



OpenShift Container Platform 4.12

Backup and restore

Backing up and restoring your OpenShift Container Platform cluster

OpenShift Container Platform 4.12 Backup and restore

Backing up and restoring your OpenShift Container Platform cluster

Legal Notice

Copyright © 2024 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, the Red Hat 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 is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

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

All other trademarks are the property of their respective owners.

Abstract

This document provides instructions for backing up your cluster's data and for recovering from various disaster scenarios.

Table of Contents

CHAPTER 1. BACKUP AND RESTORE	9
1.1. CONTROL PLANE BACKUP AND RESTORE OPERATIONS	9
1.2. APPLICATION BACKUP AND RESTORE OPERATIONS	9
1.2.1. OADP requirements	9
1.2.2. Backing up and restoring applications	10
CHAPTER 2. SHUTTING DOWN THE CLUSTER GRACEFULLY	12
2.1. PREREQUISITES	12
2.2. SHUTTING DOWN THE CLUSTER	12
2.3. ADDITIONAL RESOURCES	14
CHAPTER 3. RESTARTING THE CLUSTER GRACEFULLY	15
3.1. PREREQUISITES	15
3.2. RESTARTING THE CLUSTER	15
CHAPTER 4. OADP APPLICATION BACKUP AND RESTORE	18
4.1. INTRODUCTION TO OPENSIFT API FOR DATA PROTECTION	18
4.1.1. OpenShift API for Data Protection APIs	18
4.2. OADP RELEASE NOTES	18
4.2.1. OADP 1.3 release notes	18
4.2.1.1. OADP 1.3.1 release notes	18
4.2.1.1.1. New features	18
4.2.1.1.2. Resolved issues	18
4.2.1.1.3. Known issues	20
4.2.1.2. OADP 1.3.0 release notes	20
4.2.1.2.1. New features	20
4.2.1.2.2. Resolved issues	21
4.2.1.2.3. Known issues	22
4.2.1.2.4. Upgrade notes	22
4.2.1.2.4.1. Changes from OADP 1.2 to 1.3	22
4.2.1.2.4.2. Upgrading steps	23
4.2.1.2.4.3. Upgrading from OADP 1.2 Technology Preview Data Mover	23
4.2.1.2.4.4. Backing up the DPA configuration	24
4.2.1.2.4.5. Upgrading the OADP Operator	24
4.2.1.2.4.6. Converting DPA to the new version	24
4.2.1.2.4.7. Verifying the upgrade	25
4.2.2. OADP 1.2 release notes	27
4.2.2.1. OADP 1.2.4 release notes	27
4.2.2.1.1. Resolved issues	27
4.2.2.1.2. Known issues	27
4.2.2.2. OADP 1.2.3 release notes	28
4.2.2.2.1. New features	28
4.2.2.2.2. Resolved issues	28
4.2.2.2.3. Known issues	28
4.2.2.3. OADP 1.2.2 release notes	28
4.2.2.3.1. New features	28
4.2.2.3.2. Resolved issues	28
4.2.2.3.3. Known issues	29
4.2.2.4. OADP 1.2.1 release notes	30
4.2.2.4.1. New features	30
4.2.2.4.2. Resolved issues	30
4.2.2.4.3. Known issues	30

4.2.2.5. OADP 1.2.0 release notes	30
4.2.2.5.1. New features	31
4.2.2.5.1.1. Technical preview features	31
4.2.2.5.2. Resolved issues	31
4.2.2.5.3. Known issues	31
4.2.2.5.4. Upgrade notes	32
4.2.2.5.4.1. Changes from OADP 1.1 to 1.2	32
4.2.2.5.5. Upgrading steps	32
4.2.2.5.5.1. Backing up the DPA configuration	32
4.2.2.5.5.2. Upgrading the OADP Operator	32
4.2.2.5.5.3. Converting DPA to the new version	33
4.2.2.5.5.4. Verifying the upgrade	34
4.2.3. OADP 1.1 release notes	35
4.2.3.1. {oadp-short} 1.1.8 release notes	35
4.2.3.1.1. Known issues	35
4.2.3.2. OADP 1.1.7 release notes	35
4.2.3.2.1. Resolved issues	35
4.2.3.2.2. Known issues	35
4.2.3.3. OADP 1.1.6 release notes	36
4.2.3.3.1. Resolved issues	36
4.2.3.3.2. Known issues	36
4.2.3.4. OADP 1.1.5 release notes	36
4.2.3.4.1. New features	36
4.2.3.4.2. Resolved issues	36
4.2.3.4.3. Known issues	36
4.2.3.5. OADP 1.1.4 release notes	36
4.2.3.5.1. New features	36
4.2.3.5.2. Resolved issues	36
4.2.3.5.3. Known issues	37
4.2.3.6. OADP 1.1.3 release notes	37
4.2.3.6.1. New features	37
4.2.3.6.2. Resolved issues	38
4.2.3.6.3. Known issues	38
4.2.3.7. OADP 1.1.2 release notes	38
4.2.3.7.1. Product recommendations	38
4.2.3.7.2. Resolved issues	38
4.2.3.7.3. Known issues	38
4.2.3.8. OADP 1.1.1 release notes	39
4.2.3.8.1. Product recommendations	39
4.2.3.8.2. Known issues	39
4.3. OADP FEATURES AND PLUGINS	39
4.3.1. OADP features	39
4.3.2. OADP plugins	40
4.3.3. About OADP Velero plugins	41
4.3.3.1. Default Velero cloud provider plugins	42
4.3.3.2. Custom Velero plugins	42
4.3.3.3. Velero plugins returning "received EOF, stopping recv loop" message	43
4.3.4. Supported architectures for OADP	43
4.3.5. OADP support for IBM Power and IBM Z	43
4.3.5.1. OADP support for target backup locations using IBM Power	43
4.3.5.2. OADP testing and support for target backup locations using IBM Z	43
4.3.6. OADP plugins known issues	44
4.3.6.1. Velero plugin panics during imagestream backups due to a missing secret	44

4.3.6.1.1. Workaround to avoid the panic error	44
4.3.6.2. OpenShift ADP Controller segmentation fault	44
4.3.6.2.1. OpenShift ADP Controller segmentation fault workaround	45
4.4. INSTALLING AND CONFIGURING OADP	45
4.4.1. About installing OADP	45
4.4.1.1. AWS S3 compatible backup storage providers	46
4.4.1.1.1. Supported backup storage providers	46
4.4.1.1.2. Unsupported backup storage providers	47
4.4.1.1.3. Backup storage providers with known limitations	47
4.4.1.2. Configuring Multicloud Object Gateway (MCG) for disaster recovery on OpenShift Data Foundation	47
4.4.1.3. About OADP update channels	48
4.4.1.4. Installation of OADP on multiple namespaces	49
4.4.1.5. Velero CPU and memory requirements based on collected data	49
4.4.1.5.1. CPU and memory requirement for configurations	49
4.4.1.5.2. NodeAgent CPU for large usage	50
4.4.2. Installing the OADP Operator	51
4.4.2.1. OADP-Velero-OpenShift Container Platform version relationship	51
4.4.3. Configuring the OpenShift API for Data Protection with Amazon Web Services	52
4.4.3.1. Configuring Amazon Web Services	52
4.4.3.2. About backup and snapshot locations and their secrets	55
Backup locations	55
Snapshot locations	55
Secrets	55
4.4.3.2.1. Creating a default Secret	55
4.4.3.2.2. Creating profiles for different credentials	56
4.4.3.3. Configuring the Data Protection Application	57
4.4.3.3.1. Setting Velero CPU and memory resource allocations	57
4.4.3.3.2. Enabling self-signed CA certificates	58
4.4.3.4. Installing the Data Protection Application	59
4.4.3.4.1. Enabling CSI in the DataProtectionApplication CR	61
4.4.4. Configuring the OpenShift API for Data Protection with Microsoft Azure	62
4.4.4.1. Configuring Microsoft Azure	62
4.4.4.2. About backup and snapshot locations and their secrets	64
Backup locations	64
Snapshot locations	64
Secrets	64
4.4.4.2.1. Creating a default Secret	65
4.4.4.2.2. Creating secrets for different credentials	65
4.4.4.3. Configuring the Data Protection Application	66
4.4.4.3.1. Setting Velero CPU and memory resource allocations	67
4.4.4.3.2. Enabling self-signed CA certificates	67
4.4.4.4. Installing the Data Protection Application	68
4.4.4.4.1. Enabling CSI in the DataProtectionApplication CR	71
4.4.5. Configuring the OpenShift API for Data Protection with Google Cloud Platform	71
4.4.5.1. Configuring Google Cloud Platform	72
4.4.5.2. About backup and snapshot locations and their secrets	73
Backup locations	74
Snapshot locations	74
Secrets	74
4.4.5.2.1. Creating a default Secret	74
4.4.5.2.2. Creating secrets for different credentials	75
4.4.5.3. Configuring the Data Protection Application	76

4.4.5.3.1. Setting Velero CPU and memory resource allocations	76
4.4.5.3.2. Enabling self-signed CA certificates	77
4.4.5.4. Installing the Data Protection Application	77
4.4.5.4.1. Enabling CSI in the DataProtectionApplication CR	80
4.4.6. Configuring the OpenShift API for Data Protection with Multicloud Object Gateway	81
4.4.6.1. Retrieving Multicloud Object Gateway credentials	81
4.4.6.2. About backup and snapshot locations and their secrets	82
Backup locations	82
Snapshot locations	82
Secrets	82
4.4.6.2.1. Creating a default Secret	82
4.4.6.2.2. Creating secrets for different credentials	83
4.4.6.3. Configuring the Data Protection Application	84
4.4.6.3.1. Setting Velero CPU and memory resource allocations	84
4.4.6.3.2. Enabling self-signed CA certificates	85
4.4.6.4. Installing the Data Protection Application	86
4.4.6.4.1. Enabling CSI in the DataProtectionApplication CR	88
4.4.7. Configuring the OpenShift API for Data Protection with OpenShift Data Foundation	89
4.4.7.1. About backup and snapshot locations and their secrets	89
Backup locations	89
Snapshot locations	90
Secrets	90
4.4.7.1.1. Creating a default Secret	90
4.4.7.1.2. Creating secrets for different credentials	91
4.4.7.2. Configuring the Data Protection Application	92
4.4.7.2.1. Setting Velero CPU and memory resource allocations	92
4.4.7.2.1.1. Adjusting Ceph CPU and memory requirements based on collected data	93
4.4.7.2.1.1.1. CPU and memory requirement for configurations	93
4.4.7.2.2. Enabling self-signed CA certificates	93
4.4.7.3. Installing the Data Protection Application	94
4.4.7.3.1. Creating an Object Bucket Claim for disaster recovery on OpenShift Data Foundation	97
4.4.7.3.2. Enabling CSI in the DataProtectionApplication CR	97
4.5. UNINSTALLING OADP	98
4.5.1. Uninstalling the OpenShift API for Data Protection	98
4.6. OADP BACKING UP	98
4.6.1. Backing up applications	98
4.6.1.1. Known issues	99
4.6.2. Creating a Backup CR	99
4.6.3. Backing up persistent volumes with CSI snapshots	100
4.6.4. Backing up applications with File System Backup: Kopia or Restic	101
4.6.5. Creating backup hooks	102
4.6.6. Scheduling backups using Schedule CR	104
4.6.7. Deleting backups	105
4.6.8. About Kopia	106
4.6.8.1. OADP integration with Kopia	106
4.7. OADP RESTORING	107
4.7.1. Restoring applications	107
4.7.1.1. Creating a Restore CR	107
4.7.1.2. Creating restore hooks	109
4.8. OADP AND ROSA	111
4.8.1. Backing up applications on ROSA clusters using OADP	111
4.8.1.1. Preparing AWS credentials for OADP	112
4.8.1.2. Installing the OADP Operator and providing the IAM role	114

4.8.1.3. Example: Backing up workload on OADP ROSA STS, with an optional cleanup	119
4.8.1.3.1. Performing a backup with OADP and ROSA STS	119
4.8.1.3.2. Cleaning up a cluster after a backup with OADP and ROSA STS	122
4.9. OADP DATA MOVER	123
4.9.1. OADP Data Mover Introduction	123
4.9.1.1. OADP Data Mover prerequisites	124
4.9.2. Using Data Mover for CSI snapshots	124
4.9.3. Using OADP 1.2 Data Mover with Ceph storage	129
4.9.3.1. Prerequisites for using OADP 1.2 Data Mover with Ceph storage	130
4.9.3.2. Defining custom resources for use with OADP 1.2 Data Mover	130
4.9.3.2.1. Defining CephFS custom resources for use with OADP 1.2 Data Mover	130
4.9.3.2.2. Defining CephRBD custom resources for use with OADP 1.2 Data Mover	131
4.9.3.2.3. Defining additional custom resources for use with OADP 1.2 Data Mover	132
4.9.3.3. Backing up and restoring data using OADP 1.2 Data Mover and CephFS storage	133
4.9.3.3.1. Creating a DPA for use with CephFS storage	134
4.9.3.3.2. Backing up data using OADP 1.2 Data Mover and CephFS storage	135
4.9.3.3.3. Restoring data using OADP 1.2 Data Mover and CephFS storage	136
4.9.3.4. Backing up and restoring data using OADP 1.2 Data Mover and split volumes (CephFS and Ceph RBD)	137
4.9.3.4.1. Creating a DPA for use with split volumes	137
4.9.3.4.2. Backing up data using OADP 1.2 Data Mover and split volumes	138
4.9.3.4.3. Restoring data using OADP 1.2 Data Mover and split volumes	139
4.9.4. Cleaning up after a backup using OADP 1.1 Data Mover	140
4.9.4.1. Deleting snapshots in a bucket	140
4.9.4.2. Deleting cluster resources	140
4.9.4.2.1. Deleting cluster resources following a successful backup and restore that used Data Mover	141
4.9.4.2.2. Deleting cluster resources following a partially successful or a failed backup and restore that used Data Mover	141
4.10. OADP 1.3 DATA MOVER	142
4.10.1. About the OADP 1.3 Data Mover	142
4.10.1.1. Enabling the built-in Data Mover	143
4.10.1.2. Built-in Data Mover controller and custom resource definitions (CRDs)	143
4.10.2. Backing up and restoring CSI snapshots	143
4.10.2.1. Backing up persistent volumes with CSI snapshots	144
4.10.2.2. Restoring CSI volume snapshots	146
4.11. TROUBLESHOOTING	147
4.11.1. Downloading the Velero CLI tool	147
4.11.1.1. OADP-Velero-OpenShift Container Platform version relationship	148
4.11.2. Accessing the Velero binary in the Velero deployment in the cluster	149
4.11.3. Debugging Velero resources with the OpenShift CLI tool	149
Velero CRs	149
Velero pod logs	149
Velero pod debug logs	149
4.11.4. Debugging Velero resources with the Velero CLI tool	150
Syntax	150
Help option	150
Describe command	150
Logs command	151
4.11.5. Pods crash or restart due to lack of memory or CPU	151
4.11.5.1. Setting resource requests for a Velero pod	151
4.11.5.2. Setting resource requests for a Restic pod	152
4.11.6. Issues with Velero and admission webhooks	153
4.11.6.1. Restoring workarounds for Velero backups that use admission webhooks	153

4.11.6.1.1. Restoring Knative resources	153
4.11.6.1.2. Restoring IBM AppConnect resources	153
4.11.6.2. OADP plugins known issues	154
4.11.6.2.1. Velero plugin panics during imagestream backups due to a missing secret	154
4.11.6.2.1.1. Workaround to avoid the panic error	154
4.11.6.2.2. OpenShift ADP Controller segmentation fault	155
4.11.6.2.2.1. OpenShift ADP Controller segmentation fault workaround	155
4.11.6.3. Velero plugins returning "received EOF, stopping recv loop" message	155
4.11.7. Installation issues	155
4.11.7.1. Backup storage contains invalid directories	155
4.11.7.2. Incorrect AWS credentials	156
4.11.8. OADP Operator issues	156
4.11.8.1. OADP Operator fails silently	156
4.11.9. OADP timeouts	157
4.11.9.1. Restic timeout	158
4.11.9.2. Velero resource timeout	158
4.11.9.3. Data Mover timeout	159
4.11.9.4. CSI snapshot timeout	159
4.11.9.5. Velero default item operation timeout	160
4.11.9.6. Item operation timeout - restore	161
4.11.9.7. Item operation timeout - backup	161
4.11.10. Backup and Restore CR issues	162
4.11.10.1. Backup CR cannot retrieve volume	162
4.11.10.2. Backup CR status remains in progress	162
4.11.10.3. Backup CR status remains in PartiallyFailed	162
4.11.11. Restic issues	163
4.11.11.1. Restic permission error for NFS data volumes with root_squash enabled	163
4.11.11.2. Restic Backup CR cannot be recreated after bucket is emptied	164
4.11.12. Using the must-gather tool	164
4.11.12.1. Using must-gather with insecure TLS connections	165
4.11.12.2. Combining options when using the must-gather tool	165
4.11.13. OADP Monitoring	165
4.11.13.1. OADP monitoring setup	166
4.11.13.2. Creating OADP service monitor	167
4.11.13.3. Creating an alerting rule	168
4.11.13.4. List of available metrics	170
4.11.13.5. Viewing metrics using the Observe UI	173
4.12. APIS USED WITH OADP	173
4.12.1. Velero API	174
4.12.2. OADP API	174
4.13. ADVANCED OADP FEATURES AND FUNCTIONALITIES	179
4.13.1. Working with different Kubernetes API versions on the same cluster	179
4.13.1.1. Listing the Kubernetes API group versions on a cluster	179
4.13.1.2. About Enable API Group Versions	179
4.13.1.3. Using Enable API Group Versions	180
4.13.2. Backing up data from one cluster and restoring it to another cluster	181
4.13.2.1. About backing up data from one cluster and restoring it on another cluster	181
4.13.2.1.1. Operators	181
4.13.2.1.2. Use of Velero	181
4.13.2.2. About determining which pod volumes to back up	181
4.13.2.2.1. Limitations	182
4.13.2.2.2. Backing up pod volumes by using the opt-in method	182
4.13.2.2.3. Backing up pod volumes by using the opt-out method	183

4.13.2.3. UID and GID ranges	183
4.13.2.4. Backing up data from one cluster and restoring it to another cluster	185
4.13.3. Additional resources	185
CHAPTER 5. CONTROL PLANE BACKUP AND RESTORE	186
5.1. BACKING UP ETCD	186
5.1.1. Backing up etcd data	186
5.2. REPLACING AN UNHEALTHY ETCD MEMBER	188
5.2.1. Prerequisites	188
5.2.2. Identifying an unhealthy etcd member	188
5.2.3. Determining the state of the unhealthy etcd member	189
5.2.4. Replacing the unhealthy etcd member	190
5.2.4.1. Replacing an unhealthy etcd member whose machine is not running or whose node is not ready	191
5.2.4.2. Replacing an unhealthy etcd member whose etcd pod is crashlooping	198
5.2.4.3. Replacing an unhealthy bare metal etcd member whose machine is not running or whose node is not ready	203
5.2.5. Additional resources	214
5.3. BACKING UP AND RESTORING ETCD ON A HOSTED CLUSTER	214
5.3.1. Taking a snapshot of etcd on a hosted cluster	214
5.3.2. Restoring an etcd snapshot on a hosted cluster	216
5.3.3. Additional resources	216
5.4. DISASTER RECOVERY	216
5.4.1. About disaster recovery	216
5.4.2. Restoring to a previous cluster state	217
5.4.2.1. About restoring cluster state	217
5.4.2.2. Restoring to a previous cluster state	218
5.4.2.3. Additional resources	231
5.4.2.4. Issues and workarounds for restoring a persistent storage state	231
5.4.3. Recovering from expired control plane certificates	232
5.4.3.1. Recovering from expired control plane certificates	232
5.4.4. Disaster recovery for a hosted cluster within an AWS region	233
5.4.4.1. Example environment and context	234
5.4.4.2. Overview of the backup and restore process	236
5.4.4.3. Backing up a hosted cluster	240
5.4.4.4. Restoring a hosted cluster	245
5.4.4.5. Deleting a hosted cluster from your source management cluster	248
5.4.4.6. Running a script to back up and restore a hosted cluster	250

CHAPTER 1. BACKUP AND RESTORE

1.1. CONTROL PLANE BACKUP AND RESTORE OPERATIONS

As a cluster administrator, you might need to stop an OpenShift Container Platform cluster for a period and restart it later. Some reasons for restarting a cluster are that you need to perform maintenance on a cluster or want to reduce resource costs. In OpenShift Container Platform, you can perform a [graceful shutdown of a cluster](#) so that you can easily restart the cluster later.

You must [back up etcd data](#) before shutting down a cluster; etcd is the key-value store for OpenShift Container Platform, which persists the state of all resource objects. An etcd backup plays a crucial role in disaster recovery. In OpenShift Container Platform, you can also [replace an unhealthy etcd member](#).

When you want to get your cluster running again, [restart the cluster gracefully](#).



NOTE

A cluster's certificates expire one year after the installation date. You can shut down a cluster and expect it to restart gracefully while the certificates are still valid. Although the cluster automatically retrieves the expired control plane certificates, you must still [approve the certificate signing requests \(CSRs\)](#).

You might run into several situations where OpenShift Container Platform does not work as expected, such as:

- You have a cluster that is not functional after the restart because of unexpected conditions, such as node failure or network connectivity issues.
- You have deleted something critical in the cluster by mistake.
- You have lost the majority of your control plane hosts, leading to etcd quorum loss.

You can always recover from a disaster situation by [restoring your cluster to its previous state](#) using the saved etcd snapshots.

Additional resources

- [Quorum protection with machine lifecycle hooks](#)

1.2. APPLICATION BACKUP AND RESTORE OPERATIONS

As a cluster administrator, you can back up and restore applications running on OpenShift Container Platform by using the OpenShift API for Data Protection (OADP).

OADP backs up and restores Kubernetes resources and internal images, at the granularity of a namespace, by using the version of Velero that is appropriate for the version of OADP you install, according to the table in [Downloading the Velero CLI tool](#). OADP backs up and restores persistent volumes (PVs) by using snapshots or Restic. For details, see [OADP features](#).

1.2.1. OADP requirements

OADP has the following requirements:

- You must be logged in as a user with a **cluster-admin** role.


- You must have object storage for storing backups, such as one of the following storage types:
 - OpenShift Data Foundation
 - Amazon Web Services
 - Microsoft Azure
 - Google Cloud Platform
 - S3-compatible object storage



NOTE

If you want to use CSI backup on OCP 4.11 and later, install OADP 1.1.x.

OADP 1.0.x does not support CSI backup on OCP 4.11 and later. OADP 1.0. x includes Velero 1.7.x and expects the API group **snapshot.storage.k8s.io/v1beta1**, which is not present on OCP 4.11 and later.



IMPORTANT

The **CloudStorage** API for S3 storage is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

- To back up PVs with snapshots, you must have cloud storage that has a native snapshot API or supports Container Storage Interface (CSI) snapshots, such as the following providers:
 - Amazon Web Services
 - Microsoft Azure
 - Google Cloud Platform
 - CSI snapshot-enabled cloud storage, such as Ceph RBD or Ceph FS



NOTE

If you do not want to back up PVs by using snapshots, you can use [Restic](#), which is installed by the OADP Operator by default.

1.2.2. Backing up and restoring applications

You back up applications by creating a **Backup** custom resource (CR). See [Creating a Backup CR](#). You can configure the following backup options:

- [Creating backup hooks](#) to run commands before or after the backup operation

- [Scheduling backups](#)
- [Restic backups](#)
- You restore application backups by creating a **Restore** (CR). See [Creating a Restore CR](#).
- You can configure [restore hooks](#) to run commands in init containers or in the application container during the restore operation.

CHAPTER 2. SHUTTING DOWN THE CLUSTER GRACEFULLY

This document describes the process to gracefully shut down your cluster. You might need to temporarily shut down your cluster for maintenance reasons, or to save on resource costs.

2.1. PREREQUISITES

- Take an [etcd backup](#) prior to shutting down the cluster.



IMPORTANT

It is important to take an etcd backup before performing this procedure so that your cluster can be restored if you encounter any issues when restarting the cluster.

For example, the following conditions can cause the restarted cluster to malfunction:

- etcd data corruption during shutdown
- Node failure due to hardware
- Network connectivity issues

If your cluster fails to recover, follow the steps to [restore to a previous cluster state](#).

2.2. SHUTTING DOWN THE CLUSTER

You can shut down your cluster in a graceful manner so that it can be restarted at a later date.



NOTE

You can shut down a cluster until a year from the installation date and expect it to restart gracefully. After a year from the installation date, the cluster certificates expire.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have taken an etcd backup.

Procedure

1. If you plan to shut down the cluster for an extended period of time, determine the date that cluster certificates expire.

You must restart the cluster prior to the date that certificates expire. As the cluster restarts, the process might require you to manually approve the pending certificate signing requests (CSRs) to recover kubelet certificates.

- a. Check the expiration date for the **kube-apiserver-to-kubelet-signer** CA certificate:

```
$ oc -n openshift-kube-apiserver-operator get secret kube-apiserver-to-kubelet-signer -o jsonpath='{.metadata.annotations.auth\.openshift\.io/certificate-not-after}'
```


-

Example output

```
2023-08-05T14:37:50Z
```

- b. Check the expiration date for the kubelet certificates:
 - i. Start a debug session for a control plane node by running the following command:

```
$ oc debug node/<node_name>
```

- ii. Change your root directory to **/host** by running the following command:

```
sh-4.4# chroot /host
```

- iii. Check the kubelet client certificate expiration date by running the following command:

```
sh-5.1# openssl x509 -in /var/lib/kubelet/pki/kubelet-client-current.pem -noout -enddate
```

Example output

```
notAfter=Jun 6 10:50:07 2023 GMT
```

- iv. Check the kubelet server certificate expiration date by running the following command:

```
sh-5.1# openssl x509 -in /var/lib/kubelet/pki/kubelet-server-current.pem -noout -enddate
```

Example output

```
notAfter=Jun 6 10:50:07 2023 GMT
```

- v. Exit the debug session.
 - vi. Repeat these steps to check certificate expiration dates on all control plane nodes. To ensure that the cluster can restart gracefully, plan to restart it before the earliest certificate expiration date.
2. Shut down all of the nodes in the cluster. You can do this from your cloud provider's web console, or run the following loop:

```
$ for node in $(oc get nodes -o jsonpath='{.items[*].metadata.name}'); do oc debug node/${node} -- chroot /host shutdown -h 1; done 1
```

- 1** **-h 1** indicates how long, in minutes, this process lasts before the control-plane nodes are shut down. For large-scale clusters with 10 nodes or more, set to 10 minutes or longer to make sure all the compute nodes have time to shut down first.

Example output

```
Starting pod/ip-10-0-130-169us-east-2computeinternal-debug ...
To use host binaries, run `chroot /host`
Shutdown scheduled for Mon 2021-09-13 09:36:17 UTC, use 'shutdown -c' to cancel.
```

```
Removing debug pod ...
Starting pod/ip-10-0-150-116us-east-2computeinternal-debug ...
To use host binaries, run `chroot /host`
Shutdown scheduled for Mon 2021-09-13 09:36:29 UTC, use 'shutdown -c' to cancel.
```

Shutting down the nodes using one of these methods allows pods to terminate gracefully, which reduces the chance for data corruption.



NOTE

Adjust the shut down time to be longer for large-scale clusters:

```
$ for node in $(oc get nodes -o jsonpath='{.items[*].metadata.name}'); do oc
debug node/${node} -- chroot /host shutdown -h 10; done
```



NOTE

It is not necessary to drain control plane nodes of the standard pods that ship with OpenShift Container Platform prior to shutdown.

Cluster administrators are responsible for ensuring a clean restart of their own workloads after the cluster is restarted. If you drained control plane nodes prior to shutdown because of custom workloads, you must mark the control plane nodes as schedulable before the cluster will be functional again after restart.

3. Shut off any cluster dependencies that are no longer needed, such as external storage or an LDAP server. Be sure to consult your vendor's documentation before doing so.



IMPORTANT

If you deployed your cluster on a cloud-provider platform, do not shut down, suspend, or delete the associated cloud resources. If you delete the cloud resources of a suspended virtual machine, OpenShift Container Platform might not restore successfully.

2.3. ADDITIONAL RESOURCES

- [Restarting the cluster gracefully](#)

CHAPTER 3. RESTARTING THE CLUSTER GRACEFULLY

This document describes the process to restart your cluster after a graceful shutdown.

Even though the cluster is expected to be functional after the restart, the cluster might not recover due to unexpected conditions, for example:

- etcd data corruption during shutdown
- Node failure due to hardware
- Network connectivity issues

If your cluster fails to recover, follow the steps to [restore to a previous cluster state](#).

3.1. PREREQUISITES

- You have [gracefully shut down your cluster](#).

