizing the Web Console	Added Configuring Navigation Menus section.
Setting up the Registry → Known Issues	Added troubleshooting guidance on Image Pruning Failures .
Master and Node Configuration	Added a Audit Configuration section.
Installing → Prerequisites	Added information about disabling dnsmasq .

Affected Topic	Description of Change
Redeploying Certificates	New topic reviewing how to back up and redeploy cluster certificates using the ansible-playbook command.
Installing → Advanced Installation	Added example for a multi-master install with etcd on the same hosts.
Enabling Cluster Metrics	Added capacity planning guidance for OpenShift Container Platform metrics.
Installing → Prerequisites	Updated scale recommendations.
Installing → Advanced Installation	Updated the Multiple Masters Using HAProxy Inventory File example with guidance on applying updated node defaults.
Upgrading → Performing Manual Cluster Upgrades	Updated version numbers for image streams across the Updating the Default Image Streams and Templates section.
Persistent Storage Examples → Binding Persistent Volumes by Labels	Updated the persistentVolumeReclaimPolicy setting to retain in the glusterfs-pv.yaml example , since recycle is not supported in this case.
Persistent Storage Examples → Complete Example Using GlusterFS	Updated the GlusterFS persistent storage example to use NGNIX instead of busybox.
Configuring Pipeline Execution	Corrected instructions for enabling Jenkins auto-provision.
Configuring Routing	Changed "Native Container Routing" topic to "Configuring Routing" and added information about Configuring Route Timeouts.
Aggregating Container Logs	Added clarifying details to the Warning box in the Persistent Elasticsearch Storage section regarding the NFS workaround.
Upgrading → Performing Manual Cluster Upgrades	Added a new Update Your Configuration File section.
Setting up the Registry → Extended Registry Configuration	Emphasized the new mandatory middleware configuration entries.

Affected Topic	Description of Change
Deploying a Docker Registry	Extended the registry configuration file example within the Deploying Updated Configuration section to include the blobrepositorycachettl option.
Storage Examples → Binding Persistent Volumes by Labels	New topic providing an end-to-end example for binding persistent volume claims (PVCs) to persistent volumes (PVs) by defining labels in the PV and matching selectors in the PVC.
Persistent Storage Examples → Selector-Label Volume Binding	New topic outlining how to bind persistent volumes claims (PVCs) to persistent volumes (PVs) via selector and label attributes.
Upgrading → Blue-Green Deployments	Added new topic.
Enabling Cluster Metrics	Added additional details to the Accessing Hawkular Metrics Directly section.
Installing → Deploying a Router	Added a new Protecting Against DDoS Attacks section.
Configuring Pipeline Execution	New section.
Installing → Prerequisites	Added that the deserialization cache size can be reduced using a setting in master-config.yaml .
Aggregating Container Logs	Added information about configuration from configmaps, Fluentd, and Curator.
Installing → Deploying a Docker Registry	Edited references to oc secrets add .
Configuring Persistent Storage → Volume Security	Fixed formatting of the oc get project default -o yaml example output within the SCCs, Defaults, and Allowed Ranges section.
Configuring Authentication	Updated OAuth grant strategies information.
Installing → Deploying a Docker Registry	Updated mandatory configuration options for the registry's configuration file.
Configuring the SDN	Updated migration steps for SDN plug-ins.
Performing Manual Cluster Upgrades	Added a Warning box about excluding roles from reconciliation.

Affected Topic	Description of Change
Configuring Authentication	Added OpenID and GitLab challenge options.
Enabling Cluster Metrics	Added a new Metrics Deployer Validations section.
Upgrading → Performing Manual Cluster Upgrades	Added recent image quota restrictions.