3.2. RESTARTING THE CLUSTER

You can restart your cluster after it has been shut down gracefully.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- This procedure assumes that you gracefully shut down the cluster.

Procedure

1. Power on any cluster dependencies, such as external storage or an LDAP server.
2. Start all cluster machines.
Use the appropriate method for your cloud environment to start the machines, for example, from your cloud provider's web console.

Wait approximately 10 minutes before continuing to check the status of control plane nodes.
3. Verify that all control plane nodes are ready.

```
$ oc get nodes -l node-role.kubernetes.io/master
```

The control plane nodes are ready if the status is **Ready**, as shown in the following output:

```
NAME                                STATUS ROLES  AGE  VERSION
ip-10-0-168-251.ec2.internal        Ready  master  75m  v1.25.0
ip-10-0-170-223.ec2.internal        Ready  master  75m  v1.25.0
ip-10-0-211-16.ec2.internal         Ready  master  75m  v1.25.0
```

4. If the control plane nodes are *not* ready, then check whether there are any pending certificate signing requests (CSRs) that must be approved.
 - a. Get the list of current CSRs:

```
$ oc get csr
```

- b. Review the details of a CSR to verify that it is valid:

```
$ oc describe csr <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- c. Approve each valid CSR:

```
$ oc adm certificate approve <csr_name>
```

5. After the control plane nodes are ready, verify that all worker nodes are ready.

```
$ oc get nodes -l node-role.kubernetes.io/worker
```

The worker nodes are ready if the status is **Ready**, as shown in the following output:

```
NAME                                STATUS ROLES  AGE  VERSION
ip-10-0-179-95.ec2.internal  Ready  worker  64m  v1.25.0
ip-10-0-182-134.ec2.internal  Ready  worker  64m  v1.25.0
ip-10-0-250-100.ec2.internal  Ready  worker  64m  v1.25.0
```

6. If the worker nodes are *not* ready, then check whether there are any pending certificate signing requests (CSRs) that must be approved.

- a. Get the list of current CSRs:

```
$ oc get csr
```

- b. Review the details of a CSR to verify that it is valid:

```
$ oc describe csr <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- c. Approve each valid CSR:

```
$ oc adm certificate approve <csr_name>
```

7. Verify that the cluster started properly.

- a. Check that there are no degraded cluster Operators.

```
$ oc get clusteroperators
```

Check that there are no cluster Operators with the **DEGRADED** condition set to **True**.

```
NAME                                VERSION AVAILABLE PROGRESSING DEGRADED
SINCE
authentication                       4.12.0  True    False    False    59m
```

```

cloud-credential          4.12.0 True   False   False   85m
cluster-autoscaler       4.12.0 True   False   False   73m
config-operator          4.12.0 True   False   False   73m
console                   4.12.0 True   False   False   62m
csi-snapshot-controller  4.12.0 True   False   False   66m
dns                       4.12.0 True   False   False   76m
etcd                      4.12.0 True   False   False   76m
...

```

- b. Check that all nodes are in the **Ready** state:

```
$ oc get nodes
```

Check that the status for all nodes is **Ready**.

```

NAME                                STATUS ROLES  AGE  VERSION
ip-10-0-168-251.ec2.internal        Ready  master  82m  v1.25.0
ip-10-0-170-223.ec2.internal        Ready  master  82m  v1.25.0
ip-10-0-179-95.ec2.internal         Ready  worker  70m  v1.25.0
ip-10-0-182-134.ec2.internal        Ready  worker  70m  v1.25.0
ip-10-0-211-16.ec2.internal         Ready  master  82m  v1.25.0
ip-10-0-250-100.ec2.internal        Ready  worker  69m  v1.25.0

```

If the cluster did not start properly, you might need to restore your cluster using an etcd backup.

8. After the control plane and worker nodes are ready, mark all the nodes in the cluster as schedulable. Run the following command:

```
for node in $(oc get nodes -o jsonpath='{.items[*].metadata.name}'); do echo ${node} ; oc adm uncordon ${node} ; done
```

Additional resources

- See [Restoring to a previous cluster state](#) for how to use an etcd backup to restore if your cluster failed to recover after restarting.

CHAPTER 4. OADP APPLICATION BACKUP AND RESTORE

4.1. INTRODUCTION TO OPENSIFT API FOR DATA PROTECTION

The OpenShift API for Data Protection (OADP) product safeguards customer applications on OpenShift Container Platform. It offers comprehensive disaster recovery protection, covering OpenShift Container Platform applications, application-related cluster resources, persistent volumes, and internal images. OADP is also capable of backing up both containerized applications and virtual machines (VMs).

However, OADP does not serve as a disaster recovery solution for [etcd](#) or OpenShift Operators.

4.1.1. OpenShift API for Data Protection APIs

OpenShift API for Data Protection (OADP) provides APIs that enable multiple approaches to customizing backups and preventing the inclusion of unnecessary or inappropriate resources.

OADP provides the following APIs:

- [Backup](#)
- [Restore](#)
- [Schedule](#)
- [BackupStorageLocation](#)
- [VolumeSnapshotLocation](#)

Additional resources

- [Backing up etcd](#)

4.2. OADP RELEASE NOTES

4.2.1. OADP 1.3 release notes

The release notes for OpenShift API for Data Protection (OADP) describe new features and enhancements, deprecated features, product recommendations, known issues, and resolved issues.

4.2.1.1. OADP 1.3.1 release notes

The OpenShift API for Data Protection (OADP) 1.3.1 release notes lists new features and resolved issues.

4.2.1.1.1. New features

OADP 1.3.0 Data Mover is now fully supported

The OADP built-in Data Mover, introduced in OADP 1.3.0 as a Technology Preview, is now fully supported for both containerized and virtual machine workloads.

4.2.1.1.2. Resolved issues

{ibm-cloud-name} Object Storage is now supported as a backup storage provider

{ibm-cloud-name} Object Storage is one of the AWS S3 compatible backup storage providers, which was unsupported previously. With this update, {ibm-cloud-name} Object Storage is now supported as an AWS S3 compatible backup storage provider.

[OADP-3788](#)

OADP operator now correctly reports the missing region error

Previously, when you specified **profile:default** without specifying the **region** in the AWS Backup Storage Location (BSL) configuration, the OADP operator failed to report the **missing region** error on the Data Protection Application (DPA) custom resource (CR). This update corrects validation of DPA BSL specification for AWS. As a result, the OADP Operator reports the **missing region** error.

[OADP-3044](#)

Custom labels are not removed from the openshift-adp namespace

Previously, the **openshift-adp-controller-manager** pod would reset the labels attached to the **openshift-adp** namespace. This caused synchronization issues for applications requiring custom labels such as Argo CD, leading to improper functionality. With this update, this issue is fixed and custom labels are not removed from the **openshift-adp** namespace.

[OADP-3189](#)

OADP must-gather image collects CRDs

Previously, the OADP **must-gather** image did not collect the custom resource definitions (CRDs) shipped by OADP. Consequently, you could not use the **omg** tool to extract data in the support shell. With this fix, the **must-gather** image now collects CRDs shipped by OADP and can use the **omg** tool to extract data.

[OADP-3229](#)

Garbage collection has the correct description for the default frequency value

Previously, the **garbage-collection-frequency** field had a wrong description for the default frequency value. With this update, **garbage-collection-frequency** has a correct value of one hour for the **gc-controller** reconciliation default frequency.

[OADP-3486](#)

FIPS Mode flag is available in OperatorHub

By setting the **fips-compliant** flag to **true**, the FIPS mode flag is now added to the OADP Operator listing in OperatorHub. This feature was enabled in OADP 1.3.0 but did not show up in the Red Hat Container catalog as being FIPS enabled.

[OADP-3495](#)

CSI plugin does not panic with a nil pointer when csiSnapshotTimeout is set to a short duration

Previously, when the **csiSnapshotTimeout** parameter was set to a short duration, the CSI plugin encountered the following error: **plugin panicked: runtime error: invalid memory address or nil pointer dereference**.

With this fix, the backup fails with the following error: **Timed out awaiting reconciliation of volumesnapshot**.

[OADP-3069](#)

For a complete list of all issues resolved in this release, see the list of [OADP 1.3.1 resolved issues](#) in Jira.

4.2.1.1.3. Known issues

Backup and storage restrictions for Single-node OpenShift clusters deployed on {ibm-power-name} and {ibm-z-name} platforms

Review the following backup and storage related restrictions for Single-node OpenShift clusters that are deployed on {ibm-power-name} and {ibm-z-name} platforms:

Storage

Only NFS storage is currently compatible with single-node OpenShift clusters deployed on {ibm-power-name} and {ibm-z-name} platforms.

Backup

Only the backing up applications with File System Backup such as **kopia** and **restic** are supported for backup and restore operations.

[OADP-3787](#)

Cassandra application pods enter in the CrashLoopBackoff status after restoring OADP

After OADP restores, the Cassandra application pods might enter in the **CrashLoopBackoff** status. To work around this problem, delete the **StatefulSet** pods with any error or the **CrashLoopBackoff** state after restoring OADP. The **StatefulSet** controller recreates these pods and it runs normally.

[OADP-3767](#)

4.2.1.2. OADP 1.3.0 release notes

The OpenShift API for Data Protection (OADP) 1.3.0 release notes lists new features, resolved issues and bugs, and known issues.

4.2.1.2.1. New features

Velero built-in DataMover

OADP 1.3 includes a built-in Data Mover that you can use to move Container Storage Interface (CSI) volume snapshots to a remote object store. The built-in Data Mover allows you to restore stateful applications from the remote object store if a failure, accidental deletion, or corruption of the cluster occurs. It uses Kopia as the uploader mechanism to read the snapshot data and to write to the Unified Repository.

Velero built-in DataMover is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

Backing up applications with File System Backup: Kopia or Restic

Velero's File System Backup (FSB) supports two backup libraries: the Restic path and the Kopia path.

Velero allows users to select between the two paths.

For backup, specify the path during the installation through the **uploader-type** flag. The valid value is either **restic** or **kopia**. This field defaults to **kopia** if the value is not specified. The selection cannot be changed after the installation.

GCP Cloud authentication

Google Cloud Platform (GCP) authentication enables you to use short-lived Google credentials.

GCP with Workload Identity Federation enables you to use Identity and Access Management (IAM) to grant external identities IAM roles, including the ability to impersonate service accounts. This eliminates the maintenance and security risks associated with service account keys.

AWS ROSA STS authentication

You can use OpenShift API for Data Protection (OADP) with Red Hat OpenShift Service on AWS (ROSA) clusters to backup and restore application data.

ROSA provides seamless integration with a wide range of AWS compute, database, analytics, machine learning, networking, mobile, and other services to speed up the building and delivering of differentiating experiences to your customers.

You can subscribe to the service directly from your AWS account.

After the clusters are created, you can operate your clusters by using the OpenShift web console. The ROSA service also uses OpenShift APIs and command-line interface (CLI) tools.

4.2.1.2.2. Resolved issues

ACM applications were removed and re-created on managed clusters after restore

Applications on managed clusters were deleted and re-created upon restore activation. OpenShift API for Data Protection (OADP 1.2) backup and restore process is faster than the older versions. The OADP performance change caused this behavior when restoring ACM resources. Therefore, some resources were restored before other resources, which caused the removal of the applications from managed clusters. [OADP-2686](#)

Restic restore was partially failing due to Pod Security standard

During interoperability testing, OpenShift Container Platform 4.14 had the pod Security mode set to **enforce**, which caused the pod to be denied. This was caused due to the restore order. The pod was getting created before the security context constraints (SCC) resource, since the pod violated the **podSecurity** standard, it denied the pod. When setting the restore priority field on the Velero server, restore is successful. [OADP-2688](#)

Possible pod volume backup failure if Velero is installed in several namespaces

There was a regression in Pod Volume Backup (PVB) functionality when Velero was installed in several namespaces. The PVB controller was not properly limiting itself to PVBs in its own namespace. [OADP-2308](#)

OADP Velero plugins returning "received EOF, stopping recv loop" message

In OADP, Velero plugins were started as separate processes. When the Velero operation completes, either successfully or not, they exit. Therefore, if you see a **received EOF, stopping recv loop** messages in debug logs, it does not mean an error occurred, it means that a plugin operation has completed. [OADP-2176](#)

CVE-2023-39325 Multiple HTTP/2 enabled web servers are vulnerable to a DDoS attack (Rapid Reset Attack)

In previous releases of OADP, the HTTP/2 protocol was susceptible to a denial of service attack because request cancellation could reset multiple streams quickly. The server had to set up and tear down the streams while not hitting any server-side limit for the maximum number of active streams per connection. This resulted in a denial of service due to server resource consumption.

For more information, see [CVE-2023-39325 \(Rapid Reset Attack\)](#)

For a complete list of all issues resolved in this release, see the list of [OADP 1.3.0 resolved issues](#) in Jira.

4.2.1.2.3. Known issues

CSI plugin errors on nil pointer when csiSnapshotTimeout is set to a short duration

The CSI plugin errors on nil pointer when **csiSnapshotTimeout** is set to a short duration. Sometimes it succeeds to complete the snapshot within a short duration, but often it panics with the backup **PartiallyFailed** with the following error: **plugin panicked: runtime error: invalid memory address or nil pointer dereference.**

Backup is marked as PartiallyFailed when volumeSnapshotContent CR has an error

If any of the **VolumeSnapshotContent** CRs have an error related to removing the **VolumeSnapshotBeingCreated** annotation, it moves the backup to the **WaitingForPluginOperationsPartiallyFailed** phase. [OADP-2871](#)

Performance issues when restoring 30,000 resources for the first time

When restoring 30,000 resources for the first time, without an existing-resource-policy, it takes twice as long to restore them, than it takes during the second and third try with an existing-resource-policy set to **update**. [OADP-3071](#)

Post restore hooks might start running before Datadownload operation has released the related PV


Due to the asynchronous nature of the Data Mover operation, a post-hook might be attempted before the related pods persistent volumes (PVs) are released by the Data Mover persistent volume claim (PVC).

GCP-Workload Identity Federation VSL backup PartiallyFailed

VSL backup **PartiallyFailed** when GCP workload identity is configured on GCP.

For a complete list of all known issues in this release, see the list of [OADP 1.3.0 known issues](#) in Jira.

4.2.1.2.4. Upgrade notes



NOTE

Always upgrade to the next minor version. **Do not** skip versions. To update to a later version, upgrade only one channel at a time. For example, to upgrade from OpenShift API for Data Protection (OADP) 1.1 to 1.3, upgrade first to 1.2, and then to 1.3.

4.2.1.2.4.1. Changes from OADP 1.2 to 1.3

The Velero server has been updated from version 1.11 to 1.12.

OpenShift API for Data Protection (OADP) 1.3 uses the Velero built-in Data Mover instead of the VolumeSnapshotMover (VSM) or the Volsync Data Mover.

This changes the following:

- The **spec.features.dataMover** field and the VSM plugin are not compatible with OADP 1.3, and you must remove the configuration from the **DataProtectionApplication** (DPA) configuration.
- The Volsync Operator is no longer required for Data Mover functionality, and you can remove it.
- The custom resource definitions **volumesnapshotbackups.datamover.oadp.openshift.io** and **volumesnapshotrestores.datamover.oadp.openshift.io** are no longer required, and you can remove them.
- The secrets used for the OADP-1.2 Data Mover are no longer required, and you can remove them.

OADP 1.3 supports Kopia, which is an alternative file system backup tool to Restic.

- To employ Kopia, use the new **spec.configuration.nodeAgent** field as shown in the following example:

Example

```
spec:
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
# ...
```

- The **spec.configuration.restic** field is deprecated in OADP 1.3 and will be removed in a future version of OADP. To avoid seeing deprecation warnings, remove the **restic** key and its values, and use the following new syntax:

Example

```
spec:
  configuration:
    nodeAgent:
      enable: true
      uploaderType: restic
# ...
```



NOTE

In a future OADP release, it is planned that the **kopia** tool will become the default **uploaderType** value.

4.2.1.2.4.2. Upgrading steps

4.2.1.2.4.3. Upgrading from OADP 1.2 Technology Preview Data Mover

OpenShift API for Data Protection (OADP) 1.2 Data Mover backups **cannot** be restored with OADP 1.3. To prevent a gap in the data protection of your applications, complete the following steps before upgrading to OADP 1.3:

Procedure

1. If your cluster backups are sufficient and Container Storage Interface (CSI) storage is available, back up the applications with a CSI backup.
2. If you require off cluster backups:
 - a. Back up the applications with a file system backup that uses the **--default-volumes-to-fs-backup=true** or **backup.spec.defaultVolumesToFsBackup** options.
 - b. Back up the applications with your object storage plugins, for example, **velero-plugin-for-aws**.



NOTE

To restore OADP 1.2 Data Mover backup, you must uninstall OADP, and install and configure OADP 1.2.

4.2.1.2.4.4. Backing up the DPA configuration

You must back up your current **DataProtectionApplication** (DPA) configuration.

Procedure

- Save your current DPA configuration by running the following command:

Example

```
$ oc get dpa -n openshift-adp -o yaml > dpa.orig.backup
```

4.2.1.2.4.5. Upgrading the OADP Operator

Use the following sequence when upgrading the OpenShift API for Data Protection (OADP) Operator.

Procedure

1. Change your subscription channel for the OADP Operator from **stable-1.2** to **stable-1.3**.
2. Allow time for the Operator and containers to update and restart.

Additional resources

- [Updating installed Operators](#)

4.2.1.2.4.6. Converting DPA to the new version

If you need to move backups off cluster with the Data Mover, reconfigure the **DataProtectionApplication** (DPA) manifest as follows.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. In the **Provided APIs** section, click **View more**.
3. Click **Create instance** in the **DataProtectionApplication** box.
4. Click **YAML View** to display the current DPA parameters.

Example current DPA

```
spec:
  configuration:
    features:
      dataMover:
        enable: true
        credentialName: dm-credentials
    velero:
      defaultPlugins:
        - vsm
        - csi
        - openshift
# ...
```

5. Update the DPA parameters:
 - Remove the **features.dataMover** key and values from the DPA.
 - Remove the VolumeSnapshotMover (VSM) plugin.
 - Add the **nodeAgent** key and values.

Example updated DPA

```
spec:
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
# ...
```

6. Wait for the DPA to reconcile successfully.

4.2.1.2.4.7. Verifying the upgrade

Use the following procedure to verify the upgrade.

Procedure

1. Verify the installation by viewing the OpenShift API for Data Protection (OADP) resources by running the following command:

```
$ oc get all -n openshift-adp
```

Example output

```

NAME                                READY STATUS RESTARTS AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2   Running 0      2m8s
pod/node-agent-9cq4q                        1/1   Running 0       94s
pod/node-agent-m4lts                        1/1   Running 0       94s
pod/node-agent-pv4kr                        1/1   Running 0       95s
pod/velero-588db7f655-n842v                1/1   Running 0       95s

NAME                                TYPE          CLUSTER-IP    EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP    172.30.70.140
<none>    8443/TCP    2m8s
service/openshift-adp-velero-metrics-svc                  ClusterIP    172.30.10.0   <none>
8085/TCP    8h

NAME                                DESIRED CURRENT READY UP-TO-DATE AVAILABLE NODE
SELECTOR AGE
daemonset.apps/node-agent           3         3         3     3         3         <none>    96s

NAME                                READY UP-TO-DATE AVAILABLE AGE
deployment.apps/oadp-operator-controller-manager  1/1     1         1     2m9s
deployment.apps/velero                    1/1     1         1     96s

NAME                                DESIRED CURRENT READY AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47  1         1         1     2m9s
replicaset.apps/velero-588db7f655                1         1         1     96s

```

2. Verify that the **DataProtectionApplication** (DPA) is reconciled by running the following command:

```
$ oc get dpa dpa-sample -n openshift-adp -o jsonpath='{.status}'
```

Example output

```
{
  "conditions": [
    {
      "lastTransitionTime": "2023-10-27T01:23:57Z",
      "message": "Reconcile complete",
      "reason": "Complete",
      "status": "True",
      "type": "Reconciled"
    }
  ]
}
```

3. Verify the **type** is set to **Reconciled**.
4. Verify the backup storage location and confirm that the **PHASE** is **Available** by running the following command:

```
$ oc get backupStorageLocation -n openshift-adp
```

Example output

```

NAME          PHASE    LAST VALIDATED AGE    DEFAULT
dpa-sample-1  Available 1s      3d16h true

```

In OADP 1.3 you can start data movement off cluster per backup versus creating a **DataProtectionApplication** (DPA) configuration.

Example

```
$ velero backup create example-backup --include-namespaces mysql-persistent --snapshot-move-data=true
```

Example

```
apiVersion: velero.io/v1
kind: Backup
metadata:
  name: example-backup
  namespace: openshift-adp
spec:
  snapshotMoveData: true
  includedNamespaces:
  - mysql-persistent
  storageLocation: dpa-sample-1
  ttl: 720h0m0s
# ...
```

4.2.2. OADP 1.2 release notes

The release notes for OpenShift API for Data Protection (OADP) 1.2 describe new features and enhancements, deprecated features, product recommendations, known issues, and resolved issues.

4.2.2.1. OADP 1.2.4 release notes

OpenShift API for Data Protection (OADP) 1.2.4 is a Container Grade Only (CGO) release, released to refresh the health grades of the containers, with no changes to any code in the product itself compared to that of {oadp-short} 1.2.3.

4.2.2.1.1. Resolved issues

There are no resolved issues in {oadp-short} 1.2.4.

4.2.2.1.2. Known issues

The {oadp-short} 1.2.4 has the following known issue:

Data Protection Application (DPA) does not reconcile when the credentials secret is updated

Currently, the {oadp-short} Operator does not reconcile when you update the **cloud-credentials** secret. This occurs because there are no {oadp-short} specific labels or owner references on the **cloud-credentials** secret. If you create a **cloud-credentials** secret with incorrect credentials, such as empty data, the Operator reconciles and creates a Backup Storage Location (BSL) and registry deployment with the empty data. As a result, when you update the **cloud-credentials** secret with the correct credentials, the Operator does not immediately reconcile to catch the new credentials.

Workaround: Update to {oadp-short} 1.3.

[\(OADP-3327\)](#)

4.2.2.2. OADP 1.2.3 release notes

4.2.2.2.1. New features

There are no new features in the release of OpenShift API for Data Protection (OADP) 1.2.3.

4.2.2.2.2. Resolved issues

The following highlighted issues are resolved in OADP 1.2.3:

Multiple HTTP/2 enabled web servers are vulnerable to a DDoS attack (Rapid Reset Attack)

In previous releases of OADP 1.2, the HTTP/2 protocol was susceptible to a denial of service attack because request cancellation could reset multiple streams quickly. The server had to set up and tear down the streams while not hitting any server-side limit for the maximum number of active streams per connection. This resulted in a denial of service due to server resource consumption. For a list of all OADP issues associated with this CVE, see the following [Jira list](#).

For more information, see [CVE-2023-39325 \(Rapid Reset Attack\)](#).

For a complete list of all issues resolved in the release of OADP 1.2.3, see the list of [OADP 1.2.3 resolved issues](#) in Jira.

4.2.2.2.3. Known issues

The {oadp-short} 1.2.3 has the following known issue:

Data Protection Application (DPA) does not reconcile when the credentials secret is updated

Currently, the {oadp-short} Operator does not reconcile when you update the **cloud-credentials** secret. This occurs because there are no {oadp-short} specific labels or owner references on the **cloud-credentials** secret. If you create a **cloud-credentials** secret with incorrect credentials, such as empty data, the Operator reconciles and creates a Backup Storage Location (BSL) and registry deployment with the empty data. As a result, when you update the **cloud-credentials** secret with the correct credentials, the Operator does not immediately reconcile to catch the new credentials.

Workaround: Update to {oadp-short} 1.3.

[\(OADP-3327\)](#)

4.2.2.3. OADP 1.2.2 release notes

4.2.2.3.1. New features

There are no new features in the release of OpenShift API for Data Protection (OADP) 1.2.2.

4.2.2.3.2. Resolved issues

The following highlighted issues are resolved in OADP 1.2.2:

Restic restore partially failed due to a Pod Security standard

In previous releases of OADP 1.2, OpenShift Container Platform 4.14 enforced a pod security admission (PSA) policy that hindered the readiness of pods during a Restic restore process.

This issue has been resolved in the release of OADP 1.2.2, and also OADP 1.1.6. Therefore, it is recommended that users upgrade to these releases.

For more information, see [Restic restore partially failing on OCP 4.14 due to changed PSA policy . \(OADP-2094\)](#)

Backup of an app with internal images partially failed with plugin panicked error

In previous releases of OADP 1.2, the backup of an application with internal images partially failed with plugin panicked error returned. The backup partially fails with this error in the Velero logs:

```
time="2022-11-23T15:40:46Z" level=info msg="1 errors encountered backup up item"
backup=openshift-adp/django-persistent-67a5b83d-6b44-11ed-9cba-902e163f806c
logSource="/remote-source/velero/app/pkg/backup/backup.go:413" name=django-psql-persistent
time="2022-11-23T15:40:46Z" level=error msg="Error backing up item" backup=openshift-
adp/django-persistent-67a5b83d-6b44-11ed-9cba-902e163f8
```

This issue has been resolved in OADP 1.2.2. ([OADP-1057](#)).

ACM cluster restore was not functioning as expected due to restore order

In previous releases of OADP 1.2, ACM cluster restore was not functioning as expected due to restore order. ACM applications were removed and re-created on managed clusters after restore activation. ([OADP-2505](#))

VM's using filesystemOverhead failed when backing up and restoring due to volume size mismatch

In previous releases of OADP 1.2, due to storage provider implementation choices, whenever there was a difference between the application persistent volume claims (PVCs) storage request and the snapshot size of the same PVC, VM's using filesystemOverhead failed when backing up and restoring. This issue has been resolved in the Data Mover of OADP 1.2.2. ([OADP-2144](#))

OADP did not contain an option to set VolSync replication source prune interval

In previous releases of OADP 1.2, there was no option to set the VolSync replication source **pruneInterval**. ([OADP-2052](#))

Possible pod volume backup failure if Velero was installed in multiple namespaces

In previous releases of OADP 1.2, there was a possibility of pod volume backup failure if Velero was installed in multiple namespaces. ([OADP-2409](#))

Backup Storage Locations moved to unavailable phase when VSL uses custom secret

In previous releases of OADP 1.2, Backup Storage Locations moved to unavailable phase when Volume Snapshot Location used custom secret. ([OADP-1737](#))

For a complete list of all issues resolved in the release of OADP 1.2.2, see the list of [OADP 1.2.2 resolved issues](#) in Jira.

4.2.2.3.3. Known issues

The following issues have been highlighted as known issues in the release of OADP 1.2.2:

Must-gather command fails to remove ClusterRoleBinding resources

The `oc adm must-gather` command fails to remove `ClusterRoleBinding` resources.

The `oc adm must-gather` command fails to remove `ClusterRoleBinding` resources, which are left on cluster due to admission webhook. Therefore, requests for the removal of the `ClusterRoleBinding` resources are denied. ([OADP-27730](#))

```
admission webhook "clusterrolebindings-validation.managed.openshift.io" denied the request:
Deleting ClusterRoleBinding must-gather-p7vwj is not allowed
```

For a complete list of all known issues in this release, see the list of [OADP 1.2.2 known issues](#) in Jira.

4.2.2.4. OADP 1.2.1 release notes

4.2.2.4.1. New features

There are no new features in the release of OpenShift API for Data Protection (OADP) 1.2.1.

4.2.2.4.2. Resolved issues

For a complete list of all issues resolved in the release of OADP 1.2.1, see the list of [OADP 1.2.1 resolved issues](#) in Jira.

4.2.2.4.3. Known issues

The following issues have been highlighted as known issues in the release of OADP 1.2.1:

DataMover Restic retain and prune policies do not work as expected

The retention and prune features provided by VolSync and Restic are not working as expected. Because there is no working option to set the prune interval on VolSync replication, you have to manage and prune remotely stored backups on S3 storage outside of OADP. For more details, see:

- [OADP-2052](#)
- [OADP-2048](#)
- [OADP-2175](#)
- [OADP-1690](#)



IMPORTANT

OADP Data Mover is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

For a complete list of all known issues in this release, see the list of [OADP 1.2.1 known issues](#) in Jira.

4.2.2.5. OADP 1.2.0 release notes

The OADP 1.2.0 release notes include information about new features, bug fixes, and known issues.

4.2.2.5.1. New features

Resource timeouts

The new **resourceTimeout** option specifies the timeout duration in minutes for waiting on various Velero resources. This option applies to resources such as Velero CRD availability, **volumeSnapshot** deletion, and backup repository availability. The default duration is 10 minutes.

AWS S3 compatible backup storage providers

You can back up objects and snapshots on AWS S3 compatible providers.

4.2.2.5.1.1. Technical preview features

Data Mover

The OADP Data Mover enables you to back up Container Storage Interface (CSI) volume snapshots to a remote object store. When you enable Data Mover, you can restore stateful applications using CSI volume snapshots pulled from the object store in case of accidental cluster deletion, cluster failure, or data corruption.



IMPORTANT

OADP Data Mover is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

4.2.2.5.2. Resolved issues

For a complete list of all issues resolved in this release, see the list of [OADP 1.2.0 resolved issues](#) in Jira.

4.2.2.5.3. Known issues

The following issues have been highlighted as known issues in the release of OADP 1.2.0:

Multiple HTTP/2 enabled web servers are vulnerable to a DDoS attack (Rapid Reset Attack)

The HTTP/2 protocol is susceptible to a denial of service attack because request cancellation can reset multiple streams quickly. The server has to set up and tear down the streams while not hitting any server-side limit for the maximum number of active streams per connection. This results in a denial of service due to server resource consumption.

It is advised to upgrade to OADP 1.2.3, which resolves this issue.

For more information, see [CVE-2023-39325 \(Rapid Reset Attack\)](#).


An incorrect hostname can be created when changing a hostname in a generated route.

By default, the OpenShift Container Platform cluster makes sure that the **openshift.io/host.generated: true** annotation is turned on and fills in the field for both the routes that are generated and those that are not generated.

You cannot modify the value for the **.spec.host** field based on the base domain name of your cluster in the generated and non-generated routes.

If you modify the value for the **.spec.host** field, it is not possible to restore the default value that was generated by the OpenShift Container Platform cluster. After you restore your OpenShift Container Platform cluster, the Operator resets the value for the field.

4.2.2.5.4. Upgrade notes



NOTE

Always upgrade to the next minor version. **Do not** skip versions. To update to a later version, upgrade only one channel at a time. For example, to upgrade from OpenShift API for Data Protection (OADP) 1.1 to 1.3, upgrade first to 1.2, then to 1.3.

4.2.2.5.4.1. Changes from OADP 1.1 to 1.2

The Velero server was updated from version 1.9 to 1.11.

In OADP 1.2, the **DataProtectionApplication** (DPA) configuration **dpa.spec.configuration.velero.args** has the following changes:

- The **default-volumes-to-restic** field was renamed to **default-volumes-to-fs-backup**. If you use **dpa.spec.configuration.velero.args**, you must add it again with the new name to your DPA after upgrading OADP.
- The **restic-timeout** field was renamed to **fs-backup-timeout**. If you use **dpa.spec.configuration.velero.args**, you must add it again with the new name to your DPA after upgrading OADP.
- The **restic** daemon set was renamed to **node-agent**. OADP automatically updates the name of the daemon set.
- The custom resource definition **resticrepositories.velero.io** was renamed to **backuprepositories.velero.io**.
- The custom resource definition **resticrepositories.velero.io** can be removed from the cluster.

4.2.2.5.5. Upgrading steps

4.2.2.5.5.1. Backing up the DPA configuration

You must back up your current **DataProtectionApplication** (DPA) configuration.

Procedure

- Save your current DPA configuration by running the following command:

Example

```
$ oc get dpa -n openshift-adp -o yaml > dpa.orig.backup
```

4.2.2.5.5.2. Upgrading the OADP Operator

Use the following sequence when upgrading the OpenShift API for Data Protection (OADP) Operator.

Procedure

1. Change your subscription channel for the OADP Operator from **stable-1.1** to **stable-1.2**.
2. Allow time for the Operator and containers to update and restart.

Additional resources

- [Configuring Amazon Web Services](#)
- [Using Data Mover for CSI snapshots](#)
- [Updating installed Operators](#)

4.2.2.5.5.3. Converting DPA to the new version

If you use the fields that were updated in the **spec.configuration.velero.args** stanza, you must configure your **DataProtectionApplication** (DPA) manifest to use the new parameter names.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. Select **Provided APIs**, click **Create instance** in the **DataProtectionApplication** box.
3. Click **YAML View** to display the current DPA parameters.

Example current DPA

```
spec:
  configuration:
    velero:
      args:
        default-volumes-to-fs-backup: true
        default-restic-prune-frequency: 6000
        fs-backup-timeout: 600
# ...
```

4. Update the DPA parameters:
5. Update the DPA parameter names without changing their values:
 - a. Change the **default-volumes-to-restic** key to **default-volumes-to-fs-backup**.
 - b. Change the **default-restic-prune-frequency** key to **default-repo-maintain-frequency**.
 - c. Change the **restic-timeout** key to **fs-backup-timeout**.

.Example updated DPA

```
spec:
  configuration:
    velero:
```

```
args:
  default-volumes-to-fs-backup: true
  default-repo-maintain-frequency: 6000
  fs-backup-timeout: 600
# ...
```

6. Wait for the DPA to reconcile successfully.

4.2.2.5.5.4. Verifying the upgrade

Use the following procedure to verify the upgrade.

Procedure

1. Verify the installation by viewing the OpenShift API for Data Protection (OADP) resources by running the following command:

```
$ oc get all -n openshift-adp
```

Example output

```
NAME                                READY STATUS RESTARTS AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2   Running 0      2m8s
pod/restic-9cq4q                               1/1   Running 0      94s
pod/restic-m4lts                               1/1   Running 0      94s
pod/restic-pv4kr                               1/1   Running 0      95s
pod/velero-588db7f655-n842v                   1/1   Running 0      95s
```

```
NAME                                TYPE      CLUSTER-IP      EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP  172.30.70.140
<none>    8443/TCP  2m8s
```

```
NAME            DESIRED  CURRENT  READY  UP-TO-DATE  AVAILABLE  NODE
SELECTOR  AGE
daemonset.apps/restic  3        3        3      3           3          <none>    96s
```

```
NAME                                READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/oadp-operator-controller-manager  1/1    1            1          2m9s
deployment.apps/velero                          1/1    1            1          96s
```

```
NAME                                DESIRED  CURRENT  READY  AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47  1        1        1      2m9s
replicaset.apps/velero-588db7f655                          1        1        1      96s
```

2. Verify that the **DataProtectionApplication** (DPA) is reconciled by running the following command:

```
$ oc get dpa dpa-sample -n openshift-adp -o jsonpath='{.status}'
```

Example output

```
{"conditions":[{"lastTransitionTime":"2023-10-27T01:23:57Z","message":"Reconcile complete","reason":"Complete","status":"True","type":"Reconciled"}]}
```

- 3. Verify the **type** is set to **Reconciled**.
- 4. Verify the backup storage location and confirm that the **PHASE** is **Available** by running the following command:

```
$ oc get backupStorageLocation -n openshift-adp
```

Example output

```
NAME          PHASE    LAST VALIDATED  AGE    DEFAULT
dpa-sample-1  Available  1s              3d16h  true
```

4.2.3. OADP 1.1 release notes

The release notes for OpenShift API for Data Protection (OADP) 1.1 describe new features and enhancements, deprecated features, product recommendations, known issues, and resolved issues.

4.2.3.1. {oadp-short} 1.1.8 release notes

The OpenShift API for Data Protection (OADP) 1.1.8 release notes lists any known issues. There are no resolved issues in this release.

4.2.3.1.1. Known issues

For a complete list of all known issues in {oadp-short} 1.1.8, see the list of [OADP 1.1.8 known issues](#) in Jira.

4.2.3.2. OADP 1.1.7 release notes

The OADP 1.1.7 release notes lists any resolved issues and known issues.

4.2.3.2.1. Resolved issues

The following highlighted issues are resolved in OADP 1.1.7:

Multiple HTTP/2 enabled web servers are vulnerable to a DDoS attack (Rapid Reset Attack)

In previous releases of OADP 1.1, the HTTP/2 protocol was susceptible to a denial of service attack because request cancellation could reset multiple streams quickly. The server had to set up and tear down the streams while not hitting any server-side limit for the maximum number of active streams per connection. This resulted in a denial of service due to server resource consumption. For a list of all OADP issues associated with this CVE, see the following [Jira list](#).

For more information, see [CVE-2023-39325 \(Rapid Reset Attack\)](#).

For a complete list of all issues resolved in the release of OADP 1.1.7, see the list of [OADP 1.1.7 resolved issues](#) in Jira.

4.2.3.2.2. Known issues

There are no known issues in the release of OADP 1.1.7.

4.2.3.3. OADP 1.1.6 release notes

The OADP 1.1.6 release notes lists any new features, resolved issues and bugs, and known issues.

4.2.3.3.1. Resolved issues

Restic restore partially failing due to Pod Security standard

OCP 4.14 introduced pod security standards that meant the **privileged** profile is **enforced**. In previous releases of OADP, this profile caused the pod to receive **permission denied** errors. This issue was caused because of the restore order. The pod was created before the security context constraints (SCC) resource. As this pod violated the pod security standard, the pod was denied and subsequently failed. [OADP-2420](#)

Restore partially failing for job resource

In previous releases of OADP, the restore of job resource was partially failing in OCP 4.14. This issue was not seen in older OCP versions. The issue was caused by an additional label being to the job resource, which was not present in older OCP versions. [OADP-2530](#)

For a complete list of all issues resolved in this release, see the list of [OADP 1.1.6 resolved issues](#) in Jira.

4.2.3.3.2. Known issues

For a complete list of all known issues in this release, see the list of [OADP 1.1.6 known issues](#) in Jira.

4.2.3.4. OADP 1.1.5 release notes

The OADP 1.1.5 release notes lists any new features, resolved issues and bugs, and known issues.

4.2.3.4.1. New features

This version of OADP is a service release. No new features are added to this version.

4.2.3.4.2. Resolved issues

For a complete list of all issues resolved in this release, see the list of [OADP 1.1.5 resolved issues](#) in Jira.

4.2.3.4.3. Known issues

For a complete list of all known issues in this release, see the list of [OADP 1.1.5 known issues](#) in Jira.

4.2.3.5. OADP 1.1.4 release notes

The OADP 1.1.4 release notes lists any new features, resolved issues and bugs, and known issues.

4.2.3.5.1. New features

This version of OADP is a service release. No new features are added to this version.

4.2.3.5.2. Resolved issues

Add support for all the velero deployment server arguments

In previous releases of OADP, OADP did not facilitate the support of all the upstream Velero server arguments. This issue has been resolved in OADP 1.1.4 and all the upstream Velero server arguments are supported. [OADP-1557](#)

Data Mover can restore from an incorrect snapshot when there was more than one VSR for the restore name and pvc name

In previous releases of OADP, OADP Data Mover could restore from an incorrect snapshot if there was more than one Volume Snapshot Restore (VSR) resource in the cluster for the same Velero **restore** name and PersistentVolumeClaim (pvc) name. [OADP-1822](#)

Cloud Storage API BSLs need OwnerReference

In previous releases of OADP, ACM BackupSchedules failed validation because of a missing **OwnerReference** on Backup Storage Locations (BSLs) created with **dpa.spec.backupLocations.bucket**. [OADP-1511](#)

For a complete list of all issues resolved in this release, see the list of [OADP 1.1.4 resolved issues](#) in Jira.

4.2.3.5.3. Known issues

This release has the following known issues:

OADP backups might fail because a UID/GID range might have changed on the cluster

OADP backups might fail because a UID/GID range might have changed on the cluster where the application has been restored, with the result that OADP does not back up and restore OpenShift Container Platform UID/GID range metadata. To avoid the issue, if the backed application requires a specific UUID, ensure the range is available when restored. An additional workaround is to allow OADP to create the namespace in the restore operation.

A restoration might fail if ArgoCD is used during the process due to a label used by ArgoCD

A restoration might fail if ArgoCD is used during the process due to a label used by ArgoCD, **app.kubernetes.io/instance**. This label identifies which resources ArgoCD needs to manage, which can create a conflict with OADP's procedure for managing resources on restoration. To work around this issue, set **.spec.resourceTrackingMethod** on the ArgoCD YAML to **annotation+label** or **annotation**. If the issue continues to persist, then disable ArgoCD before beginning to restore, and enable it again when restoration is finished.

OADP Velero plugins returning "received EOF, stopping recv loop" message

Velero plugins are started as separate processes. When the Velero operation has completed, either successfully or not, they exit. Therefore if you see a **received EOF, stopping recv loop** messages in debug logs, it does not mean an error occurred. The message indicates that a plugin operation has completed. [OADP-2176](#)

For a complete list of all known issues in this release, see the list of [OADP 1.1.4 known issues](#) in Jira.

4.2.3.6. OADP 1.1.3 release notes

The OADP 1.1.3 release notes lists any new features, resolved issues and bugs, and known issues.

4.2.3.6.1. New features

This version of OADP is a service release. No new features are added to this version.

4.2.3.6.2. Resolved issues

For a complete list of all issues resolved in this release, see the list of [OADP 1.1.3 resolved issues](#) in Jira.

4.2.3.6.3. Known issues

For a complete list of all known issues in this release, see the list of [OADP 1.1.3 known issues](#) in Jira.

4.2.3.7. OADP 1.1.2 release notes

The OADP 1.1.2 release notes include product recommendations, a list of fixed bugs and descriptions of known issues.

4.2.3.7.1. Product recommendations

VolSync

To prepare for the upgrade from VolSync 0.5.1 to the latest version available from the VolSync **stable** channel, you must add this annotation in the **openshift-adp** namespace by running the following command:

```
$ oc annotate --overwrite namespace/openshift-adp volsync.backube/privileged-movers='true'
```

Velero

In this release, Velero has been upgraded from version 1.9.2 to version [1.9.5](#).

Restic

In this release, Restic has been upgraded from version 0.13.1 to version [0.14.0](#).

4.2.3.7.2. Resolved issues

The following issues have been resolved in this release:

- [OADP-1150](#)
- [OADP-290](#)
- [OADP-1056](#)

4.2.3.7.3. Known issues

This release has the following known issues:

- OADP currently does not support backup and restore of AWS EFS volumes using restic in Velero ([OADP-778](#)).
- CSI backups might fail due to a Ceph limitation of **VolumeSnapshotContent** snapshots per PVC.
You can create many snapshots of the same persistent volume claim (PVC) but cannot schedule periodic creation of snapshots:
 - For CephFS, you can create up to 100 snapshots per PVC. ([OADP-804](#))

- For RADOS Block Device (RBD), you can create up to 512 snapshots for each PVC. ([OADP-975](#))

For more information, see [Volume Snapshots](#).

4.2.3.8. OADP 1.1.1 release notes

The OADP 1.1.1 release notes include product recommendations and descriptions of known issues.

4.2.3.8.1. Product recommendations

Before you install OADP 1.1.1, it is recommended to either install VolSync 0.5.1 or to upgrade to it.

4.2.3.8.2. Known issues

This release has the following known issues:

- Multiple HTTP/2 enabled web servers are vulnerable to a DDoS attack (Rapid Reset Attack)
The HTTP/2 protocol is susceptible to a denial of service attack because request cancellation can reset multiple streams quickly. The server has to set up and tear down the streams while not hitting any server-side limit for the maximum number of active streams per connection. This results in a denial of service due to server resource consumption. For a list of all OADP issues associated with this CVE, see the following [Jira list](#).

It is advised to upgrade to OADP 1.1.7 or 1.2.3, which resolve this issue.

For more information, see [CVE-2023-39325 \(Rapid Reset Attack\)](#).

- OADP currently does not support backup and restore of AWS EFS volumes using restic in Velero ([OADP-778](#)).
- CSI backups might fail due to a Ceph limitation of **VolumeSnapshotContent** snapshots per PVC.
You can create many snapshots of the same persistent volume claim (PVC) but cannot schedule periodic creation of snapshots:
 - For CephFS, you can create up to 100 snapshots per PVC.
 - For RADOS Block Device (RBD), you can create up to 512 snapshots for each PVC. ([OADP-804](#)) and ([OADP-975](#))
For more information, see [Volume Snapshots](#).

4.3. OADP FEATURES AND PLUGINS

OpenShift API for Data Protection (OADP) features provide options for backing up and restoring applications.

The default plugins enable Velero to integrate with certain cloud providers and to back up and restore OpenShift Container Platform resources.

4.3.1. OADP features

OpenShift API for Data Protection (OADP) supports the following features:

Backup

You can use OADP to back up all applications on the OpenShift Platform, or you can filter the resources by type, namespace, or label.

OADP backs up Kubernetes objects and internal images by saving them as an archive file on object storage. OADP backs up persistent volumes (PVs) by creating snapshots with the native cloud snapshot API or with the Container Storage Interface (CSI). For cloud providers that do not support snapshots, OADP backs up resources and PV data with Restic.



NOTE

You must exclude Operators from the backup of an application for backup and restore to succeed.

Restore

You can restore resources and PVs from a backup. You can restore all objects in a backup or filter the objects by namespace, PV, or label.



NOTE

You must exclude Operators from the backup of an application for backup and restore to succeed.

Schedule

You can schedule backups at specified intervals.

Hooks

You can use hooks to run commands in a container on a pod, for example, **fsfreeze** to freeze a file system. You can configure a hook to run before or after a backup or restore. Restore hooks can run in an init container or in the application container.

4.3.2. OADP plugins

The OpenShift API for Data Protection (OADP) provides default Velero plugins that are integrated with storage providers to support backup and snapshot operations. You can create [custom plugins](#) based on the Velero plugins.

OADP also provides plugins for OpenShift Container Platform resource backups, OpenShift Virtualization resource backups, and Container Storage Interface (CSI) snapshots.

Table 4.1. OADP plugins

OADP plugin	Function	Storage location
aws	Backs up and restores Kubernetes objects.	AWS S3
	Backs up and restores volumes with snapshots.	AWS EBS
azure	Backs up and restores Kubernetes objects.	Microsoft Azure Blob storage

OADP plugin	Function	Storage location
	Backs up and restores volumes with snapshots.	Microsoft Azure Managed Disks
gcp	Backs up and restores Kubernetes objects.	Google Cloud Storage
	Backs up and restores volumes with snapshots.	Google Compute Engine Disks
openshift	Backs up and restores OpenShift Container Platform resources. ^[1]	Object store
kubevirt	Backs up and restores OpenShift Virtualization resources. ^[2]	Object store
csi	Backs up and restores volumes with CSI snapshots. ^[3]	Cloud storage that supports CSI snapshots
vsm	VolumeSnapshotMover relocates snapshots from the cluster into an object store to be used during a restore process to recover stateful applications, in situations such as cluster deletion. ^[4]	Object store

1. Mandatory.
2. Virtual machine disks are backed up with CSI snapshots or Restic.
3. The **csi** plugin uses the Kubernetes CSI snapshot API.
 - OADP 1.1 or later uses **snapshot.storage.k8s.io/v1**
 - OADP 1.0 uses **snapshot.storage.k8s.io/v1beta1**
4. OADP 1.2 only.

4.3.3. About OADP Velero plugins

You can configure two types of plugins when you install Velero:

- Default cloud provider plugins
- Custom plugins

Both types of plugin are optional, but most users configure at least one cloud provider plugin.

4.3.3.1. Default Velero cloud provider plugins

You can install any of the following default Velero cloud provider plugins when you configure the `oadp_v1alpha1_dpa.yaml` file during deployment:

- **aws** (Amazon Web Services)
- **gcp** (Google Cloud Platform)
- **azure** (Microsoft Azure)
- **openshift** (OpenShift Velero plugin)
- **csi** (Container Storage Interface)
- **kubvirt** (KubeVirt)

You specify the desired default plugins in the `oadp_v1alpha1_dpa.yaml` file during deployment.

Example file

The following `.yaml` file installs the **openshift**, **aws**, **azure**, and **gcp** plugins:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: dpa-sample
spec:
  configuration:
    velero:
      defaultPlugins:
        - openshift
        - aws
        - azure
        - gcp
```

4.3.3.2. Custom Velero plugins

You can install a custom Velero plugin by specifying the plugin **image** and **name** when you configure the `oadp_v1alpha1_dpa.yaml` file during deployment.

You specify the desired custom plugins in the `oadp_v1alpha1_dpa.yaml` file during deployment.

Example file

The following `.yaml` file installs the default **openshift**, **azure**, and **gcp** plugins and a custom plugin that has the name **custom-plugin-example** and the image **quay.io/example-repo/custom-velero-plugin**:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: dpa-sample
spec:
  configuration:
    velero:
      defaultPlugins:
```

```

- openshift
- azure
- gcp
customPlugins:
- name: custom-plugin-example
  image: quay.io/example-repo/custom-velero-plugin

```

4.3.3.3. Velero plugins returning "received EOF, stopping recv loop" message




NOTE

Velero plugins are started as separate processes. After the Velero operation has completed, either successfully or not, they exit. Receiving a **received EOF, stopping recv loop** message in the debug logs indicates that a plugin operation has completed. It does not mean that an error has occurred.

4.3.4. Supported architectures for OADP

OpenShift API for Data Protection (OADP) supports the following architectures:

- AMD64
- ARM64
- PPC64le
- s390x



NOTE

OADP 1.2.0 and later versions support the ARM64 architecture.

4.3.5. OADP support for IBM Power and IBM Z

OpenShift API for Data Protection (OADP) is platform neutral. The information that follows relates only to IBM Power and to IBM Z.

OADP 1.1.0 was tested successfully against OpenShift Container Platform 4.11 for both IBM Power and IBM Z. The sections that follow give testing and support information for OADP 1.1.0 in terms of backup locations for these systems.

4.3.5.1. OADP support for target backup locations using IBM Power

IBM Power running with OpenShift Container Platform 4.11 and 4.12, and OpenShift API for Data Protection (OADP) 1.1.2 was tested successfully against an AWS S3 backup location target. Although the test involved only an AWS S3 target, Red Hat supports running IBM Power with OpenShift Container Platform 4.11 and 4.12, and OADP 1.1.2 against all non-AWS S3 backup location targets as well.

4.3.5.2. OADP testing and support for target backup locations using IBM Z

IBM Z running with OpenShift Container Platform 4.11 and 4.12, and OpenShift API for Data Protection (OADP) 1.1.2 was tested successfully against an AWS S3 backup location target. Although the test involved only an AWS S3 target, Red Hat supports running IBM Z with OpenShift Container Platform 4.11 and 4.12, and OADP 1.1.2 against all non-AWS S3 backup location targets as well.

4.3.6. OADP plugins known issues

The following section describes known issues in OpenShift API for Data Protection (OADP) plugins:

4.3.6.1. Velero plugin panics during imagestream backups due to a missing secret

When the backup and the Backup Storage Location (BSL) are managed outside the scope of the Data Protection Application (DPA), the OADP controller, meaning the DPA reconciliation does not create the relevant **oadp-<bsl_name>-<bsl_provider>-registry-secret**.

When the backup is run, the OpenShift Velero plugin panics on the imagestream backup, with the following panic error:

```
024-02-27T10:46:50.028951744Z time="2024-02-27T10:46:50Z" level=error msg="Error backing up item"
backup=openshift-adp/<backup name> error="error executing custom action
(groupResource=imagestreams.image.openshift.io,
namespace=<BSL Name>, name=postgres): rpc error: code = Aborted desc = plugin panicked:
runtime error: index out of range with length 1, stack trace: goroutine 94...
```

4.3.6.1.1. Workaround to avoid the panic error

To avoid the Velero plugin panic error, perform the following steps:

1. Label the custom BSL with the relevant label:

```
$ oc label BackupStorageLocation <bsl_name> app.kubernetes.io/component=bsl
```

2. After the BSL is labeled, wait until the DPA reconciles.



NOTE

You can force the reconciliation by making any minor change to the DPA itself.

3. When the DPA reconciles, confirm that the relevant **oadp-<bsl_name>-<bsl_provider>-registry-secret** has been created and that the correct registry data has been populated into it:

```
$ oc -n openshift-adp get secret/oadp-<bsl_name>-<bsl_provider>-registry-secret -o json | jq
-r '.data'
```

4.3.6.2. OpenShift ADP Controller segmentation fault

If you configure a DPA with both **cloudstorage** and **restic** enabled, the **openshift-adp-controller-manager** pod crashes and restarts indefinitely until the pod fails with a crash loop segmentation fault.

You can have either **velero** or **cloudstorage** defined, because they are mutually exclusive fields.

- If you have both **velero** and **cloudstorage** defined, the **openshift-adp-controller-manager** fails.
- If you have neither **velero** nor **cloudstorage** defined, the **openshift-adp-controller-manager** fails.

For more information about this issue, see [OADP-1054](#).

4.3.6.2.1. OpenShift ADP Controller segmentation fault workaround

You must define either **velero** or **cloudstorage** when you configure a DPA. If you define both APIs in your DPA, the **openshift-adp-controller-manager** pod fails with a crash loop segmentation fault.

4.4. INSTALLING AND CONFIGURING OADP

4.4.1. About installing OADP

As a cluster administrator, you install the OpenShift API for Data Protection (OADP) by installing the OADP Operator. The OADP Operator installs [Velero 1.12](#).



NOTE

Starting from OADP 1.0.4, all OADP 1.0.z versions can only be used as a dependency of the MTC Operator and are not available as a standalone Operator.

To back up Kubernetes resources and internal images, you must have object storage as a backup location, such as one of the following storage types:

- [Amazon Web Services](#)
- [Microsoft Azure](#)
- [Google Cloud Platform](#)
- [Multicloud Object Gateway](#)
- AWS S3 compatible object storage, such as Multicloud Object Gateway or MinIO

You can configure multiple backup storage locations within the same namespace for each individual OADP deployment.



NOTE

Unless specified otherwise, "NooBaa" refers to the open source project that provides lightweight object storage, while "Multicloud Object Gateway (MCG)" refers to the Red Hat distribution of NooBaa.

For more information on the MCG, see [Accessing the Multicloud Object Gateway with your applications](#).



IMPORTANT

The **CloudStorage** API, which automates the creation of a bucket for object storage, is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can back up persistent volumes (PVs) by using snapshots or Restic.

To back up PVs with snapshots, you must have a cloud provider that supports either a native snapshot API or Container Storage Interface (CSI) snapshots, such as one of the following cloud providers:

- [Amazon Web Services](#)
- [Microsoft Azure](#)
- [Google Cloud Platform](#)
- CSI snapshot-enabled cloud provider, such as [OpenShift Data Foundation](#)



NOTE

If you want to use CSI backup on OCP 4.11 and later, install OADP 1.1.x.

OADP 1.0.x does not support CSI backup on OCP 4.11 and later. OADP 1.0. x includes Velero 1.7.x and expects the API group **snapshot.storage.k8s.io/v1beta1**, which is not present on OCP 4.11 and later.

If your cloud provider does not support snapshots or if your storage is NFS, you can back up applications with [Restic backups](#) on object storage.

You create a default **Secret** and then you install the Data Protection Application.

4.4.1.1. AWS S3 compatible backup storage providers

OADP is compatible with many object storage providers for use with different backup and snapshot operations. Several object storage providers are fully supported, several are unsupported but known to work, and some have known limitations.

4.4.1.1.1. Supported backup storage providers

The following AWS S3 compatible object storage providers are fully supported by OADP through the AWS plugin for use as backup storage locations:

- MinIO
- Multicloud Object Gateway (MCG)
- Amazon Web Services (AWS) S3

**NOTE**

The following compatible object storage providers are supported and have their own Velero object store plugins:

- Google Cloud Platform (GCP)
- Microsoft Azure

4.4.1.1.2. Unsupported backup storage providers

The following AWS S3 compatible object storage providers, are known to work with Velero through the AWS plugin, for use as backup storage locations, however, they are unsupported and have not been tested by Red Hat:

- IBM Cloud
- Oracle Cloud
- DigitalOcean
- NooBaa, unless installed using Multicloud Object Gateway (MCG)
- Tencent Cloud
- Ceph RADOS v12.2.7
- Quobyte
- Cloudian HyperStore

**NOTE**

Unless specified otherwise, "NooBaa" refers to the open source project that provides lightweight object storage, while "Multicloud Object Gateway (MCG)" refers to the Red Hat distribution of NooBaa.

For more information on the MCG, see [Accessing the Multicloud Object Gateway with your applications](#).

4.4.1.1.3. Backup storage providers with known limitations

The following AWS S3 compatible object storage providers are known to work with Velero through the AWS plugin with a limited feature set:

- Swift - It works for use as a backup storage location for backup storage, but is not compatible with Restic for filesystem-based volume backup and restore.

4.4.1.2. Configuring Multicloud Object Gateway (MCG) for disaster recovery on OpenShift Data Foundation

If you use cluster storage for your MCG bucket **backupStorageLocation** on OpenShift Data Foundation, configure MCG as an external object store.

**WARNING**

Failure to configure MCG as an external object store might lead to backups not being available.

**NOTE**

Unless specified otherwise, "NooBaa" refers to the open source project that provides lightweight object storage, while "Multicloud Object Gateway (MCG)" refers to the Red Hat distribution of NooBaa.

For more information on the MCG, see [Accessing the Multicloud Object Gateway with your applications](#).

Procedure

- Configure MCG as an external object store as described in [Adding storage resources for hybrid or Multicloud](#).

Additional resources

- [Overview of backup and snapshot locations in the Velero documentation](#)

4.4.1.3. About OADP update channels

When you install an OADP Operator, you choose an *update channel*. This channel determines which upgrades to the OADP Operator and to Velero you receive. You can switch channels at any time.

The following update channels are available:

- The **stable** channel is now deprecated. The **stable** channel contains the patches (z-stream updates) of OADP **ClusterServiceVersion** for **oadp.v1.1.z** and older versions from **oadp.v1.0.z**.
- The **stable-1.0** channel contains **oadp.v1.0.z**, the most recent OADP 1.0 **ClusterServiceVersion**.
- The **stable-1.1** channel contains **oadp.v1.1.z**, the most recent OADP 1.1 **ClusterServiceVersion**.
- The **stable-1.2** channel contains **oadp.v1.2.z**, the most recent OADP 1.2 **ClusterServiceVersion**.
- The **stable-1.3** channel contains **oadp.v1.3.z**, the most recent OADP 1.3 **ClusterServiceVersion**.

Which update channel is right for you?

- The **stable** channel is now deprecated. If you are already using the stable channel, you will continue to get updates from **oadp.v1.1.z**.

- Choose the **stable-1.y** update channel to install OADP 1.y and to continue receiving patches for it. If you choose this channel, you will receive all z-stream patches for version 1.y.z.

When must you switch update channels?

- If you have OADP 1.y installed, and you want to receive patches only for that y-stream, you must switch from the **stable** update channel to the **stable-1.y** update channel. You will then receive all z-stream patches for version 1.y.z.
- If you have OADP 1.0 installed, want to upgrade to OADP 1.1, and then receive patches only for OADP 1.1, you must switch from the **stable-1.0** update channel to the **stable-1.1** update channel. You will then receive all z-stream patches for version 1.1.z.
- If you have OADP 1.y installed, with y greater than 0, and want to switch to OADP 1.0, you must *uninstall* your OADP Operator and then reinstall it using the **stable-1.0** update channel. You will then receive all z-stream patches for version 1.0.z.



NOTE

You cannot switch from OADP 1.y to OADP 1.0 by switching update channels. You must uninstall the Operator and then reinstall it.

4.4.1.4. Installation of OADP on multiple namespaces

You can install OADP into multiple namespaces on the same cluster so that multiple project owners can manage their own OADP instance. This use case has been validated with Restic and CSI.

You install each instance of OADP as specified by the per-platform procedures contained in this document with the following additional requirements:

- All deployments of OADP on the same cluster must be the same version, for example, 1.1.4. Installing different versions of OADP on the same cluster is **not** supported.
- Each individual deployment of OADP must have a unique set of credentials and at least one **BackupStorageLocation** configuration. You can also use multiple **BackupStorageLocation** configurations within the same namespace.
- By default, each OADP deployment has cluster-level access across namespaces. OpenShift Container Platform administrators need to review security and RBAC settings carefully and make any necessary changes to them to ensure that each OADP instance has the correct permissions.

Additional resources

- [Cluster service version](#)

4.4.1.5. Velero CPU and memory requirements based on collected data

The following recommendations are based on observations of performance made in the scale and performance lab. The backup and restore resources can be impacted by the type of plugin, the amount of resources required by that backup or restore, and the respective data contained in the persistent volumes (PVs) related to those resources.

4.4.1.5.1. CPU and memory requirement for configurations

Configuration types	[1] Average usage	[2] Large usage	resourceTimeouts
CSI	Velero: CPU- Request 200m, Limits 1000m Memory - Request 256Mi, Limits 1024Mi	Velero: CPU- Request 200m, Limits 2000m Memory- Request 256Mi, Limits 2048Mi	N/A
Restic	[3] Restic: CPU- Request 1000m, Limits 2000m Memory - Request 16Gi, Limits 32Gi	[4] Restic: CPU - Request 2000m, Limits 8000m Memory - Request 16Gi, Limits 40Gi	900m
[5] DataMover	N/A	N/A	10m - average usage 60m - large usage

1. Average usage - use these settings for most usage situations.
2. Large usage - use these settings for large usage situations, such as a large PV (500GB Usage), multiple namespaces (100+), or many pods within a single namespace (2000 pods+), and for optimal performance for backup and restore involving large datasets.
3. Restic resource usage corresponds to the amount of data, and type of data. For example, many small files or large amounts of data can cause Restic to utilize large amounts of resources. The [Velero](#) documentation references 500m as a supplied default, for most of our testing we found 200m request suitable with 1000m limit. As cited in the Velero documentation, exact CPU and memory usage is dependent on the scale of files and directories, in addition to environmental limitations.
4. Increasing the CPU has a significant impact on improving backup and restore times.
5. DataMover - DataMover default resourceTimeout is 10m. Our tests show that for restoring a large PV (500GB usage), it is required to increase the resourceTimeout to 60m.



NOTE

The resource requirements listed throughout the guide are for average usage only. For large usage, adjust the settings as described in the table above.

4.4.15.2. NodeAgent CPU for large usage

Testing shows that increasing **NodeAgent** CPU can significantly improve backup and restore times when using OpenShift API for Data Protection (OADP).



IMPORTANT

It is not recommended to use Kopia without limits in production environments on nodes running production workloads due to Kopia's aggressive consumption of resources. However, running Kopia with limits that are too low results in CPU limiting and slow backups and restore situations. Testing showed that running Kopia with 20 cores and 32 Gi memory supported backup and restore operations of over 100 GB of data, multiple namespaces, or over 2000 pods in a single namespace.

Testing detected no CPU limiting or memory saturation with these resource specifications.

You can set these limits in Ceph MDS pods by following the procedure in [Changing the CPU and memory resources on the rook-ceph pods](#).

You need to add the following lines to the storage cluster Custom Resource (CR) to set the limits:

```
resources:
  mds:
    limits:
      cpu: "3"
      memory: 128Gi
    requests:
      cpu: "3"
      memory: 8Gi
```

4.4.2. Installing the OADP Operator

You can install the OpenShift API for Data Protection (OADP) Operator on OpenShift Container Platform 4.12 by using Operator Lifecycle Manager (OLM).

The OADP Operator installs [Velero 1.12](#).

Prerequisites

- You must be logged in as a user with **cluster-admin** privileges.

Procedure

- In the OpenShift Container Platform web console, click **Operators** → **OperatorHub**.
- Use the **Filter by keyword** field to find the **OADP Operator**.
- Select the **OADP Operator** and click **Install**.
- Click **Install** to install the Operator in the **openshift-adp** project.
- Click **Operators** → **Installed Operators** to verify the installation.

4.4.2.1. OADP-Velero-OpenShift Container Platform version relationship

OADP version	Velero version	OpenShift Container Platform version
1.1.0	1.9	4.9 and later
1.1.1	1.9	4.9 and later
1.1.2	1.9	4.9 and later
1.1.3	1.9	4.9 and later
1.1.4	1.9	4.9 and later
1.1.5	1.9	4.9 and later
1.1.6	1.9	4.11 and later
1.1.7	1.9	4.11 and later
1.2.0	1.11	4.11 and later
1.2.1	1.11	4.11 and later
1.2.2	1.11	4.11 and later
1.2.3	1.11	4.11 and later
1.3.0	1.12	4.12 and later

4.4.3. Configuring the OpenShift API for Data Protection with Amazon Web Services

You install the OpenShift API for Data Protection (OADP) with Amazon Web Services (AWS) by installing the OADP Operator. The Operator installs [Velero 1.12](#).



NOTE

Starting from OADP 1.0.4, all OADP 1.0.z versions can only be used as a dependency of the MTC Operator and are not available as a standalone Operator.

You configure AWS for Velero, create a default **Secret**, and then install the Data Protection Application. For more details, see [Installing the OADP Operator](#).

To install the OADP Operator in a restricted network environment, you must first disable the default OperatorHub sources and mirror the Operator catalog. See [Using Operator Lifecycle Manager on restricted networks](#) for details.

4.4.3.1. Configuring Amazon Web Services

You configure Amazon Web Services (AWS) for the OpenShift API for Data Protection (OADP).

Prerequisites

- You must have the [AWS CLI](#) installed.

Procedure

- Set the **BUCKET** variable:

```
$ BUCKET=<your_bucket>
```

- Set the **REGION** variable:

```
$ REGION=<your_region>
```

- Create an AWS S3 bucket:

```
$ aws s3api create-bucket \
  --bucket $BUCKET \
  --region $REGION \
  --create-bucket-configuration LocationConstraint=$REGION ❶
```

- ❶ **us-east-1** does not support a **LocationConstraint**. If your region is **us-east-1**, omit **--create-bucket-configuration LocationConstraint=\$REGION**.

- Create an IAM user:

```
$ aws iam create-user --user-name velero ❶
```

- ❶ If you want to use Velero to back up multiple clusters with multiple S3 buckets, create a unique user name for each cluster.

- Create a **velero-policy.json** file:

```
$ cat > velero-policy.json <<EOF
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeVolumes",
        "ec2:DescribeSnapshots",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:CreateSnapshot",
        "ec2>DeleteSnapshot"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
```

```

        "s3:DeleteObject",
        "s3:PutObject",
        "s3:AbortMultipartUpload",
        "s3:ListMultipartUploadParts"
    ],
    "Resource": [
        "arn:aws:s3:::${BUCKET}/*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation",
        "s3:ListBucketMultipartUploads"
    ],
    "Resource": [
        "arn:aws:s3:::${BUCKET}"
    ]
  }
]
}
EOF

```

6. Attach the policies to give the **velero** user the minimum necessary permissions:

```

$ aws iam put-user-policy \
  --user-name velero \
  --policy-name velero \
  --policy-document file://velero-policy.json

```

7. Create an access key for the **velero** user:

```

$ aws iam create-access-key --user-name velero

```

Example output

```

{
  "AccessKey": {
    "UserName": "velero",
    "Status": "Active",
    "CreateDate": "2017-07-31T22:24:41.576Z",
    "SecretAccessKey": <AWS_SECRET_ACCESS_KEY>,
    "AccessKeyId": <AWS_ACCESS_KEY_ID>
  }
}

```

8. Create a **credentials-velero** file:

```

$ cat << EOF > ./credentials-velero
[default]
aws_access_key_id=<AWS_ACCESS_KEY_ID>
aws_secret_access_key=<AWS_SECRET_ACCESS_KEY>
EOF

```

You use the **credentials-velero** file to create a **Secret** object for AWS before you install the Data Protection Application.

4.4.3.2. About backup and snapshot locations and their secrets

You specify backup and snapshot locations and their secrets in the **DataProtectionApplication** custom resource (CR).

Backup locations

You specify S3-compatible object storage, such as Multicloud Object Gateway or MinIO, as a backup location.

Velero backs up OpenShift Container Platform resources, Kubernetes objects, and internal images as an archive file on object storage.

Snapshot locations

If you use your cloud provider's native snapshot API to back up persistent volumes, you must specify the cloud provider as the snapshot location.

If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a **VolumeSnapshotClass** CR to register the CSI driver.

If you use Restic, you do not need to specify a snapshot location because Restic backs up the file system on object storage.

Secrets

If the backup and snapshot locations use the same credentials or if you do not require a snapshot location, you create a default **Secret**.

If the backup and snapshot locations use different credentials, you create two secret objects:

- Custom **Secret** for the backup location, which you specify in the **DataProtectionApplication** CR.
- Default **Secret** for the snapshot location, which is not referenced in the **DataProtectionApplication** CR.



IMPORTANT

The Data Protection Application requires a default **Secret**. Otherwise, the installation will fail.

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file.

4.4.3.2.1. Creating a default Secret

You create a default **Secret** if your backup and snapshot locations use the same credentials or if you do not require a snapshot location.

The default name of the **Secret** is **cloud-credentials**.



NOTE

The **DataProtectionApplication** custom resource (CR) requires a default **Secret**. Otherwise, the installation will fail. If the name of the backup location **Secret** is not specified, the default name is used.

If you do not want to use the backup location credentials during the installation, you can create a **Secret** with the default name by using an empty **credentials-velero** file.

Prerequisites

- Your object storage and cloud storage, if any, must use the same credentials.
- You must configure object storage for Velero.
- You must create a **credentials-velero** file for the object storage in the appropriate format.

Procedure

- Create a **Secret** with the default name:

```
$ oc create secret generic cloud-credentials -n openshift-adp --from-file cloud=credentials-velero
```

The **Secret** is referenced in the **spec.backupLocations.credential** block of the **DataProtectionApplication** CR when you install the Data Protection Application.

4.4.3.2.2. Creating profiles for different credentials

If your backup and snapshot locations use different credentials, you create separate profiles in the **credentials-velero** file.

Then, you create a **Secret** object and specify the profiles in the **DataProtectionApplication** custom resource (CR).

Procedure

1. Create a **credentials-velero** file with separate profiles for the backup and snapshot locations, as in the following example:

```
[backupStorage]
aws_access_key_id=<AWS_ACCESS_KEY_ID>
aws_secret_access_key=<AWS_SECRET_ACCESS_KEY>

[volumeSnapshot]
aws_access_key_id=<AWS_ACCESS_KEY_ID>
aws_secret_access_key=<AWS_SECRET_ACCESS_KEY>
```

2. Create a **Secret** object with the **credentials-velero** file:

```
$ oc create secret generic cloud-credentials -n openshift-adp --from-file cloud=credentials-velero 1
```

3. Add the profiles to the **DataProtectionApplication** CR, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  ...
  backupLocations:
  - name: default
    velero:
      provider: aws
      default: true
      objectStorage:
        bucket: <bucket_name>
        prefix: <prefix>
      config:
        region: us-east-1
        profile: "backupStorage"
      credential:
        key: cloud
        name: cloud-credentials
  snapshotLocations:
  - velero:
      provider: aws
      config:
        region: us-west-2
        profile: "volumeSnapshot"

```

4.4.3.3. Configuring the Data Protection Application

You can configure the Data Protection Application by setting Velero resource allocations or enabling self-signed CA certificates.

4.4.3.3.1. Setting Velero CPU and memory resource allocations

You set the CPU and memory resource allocations for the **Velero** pod by editing the **DataProtectionApplication** custom resource (CR) manifest.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the values in the **spec.configuration.velero.podConfig.ResourceAllocations** block of the **DataProtectionApplication** CR manifest, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  configuration:

```

```

velero:
  podConfig:
    nodeSelector: <node selector> 1
    resourceAllocations: 2
    limits:
      cpu: "1"
      memory: 1024Mi
    requests:
      cpu: 200m
      memory: 256Mi

```

1 1 Specify the node selector to be supplied to Velero podSpec.

2 The **resourceAllocations** listed are for average usage.

4.4.3.3.2. Enabling self-signed CA certificates

You must enable a self-signed CA certificate for object storage by editing the **DataProtectionApplication** custom resource (CR) manifest to prevent a **certificate signed by unknown authority** error.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the **spec.backupLocations.velero.objectStorage.caCert** parameter and **spec.backupLocations.velero.config** parameters of the **DataProtectionApplication** CR manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  backupLocations:
  - name: default
    velero:
      provider: aws
      default: true
      objectStorage:
        bucket: <bucket>
        prefix: <prefix>
        caCert: <base64_encoded_cert_string> 1
      config:
        insecureSkipTLSVerify: "false" 2
  ...

```

1 Specify the Base64-encoded CA certificate string.

2 The **insecureSkipTLSVerify** configuration can be set to either **"true"** or **"false"**. If set to **"true"**, SSL/TLS security is disabled. If set to **"false"**, SSL/TLS security is enabled.

4.4.3.4. Installing the Data Protection Application

You install the Data Protection Application (DPA) by creating an instance of the **DataProtectionApplication** API.

Prerequisites

- You must install the OADP Operator.
- You must configure object storage as a backup location.
- If you use snapshots to back up PVs, your cloud provider must support either a native snapshot API or Container Storage Interface (CSI) snapshots.
- If the backup and snapshot locations use the same credentials, you must create a **Secret** with the default name, **cloud-credentials**.
- If the backup and snapshot locations use different credentials, you must create a **Secret** with the default name, **cloud-credentials**, which contains separate profiles for the backup and snapshot location credentials.



NOTE

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file. If there is no default **Secret**, the installation will fail.



NOTE

Velero creates a secret named **velero-repo-credentials** in the OADP namespace, which contains a default backup repository password. You can update the secret with your own password encoded as base64 **before** you run your first backup targeted to the backup repository. The value of the key to update is **Data[repository-password]**.

After you create your DPA, the first time that you run a backup targeted to the backup repository, Velero creates a backup repository whose secret is **velero-repo-credentials**, which contains either the default password or the one you replaced it with. If you update the secret password **after** the first backup, the new password will not match the password in **velero-repo-credentials**, and therefore, Velero will not be able to connect with the older backups.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. Under **Provided APIs**, click **Create instance** in the **DataProtectionApplication** box.
3. Click **YAML View** and update the parameters of the **DataProtectionApplication** manifest:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
```

```

spec:
  configuration:
    velero:
      defaultPlugins:
        - openshift 1
        - aws
      resourceTimeout: 10m 2
    restic:
      enable: true 3
      podConfig:
        nodeSelector: <node_selector> 4
  backupLocations:
    - name: default
      velero:
        provider: aws
        default: true
        objectStorage:
          bucket: <bucket_name> 5
          prefix: <prefix> 6
        config:
          region: <region>
          profile: "default"
        credential:
          key: cloud
          name: cloud-credentials 7
  snapshotLocations: 8
    - velero:
        provider: aws
        config:
          region: <region> 9
          profile: "default"

```

- 1 The **openshift** plugin is mandatory.
- 2 Specify how many minutes to wait for several Velero resources before timeout occurs, such as Velero CRD availability, volumeSnapshot deletion, and backup repository availability. The default is 10m.
- 3 Set this value to **false** if you want to disable the Restic installation. Restic deploys a daemon set, which means that Restic pods run on each working node. In OADP version 1.2 and later, you can configure Restic for backups by adding **spec.defaultVolumesToFsBackup: true** to the **Backup** CR. In OADP version 1.1, add **spec.defaultVolumesToRestic: true** to the **Backup** CR.
- 4 Specify on which nodes Restic is available. By default, Restic runs on all nodes.
- 5 Specify a bucket as the backup storage location. If the bucket is not a dedicated bucket for Velero backups, you must specify a prefix.
- 6 Specify a prefix for Velero backups, for example, **velero**, if the bucket is used for multiple purposes.
- 7 Specify the name of the **Secret** object that you created. If you do not specify this value, the default name, **cloud-credentials**, is used. If you specify a custom name, the custom name is used for the backup location.

- 8 Specify a snapshot location, unless you use CSI snapshots or Restic to back up PVs.
- 9 The snapshot location must be in the same region as the PVs.

4. Click **Create**.

5. Verify the installation by viewing the OADP resources:

```
$ oc get all -n openshift-adp
```

Example output

```
NAME                                READY STATUS  RESTARTS  AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2  Running  0        2m8s
pod/restic-9cq4q                                1/1  Running  0        94s
pod/restic-m4lts                                1/1  Running  0        94s
pod/restic-pv4kr                                1/1  Running  0        95s
pod/velero-588db7f655-n842v                    1/1  Running  0        95s
```

```
NAME                                TYPE      CLUSTER-IP  EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP  172.30.70.140
<none>      8443/TCP  2m8s
```

```
NAME          DESIRED  CURRENT  READY  UP-TO-DATE  AVAILABLE  NODE
SELECTOR  AGE
daemonset.apps/restic  3        3        3      3           3          <none>      96s
```

```
NAME                                READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/oadp-operator-controller-manager  1/1    1           1          2m9s
deployment.apps/velero                          1/1    1           1          96s
```

```
NAME                                DESIRED  CURRENT  READY  AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47  1        1        1      2m9s
replicaset.apps/velero-588db7f655                          1        1        1      96s
```

4.4.3.4.1. Enabling CSI in the DataProtectionApplication CR

You enable the Container Storage Interface (CSI) in the **DataProtectionApplication** custom resource (CR) in order to back up persistent volumes with CSI snapshots.

Prerequisites

- The cloud provider must support CSI snapshots.

Procedure

- Edit the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
spec:
```

```
configuration:
  velero:
    defaultPlugins:
      - openshift
      - csi 1
```

- 1** Add the **csi** default plugin.

4.4.4. Configuring the OpenShift API for Data Protection with Microsoft Azure

You install the OpenShift API for Data Protection (OADP) with Microsoft Azure by installing the OADP Operator. The Operator installs [Velero 1.12](#).



NOTE

Starting from OADP 1.0.4, all OADP 1.0.z versions can only be used as a dependency of the MTC Operator and are not available as a standalone Operator.

You configure Azure for Velero, create a default **Secret**, and then install the Data Protection Application. For more details, see [Installing the OADP Operator](#).

To install the OADP Operator in a restricted network environment, you must first disable the default OperatorHub sources and mirror the Operator catalog. See [Using Operator Lifecycle Manager on restricted networks](#) for details.

4.4.4.1. Configuring Microsoft Azure

You configure a Microsoft Azure for the OpenShift API for Data Protection (OADP).

Prerequisites

- You must have the [Azure CLI](#) installed.

Procedure

1. Log in to Azure:

```
$ az login
```

2. Set the **AZURE_RESOURCE_GROUP** variable:

```
$ AZURE_RESOURCE_GROUP=Velero_Backups
```

3. Create an Azure resource group:

```
$ az group create -n $AZURE_RESOURCE_GROUP --location CentralUS 1
```

- 1** Specify your location.

4. Set the **AZURE_STORAGE_ACCOUNT_ID** variable:

```
$ AZURE_STORAGE_ACCOUNT_ID=$(velero$(uuidgen | cut -d '-' -f5 | tr 'A-Z' 'a-z'))
```

5. Create an Azure storage account:

```
$ az storage account create \
  --name $AZURE_STORAGE_ACCOUNT_ID \
  --resource-group $AZURE_RESOURCE_GROUP \
  --sku Standard_GRS \
  --encryption-services blob \
  --https-only true \
  --kind BlobStorage \
  --access-tier Hot
```

6. Set the **BLOB_CONTAINER** variable:

```
$ BLOB_CONTAINER=velero
```

7. Create an Azure Blob storage container:

```
$ az storage container create \
  -n $BLOB_CONTAINER \
  --public-access off \
  --account-name $AZURE_STORAGE_ACCOUNT_ID
```

8. Obtain the storage account access key:

```
$ AZURE_STORAGE_ACCOUNT_ACCESS_KEY=$(az storage account keys list \
  --account-name $AZURE_STORAGE_ACCOUNT_ID \
  --query "[?keyName == 'key1'].value" -o tsv)
```

9. Create a custom role that has the minimum required permissions:

```
AZURE_ROLE=Velero
az role definition create --role-definition '{
  "Name": "$AZURE_ROLE",
  "Description": "Velero related permissions to perform backups, restores and deletions",
  "Actions": [
    "Microsoft.Compute/disks/read",
    "Microsoft.Compute/disks/write",
    "Microsoft.Compute/disks/endGetAccess/action",
    "Microsoft.Compute/disks/beginGetAccess/action",
    "Microsoft.Compute/snapshots/read",
    "Microsoft.Compute/snapshots/write",
    "Microsoft.Compute/snapshots/delete",
    "Microsoft.Storage/storageAccounts/listkeys/action",
    "Microsoft.Storage/storageAccounts/regeneratekey/action"
  ],
  "AssignableScopes": ["/subscriptions/$AZURE_SUBSCRIPTION_ID"]
}'
```

10. Create a **credentials-velero** file:

```
$ cat << EOF > ./credentials-velero
```

```

AZURE_SUBSCRIPTION_ID=${AZURE_SUBSCRIPTION_ID}
AZURE_TENANT_ID=${AZURE_TENANT_ID}
AZURE_CLIENT_ID=${AZURE_CLIENT_ID}
AZURE_CLIENT_SECRET=${AZURE_CLIENT_SECRET}
AZURE_RESOURCE_GROUP=${AZURE_RESOURCE_GROUP}
AZURE_STORAGE_ACCOUNT_ACCESS_KEY=${AZURE_STORAGE_ACCOUNT_ACCESS_KEY} ❶
AZURE_CLOUD_NAME=AzurePublicCloud
EOF

```

- ❶ Mandatory. You cannot back up internal images if the **credentials-velero** file contains only the service principal credentials.

You use the **credentials-velero** file to create a **Secret** object for Azure before you install the Data Protection Application.

4.4.4.2. About backup and snapshot locations and their secrets

You specify backup and snapshot locations and their secrets in the **DataProtectionApplication** custom resource (CR).

Backup locations

You specify S3-compatible object storage, such as Multicloud Object Gateway or MinIO, as a backup location.

Velero backs up OpenShift Container Platform resources, Kubernetes objects, and internal images as an archive file on object storage.

Snapshot locations

If you use your cloud provider's native snapshot API to back up persistent volumes, you must specify the cloud provider as the snapshot location.

If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a **VolumeSnapshotClass** CR to register the CSI driver.

If you use Restic, you do not need to specify a snapshot location because Restic backs up the file system on object storage.

Secrets

If the backup and snapshot locations use the same credentials or if you do not require a snapshot location, you create a default **Secret**.

If the backup and snapshot locations use different credentials, you create two secret objects:

- Custom **Secret** for the backup location, which you specify in the **DataProtectionApplication** CR.
- Default **Secret** for the snapshot location, which is not referenced in the **DataProtectionApplication** CR.



IMPORTANT

The Data Protection Application requires a default **Secret**. Otherwise, the installation will fail.

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file.

4.4.4.2.1. Creating a default Secret

You create a default **Secret** if your backup and snapshot locations use the same credentials or if you do not require a snapshot location.

The default name of the **Secret** is **cloud-credentials-azure**.



NOTE

The **DataProtectionApplication** custom resource (CR) requires a default **Secret**. Otherwise, the installation will fail. If the name of the backup location **Secret** is not specified, the default name is used.

If you do not want to use the backup location credentials during the installation, you can create a **Secret** with the default name by using an empty **credentials-velero** file.

Prerequisites

- Your object storage and cloud storage, if any, must use the same credentials.
- You must configure object storage for Velero.
- You must create a **credentials-velero** file for the object storage in the appropriate format.

Procedure

- Create a **Secret** with the default name:

```
$ oc create secret generic cloud-credentials-azure -n openshift-adp --from-file
cloud=credentials-velero
```

The **Secret** is referenced in the **spec.backupLocations.credential** block of the **DataProtectionApplication** CR when you install the Data Protection Application.

4.4.4.2.2. Creating secrets for different credentials

If your backup and snapshot locations use different credentials, you must create two **Secret** objects:

- Backup location **Secret** with a custom name. The custom name is specified in the **spec.backupLocations** block of the **DataProtectionApplication** custom resource (CR).
- Snapshot location **Secret** with the default name, **cloud-credentials-azure**. This **Secret** is not specified in the **DataProtectionApplication** CR.

Procedure

1. Create a **credentials-velero** file for the object storage in the appropriate format.

1. Create a **credentials-velero** file for the snapshot location in the appropriate format for your cloud provider.
2. Create a **Secret** for the snapshot location with the default name:

```
$ oc create secret generic cloud-credentials-azure -n openshift-adp --from-file
cloud=credentials-velero
```

3. Create a **credentials-velero** file for the backup location in the appropriate format for your object storage.
4. Create a **Secret** for the backup location with a custom name:

```
$ oc create secret generic <custom_secret> -n openshift-adp --from-file cloud=credentials-
velero
```

5. Add the **Secret** with the custom name to the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  ...
  backupLocations:
    - velero:
      config:
        resourceGroup: <azure_resource_group>
        storageAccount: <azure_storage_account_id>
        subscriptionId: <azure_subscription_id>
        storageAccountKeyEnvVar: AZURE_STORAGE_ACCOUNT_ACCESS_KEY
      credential:
        key: cloud
        name: <custom_secret> 1
      provider: azure
      default: true
      objectStorage:
        bucket: <bucket_name>
        prefix: <prefix>
  snapshotLocations:
    - velero:
      config:
        resourceGroup: <azure_resource_group>
        subscriptionId: <azure_subscription_id>
        incremental: "true"
      provider: azure
```

- 1 Backup location **Secret** with custom name.

4.4.4.3. Configuring the Data Protection Application

You can configure the Data Protection Application by setting Velero resource allocations or enabling self-signed CA certificates.

4.4.4.3.1. Setting Velero CPU and memory resource allocations

You set the CPU and memory resource allocations for the **Velero** pod by editing the **DataProtectionApplication** custom resource (CR) manifest.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the values in the **spec.configuration.velero.podConfig.ResourceAllocations** block of the **DataProtectionApplication** CR manifest, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  configuration:
    velero:
      podConfig:
        nodeSelector: <node selector> 1
        resourceAllocations: 2
          limits:
            cpu: "1"
            memory: 1024Mi
          requests:
            cpu: 200m
            memory: 256Mi
```

1 Specify the node selector to be supplied to Velero podSpec.

2 The **resourceAllocations** listed are for average usage.

4.4.4.3.2. Enabling self-signed CA certificates

You must enable a self-signed CA certificate for object storage by editing the **DataProtectionApplication** custom resource (CR) manifest to prevent a **certificate signed by unknown authority** error.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the **spec.backupLocations.velero.objectStorage.caCert** parameter and **spec.backupLocations.velero.config** parameters of the **DataProtectionApplication** CR manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  backupLocations:
    - name: default
      velero:
        provider: aws
        default: true
        objectStorage:
          bucket: <bucket>
          prefix: <prefix>
          caCert: <base64_encoded_cert_string> ❶
        config:
          insecureSkipTLSVerify: "false" ❷
  ...

```

- ❶ Specify the Base64-encoded CA certificate string.
- ❷ The **insecureSkipTLSVerify** configuration can be set to either **"true"** or **"false"**. If set to **"true"**, SSL/TLS security is disabled. If set to **"false"**, SSL/TLS security is enabled.

4.4.4.4. Installing the Data Protection Application

You install the Data Protection Application (DPA) by creating an instance of the **DataProtectionApplication** API.

Prerequisites

- You must install the OADP Operator.
- You must configure object storage as a backup location.
- If you use snapshots to back up PVs, your cloud provider must support either a native snapshot API or Container Storage Interface (CSI) snapshots.
- If the backup and snapshot locations use the same credentials, you must create a **Secret** with the default name, **cloud-credentials-azure**.
- If the backup and snapshot locations use different credentials, you must create two **Secrets**:
 - **Secret** with a custom name for the backup location. You add this **Secret** to the **DataProtectionApplication** CR.
 - **Secret** with the default name, **cloud-credentials-azure**, for the snapshot location. This **Secret** is not referenced in the **DataProtectionApplication** CR.

**NOTE**

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file. If there is no default **Secret**, the installation will fail.

**NOTE**

Velero creates a secret named **velero-repo-credentials** in the OADP namespace, which contains a default backup repository password. You can update the secret with your own password encoded as base64 **before** you run your first backup targeted to the backup repository. The value of the key to update is **Data[repository-password]**.

After you create your DPA, the first time that you run a backup targeted to the backup repository, Velero creates a backup repository whose secret is **velero-repo-credentials**, which contains either the default password or the one you replaced it with. If you update the secret password **after** the first backup, the new password will not match the password in **velero-repo-credentials**, and therefore, Velero will not be able to connect with the older backups.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. Under **Provided APIs**, click **Create instance** in the **DataProtectionApplication** box.
3. Click **YAML View** and update the parameters of the **DataProtectionApplication** manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  configuration:
    velero:
      defaultPlugins:
        - azure
        - openshift 1
      resourceTimeout: 10m 2
    restic:
      enable: true 3
      podConfig:
        nodeSelector: <node_selector> 4
  backupLocations:
    - velero:
      config:
        resourceGroup: <azure_resource_group> 5
        storageAccount: <azure_storage_account_id> 6
        subscriptionId: <azure_subscription_id> 7
        storageAccountKeyEnvVar: AZURE_STORAGE_ACCOUNT_ACCESS_KEY
      credential:
        key: cloud

```

```

    name: cloud-credentials-azure 8
    provider: azure
    default: true
    objectStorage:
      bucket: <bucket_name> 9
      prefix: <prefix> 10
snapshotLocations: 11
- velero:
  config:
    resourceGroup: <azure_resource_group>
    subscriptionId: <azure_subscription_id>
    incremental: "true"
  name: default
  provider: azure

```

- 1** The **openshift** plugin is mandatory.
- 2** Specify how many minutes to wait for several Velero resources before timeout occurs, such as Velero CRD availability, volumeSnapshot deletion, and backup repository availability. The default is 10m.
- 3** Set this value to **false** if you want to disable the Restic installation. Restic deploys a daemon set, which means that Restic pods run on each working node. In OADP version 1.2 and later, you can configure Restic for backups by adding **spec.defaultVolumesToFsBackup: true** to the **Backup** CR. In OADP version 1.1, add **spec.defaultVolumesToRestic: true** to the **Backup** CR.
- 4** Specify on which nodes Restic is available. By default, Restic runs on all nodes.
- 5** Specify the Azure resource group.
- 6** Specify the Azure storage account ID.
- 7** Specify the Azure subscription ID.
- 8** If you do not specify this value, the default name, **cloud-credentials-azure**, is used. If you specify a custom name, the custom name is used for the backup location.
- 9** Specify a bucket as the backup storage location. If the bucket is not a dedicated bucket for Velero backups, you must specify a prefix.
- 10** Specify a prefix for Velero backups, for example, **velero**, if the bucket is used for multiple purposes.
- 11** You do not need to specify a snapshot location if you use CSI snapshots or Restic to back up PVs.

4. Click **Create**.

5. Verify the installation by viewing the OADP resources:

```
$ oc get all -n openshift-adp
```

Example output

```

NAME                                READY STATUS RESTARTS AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2   Running 0      2m8s
pod/restic-9cq4q                                1/1   Running 0      94s
pod/restic-m4lts                                1/1   Running 0      94s
pod/restic-pv4kr                                1/1   Running 0      95s
pod/velero-588db7f655-n842v                    1/1   Running 0      95s

```

```

NAME                                TYPE          CLUSTER-IP    EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP    172.30.70.140
<none>      8443/TCP    2m8s

```

```

NAME            DESIRED CURRENT READY UP-TO-DATE AVAILABLE NODE
SELECTOR AGE
daemonset.apps/restic 3      3      3      3      3      <none> 96s

```

```

NAME                                READY UP-TO-DATE AVAILABLE AGE
deployment.apps/oadp-operator-controller-manager  1/1   1      1      2m9s
deployment.apps/velero                          1/1   1      1      96s

```

```

NAME                                DESIRED CURRENT READY AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47  1      1      1      2m9s
replicaset.apps/velero-588db7f655                          1      1      1      96s

```

4.4.4.4.1. Enabling CSI in the DataProtectionApplication CR

You enable the Container Storage Interface (CSI) in the **DataProtectionApplication** custom resource (CR) in order to back up persistent volumes with CSI snapshots.

Prerequisites

- The cloud provider must support CSI snapshots.

Procedure

- Edit the **DataProtectionApplication** CR, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
spec:
  configuration:
    velero:
      defaultPlugins:
        - openshift
        - csi 1

```

- 1 Add the **csi** default plugin.

4.4.5. Configuring the OpenShift API for Data Protection with Google Cloud Platform

You install the OpenShift API for Data Protection (OADP) with Google Cloud Platform (GCP) by installing the OADP Operator. The Operator installs [Velero 1.12](#).



NOTE

Starting from OADP 1.0.4, all OADP 1.0.z versions can only be used as a dependency of the MTC Operator and are not available as a standalone Operator.

You configure GCP for Velero, create a default **Secret**, and then install the Data Protection Application. For more details, see [Installing the OADP Operator](#).

To install the OADP Operator in a restricted network environment, you must first disable the default OperatorHub sources and mirror the Operator catalog. See [Using Operator Lifecycle Manager on restricted networks](#) for details.

4.4.5.1. Configuring Google Cloud Platform

You configure Google Cloud Platform (GCP) for the OpenShift API for Data Protection (OADP).

Prerequisites

- You must have the **gcloud** and **gsutil** CLI tools installed. See the [Google cloud documentation](#) for details.

Procedure

1. Log in to GCP:

```
$ gcloud auth login
```

2. Set the **BUCKET** variable:

```
$ BUCKET=<bucket> 1
```

- 1** Specify your bucket name.

3. Create the storage bucket:

```
$ gsutil mb gs://$BUCKET/
```

4. Set the **PROJECT_ID** variable to your active project:

```
$ PROJECT_ID=$(gcloud config get-value project)
```

5. Create a service account:

```
$ gcloud iam service-accounts create velero \  
--display-name "Velero service account"
```

6. List your service accounts:

```
$ gcloud iam service-accounts list
```

7. Set the **SERVICE_ACCOUNT_EMAIL** variable to match its **email** value:

```
$ SERVICE_ACCOUNT_EMAIL=$(gcloud iam service-accounts list \
  --filter="displayName:Velero service account" \
  --format 'value(email)')
```

8. Attach the policies to give the **velero** user the minimum necessary permissions:

```
$ ROLE_PERMISSIONS=(
  compute.disks.get
  compute.disks.create
  compute.disks.createSnapshot
  compute.snapshots.get
  compute.snapshots.create
  compute.snapshots.useReadOnly
  compute.snapshots.delete
  compute.zones.get
  storage.objects.create
  storage.objects.delete
  storage.objects.get
  storage.objects.list
  iam.serviceAccounts.signBlob
)
```

9. Create the **velero.server** custom role:

```
$ gcloud iam roles create velero.server \
  --project $PROJECT_ID \
  --title "Velero Server" \
  --permissions "${IFS=","; echo "${ROLE_PERMISSIONS[*]}")"
```

10. Add IAM policy binding to the project:

```
$ gcloud projects add-iam-policy-binding $PROJECT_ID \
  --member serviceAccount:$SERVICE_ACCOUNT_EMAIL \
  --role projects/$PROJECT_ID/roles/velero.server
```

11. Update the IAM service account:

```
$ gsutil iam ch serviceAccount:$SERVICE_ACCOUNT_EMAIL:objectAdmin gs://${BUCKET}
```

12. Save the IAM service account keys to the **credentials-velero** file in the current directory:

```
$ gcloud iam service-accounts keys create credentials-velero \
  --iam-account $SERVICE_ACCOUNT_EMAIL
```

You use the **credentials-velero** file to create a **Secret** object for GCP before you install the Data Protection Application.

4.4.5.2. About backup and snapshot locations and their secrets

You specify backup and snapshot locations and their secrets in the **DataProtectionApplication** custom resource (CR).

Backup locations

You specify S3-compatible object storage, such as Multicloud Object Gateway or MinIO, as a backup location.

Velero backs up OpenShift Container Platform resources, Kubernetes objects, and internal images as an archive file on object storage.

Snapshot locations

If you use your cloud provider's native snapshot API to back up persistent volumes, you must specify the cloud provider as the snapshot location.

If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a **VolumeSnapshotClass** CR to register the CSI driver.

If you use Restic, you do not need to specify a snapshot location because Restic backs up the file system on object storage.

Secrets

If the backup and snapshot locations use the same credentials or if you do not require a snapshot location, you create a default **Secret**.

If the backup and snapshot locations use different credentials, you create two secret objects:

- Custom **Secret** for the backup location, which you specify in the **DataProtectionApplication** CR.
- Default **Secret** for the snapshot location, which is not referenced in the **DataProtectionApplication** CR.



IMPORTANT

The Data Protection Application requires a default **Secret**. Otherwise, the installation will fail.

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file.

4.4.5.2.1. Creating a default Secret

You create a default **Secret** if your backup and snapshot locations use the same credentials or if you do not require a snapshot location.

The default name of the **Secret** is **cloud-credentials-gcp**.



NOTE

The **DataProtectionApplication** custom resource (CR) requires a default **Secret**. Otherwise, the installation will fail. If the name of the backup location **Secret** is not specified, the default name is used.

If you do not want to use the backup location credentials during the installation, you can create a **Secret** with the default name by using an empty **credentials-velero** file.

Prerequisites

- Your object storage and cloud storage, if any, must use the same credentials.

- You must configure object storage for Velero.
- You must create a **credentials-velero** file for the object storage in the appropriate format.

Procedure

- Create a **Secret** with the default name:

```
$ oc create secret generic cloud-credentials-gcp -n openshift-adp --from-file
cloud=credentials-velero
```

The **Secret** is referenced in the **spec.backupLocations.credential** block of the **DataProtectionApplication** CR when you install the Data Protection Application.

4.4.5.2.2. Creating secrets for different credentials

If your backup and snapshot locations use different credentials, you must create two **Secret** objects:

- Backup location **Secret** with a custom name. The custom name is specified in the **spec.backupLocations** block of the **DataProtectionApplication** custom resource (CR).
- Snapshot location **Secret** with the default name, **cloud-credentials-gcp**. This **Secret** is not specified in the **DataProtectionApplication** CR.

Procedure

1. Create a **credentials-velero** file for the snapshot location in the appropriate format for your cloud provider.
2. Create a **Secret** for the snapshot location with the default name:

```
$ oc create secret generic cloud-credentials-gcp -n openshift-adp --from-file
cloud=credentials-velero
```

3. Create a **credentials-velero** file for the backup location in the appropriate format for your object storage.
4. Create a **Secret** for the backup location with a custom name:

```
$ oc create secret generic <custom_secret> -n openshift-adp --from-file cloud=credentials-
velero
```

5. Add the **Secret** with the custom name to the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  ...
  backupLocations:
    - velero:
```

```

provider: gcp
default: true
credential:
  key: cloud
  name: <custom_secret> 1
objectStorage:
  bucket: <bucket_name>
  prefix: <prefix>
snapshotLocations:
- velero:
  provider: gcp
  default: true
  config:
    project: <project>
    snapshotLocation: us-west1

```

- 1 Backup location **Secret** with custom name.

4.4.5.3. Configuring the Data Protection Application

You can configure the Data Protection Application by setting Velero resource allocations or enabling self-signed CA certificates.

4.4.5.3.1. Setting Velero CPU and memory resource allocations

You set the CPU and memory resource allocations for the **Velero** pod by editing the **DataProtectionApplication** custom resource (CR) manifest.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the values in the **spec.configuration.velero.podConfig.ResourceAllocations** block of the **DataProtectionApplication** CR manifest, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  configuration:
    velero:
      podConfig:
        nodeSelector: <node selector> 1
        resourceAllocations: 2
          limits:
            cpu: "1"
            memory: 1024Mi
          requests:
            cpu: 200m
            memory: 256Mi

```


-
- 1 Specify the node selector to be supplied to Velero podSpec.
- 2 The **resourceAllocations** listed are for average usage.

4.4.5.3.2. Enabling self-signed CA certificates

You must enable a self-signed CA certificate for object storage by editing the **DataProtectionApplication** custom resource (CR) manifest to prevent a **certificate signed by unknown authority** error.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the **spec.backupLocations.velero.objectStorage.caCert** parameter and **spec.backupLocations.velero.config** parameters of the **DataProtectionApplication** CR manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  backupLocations:
    - name: default
      velero:
        provider: aws
        default: true
        objectStorage:
          bucket: <bucket>
          prefix: <prefix>
          caCert: <base64_encoded_cert_string> 1
        config:
          insecureSkipTLSVerify: "false" 2
  ...

```

- 1 Specify the Base64-encoded CA certificate string.
- 2 The **insecureSkipTLSVerify** configuration can be set to either **"true"** or **"false"**. If set to **"true"**, SSL/TLS security is disabled. If set to **"false"**, SSL/TLS security is enabled.

4.4.5.4. Installing the Data Protection Application

You install the Data Protection Application (DPA) by creating an instance of the **DataProtectionApplication** API.

Prerequisites

- You must install the OADP Operator.

- You must configure object storage as a backup location.
- If you use snapshots to back up PVs, your cloud provider must support either a native snapshot API or Container Storage Interface (CSI) snapshots.
- If the backup and snapshot locations use the same credentials, you must create a **Secret** with the default name, **cloud-credentials-gcp**.
- If the backup and snapshot locations use different credentials, you must create two **Secrets**:
 - **Secret** with a custom name for the backup location. You add this **Secret** to the **DataProtectionApplication** CR.
 - **Secret** with the default name, **cloud-credentials-gcp**, for the snapshot location. This **Secret** is not referenced in the **DataProtectionApplication** CR.



NOTE

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file. If there is no default **Secret**, the installation will fail.



NOTE

Velero creates a secret named **velero-repo-credentials** in the OADP namespace, which contains a default backup repository password. You can update the secret with your own password encoded as base64 **before** you run your first backup targeted to the backup repository. The value of the key to update is **Data[repository-password]**.

After you create your DPA, the first time that you run a backup targeted to the backup repository, Velero creates a backup repository whose secret is **velero-repo-credentials**, which contains either the default password or the one you replaced it with. If you update the secret password **after** the first backup, the new password will not match the password in **velero-repo-credentials**, and therefore, Velero will not be able to connect with the older backups.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. Under **Provided APIs**, click **Create instance** in the **DataProtectionApplication** box.
3. Click **YAML View** and update the parameters of the **DataProtectionApplication** manifest:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  configuration:
    velero:
      defaultPlugins:
        - gcp
```

```

- openshift 1
resourceTimeout: 10m 2
restic:
  enable: true 3
  podConfig:
    nodeSelector: <node_selector> 4
backupLocations:
- velero:
  provider: gcp
  default: true
  credential:
    key: cloud
    name: cloud-credentials-gcp 5
  objectStorage:
    bucket: <bucket_name> 6
    prefix: <prefix> 7
snapshotLocations: 8
- velero:
  provider: gcp
  default: true
  config:
    project: <project>
    snapshotLocation: us-west1 9

```

- 1** The **openshift** plugin is mandatory.
- 2** Specify how many minutes to wait for several Velero resources before timeout occurs, such as Velero CRD availability, volumeSnapshot deletion, and backup repository availability. The default is 10m.
- 3** Set this value to **false** if you want to disable the Restic installation. Restic deploys a daemon set, which means that Restic pods run on each working node. In OADP version 1.2 and later, you can configure Restic for backups by adding **spec.defaultVolumesToFsBackup: true** to the **Backup** CR. In OADP version 1.1, add **spec.defaultVolumesToRestic: true** to the **Backup** CR.
- 4** Specify on which nodes Restic is available. By default, Restic runs on all nodes.
- 5** If you do not specify this value, the default name, **cloud-credentials-gcp**, is used. If you specify a custom name, the custom name is used for the backup location.
- 6** Specify a bucket as the backup storage location. If the bucket is not a dedicated bucket for Velero backups, you must specify a prefix.
- 7** Specify a prefix for Velero backups, for example, **velero**, if the bucket is used for multiple purposes.
- 8** Specify a snapshot location, unless you use CSI snapshots or Restic to back up PVs.
- 9** The snapshot location must be in the same region as the PVs.

4. Click **Create**.

5. Verify the installation by viewing the OADP resources:

■

```
$ oc get all -n openshift-adp
```

Example output

```
NAME                                READY STATUS  RESTARTS  AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2   Running  0         2m8s
pod/restic-9cq4q                                1/1   Running  0         94s
pod/restic-m4lts                                1/1   Running  0         94s
pod/restic-pv4kr                                1/1   Running  0         95s
pod/velero-588db7f655-n842v                    1/1   Running  0         95s
```

```
NAME                                TYPE          CLUSTER-IP      EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP  172.30.70.140
<none>    8443/TCP  2m8s
```

```
NAME          DESIRED  CURRENT  READY  UP-TO-DATE  AVAILABLE  NODE
SELECTOR  AGE
daemonset.apps/restic  3        3        3      3           3          <none>    96s
```

```
NAME                                READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/oadp-operator-controller-manager  1/1    1           1         2m9s
deployment.apps/velero                          1/1    1           1         96s
```

```
NAME                                DESIRED  CURRENT  READY  AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47  1        1        1      2m9s
replicaset.apps/velero-588db7f655                          1        1        1      96s
```

4.4.5.4.1. Enabling CSI in the DataProtectionApplication CR

You enable the Container Storage Interface (CSI) in the **DataProtectionApplication** custom resource (CR) in order to back up persistent volumes with CSI snapshots.

Prerequisites

- The cloud provider must support CSI snapshots.

Procedure

- Edit the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
spec:
  configuration:
    velero:
      defaultPlugins:
        - openshift
        - csi 1
```

- Add the **csi** default plugin.

4.4.6. Configuring the OpenShift API for Data Protection with Multicloud Object Gateway

You install the OpenShift API for Data Protection (OADP) with Multicloud Object Gateway (MCG) by installing the OADP Operator. The Operator installs [Velero 1.12](#).



NOTE

Starting from OADP 1.0.4, all OADP 1.0.z versions can only be used as a dependency of the MTC Operator and are not available as a standalone Operator.

You configure [Multicloud Object Gateway](#) as a backup location. MCG is a component of OpenShift Data Foundation. You configure MCG as a backup location in the **DataProtectionApplication** custom resource (CR).



IMPORTANT

The **CloudStorage** API, which automates the creation of a bucket for object storage, is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You create a **Secret** for the backup location and then you install the Data Protection Application. For more details, see [Installing the OADP Operator](#).

To install the OADP Operator in a restricted network environment, you must first disable the default OperatorHub sources and mirror the Operator catalog. For details, see [Using Operator Lifecycle Manager on restricted networks](#).

4.4.6.1. Retrieving Multicloud Object Gateway credentials

You must retrieve the Multicloud Object Gateway (MCG) credentials in order to create a **Secret** custom resource (CR) for the OpenShift API for Data Protection (OADP).

MCG is a component of OpenShift Data Foundation.

Prerequisites

- You must deploy OpenShift Data Foundation by using the appropriate [OpenShift Data Foundation deployment guide](#).

Procedure

- Obtain the S3 endpoint, **AWS_ACCESS_KEY_ID**, and **AWS_SECRET_ACCESS_KEY** by running the [describe command](#) on the **NooBaa** custom resource.
- Create a **credentials-velero** file:

```
$ cat << EOF > ./credentials-velero
```

```
[default]
aws_access_key_id=<AWS_ACCESS_KEY_ID>
aws_secret_access_key=<AWS_SECRET_ACCESS_KEY>
EOF
```

You use the **credentials-velero** file to create a **Secret** object when you install the Data Protection Application.

4.4.6.2. About backup and snapshot locations and their secrets

You specify backup and snapshot locations and their secrets in the **DataProtectionApplication** custom resource (CR).

Backup locations

You specify S3-compatible object storage, such as Multicloud Object Gateway or MinIO, as a backup location.

Velero backs up OpenShift Container Platform resources, Kubernetes objects, and internal images as an archive file on object storage.

Snapshot locations

If you use your cloud provider's native snapshot API to back up persistent volumes, you must specify the cloud provider as the snapshot location.

If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a **VolumeSnapshotClass** CR to register the CSI driver.

If you use Restic, you do not need to specify a snapshot location because Restic backs up the file system on object storage.

Secrets

If the backup and snapshot locations use the same credentials or if you do not require a snapshot location, you create a default **Secret**.

If the backup and snapshot locations use different credentials, you create two secret objects:

- Custom **Secret** for the backup location, which you specify in the **DataProtectionApplication** CR.
- Default **Secret** for the snapshot location, which is not referenced in the **DataProtectionApplication** CR.



IMPORTANT

The Data Protection Application requires a default **Secret**. Otherwise, the installation will fail.

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file.

4.4.6.2.1. Creating a default Secret

You create a default **Secret** if your backup and snapshot locations use the same credentials or if you do not require a snapshot location.

The default name of the **Secret** is **cloud-credentials**.



NOTE

The **DataProtectionApplication** custom resource (CR) requires a default **Secret**. Otherwise, the installation will fail. If the name of the backup location **Secret** is not specified, the default name is used.

If you do not want to use the backup location credentials during the installation, you can create a **Secret** with the default name by using an empty **credentials-velero** file.

Prerequisites

- Your object storage and cloud storage, if any, must use the same credentials.
- You must configure object storage for Velero.
- You must create a **credentials-velero** file for the object storage in the appropriate format.

Procedure

- Create a **Secret** with the default name:

```
$ oc create secret generic cloud-credentials -n openshift-adp --from-file cloud=credentials-velero
```

The **Secret** is referenced in the **spec.backupLocations.credential** block of the **DataProtectionApplication** CR when you install the Data Protection Application.

4.4.6.2.2. Creating secrets for different credentials

If your backup and snapshot locations use different credentials, you must create two **Secret** objects:

- Backup location **Secret** with a custom name. The custom name is specified in the **spec.backupLocations** block of the **DataProtectionApplication** custom resource (CR).
- Snapshot location **Secret** with the default name, **cloud-credentials**. This **Secret** is not specified in the **DataProtectionApplication** CR.

Procedure

1. Create a **credentials-velero** file for the snapshot location in the appropriate format for your cloud provider.
2. Create a **Secret** for the snapshot location with the default name:

```
$ oc create secret generic cloud-credentials -n openshift-adp --from-file cloud=credentials-velero
```

3. Create a **credentials-velero** file for the backup location in the appropriate format for your object storage.
4. Create a **Secret** for the backup location with a custom name:

```
$ oc create secret generic <custom_secret> -n openshift-adp --from-file cloud=credentials-velero
```

5. Add the **Secret** with the custom name to the **DataProtectionApplication** CR, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  ...
  backupLocations:
    - velero:
      config:
        profile: "default"
        region: minio
        s3Url: <url>
        insecureSkipTLSVerify: "true"
        s3ForcePathStyle: "true"
      provider: aws
      default: true
      credential:
        key: cloud
        name: <custom_secret> 1
      objectStorage:
        bucket: <bucket_name>
        prefix: <prefix>

```

- 1 Backup location **Secret** with custom name.

4.4.6.3. Configuring the Data Protection Application

You can configure the Data Protection Application by setting Velero resource allocations or enabling self-signed CA certificates.

4.4.6.3.1. Setting Velero CPU and memory resource allocations

You set the CPU and memory resource allocations for the **Velero** pod by editing the **DataProtectionApplication** custom resource (CR) manifest.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the values in the **spec.configuration.velero.podConfig.ResourceAllocations** block of the **DataProtectionApplication** CR manifest, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...

```



```

configuration:
  velero:
    podConfig:
      nodeSelector: <node selector> 1
      resourceAllocations: 2
      limits:
        cpu: "1"
        memory: 1024Mi
      requests:
        cpu: 200m
        memory: 256Mi

```

1 Specify the node selector to be supplied to Velero podSpec.

2 The **resourceAllocations** listed are for average usage.

4.4.6.3.2. Enabling self-signed CA certificates

You must enable a self-signed CA certificate for object storage by editing the **DataProtectionApplication** custom resource (CR) manifest to prevent a **certificate signed by unknown authority** error.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the **spec.backupLocations.velero.objectStorage.caCert** parameter and **spec.backupLocations.velero.config** parameters of the **DataProtectionApplication** CR manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  backupLocations:
    - name: default
      velero:
        provider: aws
        default: true
        objectStorage:
          bucket: <bucket>
          prefix: <prefix>
          caCert: <base64_encoded_cert_string> 1
        config:
          insecureSkipTLSVerify: "false" 2
  ...

```

1 Specify the Base64-encoded CA certificate string.

2

The **insecureSkipTLSVerify** configuration can be set to either **"true"** or **"false"**. If set to **"true"**, SSL/TLS security is disabled. If set to **"false"**, SSL/TLS security is enabled.

4.4.6.4. Installing the Data Protection Application

You install the Data Protection Application (DPA) by creating an instance of the **DataProtectionApplication** API.

Prerequisites

- You must install the OADP Operator.
- You must configure object storage as a backup location.
- If you use snapshots to back up PVs, your cloud provider must support either a native snapshot API or Container Storage Interface (CSI) snapshots.
- If the backup and snapshot locations use the same credentials, you must create a **Secret** with the default name, **cloud-credentials**.
- If the backup and snapshot locations use different credentials, you must create two **Secrets**:
 - **Secret** with a custom name for the backup location. You add this **Secret** to the **DataProtectionApplication** CR.
 - **Secret** with the default name, **cloud-credentials**, for the snapshot location. This **Secret** is not referenced in the **DataProtectionApplication** CR.



NOTE

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file. If there is no default **Secret**, the installation will fail.



NOTE

Velero creates a secret named **velero-repo-credentials** in the OADP namespace, which contains a default backup repository password. You can update the secret with your own password encoded as base64 **before** you run your first backup targeted to the backup repository. The value of the key to update is **Data[repository-password]**.

After you create your DPA, the first time that you run a backup targeted to the backup repository, Velero creates a backup repository whose secret is **velero-repo-credentials**, which contains either the default password or the one you replaced it with. If you update the secret password **after** the first backup, the new password will not match the password in **velero-repo-credentials**, and therefore, Velero will not be able to connect with the older backups.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. Under **Provided APIs**, click **Create instance** in the **DataProtectionApplication** box.

3. Click **YAML View** and update the parameters of the **DataProtectionApplication** manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  configuration:
    velero:
      defaultPlugins:
        - aws
        - openshift 1
      resourceTimeout: 10m 2
    restic:
      enable: true 3
      podConfig:
        nodeSelector: <node_selector> 4
  backupLocations:
    - velero:
      config:
        profile: "default"
        region: minio
        s3Url: <url> 5
        insecureSkipTLSVerify: "true"
        s3ForcePathStyle: "true"
      provider: aws
      default: true
      credential:
        key: cloud
        name: cloud-credentials 6
      objectStorage:
        bucket: <bucket_name> 7
        prefix: <prefix> 8

```

- 1 The **openshift** plugin is mandatory.
- 2 Specify how many minutes to wait for several Velero resources before timeout occurs, such as Velero CRD availability, volumeSnapshot deletion, and backup repository availability. The default is 10m.
- 3 Set this value to **false** if you want to disable the Restic installation. Restic deploys a daemon set, which means that Restic pods run on each working node. In OADP version 1.2 and later, you can configure Restic for backups by adding **spec.defaultVolumesToFsBackup: true** to the **Backup** CR. In OADP version 1.1, add **spec.defaultVolumesToRestic: true** to the **Backup** CR.
- 4 Specify on which nodes Restic is available. By default, Restic runs on all nodes.
- 5 Specify the URL of the S3 endpoint.
- 6 If you do not specify this value, the default name, **cloud-credentials**, is used. If you specify a custom name, the custom name is used for the backup location.
- 7

Specify a bucket as the backup storage location. If the bucket is not a dedicated bucket for Velero backups, you must specify a prefix.

- 8 Specify a prefix for Velero backups, for example, **velero**, if the bucket is used for multiple purposes.

4. Click **Create**.
5. Verify the installation by viewing the OADP resources:

```
$ oc get all -n openshift-adp
```

Example output

```
NAME                                READY STATUS  RESTARTS  AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2   Running  0         2m8s
pod/restic-9cq4q                                1/1   Running  0         94s
pod/restic-m4lts                                1/1   Running  0         94s
pod/restic-pv4kr                                1/1   Running  0         95s
pod/velero-588db7f655-n842v                    1/1   Running  0         95s
```

```
NAME                                TYPE      CLUSTER-IP    EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP  172.30.70.140
<none>      8443/TCP  2m8s
```

```
NAME          DESIRED  CURRENT  READY  UP-TO-DATE  AVAILABLE  NODE
SELECTOR  AGE
daemonset.apps/restic  3        3        3      3           3          <none>      96s
```

```
NAME                                READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/oadp-operator-controller-manager  1/1    1           1          2m9s
deployment.apps/velero                          1/1    1           1          96s
```

```
NAME                                DESIRED  CURRENT  READY  AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47  1        1        1      2m9s
replicaset.apps/velero-588db7f655                        1        1        1      96s
```

4.4.6.4.1. Enabling CSI in the DataProtectionApplication CR

You enable the Container Storage Interface (CSI) in the **DataProtectionApplication** custom resource (CR) in order to back up persistent volumes with CSI snapshots.

Prerequisites

- The cloud provider must support CSI snapshots.

Procedure

- Edit the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
```

```
spec:
  configuration:
    velero:
      defaultPlugins:
        - openshift
        - csi 1
```

- 1** Add the **csi** default plugin.

Additional resources

- [Performance tuning guide for Multicloud Object Gateway](#).

4.4.7. Configuring the OpenShift API for Data Protection with OpenShift Data Foundation

You install the OpenShift API for Data Protection (OADP) with OpenShift Data Foundation by installing the OADP Operator and configuring a backup location and a snapshot location. Then, you install the Data Protection Application.



NOTE

Starting from OADP 1.0.4, all OADP 1.0.z versions can only be used as a dependency of the MTC Operator and are not available as a standalone Operator.

You can configure [Multicloud Object Gateway](#) or any S3-compatible object storage as a backup location.



IMPORTANT

The **CloudStorage** API, which automates the creation of a bucket for object storage, is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You create a **Secret** for the backup location and then you install the Data Protection Application. For more details, see [Installing the OADP Operator](#).

To install the OADP Operator in a restricted network environment, you must first disable the default OperatorHub sources and mirror the Operator catalog. For details, see [Using Operator Lifecycle Manager on restricted networks](#).

4.4.7.1. About backup and snapshot locations and their secrets

You specify backup and snapshot locations and their secrets in the **DataProtectionApplication** custom resource (CR).

Backup locations

You specify S3-compatible object storage, such as Multicloud Object Gateway or MinIO, as a backup location.

Velero backs up OpenShift Container Platform resources, Kubernetes objects, and internal images as an archive file on object storage.

Snapshot locations

If you use your cloud provider's native snapshot API to back up persistent volumes, you must specify the cloud provider as the snapshot location.

If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a **VolumeSnapshotClass** CR to register the CSI driver.

If you use Restic, you do not need to specify a snapshot location because Restic backs up the file system on object storage.

Secrets

If the backup and snapshot locations use the same credentials or if you do not require a snapshot location, you create a default **Secret**.

If the backup and snapshot locations use different credentials, you create two secret objects:

- Custom **Secret** for the backup location, which you specify in the **DataProtectionApplication** CR.
- Default **Secret** for the snapshot location, which is not referenced in the **DataProtectionApplication** CR.



IMPORTANT

The Data Protection Application requires a default **Secret**. Otherwise, the installation will fail.

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file.

Additional resources

- [Creating an Object Bucket Claim using the OpenShift Web Console](#) .

4.4.7.1.1. Creating a default Secret

You create a default **Secret** if your backup and snapshot locations use the same credentials or if you do not require a snapshot location.

The default name of the **Secret** is **cloud-credentials**.



NOTE

The **DataProtectionApplication** custom resource (CR) requires a default **Secret**. Otherwise, the installation will fail. If the name of the backup location **Secret** is not specified, the default name is used.

If you do not want to use the backup location credentials during the installation, you can create a **Secret** with the default name by using an empty **credentials-velero** file.

Prerequisites

- Your object storage and cloud storage, if any, must use the same credentials.
- You must configure object storage for Velero.
- You must create a **credentials-velero** file for the object storage in the appropriate format.

Procedure

- Create a **Secret** with the default name:

```
$ oc create secret generic cloud-credentials -n openshift-adp --from-file cloud=credentials-velero
```

The **Secret** is referenced in the **spec.backupLocations.credential** block of the **DataProtectionApplication** CR when you install the Data Protection Application.

4.4.7.1.2. Creating secrets for different credentials

If your backup and snapshot locations use different credentials, you must create two **Secret** objects:

- Backup location **Secret** with a custom name. The custom name is specified in the **spec.backupLocations** block of the **DataProtectionApplication** custom resource (CR).
- Snapshot location **Secret** with the default name, **cloud-credentials**. This **Secret** is not specified in the **DataProtectionApplication** CR.

Procedure

1. Create a **credentials-velero** file for the snapshot location in the appropriate format for your cloud provider.
2. Create a **Secret** for the snapshot location with the default name:

```
$ oc create secret generic cloud-credentials -n openshift-adp --from-file cloud=credentials-velero
```

3. Create a **credentials-velero** file for the backup location in the appropriate format for your object storage.
4. Create a **Secret** for the backup location with a custom name:

```
$ oc create secret generic <custom_secret> -n openshift-adp --from-file cloud=credentials-velero
```

5. Add the **Secret** with the custom name to the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
```

```

...
backupLocations:
- velero:
  config:
    profile: "default"
    region: minio
    s3Url: <url>
    insecureSkipTLSVerify: "true"
    s3ForcePathStyle: "true"
  provider: gcp
  default: true
  credential:
    key: cloud
    name: <custom_secret> ❶
  objectStorage:
    bucket: <bucket_name>
    prefix: <prefix>

```

- ❶ Backup location **Secret** with custom name.

4.4.7.2. Configuring the Data Protection Application

You can configure the Data Protection Application by setting Velero resource allocations or enabling self-signed CA certificates.

4.4.7.2.1. Setting Velero CPU and memory resource allocations

You set the CPU and memory resource allocations for the **Velero** pod by editing the **DataProtectionApplication** custom resource (CR) manifest.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the values in the **spec.configuration.velero.podConfig.ResourceAllocations** block of the **DataProtectionApplication** CR manifest, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
...
configuration:
  velero:
    podConfig:
      nodeSelector: <node selector> ❶
      resourceAllocations: ❷
      limits:
        cpu: "1"
        memory: 1024Mi

```



```
requests:
  cpu: 200m
  memory: 256Mi
```

- 1 Specify the node selector to be supplied to Velero podSpec.
- 2 The **resourceAllocations** listed are for average usage.

4.4.7.2.1.1. Adjusting Ceph CPU and memory requirements based on collected data

The following recommendations are based on observations of performance made in the scale and performance lab. The changes are specifically related to {odf-first}. If working with {odf-short}, consult the appropriate tuning guides for official recommendations.

4.4.7.2.1.1.1. CPU and memory requirement for configurations

Backup and restore operations require large amounts of CephFS **PersistentVolumes** (PVs). To avoid Ceph MDS pods restarting with an **out-of-memory** (OOM) error, the following configuration is suggested:

Configuration types	Request	Max limit
CPU	Request changed to 3	Max limit to 3
Memory	Request changed to 8 Gi	Max limit to 128 Gi

4.4.7.2.2. Enabling self-signed CA certificates

You must enable a self-signed CA certificate for object storage by editing the **DataProtectionApplication** custom resource (CR) manifest to prevent a **certificate signed by unknown authority** error.

Prerequisites

- You must have the OpenShift API for Data Protection (OADP) Operator installed.

Procedure

- Edit the **spec.backupLocations.velero.objectStorage.caCert** parameter and **spec.backupLocations.velero.config** parameters of the **DataProtectionApplication** CR manifest:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
spec:
  ...
  backupLocations:
    - name: default
      velero:
        provider: aws
```

```

default: true
objectStorage:
  bucket: <bucket>
  prefix: <prefix>
  caCert: <base64_encoded_cert_string> 1
config:
  insecureSkipTLSVerify: "false" 2
...

```

- 1 Specify the Base64-encoded CA certificate string.
- 2 The **`insecureSkipTLSVerify`** configuration can be set to either **"true"** or **"false"**. If set to **"true"**, SSL/TLS security is disabled. If set to **"false"**, SSL/TLS security is enabled.

4.4.7.3. Installing the Data Protection Application

You install the Data Protection Application (DPA) by creating an instance of the **DataProtectionApplication** API.

Prerequisites

- You must install the OADP Operator.
- You must configure object storage as a backup location.
- If you use snapshots to back up PVs, your cloud provider must support either a native snapshot API or Container Storage Interface (CSI) snapshots.
- If the backup and snapshot locations use the same credentials, you must create a **Secret** with the default name, **cloud-credentials**.
- If the backup and snapshot locations use different credentials, you must create two **Secrets**:
 - **Secret** with a custom name for the backup location. You add this **Secret** to the **DataProtectionApplication** CR.
 - **Secret** with the default name, **cloud-credentials**, for the snapshot location. This **Secret** is not referenced in the **DataProtectionApplication** CR.



NOTE

If you do not want to specify backup or snapshot locations during the installation, you can create a default **Secret** with an empty **credentials-velero** file. If there is no default **Secret**, the installation will fail.



NOTE

Velero creates a secret named **velero-repo-credentials** in the OADP namespace, which contains a default backup repository password. You can update the secret with your own password encoded as base64 **before** you run your first backup targeted to the backup repository. The value of the key to update is **Data[repository-password]**.

After you create your DPA, the first time that you run a backup targeted to the backup repository, Velero creates a backup repository whose secret is **velero-repo-credentials**, which contains either the default password or the one you replaced it with. If you update the secret password **after** the first backup, the new password will not match the password in **velero-repo-credentials**, and therefore, Velero will not be able to connect with the older backups.

Procedure

1. Click **Operators** → **Installed Operators** and select the OADP Operator.
2. Under **Provided APIs**, click **Create instance** in the **DataProtectionApplication** box.
3. Click **YAML View** and update the parameters of the **DataProtectionApplication** manifest:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_sample>
  namespace: openshift-adp
spec:
  configuration:
    velero:
      defaultPlugins:
        - aws
        - openshift ❶
      resourceTimeout: 10m ❷
    restic:
      enable: true ❸
      podConfig:
        nodeSelector: <node_selector> ❹
  backupLocations:
    - velero:
      config:
        profile: "default"
        region: minio
        s3Url: <url> ❺
        insecureSkipTLSVerify: "true"
        s3ForcePathStyle: "true"
      provider: gcp
      default: true
      credential:
        key: cloud
        name: cloud-credentials ❻

```

```
objectStorage:
  bucket: <bucket_name> 7
  prefix: <prefix> 8
```

- 1 The **openshift** plugin is mandatory.
- 2 Specify how many minutes to wait for several Velero resources before timeout occurs, such as Velero CRD availability, volumeSnapshot deletion, and backup repository availability. The default is 10m.
- 3 Set this value to **false** if you want to disable the Restic installation. Restic deploys a daemon set, which means that Restic pods run on each working node. In OADP version 1.2 and later, you can configure Restic for backups by adding **spec.defaultVolumesToFsBackup: true** to the **Backup** CR. In OADP version 1.1, add **spec.defaultVolumesToRestic: true** to the **Backup** CR.
- 4 Specify on which nodes Restic is available. By default, Restic runs on all nodes.
- 5 Specify the URL of the S3 endpoint.
- 6 If you do not specify this value, the default name, **cloud-credentials**, is used. If you specify a custom name, the custom name is used for the backup location.
- 7 Specify a bucket as the backup storage location. If the bucket is not a dedicated bucket for Velero backups, you must specify a prefix.
- 8 Specify a prefix for Velero backups, for example, **velero**, if the bucket is used for multiple purposes.

4. Click **Create**.

5. Verify the installation by viewing the OADP resources:

```
$ oc get all -n openshift-adp
```

Example output

```
NAME                                READY STATUS  RESTARTS  AGE
pod/oadp-operator-controller-manager-67d9494d47-6l8z8  2/2   Running  0         2m8s
pod/restic-9cq4q                                1/1   Running  0         94s
pod/restic-m4lts                                1/1   Running  0         94s
pod/restic-pv4kr                                1/1   Running  0         95s
pod/velero-588db7f655-n842v                    1/1   Running  0         95s
```

```
NAME                                TYPE      CLUSTER-IP      EXTERNAL-IP
PORT(S)  AGE
service/oadp-operator-controller-manager-metrics-service  ClusterIP  172.30.70.140
<none>    8443/TCP  2m8s
```

```
NAME          DESIRED  CURRENT  READY  UP-TO-DATE  AVAILABLE  NODE
SELECTOR  AGE
daemonset.apps/restic  3        3        3      3           3          <none>    96s
```

```
NAME                                READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/oadp-operator-controller-manager  1/1    1           1          2m9s
```

deployment.apps/velero	1/1	1	1	96s
NAME	DESIRED	CURRENT	READY	AGE
replicaset.apps/oadp-operator-controller-manager-67d9494d47	1	1	1	2m9s
replicaset.apps/velero-588db7f655	1	1	1	96s

4.4.7.3.1. Creating an Object Bucket Claim for disaster recovery on OpenShift Data Foundation

If you use cluster storage for your Multicloud Object Gateway (MCG) bucket **backupStorageLocation** on OpenShift Data Foundation, create an Object Bucket Claim (OBC) using the OpenShift Web Console.



WARNING

Failure to configure an Object Bucket Claim (OBC) might lead to backups not being available.



NOTE

Unless specified otherwise, "NooBaa" refers to the open source project that provides lightweight object storage, while "Multicloud Object Gateway (MCG)" refers to the Red Hat distribution of NooBaa.

For more information on the MCG, see [Accessing the Multicloud Object Gateway with your applications](#).

Procedure

- Create an Object Bucket Claim (OBC) using the OpenShift web console as described in [Creating an Object Bucket Claim using the OpenShift Web Console](#).

4.4.7.3.2. Enabling CSI in the DataProtectionApplication CR

You enable the Container Storage Interface (CSI) in the **DataProtectionApplication** custom resource (CR) in order to back up persistent volumes with CSI snapshots.

Prerequisites

- The cloud provider must support CSI snapshots.

Procedure

- Edit the **DataProtectionApplication** CR, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
spec:
  configuration:
```

```

velero:
  defaultPlugins:
    - openshift
    - csi 1

```

- 1** Add the **csi** default plugin.

4.5. UNINSTALLING OADP

4.5.1. Uninstalling the OpenShift API for Data Protection

You uninstall the OpenShift API for Data Protection (OADP) by deleting the OADP Operator. See [Deleting Operators from a cluster](#) for details.

4.6. OADP BACKING UP

4.6.1. Backing up applications

You back up applications by creating a **Backup** custom resource (CR). See [Creating a Backup CR](#).

- The **Backup** CR creates backup files for Kubernetes resources and internal images on S3 object storage.
- If your cloud provider has a native snapshot API or supports CSI snapshots, the **Backup** CR backs up persistent volumes (PVs) by creating snapshots. For more information about working with CSI snapshots, see [Backing up persistent volumes with CSI snapshots](#).

For more information about CSI volume snapshots, see [CSI volume snapshots](#).



IMPORTANT

The **CloudStorage** API for S3 storage is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

- If your cloud provider does not support snapshots or if your applications are on NFS data volumes, you can create backups by using Kopia or Restic. See [Backing up applications with File System Backup: Kopia or Restic](#).



IMPORTANT

The OpenShift API for Data Protection (OADP) does not support backing up volume snapshots that were created by other software.

You can create backup hooks to run commands before or after the backup operation. See [Creating backup hooks](#).

You can schedule backups by creating a **Schedule** CR instead of a **Backup** CR. See [Scheduling backups using Schedule CR](#)].

4.6.1.1. Known issues

OpenShift Container Platform 4.14 enforces a pod security admission (PSA) policy that can hinder the readiness of pods during a Restic restore process.

This issue has been resolved in the OADP 1.1.6 and OADP 1.2.2 releases, therefore it is recommended that users upgrade to these releases.

For more information, see [Restic restore partially failing on OCP 4.14 due to changed PSA policy](#) .

Additional resources

- [Installing Operators on clusters for administrators](#)
- [Installing Operators in namespaces for non-administrators](#)

4.6.2. Creating a Backup CR

You back up Kubernetes images, internal images, and persistent volumes (PVs) by creating a **Backup** custom resource (CR).

Prerequisites

- You must install the OpenShift API for Data Protection (OADP) Operator.
- The **DataProtectionApplication** CR must be in a **Ready** state.
- Backup location prerequisites:
 - You must have S3 object storage configured for Velero.
 - You must have a backup location configured in the **DataProtectionApplication** CR.
- Snapshot location prerequisites:
 - Your cloud provider must have a native snapshot API or support Container Storage Interface (CSI) snapshots.
 - For CSI snapshots, you must create a **VolumeSnapshotClass** CR to register the CSI driver.
 - You must have a volume location configured in the **DataProtectionApplication** CR.

Procedure

1. Retrieve the **backupStorageLocations** CRs by entering the following command:

```
$ oc get backupStorageLocations -n openshift-adp
```

Example output

-

NAMESPACE	NAME	PHASE	LAST VALIDATED	AGE	DEFAULT
openshift-adp	velero-sample-1	Available	11s	31m	

2. Create a **Backup** CR, as in the following example:

```

apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup>
  labels:
    velero.io/storage-location: default
  namespace: openshift-adp
spec:
  hooks: {}
  includedNamespaces:
    - <namespace> 1
  includedResources: [] 2
  excludedResources: [] 3
  storageLocation: <velero-sample-1> 4
  ttl: 720h0m0s
  labelSelector: 5
  matchLabels:
    app: <label_1>
    app: <label_2>
    app: <label_3>
  orLabelSelectors: 6
  - matchLabels:
    app: <label_1>
    app: <label_2>
    app: <label_3>

```

- 1** Specify an array of namespaces to back up.
- 2** Optional: Specify an array of resources to include in the backup. Resources might be shortcuts (for example, 'po' for 'pods') or fully-qualified. If unspecified, all resources are included.
- 3** Optional: Specify an array of resources to exclude from the backup. Resources might be shortcuts (for example, 'po' for 'pods') or fully-qualified.
- 4** Specify the name of the **backupStorageLocations** CR.
- 5** Map of {key,value} pairs of backup resources that have **all** the specified labels.
- 6** Map of {key,value} pairs of backup resources that have **one or more** of the specified labels.

3. Verify that the status of the **Backup** CR is **Completed**:

```
$ oc get backup -n openshift-adp <backup> -o jsonpath='{.status.phase}'
```

4.6.3. Backing up persistent volumes with CSI snapshots

You back up persistent volumes with Container Storage Interface (CSI) snapshots by editing the **VolumeSnapshotClass** custom resource (CR) of the cloud storage before you create the **Backup** CR, see [CSI volume snapshots](#).

For more information, see [Creating a Backup CR](#).

Prerequisites

- The cloud provider must support CSI snapshots.
- You must enable CSI in the **DataProtectionApplication** CR.

Procedure

- Add the **metadata.labels.velero.io/csi-volumesnapshot-class: "true"** key-value pair to the **VolumeSnapshotClass** CR:

```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: <volume_snapshot_class_name>
  labels:
    velero.io/csi-volumesnapshot-class: "true"
driver: <csi_driver>
deletionPolicy: Retain
```

You can now create a **Backup** CR.

4.6.4. Backing up applications with File System Backup: Kopia or Restic

You can use OADP to back up and restore Kubernetes volumes attached to pods from the file system of the volumes. This process is called File System Backup (FSB) or Pod Volume Backup (PVB). It is accomplished by using modules from the open source backup tools Restic or Kopia.

If your cloud provider does not support snapshots or if your applications are on NFS data volumes, you can create backups by using FSB.



NOTE

[Restic](#) is installed by the OADP Operator by default. If you prefer, you can install [Kopia](#) instead.

FSB integration with OADP provides a solution for backing up and restoring almost any type of Kubernetes volumes. This integration is an additional capability of OADP and is not a replacement for existing functionality.

You back up Kubernetes resources, internal images, and persistent volumes with Kopia or Restic by editing the **Backup** custom resource (CR).


You do not need to specify a snapshot location in the **DataProtectionApplication** CR.

**NOTE**

In OADP version 1.3 and later, you can use either Kopia or Restic for backing up applications.

For the Built-in DataMover, you must use Kopia.

In OADP version 1.2 and earlier, you can only use Restic for backing up applications.

**IMPORTANT**

FSB does not support backing up **hostPath** volumes. For more information, see [FSB limitations](#).

Prerequisites

- You must install the OpenShift API for Data Protection (OADP) Operator.
- You must not disable the default **nodeAgent** installation by setting **spec.configuration.nodeAgent.enable** to **false** in the **DataProtectionApplication** CR.
- You must select Kopia or Restic as the uploader by setting **spec.configuration.nodeAgent.uploaderType** to **kopia** or **restic** in the **DataProtectionApplication** CR.
- The **DataProtectionApplication** CR must be in a **Ready** state.

Procedure

- Create the **Backup** CR, as in the following example:

```
apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup>
  labels:
    velero.io/storage-location: default
  namespace: openshift-adp
spec:
  defaultVolumesToFsBackup: true 1
  ...
```

- 1** In OADP version 1.2 and later, add the **defaultVolumesToFsBackup: true** setting within the **spec** block. In OADP version 1.1, add **defaultVolumesToRestic: true**.

4.6.5. Creating backup hooks

When performing a backup, it is possible to specify one or more commands to execute in a container within a pod, based on the pod being backed up.

The commands can be configured to performed before any custom action processing (*Pre* hooks), or after all custom actions have been completed and any additional items specified by the custom action have been backed up (*Post* hooks).

You create backup hooks to run commands in a container in a pod by editing the **Backup** custom resource (CR).

Procedure

- Add a hook to the **spec.hooks** block of the **Backup** CR, as in the following example:

```

apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup>
  namespace: openshift-adp
spec:
  hooks:
    resources:
      - name: <hook_name>
        includedNamespaces:
          - <namespace> 1
        excludedNamespaces: 2
          - <namespace>
        includedResources: []
        - pods 3
        excludedResources: [] 4
        labelSelector: 5
          matchLabels:
            app: velero
            component: server
        pre: 6
          - exec:
              container: <container> 7
              command:
                - /bin/uname 8
                - -a
              onError: Fail 9
              timeout: 30s 10
        post: 11
  ...

```

- 1 Optional: You can specify namespaces to which the hook applies. If this value is not specified, the hook applies to all namespaces.
- 2 Optional: You can specify namespaces to which the hook does not apply.
- 3 Currently, pods are the only supported resource that hooks can apply to.
- 4 Optional: You can specify resources to which the hook does not apply.
- 5 Optional: This hook only applies to objects matching the label. If this value is not specified, the hook applies to all objects.
- 6 Array of hooks to run before the backup.
- 7 Optional: If the container is not specified, the command runs in the first container in the pod.

- 8 This is the entry point for the **init** container being added.
- 9 Allowed values for error handling are **Fail** and **Continue**. The default is **Fail**.
- 10 Optional: How long to wait for the commands to run. The default is **30s**.
- 11 This block defines an array of hooks to run after the backup, with the same parameters as the pre-backup hooks.

4.6.6. Scheduling backups using Schedule CR

The schedule operation allows you to create a backup of your data at a particular time, specified by a Cron expression.

You schedule backups by creating a **Schedule** custom resource (CR) instead of a **Backup** CR.



WARNING

Leave enough time in your backup schedule for a backup to finish before another backup is created.

For example, if a backup of a namespace typically takes 10 minutes, do not schedule backups more frequently than every 15 minutes.

Prerequisites

- You must install the OpenShift API for Data Protection (OADP) Operator.
- The **DataProtectionApplication** CR must be in a **Ready** state.

Procedure

1. Retrieve the **backupStorageLocations** CRs:

```
$ oc get backupStorageLocations -n openshift-adp
```

Example output

```

NAMESPACE   NAME           PHASE   LAST VALIDATED  AGE   DEFAULT
openshift-adp  velero-sample-1  Available  11s             31m

```

2. Create a **Schedule** CR, as in the following example:

```
$ cat << EOF | oc apply -f -
apiVersion: velero.io/v1
kind: Schedule
metadata:
  name: <schedule>
  namespace: openshift-adp
```

```
spec:
  schedule: 0 7 * * * 1
  template:
    hooks: {}
    includedNamespaces:
      - <namespace> 2
    storageLocation: <velero-sample-1> 3
    defaultVolumesToFsBackup: true 4
    ttl: 720h0m0s
EOF
```

- 1 **cron** expression to schedule the backup, for example, `0 7 * * *` to perform a backup every day at 7:00.



NOTE

To schedule a backup at specific intervals, enter the **<duration_in_minutes>** in the following format:

```
schedule: "*/10 * * * *"
```

Enter the minutes value between quotation marks (" ").

- 2 Array of namespaces to back up.
- 3 Name of the **backupStorageLocations** CR.
- 4 Optional: In OADP version 1.2 and later, add the **defaultVolumesToFsBackup: true** key-value pair to your configuration when performing backups of volumes with Restic. In OADP version 1.1, add the **defaultVolumesToRestic: true** key-value pair when you back up volumes with Restic.

1. Verify that the status of the **Schedule** CR is **Completed** after the scheduled backup runs:

```
$ oc get schedule -n openshift-adp <schedule> -o jsonpath='{.status.phase}'
```

4.6.7. Deleting backups

You can remove backup files by deleting the **Backup** custom resource (CR).



WARNING

After you delete the **Backup** CR and the associated object storage data, you cannot recover the deleted data.

Prerequisites

- You created a **Backup** CR.

- You know the name of the **Backup** CR and the namespace that contains it.
- You downloaded the Velero CLI tool.
- You can access the Velero binary in your cluster.

Procedure

- Choose one of the following actions to delete the **Backup** CR:
 - To delete the **Backup** CR and keep the associated object storage data, run the following command:

```
$ oc delete backup <backup_CR_name> -n <velero_namespace>
```

- To delete the **Backup** CR and delete the associated object storage data, run the following command:

```
$ velero backup delete <backup_CR_name> -n <velero_namespace>
```

Where:

<backup_CR_name>

The name of the **Backup** custom resource.

<velero_namespace>

The namespace that contains the **Backup** custom resource.

4.6.8. About Kopia

Kopia is a fast and secure open-source backup and restore tool that allows you to create encrypted snapshots of your data and save the snapshots to remote or cloud storage of your choice.

Kopia supports network and local storage locations, and many cloud or remote storage locations, including:

- Amazon S3 and any cloud storage that is compatible with S3
- Azure Blob Storage
- Google Cloud Storage platform

Kopia uses content-addressable storage for snapshots:

- Snapshots are always incremental; data that is already included in previous snapshots is not re-uploaded to the repository. A file is only uploaded to the repository again if it is modified.
- Stored data is deduplicated; if multiple copies of the same file exist, only one of them is stored.
- If files are moved or renamed, Kopia can recognize that they have the same content and does not upload them again.

4.6.8.1. OADP integration with Kopia

OADP 1.3 supports Kopia as the backup mechanism for pod volume backup in addition to Restic. You

must choose one or the other at installation by setting the **uploaderType** field in the **DataProtectionApplication** custom resource (CR). The possible values are **restic** or **kopia**. If you do not specify an **uploaderType**, OADP 1.3 defaults to using Kopia as the backup mechanism. The data is written to and read from a unified repository.

The following example shows a **DataProtectionApplication** CR configured for using Kopia:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: dpa-sample
spec:
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
# ...
```

4.7. OADP RESTORING

4.7.1. Restoring applications

You restore application backups by creating a **Restore** custom resource (CR). See [Creating a Restore CR](#).

You can create restore hooks to run commands in a container in a pod by editing the **Restore** CR. See [Creating restore hooks](#).

4.7.1.1. Creating a Restore CR

You restore a **Backup** custom resource (CR) by creating a **Restore** CR.

Prerequisites

- You must install the OpenShift API for Data Protection (OADP) Operator.
- The **DataProtectionApplication** CR must be in a **Ready** state.
- You must have a Velero **Backup** CR.
- The persistent volume (PV) capacity must match the requested size at backup time. Adjust the requested size if needed.

Procedure

1. Create a **Restore** CR, as in the following example:

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: <restore>
  namespace: openshift-adp
spec:
  backupName: <backup> 1
```

```
includedResources: [] 2
excludedResources:
- nodes
- events
- events.events.k8s.io
- backups.velero.io
- restores.velero.io
- resticrepositories.velero.io
restorePVs: true 3
```

- 1** Name of the **Backup** CR.
- 2** Optional: Specify an array of resources to include in the restore process. Resources might be shortcuts (for example, **po** for **pods**) or fully-qualified. If unspecified, all resources are included.
- 3** Optional: The **restorePVs** parameter can be set to **false** to turn off restore of **PersistentVolumes** from **VolumeSnapshot** of Container Storage Interface (CSI) snapshots or from native snapshots when **VolumeSnapshotLocation** is configured.

2. Verify that the status of the **Restore** CR is **Completed** by entering the following command:

```
$ oc get restore -n openshift-adp <restore> -o jsonpath='{.status.phase}'
```

3. Verify that the backup resources have been restored by entering the following command:

```
$ oc get all -n <namespace> 1
```

- 1** Namespace that you backed up.

4. If you use Restic to restore **DeploymentConfig** objects or if you use post-restore hooks, run the **dc-restic-post-restore.sh** cleanup script by entering the following command:

```
$ bash dc-restic-post-restore.sh <restore-name>
```



NOTE

During the restore process, the OADP Velero plug-ins scale down the **DeploymentConfig** objects and restore the pods as standalone pods. This is done to prevent the cluster from deleting the restored **DeploymentConfig** pods immediately on restore and to allow Restic and post-restore hooks to complete their actions on the restored pods. The cleanup script shown below removes these disconnected pods and scales any **DeploymentConfig** objects back up to the appropriate number of replicas.

Example 4.1. dc-restic-post-restore.sh cleanup script

```
#!/bin/bash
set -e

# if sha256sum exists, use it to check the integrity of the file
if command -v sha256sum >/dev/null 2>&1; then
```



```

CHECKSUM_CMD="sha256sum"
else
CHECKSUM_CMD="shasum -a 256"
fi

label_name () {
  if [ "${#1}" -le "63" ]; then
echo $1
return
  fi
  sha=$(echo -n $1|$CHECKSUM_CMD)
  echo "${1:0:57}${sha:0:6}"
}

OADP_NAMESPACE=${OADP_NAMESPACE:=openshift-adp}

if [[ $# -ne 1 ]]; then
  echo "usage: ${BASH_SOURCE} restore-name"
  exit 1
fi

echo using OADP Namespace $OADP_NAMESPACE
echo restore: $1

label=$(label_name $1)
echo label: $label

echo Deleting disconnected restore pods
oc delete pods -l oadp.openshift.io/disconnected-from-dc=$label

for dc in $(oc get dc --all-namespaces -l oadp.openshift.io/replicas-modified=$label -o
jsonpath='{range .items[*]}.{.metadata.namespace},"",{.metadata.name},"'
{.metadata.annotations.oadp\.openshift\.io/original-replicas},"'
{.metadata.annotations.oadp\.openshift\.io/original-paused}{"\n"}')
do
  IFS=' ' read -ra dc_arr <<< "$dc"
  if [ ${#dc_arr[0]} -gt 0 ]; then
echo Found deployment ${dc_arr[0]}/${dc_arr[1]}, setting replicas: ${dc_arr[2]}, paused:
${dc_arr[3]}
cat <<EOF | oc patch dc -n ${dc_arr[0]} ${dc_arr[1]} --patch-file /dev/stdin
spec:
  replicas: ${dc_arr[2]}
  paused: ${dc_arr[3]}
EOF
  fi
done

```

4.7.1.2. Creating restore hooks

You create restore hooks to run commands in a container in a pod by editing the **Restore** custom resource (CR).

You can create two types of restore hooks:

- An **init** hook adds an init container to a pod to perform setup tasks before the application container starts.
If you restore a Restic backup, the **restic-wait** init container is added before the restore hook init container.
- An **exec** hook runs commands or scripts in a container of a restored pod.

Procedure

- Add a hook to the **spec.hooks** block of the **Restore** CR, as in the following example:

```

apiVersion: velero.io/v1
kind: Restore
metadata:
  name: <restore>
  namespace: openshift-adp
spec:
  hooks:
    resources:
      - name: <hook_name>
        includedNamespaces:
          - <namespace> 1
        excludedNamespaces:
          - <namespace>
        includedResources:
          - pods 2
        excludedResources: []
        labelSelector: 3
          matchLabels:
            app: velero
            component: server
        postHooks:
          - init:
              initContainers:
                - name: restore-hook-init
                  image: alpine:latest
                  volumeMounts:
                    - mountPath: /restores/pvc1-vm
                      name: pvc1-vm
                  command:
                    - /bin/ash
                    - -c
                  timeout: 4
              - exec:
                  container: <container> 5
                  command:
                    - /bin/bash 6
                    - -c
                    - "psql < /backup/backup.sql"
                  waitTimeout: 5m 7
                  execTimeout: 1m 8
                  onError: Continue 9

```

- 1 Optional: Array of namespaces to which the hook applies. If this value is not specified, the hook applies to all namespaces.
- 2 Currently, pods are the only supported resource that hooks can apply to.
- 3 Optional: This hook only applies to objects matching the label selector.
- 4 Optional: Timeout specifies the maximum length of time Velero waits for **initContainers** to complete.
- 5 Optional: If the container is not specified, the command runs in the first container in the pod.
- 6 This is the entrypoint for the init container being added.
- 7 Optional: How long to wait for a container to become ready. This should be long enough for the container to start and for any preceding hooks in the same container to complete. If not set, the restore process waits indefinitely.
- 8 Optional: How long to wait for the commands to run. The default is **30s**.
- 9 Allowed values for error handling are **Fail** and **Continue**:
 - **Continue**: Only command failures are logged.
 - **Fail**: No more restore hooks run in any container in any pod. The status of the **Restore** CR will be **PartiallyFailed**.

4.8. OADP AND ROSA

4.8.1. Backing up applications on ROSA clusters using OADP

You can use OpenShift API for Data Protection (OADP) with Red Hat OpenShift Service on AWS (ROSA) clusters to back up and restore application data.

ROSA is a fully-managed, turnkey application platform that allows you to deliver value to your customers by building and deploying applications.

ROSA provides seamless integration with a wide range of Amazon Web Services (AWS) compute, database, analytics, machine learning, networking, mobile, and other services to speed up the building and delivery of differentiating experiences to your customers.

You can subscribe to the service directly from your AWS account.

After you create your clusters, you can operate your clusters with the OpenShift Container Platform web console or through [Red Hat OpenShift Cluster Manager](#). You can also use ROSA with OpenShift APIs and command-line interface (CLI) tools.

For additional information about ROSA installation, see [Installing Red Hat OpenShift Service on AWS \(ROSA\) interactive walkthrough](#).

Before installing OpenShift API for Data Protection (OADP), you must set up role and policy credentials for OADP so that it can use the Amazon Web Services API.

This process is performed in the following two stages:

1. Prepare AWS credentials
2. Install the OADP Operator and give it an IAM role

4.8.1.1. Preparing AWS credentials for OADP

An Amazon Web Services account must be prepared and configured to accept an OpenShift API for Data Protection (OADP) installation.

Procedure

1. Create the following environment variables by running the following commands:



IMPORTANT

Change the cluster name to match your ROSA cluster, and ensure you are logged into the cluster as an administrator. Ensure that all fields are outputted correctly before continuing.

```
$ export CLUSTER_NAME=my-cluster 1
export ROSA_CLUSTER_ID=$(rosa describe cluster -c ${CLUSTER_NAME} --output json |
jq -r .id)
export REGION=$(rosa describe cluster -c ${CLUSTER_NAME} --output json | jq -r
.region.id)
export OIDC_ENDPOINT=$(oc get authentication.config.openshift.io cluster -o
jsonpath='{.spec.serviceAccountIssuer}' | sed 's|^https://|')
export AWS_ACCOUNT_ID=$(aws sts get-caller-identity --query Account --output text)
export CLUSTER_VERSION=$(rosa describe cluster -c ${CLUSTER_NAME} -o json | jq -r
.version.raw_id | cut -f -2 -d '.')
export ROLE_NAME="${CLUSTER_NAME}-openshift-oadp-aws-cloud-credentials"
export SCRATCH="/tmp/${CLUSTER_NAME}/oadp"
mkdir -p ${SCRATCH}
echo "Cluster ID: ${ROSA_CLUSTER_ID}, Region: ${REGION}, OIDC Endpoint:
${OIDC_ENDPOINT}, AWS Account ID: ${AWS_ACCOUNT_ID}"
```

- 1 Replace **my-cluster** with your ROSA cluster name.

2. On the AWS account, create an IAM policy to allow access to AWS S3:
 - a. Check to see if the policy exists by running the following command:

```
$ POLICY_ARN=$(aws iam list-policies --query "Policies[?
PolicyName=='RosaOadpVer1'].{ARN:Arn}" --output text) 1
```

- 1 Replace **RosaOadp** with your policy name.

- b. Enter the following command to create the policy JSON file and then create the policy in ROSA:

**NOTE**

If the policy ARN is not found, the command creates the policy. If the policy ARN already exists, the **if** statement intentionally skips the policy creation.

```
$ if [[ -z "${POLICY_ARN}" ]]; then
  cat << EOF > ${SCRATCH}/policy.json 1
  {
  "Version": "2012-10-17",
  "Statement": [
  {
  "Effect": "Allow",
  "Action": [
  "s3:CreateBucket",
  "s3:DeleteBucket",
  "s3:PutBucketTagging",
  "s3:GetBucketTagging",
  "s3:PutEncryptionConfiguration",
  "s3:GetEncryptionConfiguration",
  "s3:PutLifecycleConfiguration",
  "s3:GetLifecycleConfiguration",
  "s3:GetBucketLocation",
  "s3:ListBucket",
  "s3:GetObject",
  "s3:PutObject",
  "s3:DeleteObject",
  "s3:ListBucketMultipartUploads",
  "s3:AbortMultipartUploads",
  "s3:ListMultipartUploadParts",
  "s3:DescribeSnapshots",
  "ec2:DescribeVolumes",
  "ec2:DescribeVolumeAttribute",
  "ec2:DescribeVolumesModifications",
  "ec2:DescribeVolumeStatus",
  "ec2:CreateTags",
  "ec2:CreateVolume",
  "ec2:CreateSnapshot",
  "ec2>DeleteSnapshot"
  ],
  "Resource": "*"
  }
  ]
  }
  EOF
```

```
POLICY_ARN=$(aws iam create-policy --policy-name "RosaOadpVer1" \
--policy-document file:///${SCRATCH}/policy.json --query Policy.Arn \
--tags Key=rosa_openshift_version,Value=${CLUSTER_VERSION}
Key=rosa_role_prefix,Value=ManagedOpenShift
Key=operator_namespace,Value=openshift-oadp Key=operator_name,Value=openshift-
oadp \
--output text)
fi
```

1 **SCRATCH** is a name for a temporary directory created for the environment variables.

- c. View the policy ARN by running the following command:

```
$ echo ${POLICY_ARN}
```

3. Create an IAM role trust policy for the cluster:

- a. Create the trust policy file by running the following command:

```
$ cat <<EOF > ${SCRATCH}/trust-policy.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::${AWS_ACCOUNT_ID}:oidc-provider/${OIDC_ENDPOINT}"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "${OIDC_ENDPOINT}:sub": [
            "system:serviceaccount:openshift-adp:openshift-adp-controller-manager",
            "system:serviceaccount:openshift-adp:velero"
          ]
        }
      }
    }
  ]
}
EOF
```

- b. Create the role by running the following command:

```
$ ROLE_ARN=$(aws iam create-role --role-name \
  "${ROLE_NAME}" \
  --assume-role-policy-document file://${SCRATCH}/trust-policy.json \
  --tags Key=rosa_cluster_id,Value=${ROSA_CLUSTER_ID} \
  Key=rosa_openshift_version,Value=${CLUSTER_VERSION} \
  Key=rosa_role_prefix,Value=ManagedOpenShift \
  Key=operator_namespace,Value=openshift-adp Key=operator_name,Value=openshift-oadp \
  --query Role.Arn --output text)
```

- c. View the role ARN by running the following command:

```
$ echo ${ROLE_ARN}
```

4. Attach the IAM policy to the IAM role by running the following command:

```
$ aws iam attach-role-policy --role-name "${ROLE_NAME}" \
  --policy-arn ${POLICY_ARN}
```

4.8.1.2. Installing the OADP Operator and providing the IAM role

AWS Security Token Service (AWS STS) is a global web service that provides short-term credentials for IAM or federated users. OpenShift Container Platform (ROSA) with STS is the recommended credential

mode for ROSA clusters. This document describes how to install OpenShift API for Data Protection (OADP) on ROSA with AWS STS.



IMPORTANT

Restic and Kopia are not supported in the OADP on ROSA with AWS STS environment. Verify that the Restic and Kopia node agent is disabled. For backing up volumes, OADP on ROSA with AWS STS supports only native snapshots and Container Storage Interface (CSI) snapshots.



IMPORTANT

In an Amazon ROSA cluster that uses STS authentication, restoring backed-up data in a different AWS region is not supported.

The Data Mover feature is not currently supported in ROSA clusters. You can use native AWS S3 tools for moving data.

Prerequisites

- An OpenShift Container Platform ROSA cluster with the required access and tokens. For instructions, see the previous procedure *Preparing AWS credentials for OADP*. If you plan to use two different clusters for backing up and restoring, you must prepare AWS credentials, including **ROLE_ARN**, for each cluster.

Procedure

1. Create an OpenShift Container Platform secret from your AWS token file by entering the following commands:

- a. Create the credentials file:

```
$ cat <<EOF > ${SCRATCH}/credentials
[default]
role_arn = ${ROLE_ARN}
web_identity_token_file = /var/run/secrets/openshift/serviceaccount/token
EOF
```

- b. Create a namespace for OADP:

```
$ oc create namespace openshift-adp
```

- c. Create the OpenShift Container Platform secret:

```
$ oc -n openshift-adp create secret generic cloud-credentials \
--from-file=${SCRATCH}/credentials
```

**NOTE**

In OpenShift Container Platform versions 4.14 and later, the OADP Operator supports a new standardized STS workflow through the Operator Lifecycle Manager (OLM) and Cloud Credentials Operator (CCO). In this workflow, you do not need to create the above secret, you only need to supply the role ARN during the installation of OLM-managed operators using the OpenShift Container Platform web console, for more information see *Installing from OperatorHub using the web console*.

The preceding secret is created automatically by CCO.

2. Install the OADP Operator:
 - a. In the OpenShift Container Platform web console, browse to **Operators** → **OperatorHub**.
 - b. Search for the **OADP Operator**.
 - c. In the **role_ARN** field, paste the `role_arn` that you created previously and click **Install**.
3. Create AWS cloud storage using your AWS credentials by entering the following command:

```
$ cat << EOF | oc create -f -
apiVersion: oadp.openshift.io/v1alpha1
kind: CloudStorage
metadata:
  name: ${CLUSTER_NAME}-oadp
  namespace: openshift-adp
spec:
  creationSecret:
    key: credentials
    name: cloud-credentials
  enableSharedConfig: true
  name: ${CLUSTER_NAME}-oadp
  provider: aws
  region: $REGION
EOF
```

4. Check your application's storage default storage class by entering the following command:

```
$ oc get pvc -n <namespace>
```

Example output

NAME	STATUS	VOLUME	CAPACITY	ACCESS	MODES
STORAGECLASS	AGE				
applog	Bound	pvc-351791ae-b6ab-4e8b-88a4-30f73caf5ef8	1Gi	RWO	gp3-
csi	4d19h				
mysql	Bound	pvc-16b8e009-a20a-4379-accb-bc81fedd0621	1Gi	RWO	gp3-
csi	4d19h				

5. Get the storage class by running the following command:

```
$ oc get storageclass
```


Example output

NAME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDINGMODE	ALLOWVOLUMEEXPANSION	AGE
gp2	kubernetes.io/aws-efs	Delete	WaitForFirstConsumer	true	4d21h
gp2-csi	ebs.csi.aws.com	Delete	WaitForFirstConsumer	true	4d21h
gp3	ebs.csi.aws.com	Delete	WaitForFirstConsumer	true	4d21h
gp3-csi (default)	ebs.csi.aws.com	Delete	WaitForFirstConsumer	true	4d21h

**NOTE**

The following storage classes will work:

- gp3-csi
- gp2-csi
- gp3
- gp2

If the application or applications that are being backed up are all using persistent volumes (PVs) with Container Storage Interface (CSI), it is advisable to include the CSI plugin in the OADP DPA configuration.

6. Create the **DataProtectionApplication** resource to configure the connection to the storage where the backups and volume snapshots are stored:
 - a. If you are using only CSI volumes, deploy a Data Protection Application by entering the following command:

```
$ cat << EOF | oc create -f -
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: ${CLUSTER_NAME}-dpa
  namespace: openshift-adp
spec:
  backupImages: true 1
  features:
    dataMover:
      enable: false
  backupLocations:
  - bucket:
      cloudStorageRef:
        name: ${CLUSTER_NAME}-oadp
      credential:
        key: credentials
        name: cloud-credentials
      prefix: velero
      default: true
```

```

    config:
      region: ${REGION}
  configuration:
    velero:
      defaultPlugins:
        - openshift
        - aws
        - csi
    restic:
      enable: false
EOF

```

- 1 ROSA supports internal image backup. Set this field to **false** if you do not want to use image backup.

- a. If you are using CSI or non-CSI volumes, deploy a Data Protection Application by entering the following command:

```

$ cat << EOF | oc create -f -
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: ${CLUSTER_NAME}-dpa
  namespace: openshift-adp
spec:
  backupImages: true 1
  features:
    dataMover:
      enable: false
  backupLocations:
  - bucket:
      cloudStorageRef:
        name: ${CLUSTER_NAME}-oadp
      credential:
        key: credentials
        name: cloud-credentials
      prefix: velero
      default: true
      config:
        region: ${REGION}
  configuration:
    velero:
      defaultPlugins:
        - openshift
        - aws
    nodeAgent: 2
      enable: false
      uploaderType: restic
  snapshotLocations:
  - velero:
      config:
        credentialsFile: /tmp/credentials/openshift-adp/cloud-credentials-credentials 3
        enableSharedConfig: "true" 4
        profile: default 5

```

```

region: ${REGION} 6
provider: aws
EOF

```

- 1 ROSA supports internal image backup. Set this field to false if you do not want to use image backup.
- 2 See the following note.
- 3 The **credentialsFile** field is the mounted location of the bucket credential on the pod.
- 4 The **enableSharedConfig** field allows the **snapshotLocations** to share or reuse the credential defined for the bucket.
- 5 Use the profile name set in the AWS credentials file.
- 6 Specify **region** as your AWS region. This must be the same as the cluster region.

You are now ready to back up and restore OpenShift Container Platform applications, as described in *Backing up applications*.

NOTE

The **enable** parameter of **restic** is set to **false** in this configuration, because OADP does not support Restic in ROSA environments.

If you use OADP 1.2, replace this configuration:

```

nodeAgent:
  enable: false
  uploaderType: restic

```

with the following configuration:

```

restic:
  enable: false

```

NOTE

If you want to use two different clusters for backing up and restoring, the two clusters must have the same AWS S3 storage names in both the cloud storage CR and the OADP **DataProtectionApplication** configuration.

Additional resources

- [Installing from OperatorHub using the web console](#) .
- [Backing up applications](#)

4.8.1.3. Example: Backing up workload on OADP ROSA STS, with an optional cleanup

4.8.1.3.1. Performing a backup with OADP and ROSA STS

The following example **hello-world** application has no persistent volumes (PVs) attached. Perform a backup with OpenShift API for Data Protection (OADP) with Red Hat OpenShift Service on AWS (ROSA) STS.

Either Data Protection Application (DPA) configuration will work.

1. Create a workload to back up by running the following commands:

```
$ oc create namespace hello-world
```

```
$ oc new-app -n hello-world --image=docker.io/openshift/hello-openshift
```

2. Expose the route by running the following command:

```
$ oc expose service/hello-openshift -n hello-world
```

3. Check that the application is working by running the following command:

```
$ curl `oc get route/hello-openshift -n hello-world -o jsonpath='{.spec.host}'`
```

Example output

```
Hello OpenShift!
```

4. Back up the workload by running the following command:

```
$ cat << EOF | oc create -f -
  apiVersion: velero.io/v1
  kind: Backup
  metadata:
    name: hello-world
    namespace: openshift-adp
  spec:
    includedNamespaces:
      - hello-world
    storageLocation: ${CLUSTER_NAME}-dpa-1
    ttl: 720h0m0s
EOF
```

5. Wait until the backup is completed and then run the following command:

```
$ watch "oc -n openshift-adp get backup hello-world -o json | jq .status"
```

Example output

```
{
  "completionTimestamp": "2022-09-07T22:20:44Z",
  "expiration": "2022-10-07T22:20:22Z",
  "formatVersion": "1.1.0",
  "phase": "Completed",
  "progress": {
    "itemsBackedUp": 58,
    "totalItems": 58
  }
}
```

```

    },
    "startTimestamp": "2022-09-07T22:20:22Z",
    "version": 1
  }

```

6. Delete the demo workload by running the following command:

```
$ oc delete ns hello-world
```

7. Restore the workload from the backup by running the following command:

```

$ cat << EOF | oc create -f -
  apiVersion: velero.io/v1
  kind: Restore
  metadata:
    name: hello-world
    namespace: openshift-adp
  spec:
    backupName: hello-world
EOF

```

8. Wait for the Restore to finish by running the following command:

```
$ watch "oc -n openshift-adp get restore hello-world -o json | jq .status"
```

Example output

```

{
  "completionTimestamp": "2022-09-07T22:25:47Z",
  "phase": "Completed",
  "progress": {
    "itemsRestored": 38,
    "totalItems": 38
  },
  "startTimestamp": "2022-09-07T22:25:28Z",
  "warnings": 9
}

```

9. Check that the workload is restored by running the following command:

```
$ oc -n hello-world get pods
```

Example output

```

NAME                                READY STATUS RESTARTS AGE
hello-openshift-9f885f7c6-kdjjp  1/1   Running 0      90s


```

10. Check the JSONPath by running the following command:

```
$ curl `oc get route/hello-openshift -n hello-world -o jsonpath='{.spec.host}'`
```

Example output

Hello OpenShift!



NOTE

For troubleshooting tips, see the OADP team's [troubleshooting documentation](#).

4.8.1.3.2. Cleaning up a cluster after a backup with OADP and ROSA STS

If you need to uninstall the OpenShift API for Data Protection (OADP) Operator together with the backups and the S3 bucket from this example, follow these instructions.

Procedure

1. Delete the workload by running the following command:

```
$ oc delete ns hello-world
```

2. Delete the Data Protection Application (DPA) by running the following command:

```
$ oc -n openshift-adp delete dpa ${CLUSTER_NAME}-dpa
```

3. Delete the cloud storage by running the following command:

```
$ oc -n openshift-adp delete cloudstorage ${CLUSTER_NAME}-oadp
```



WARNING

If this command hangs, you might need to delete the finalizer by running the following command:

```
$ oc -n openshift-adp patch cloudstorage ${CLUSTER_NAME}-oadp -p '{"metadata":{"finalizers":null}}' --type=merge
```

4. If the Operator is no longer required, remove it by running the following command:

```
$ oc -n openshift-adp delete subscription oadp-operator
```

5. Remove the namespace from the Operator:

```
$ oc delete ns openshift-adp
```

6. If the backup and restore resources are no longer required, remove them from the cluster by running the following command:

```
$ oc delete backup hello-world
```

- To delete backup, restore and remote objects in AWS S3 run the following command:

```
$ velero backup delete hello-world
```

- If you no longer need the Custom Resource Definitions (CRD), remove them from the cluster by running the following command:

```
$ for CRD in `oc get crds | grep velero | awk '{print $1}'`; do oc delete crd $CRD; done
```

- Delete the AWS S3 bucket by running the following commands:

```
$ aws s3 rm s3://${CLUSTER_NAME}-oadp --recursive
```

```
$ aws s3api delete-bucket --bucket ${CLUSTER_NAME}-oadp
```

- Detach the policy from the role by running the following command:

```
$ aws iam detach-role-policy --role-name "${ROLE_NAME}" --policy-arn "${POLICY_ARN}"
```

- Delete the role by running the following command:

```
$ aws iam delete-role --role-name "${ROLE_NAME}"
```

4.9. OADP DATA MOVER

4.9.1. OADP Data Mover Introduction

OADP Data Mover allows you to restore stateful applications from the store if a failure, accidental deletion, or corruption of the cluster occurs.



NOTE

The OADP 1.1 Data Mover is a Technology Preview feature.

The OADP 1.2 Data Mover has significantly improved features and performances, but is still a Technology Preview feature.



IMPORTANT

The OADP Data Mover is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

- You can use OADP Data Mover to back up Container Storage Interface (CSI) volume snapshots to a remote object store. See [Using Data Mover for CSI snapshots](#).

- You can use OADP 1.2 Data Mover to backup and restore application data for clusters that use CephFS, CephRBD, or both. See [Using OADP 1.2 Data Mover with Ceph storage](#) .
- You must perform a data cleanup after you perform a backup, if you are using OADP 1.1 Data Mover. See [Cleaning up after a backup using OADP 1.1 Data Mover](#) .



NOTE

Post-migration hooks are not likely to work well with the OADP 1.3 Data Mover.

The OADP 1.1 and OADP 1.2 Data Movers use synchronous processes to back up and restore application data. Because the processes are synchronous, users can be sure that any post-restore hooks start only after the persistent volumes (PVs) of the related pods are released by the persistent volume claim (PVC) of the Data Mover.

However, the OADP 1.3 Data Mover uses an asynchronous process. As a result of this difference in sequencing, a post-restore hook might be called before the related PVs were released by the PVC of the Data Mover. If this happens, the pod remains in **Pending** status and cannot run the hook. The hook attempt might time out before the pod is released, leading to a **PartiallyFailed** restore operation.

4.9.1.1. OADP Data Mover prerequisites

- You have a stateful application running in a separate namespace.
- You have installed the OADP Operator by using Operator Lifecycle Manager (OLM).
- You have created an appropriate **VolumeSnapshotClass** and **StorageClass**.
- You have installed the VolSync operator using OLM.

4.9.2. Using Data Mover for CSI snapshots

The OADP Data Mover enables customers to back up Container Storage Interface (CSI) volume snapshots to a remote object store. When Data Mover is enabled, you can restore stateful applications, using CSI volume snapshots pulled from the object store if a failure, accidental deletion, or corruption of the cluster occurs.

The Data Mover solution uses the Restic option of VolSync.

Data Mover supports backup and restore of CSI volume snapshots only.

In OADP 1.2 Data Mover **VolumeSnapshotBackups** (VSBs) and **VolumeSnapshotRestores** (VSRs) are queued using the VolumeSnapshotMover (VSM). The VSM's performance is improved by specifying a concurrent number of VSBs and VSRs simultaneously **InProgress**. After all async plugin operations are complete, the backup is marked as complete.



NOTE

The OADP 1.1 Data Mover is a Technology Preview feature.

The OADP 1.2 Data Mover has significantly improved features and performances, but is still a Technology Preview feature.



IMPORTANT

The OADP Data Mover is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).



NOTE

Red Hat recommends that customers who use OADP 1.2 Data Mover in order to back up and restore ODF CephFS volumes, upgrade or install OpenShift Container Platform version 4.12 or later for improved performance. OADP Data Mover can leverage CephFS shallow volumes in OpenShift Container Platform version 4.12 or later, which based on our testing, can improve the performance of backup times.

- [CephFS ROX details](#)

Prerequisites

- You have verified that the **StorageClass** and **VolumeSnapshotClass** custom resources (CRs) support CSI.
- You have verified that only one **VolumeSnapshotClass** CR has the annotation **snapshot.storage.kubernetes.io/is-default-class: "true"**.



NOTE

In OpenShift Container Platform version 4.12 or later, verify that this is the only default **VolumeSnapshotClass**.

- You have verified that **deletionPolicy** of the **VolumeSnapshotClass** CR is set to **Retain**.
- You have verified that only one **StorageClass** CR has the annotation **storageclass.kubernetes.io/is-default-class: "true"**.
- You have included the label **velero.io/csi-volumesnapshot-class: "true"** in your **VolumeSnapshotClass** CR.
- You have verified that the **OADP namespace** has the annotation **oc annotate --overwrite namespace/openshift-adp volsync.backube/privileged-movers="true"**.



NOTE

In OADP 1.1 the above setting is mandatory.

In OADP 1.2 the **privileged-movers** setting is not required in most scenarios. The restoring container permissions should be adequate for the Volsync copy. In some user scenarios, there may be permission errors that the **privileged-mover=true** setting should resolve.

- You have installed the VolSync Operator by using the Operator Lifecycle Manager (OLM).

**NOTE**

The VolSync Operator is required for using OADP Data Mover.

- You have installed the OADP operator by using OLM.

Procedure

1. Configure a Restic secret by creating a **.yaml** file as following:

```
apiVersion: v1
kind: Secret
metadata:
  name: <secret_name>
  namespace: openshift-adp
type: Opaque
stringData:
  RESTIC_PASSWORD: <secure_restic_password>
```

**NOTE**

By default, the Operator looks for a secret named **dm-credential**. If you are using a different name, you need to specify the name through a Data Protection Application (DPA) CR using **dpa.spec.features.dataMover.credentialName**.

2. Create a DPA CR similar to the following example. The default plugins include CSI.

Example Data Protection Application (DPA) CR

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: velero-sample
  namespace: openshift-adp
spec:
  backupLocations:
    - velero:
        config:
          profile: default
          region: us-east-1
        credential:
          key: cloud
          name: cloud-credentials
        default: true
        objectStorage:
          bucket: <bucket_name>
          prefix: <bucket-prefix>
          provider: aws
  configuration:
    restic:
      enable: <true_or_false>
    velero:
      itemOperationSyncFrequency: "10s"
  defaultPlugins:
```

```

- openshift
- aws
- csi
- vsm 1
features:
  dataMover:
    credentialName: restic-secret
    enable: true
    maxConcurrentBackupVolumes: "3" 2
    maxConcurrentRestoreVolumes: "3" 3
    pruneInterval: "14" 4
    volumeOptions: 5
    sourceVolumeOptions:
      accessMode: ReadOnlyMany
      cacheAccessMode: ReadWriteOnce
      cacheCapacity: 2Gi
    destinationVolumeOptions:
      storageClass: other-storageclass-name
      cacheAccessMode: ReadWriteMany
  snapshotLocations:
    - velero:
      config:
        profile: default
        region: us-west-2
      provider: aws

```

- 1** OADP 1.2 only.
- 2** OADP 1.2 only. Optional: Specify the upper limit of the number of snapshots allowed to be queued for backup. The default value is 10.
- 3** OADP 1.2 only. Optional: Specify the upper limit of the number of snapshots allowed to be queued for restore. The default value is 10.
- 4** OADP 1.2 only. Optional: Specify the number of days, between running Restic pruning on the repository. The prune operation repacks the data to free space, but it can also generate significant I/O traffic as a part of the process. Setting this option allows a trade-off between storage consumption, from no longer referenced data, and access costs.
- 5** OADP 1.2 only. Optional: Specify VolumeSync volume options for backup and restore.

The OADP Operator installs two custom resource definitions (CRDs), **VolumeSnapshotBackup** and **VolumeSnapshotRestore**.

Example VolumeSnapshotBackup CRD

```

apiVersion: datamover.oadp.openshift.io/v1alpha1
kind: VolumeSnapshotBackup
metadata:
  name: <vsb_name>
  namespace: <namespace_name> 1
spec:
  volumeSnapshotContent:
    name: <snapcontent_name>

```

```
protectedNamespace: <adp_namespace> 2
resticSecretRef:
  name: <restic_secret_name>
```

- 1 Specify the namespace where the volume snapshot exists.
- 2 Specify the namespace where the OADP Operator is installed. The default is **openshift-adp**.

Example VolumeSnapshotRestore CRD

```
apiVersion: datamover.oadp.openshift.io/v1alpha1
kind: VolumeSnapshotRestore
metadata:
  name: <vsr_name>
  namespace: <namespace_name> 1
spec:
  protectedNamespace: <protected_ns> 2
  resticSecretRef:
    name: <restic_secret_name>
  volumeSnapshotMoverBackupRef:
    sourcePVCDData:
      name: <source_pvc_name>
      size: <source_pvc_size>
    resticrepository: <your_rectic_repo>
    volumeSnapshotClassName: <vsclass_name>
```

- 1 Specify the namespace where the volume snapshot exists.
- 2 Specify the namespace where the OADP Operator is installed. The default is **openshift-adp**.

3. You can back up a volume snapshot by performing the following steps:

a. Create a backup CR:

```
apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup_name>
  namespace: <protected_ns> 1
spec:
  includedNamespaces:
    - <app_ns> 2
  storageLocation: velero-sample-1
```

- 1 Specify the namespace where the Operator is installed. The default namespace is **openshift-adp**.
- 2 Specify the application namespace or namespaces to be backed up.

- b. Wait up to 10 minutes and check whether the **VolumeSnapshotBackup** CR status is **Completed** by entering the following commands:

```
$ oc get vsb -n <app_ns>
```

```
$ oc get vsb <vsb_name> -n <app_ns> -o jsonpath="{.status.phase}"
```

A snapshot is created in the object store was configured in the DPA.



NOTE

If the status of the **VolumeSnapshotBackup** CR becomes **Failed**, refer to the Velero logs for troubleshooting.

4. You can restore a volume snapshot by performing the following steps:
- Delete the application namespace and the **VolumeSnapshotContent** that was created by the Velero CSI plugin.
 - Create a **Restore** CR and set **restorePVs** to **true**.

Example Restore CR


```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: <restore_name>
  namespace: <protected_ns>
spec:
  backupName: <previous_backup_name>
  restorePVs: true
```

- c. Wait up to 10 minutes and check whether the **VolumeSnapshotRestore** CR status is **Completed** by entering the following command:

```
$ oc get vsr -n <app_ns>
```

```
$ oc get vsr <vsr_name> -n <app_ns> -o jsonpath="{.status.phase}"
```

- d. Check whether your application data and resources have been restored.



NOTE

If the status of the **VolumeSnapshotRestore** CR becomes 'Failed', refer to the Velero logs for troubleshooting.

4.9.3. Using OADP 1.2 Data Mover with Ceph storage

You can use OADP 1.2 Data Mover to backup and restore application data for clusters that use CephFS, CephRBD, or both.

OADP 1.2 Data Mover leverages Ceph features that support large-scale environments. One of these is the shallow copy method, which is available for OpenShift Container Platform 4.12 and later. This feature

supports backing up and restoring **StorageClass** and **AccessMode** resources other than what is found on the source persistent volume claim (PVC).



IMPORTANT

The CephFS shallow copy feature is a back up feature. It is not part of restore operations.

4.9.3.1. Prerequisites for using OADP 1.2 Data Mover with Ceph storage

The following prerequisites apply to all back up and restore operations of data using OpenShift API for Data Protection (OADP) 1.2 Data Mover in a cluster that uses Ceph storage:

- You have installed OpenShift Container Platform 4.12 or later.
- You have installed the OADP Operator.
- You have created a secret **cloud-credentials** in the namespace **openshift-adp**.
- You have installed Red Hat OpenShift Data Foundation.
- You have installed the latest VolSync Operator by using Operator Lifecycle Manager.

4.9.3.2. Defining custom resources for use with OADP 1.2 Data Mover

When you install Red Hat OpenShift Data Foundation, it automatically creates default CephFS and a CephRBD **StorageClass** and **VolumeSnapshotClass** custom resources (CRs). You must define these CRs for use with OpenShift API for Data Protection (OADP) 1.2 Data Mover.

After you define the CRs, you must make several other changes to your environment before you can perform your back up and restore operations.

4.9.3.2.1. Defining CephFS custom resources for use with OADP 1.2 Data Mover

When you install Red Hat OpenShift Data Foundation, it automatically creates a default CephFS **StorageClass** custom resource (CR) and a default CephFS **VolumeSnapshotClass** CR. You can define these CRs for use with OpenShift API for Data Protection (OADP) 1.2 Data Mover.

Procedure

1. Define the **VolumeSnapshotClass** CR as in the following example:

Example VolumeSnapshotClass CR

```
apiVersion: snapshot.storage.k8s.io/v1
deletionPolicy: Retain 1
driver: openshift-storage.cephfs.csi.ceph.com
kind: VolumeSnapshotClass
metadata:
  annotations:
    snapshot.storage.kubernetes.io/is-default-class: true 2
  labels:
    velero.io/csi-volumesnapshot-class: true 3
name: ocs-storagecluster-cephfsplugin-snapclass
parameters:
```

```
clusterID: openshift-storage
csi.storage.k8s.io/snapshotter-secret-name: rook-csi-cephfs-provisioner
csi.storage.k8s.io/snapshotter-secret-namespace: openshift-storage
```

- 1 Must be set to **Retain**.
- 2 Must be set to **true**.
- 3 Must be set to **true**.

2. Define the **StorageClass** CR as in the following example:

Example StorageClass CR

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ocs-storagecluster-cephfs
  annotations:
    description: Provides RWO and RWX Filesystem volumes
    storageclass.kubernetes.io/is-default-class: true 1
provisioner: openshift-storage.cephfs.csi.ceph.com
parameters:
  clusterID: openshift-storage
  csi.storage.k8s.io/controller-expand-secret-name: rook-csi-cephfs-provisioner
  csi.storage.k8s.io/controller-expand-secret-namespace: openshift-storage
  csi.storage.k8s.io/node-stage-secret-name: rook-csi-cephfs-node
  csi.storage.k8s.io/node-stage-secret-namespace: openshift-storage
  csi.storage.k8s.io/provisioner-secret-name: rook-csi-cephfs-provisioner
  csi.storage.k8s.io/provisioner-secret-namespace: openshift-storage
  fsName: ocs-storagecluster-cephfilesystem
reclaimPolicy: Delete
allowVolumeExpansion: true
volumeBindingMode: Immediate
```

- 1 Must be set to **true**.

4.9.3.2.2. Defining CephRBD custom resources for use with OADP 1.2 Data Mover

When you install Red Hat OpenShift Data Foundation, it automatically creates a default CephRBD **StorageClass** custom resource (CR) and a default CephRBD **VolumeSnapshotClass** CR. You can define these CRs for use with OpenShift API for Data Protection (OADP) 1.2 Data Mover.

Procedure

1. Define the **VolumeSnapshotClass** CR as in the following example:

Example VolumeSnapshotClass CR

```
apiVersion: snapshot.storage.k8s.io/v1
deletionPolicy: Retain 1
driver: openshift-storage.rbd.csi.ceph.com
kind: VolumeSnapshotClass
```

```

metadata:
  labels:
    velero.io/csi-volumesnapshot-class: true 2
  name: ocs-storagecluster-rbdplugin-snapclass
parameters:
  clusterID: openshift-storage
  csi.storage.k8s.io/snapshotter-secret-name: rook-csi-rbd-provisioner
  csi.storage.k8s.io/snapshotter-secret-namespace: openshift-storage

```

1 Must be set to **Retain**.

2 Must be set to **true**.

2. Define the **StorageClass** CR as in the following example:

Example StorageClass CR

```

kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ocs-storagecluster-ceph-rbd
  annotations:
    description: 'Provides RWO Filesystem volumes, and RWO and RWX Block volumes'
provisioner: openshift-storage.rbd.csi.ceph.com
parameters:
  csi.storage.k8s.io/fstype: ext4
  csi.storage.k8s.io/provisioner-secret-namespace: openshift-storage
  csi.storage.k8s.io/provisioner-secret-name: rook-csi-rbd-provisioner
  csi.storage.k8s.io/node-stage-secret-name: rook-csi-rbd-node
  csi.storage.k8s.io/controller-expand-secret-name: rook-csi-rbd-provisioner
  imageFormat: '2'
  clusterID: openshift-storage
  imageFeatures: layering
  csi.storage.k8s.io/controller-expand-secret-namespace: openshift-storage
  pool: ocs-storagecluster-cephblockpool
  csi.storage.k8s.io/node-stage-secret-namespace: openshift-storage
  reclaimPolicy: Delete
  allowVolumeExpansion: true
  volumeBindingMode: Immediate

```

4.9.3.2.3. Defining additional custom resources for use with OADP 1.2 Data Mover

After you redefine the default **StorageClass** and CephRBD **VolumeSnapshotClass** custom resources (CRs), you must create the following CRs:

- A CephFS **StorageClass** CR defined to use the shallow copy feature
- A Restic **Secret** CR

Procedure

1. Create a CephFS **StorageClass** CR and set the **backingSnapshot** parameter set to **true** as in the following example:

Example CephFS StorageClass CR with backingSnapshot set to true

```

kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ocs-storagecluster-cephfs-shallow
  annotations:
    description: Provides RWO and RWX Filesystem volumes
    storageclass.kubernetes.io/is-default-class: false
provisioner: openshift-storage.cephfs.csi.ceph.com
parameters:
  csi.storage.k8s.io/provisioner-secret-namespace: openshift-storage
  csi.storage.k8s.io/provisioner-secret-name: rook-csi-cephfs-provisioner
  csi.storage.k8s.io/node-stage-secret-name: rook-csi-cephfs-node
  csi.storage.k8s.io/controller-expand-secret-name: rook-csi-cephfs-provisioner
  clusterID: openshift-storage
  fsName: ocs-storagecluster-cephfilesystem
  csi.storage.k8s.io/controller-expand-secret-namespace: openshift-storage
  backingSnapshot: true 1
  csi.storage.k8s.io/node-stage-secret-namespace: openshift-storage
  reclaimPolicy: Delete
  allowVolumeExpansion: true
  volumeBindingMode: Immediate

```

- 1** Must be set to **true**.



IMPORTANT

Ensure that the CephFS **VolumeSnapshotClass** and **StorageClass** CRs have the same value for **provisioner**.

2. Configure a Restic **Secret** CR as in the following example:

Example Restic Secret CR

```

apiVersion: v1
kind: Secret
metadata:
  name: <secret_name>
  namespace: <namespace>
type: Opaque
stringData:
  RESTIC_PASSWORD: <restic_password>

```

4.9.3.3. Backing up and restoring data using OADP 1.2 Data Mover and CephFS storage

You can use OpenShift API for Data Protection (OADP) 1.2 Data Mover to back up and restore data using CephFS storage by enabling the shallow copy feature of CephFS.

Prerequisites

- A stateful application is running in a separate namespace with persistent volume claims (PVCs) using CephFS as the provisioner.
- The **StorageClass** and **VolumeSnapshotClass** custom resources (CRs) are defined for CephFS and OADP 1.2 Data Mover.
- There is a secret **cloud-credentials** in the **openshift-adp** namespace.

4.9.3.3.1. Creating a DPA for use with CephFS storage

You must create a Data Protection Application (DPA) CR before you use the OpenShift API for Data Protection (OADP) 1.2 Data Mover to back up and restore data using CephFS storage.

Procedure

1. Verify that the **deletionPolicy** field of the **VolumeSnapshotClass** CR is set to **Retain** by running the following command:

```
$ oc get volumesnapshotclass -A -o jsonpath='{range .items[*]}{"Name: "}{.metadata.name} {" "}"{"Retention Policy: "}{.deletionPolicy}"{"\n"}{"end}'
```

2. Verify that the labels of the **VolumeSnapshotClass** CR are set to **true** by running the following command:

```
$ oc get volumesnapshotclass -A -o jsonpath='{range .items[*]}{"Name: "}{.metadata.name} {" "}"{"labels: "}{.metadata.labels}"{"\n"}{"end}'
```

3. Verify that the **storageclass.kubernetes.io/is-default-class** annotation of the **StorageClass** CR is set to **true** by running the following command:

```
$ oc get storageClass -A -o jsonpath='{range .items[*]}{"Name: "}{.metadata.name}" {" "}"{"annotations: "}{.metadata.annotations}"{"\n"}{"end}'
```

4. Create a Data Protection Application (DPA) CR similar to the following example:

Example DPA CR

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: velero-sample
  namespace: openshift-adp
spec:
  backupLocations:
    - velero:
        config:
          profile: default
          region: us-east-1
        credential:
          key: cloud
          name: cloud-credentials
        default: true
        objectStorage:
          bucket: <my_bucket>
```

```

    prefix: velero
    provider: aws
  configuration:
    restic:
      enable: false ❶
    velero:
      defaultPlugins:
        - openshift
        - aws
        - csi
        - vsm
  features:
    dataMover:
      credentialName: <restic_secret_name> ❷
      enable: true ❸
      volumeOptionsForStorageClasses: ❹
        ocs-storagecluster-cephfs:
          sourceVolumeOptions:
            accessMode: ReadOnlyMany
            cacheAccessMode: ReadWriteMany
            cacheStorageClassName: ocs-storagecluster-cephfs
            storageClassName: ocs-storagecluster-cephfs-shallow

```

- ❶ There is no default value for the **enable** field. Valid values are **true** or **false**.
- ❷ Use the Restic **Secret** that you created when you prepared your environment for working with OADP 1.2 Data Mover and Ceph. If you do not use your Restic **Secret**, the CR uses the default value **dm-credential** for this parameter.
- ❸ There is no default value for the **enable** field. Valid values are **true** or **false**.
- ❹ Optional parameter. You can define a different set of **VolumeOptionsForStorageClass** labels for each **storageClass** volume. This configuration provides a backup for volumes with different providers. The optional **VolumeOptionsForStorageClass** parameter is typically used with CephFS but can be used for any storage type.

4.9.3.3.2. Backing up data using OADP 1.2 Data Mover and CephFS storage

You can use OpenShift API for Data Protection (OADP) 1.2 Data Mover to back up data using CephFS storage by enabling the shallow copy feature of CephFS storage.

Procedure

1. Create a **Backup** CR as in the following example:

Example Backup CR

```

apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup_name>
  namespace: <protected_ns>
spec:

```

```
includedNamespaces:
- <app_ns>
storageLocation: velero-sample-1
```

2. Monitor the progress of the **VolumeSnapshotBackup** CRs by completing the following steps:
 - a. To check the progress of all the **VolumeSnapshotBackup** CRs, run the following command:

```
$ oc get vsb -n <app_ns>
```
 - b. To check the progress of a specific **VolumeSnapshotBackup** CR, run the following command:

```
$ oc get vsb <vsb_name> -n <app_ns> -ojsonpath="{.status.phase}"`
```
3. Wait several minutes until the **VolumeSnapshotBackup** CR has the status **Completed**.
4. Verify that there is at least one snapshot in the object store that is given in the Restic **Secret**. You can check for this snapshot in your targeted **BackupStorageLocation** storage provider that has a prefix of **/<OADP_namespace>**.

4.9.3.3.3. Restoring data using OADP 1.2 Data Mover and CephFS storage

You can use OpenShift API for Data Protection (OADP) 1.2 Data Mover to restore data using CephFS storage if the shallow copy feature of CephFS storage was enabled for the back up procedure. The shallow copy feature is not used in the restore procedure.

Procedure

1. Delete the application namespace by running the following command:

```
$ oc delete vsb -n <app_namespace> --all
```

2. Delete any **VolumeSnapshotContent** CRs that were created during backup by running the following command:

```
$ oc delete volumesnapshotcontent --all
```

3. Create a **Restore** CR as in the following example:

Example Restore CR

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: <restore_name>
  namespace: <protected_ns>
spec:
  backupName: <previous_backup_name>
```

4. Monitor the progress of the **VolumeSnapshotRestore** CRs by doing the following:

- a. To check the progress of all the **VolumeSnapshotRestore** CRs, run the following command:

```
$ oc get vsr -n <app_ns>
```

- b. To check the progress of a specific **VolumeSnapshotRestore** CR, run the following command:

```
$ oc get vsr <vsr_name> -n <app_ns> -ojsonpath="{.status.phase}"
```

5. Verify that your application data has been restored by running the following command:

```
$ oc get route <route_name> -n <app_ns> -ojsonpath="{.spec.host}"
```

4.9.3.4. Backing up and restoring data using OADP 1.2 Data Mover and split volumes (CephFS and Ceph RBD)

You can use OpenShift API for Data Protection (OADP) 1.2 Data Mover to back up and restore data in an environment that has *split volumes*, that is, an environment that uses both CephFS and CephRBD.

Prerequisites

- A stateful application is running in a separate namespace with persistent volume claims (PVCs) using CephFS as the provisioner.
- The **StorageClass** and **VolumeSnapshotClass** custom resources (CRs) are defined for CephFS and OADP 1.2 Data Mover.
- There is a secret **cloud-credentials** in the **openshift-adp** namespace.

4.9.3.4.1. Creating a DPA for use with split volumes

You must create a Data Protection Application (DPA) CR before you use the OpenShift API for Data Protection (OADP) 1.2 Data Mover to back up and restore data using split volumes.

Procedure

- Create a Data Protection Application (DPA) CR as in the following example:

Example DPA CR for environment with split volumes

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: velero-sample
  namespace: openshift-adp
spec:
  backupLocations:
  - velero:
      config:
        profile: default
        region: us-east-1
      credential:
        key: cloud
```

```

    name: cloud-credentials
    default: true
    objectStorage:
      bucket: <my-bucket>
      prefix: velero
      provider: aws
  configuration:
    restic:
      enable: false
    velero:
      defaultPlugins:
        - openshift
        - aws
        - csi
        - vsm
  features:
    dataMover:
      credentialName: <restic_secret_name> 1
      enable: true
      volumeOptionsForStorageClasses: 2
        ocs-storagecluster-cephfs:
          sourceVolumeOptions:
            accessMode: ReadOnlyMany
            cacheAccessMode: ReadWriteMany
            cacheStorageClassName: ocs-storagecluster-cephfs
            storageClassName: ocs-storagecluster-cephfs-shallow
        ocs-storagecluster-ceph-rbd:
          sourceVolumeOptions:
            storageClassName: ocs-storagecluster-ceph-rbd
            cacheStorageClassName: ocs-storagecluster-ceph-rbd
          destinationVolumeOptions:
            storageClassName: ocs-storagecluster-ceph-rbd
            cacheStorageClassName: ocs-storagecluster-ceph-rbd

```

- 1 Use the Restic **Secret** that you created when you prepared your environment for working with OADP 1.2 Data Mover and Ceph. If you do not, then the CR will use the default value **dm-credential** for this parameter.
- 2 A different set of **VolumeOptionsForStorageClass** labels can be defined for each **storageClass** volume, thus allowing a backup to volumes with different providers. The **VolumeOptionsForStorageClass** parameter is meant for use with CephFS. However, the optional **VolumeOptionsForStorageClass** parameter could be used for any storage type.

4.9.3.4.2. Backing up data using OADP 1.2 Data Mover and split volumes

You can use OpenShift API for Data Protection (OADP) 1.2 Data Mover to back up data in an environment that has split volumes.

Procedure

1. Create a **Backup** CR as in the following example:

Example Backup CR

```

apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup_name>
  namespace: <protected_ns>
spec:
  includedNamespaces:
  - <app_ns>
  storageLocation: velero-sample-1

```

2. Monitor the progress of the **VolumeSnapshotBackup** CRs by completing the following steps:

- a. To check the progress of all the **VolumeSnapshotBackup** CRs, run the following command:

```
$ oc get vsb -n <app_ns>
```

- b. To check the progress of a specific **VolumeSnapshotBackup** CR, run the following command:

```
$ oc get vsb <vsb_name> -n <app_ns> -ojsonpath="{.status.phase}"`
```

3. Wait several minutes until the **VolumeSnapshotBackup** CR has the status **Completed**.

4. Verify that there is at least one snapshot in the object store that is given in the Restic **Secret**. You can check for this snapshot in your targeted **BackupStorageLocation** storage provider that has a prefix of **/<OADP_namespace>**.

4.9.3.4.3. Restoring data using OADP 1.2 Data Mover and split volumes

You can use OpenShift API for Data Protection (OADP) 1.2 Data Mover to restore data in an environment that has split volumes, if the shallow copy feature of CephFS storage was enabled for the back up procedure. The shallow copy feature is not used in the restore procedure.

Procedure

1. Delete the application namespace by running the following command:

```
$ oc delete vsb -n <app_namespace> --all
```

2. Delete any **VolumeSnapshotContent** CRs that were created during backup by running the following command:

```
$ oc delete volumesnapshotcontent --all
```

3. Create a **Restore** CR as in the following example:

Example Restore CR

```

apiVersion: velero.io/v1
kind: Restore
metadata:
  name: <restore_name>

```

```
namespace: <protected_ns>
spec:
  backupName: <previous_backup_name>
```

4. Monitor the progress of the **VolumeSnapshotRestore** CRs by doing the following:

- a. To check the progress of all the **VolumeSnapshotRestore** CRs, run the following command:

```
$ oc get vsr -n <app_ns>
```

- b. To check the progress of a specific **VolumeSnapshotRestore** CR, run the following command:

```
$ oc get vsr <vsr_name> -n <app_ns> -ojsonpath="{.status.phase}"
```

5. Verify that your application data has been restored by running the following command:

```
$ oc get route <route_name> -n <app_ns> -ojsonpath="{.spec.host}"
```

4.9.4. Cleaning up after a backup using OADP 1.1 Data Mover

For OADP 1.1 Data Mover, you must perform a data cleanup after you perform a backup.

The cleanup consists of deleting the following resources:

- Snapshots in a bucket
- Cluster resources
- Volume snapshot backups (VSBs) after a backup procedure that is either run by a schedule or is run repetitively

4.9.4.1. Deleting snapshots in a bucket

OADP 1.1 Data Mover might leave one or more snapshots in a bucket after a backup. You can either delete all the snapshots or delete individual snapshots.

Procedure

- To delete all snapshots in your bucket, delete the `/<protected_namespace>` folder that is specified in the Data Protection Application (DPA) `.spec.backupLocation.objectStorage.bucket` resource.
- To delete an individual snapshot:
 1. Browse to the `/<protected_namespace>` folder that is specified in the DPA `.spec.backupLocation.objectStorage.bucket` resource.
 2. Delete the appropriate folders that are prefixed with `/<volumeSnapshotContent name>-pvc` where `<VolumeSnapshotContent name>` is the **VolumeSnapshotContent** created by Data Mover per PVC.

4.9.4.2. Deleting cluster resources

OADP 1.1 Data Mover might leave cluster resources whether or not it successfully backs up your container storage interface (CSI) volume snapshots to a remote object store.

4.9.4.2.1. Deleting cluster resources following a successful backup and restore that used Data Mover

You can delete any **VolumeSnapshotBackup** or **VolumeSnapshotRestore** CRs that remain in your application namespace after a successful backup and restore where you used Data Mover.

Procedure

1. Delete cluster resources that remain on the application namespace, the namespace with the application PVCs to backup and restore, after a backup where you use Data Mover:

```
$ oc delete vsb -n <app_namespace> --all
```

2. Delete cluster resources that remain after a restore where you use Data Mover:

```
$ oc delete vsr -n <app_namespace> --all
```

3. If needed, delete any **VolumeSnapshotContent** resources that remain after a backup and restore where you use Data Mover:

```
$ oc delete volumesnapshotcontent --all
```

4.9.4.2.2. Deleting cluster resources following a partially successful or a failed backup and restore that used Data Mover

If your backup and restore operation that uses Data Mover either fails or only partially succeeds, you must clean up any **VolumeSnapshotBackup** (VSB) or **VolumeSnapshotRestore** custom resource definitions (CRDs) that exist in the application namespace, and clean up any extra resources created by these controllers.

Procedure

1. Clean up cluster resources that remain after a backup operation where you used Data Mover by entering the following commands:
 - a. Delete VSB CRDs on the application namespace, the namespace with the application PVCs to backup and restore:

```
$ oc delete vsb -n <app_namespace> --all
```

- b. Delete **VolumeSnapshot** CRs:

```
$ oc delete volumesnapshot -A --all
```

- c. Delete **VolumeSnapshotContent** CRs:

```
$ oc delete volumesnapshotcontent --all
```

- d. Delete any PVCs on the protected namespace, the namespace the Operator is installed on.

```
$ oc delete pvc -n <protected_namespace> --all
```

- e. Delete any **ReplicationSource** resources on the namespace.

```
$ oc delete replicationsource -n <protected_namespace> --all
```

2. Clean up cluster resources that remain after a restore operation using Data Mover by entering the following commands:

- a. Delete VSR CRDs:

```
$ oc delete vsr -n <app-ns> --all
```

- b. Delete **VolumeSnapshot** CRs:

```
$ oc delete volumesnapshot -A --all
```

- c. Delete **VolumeSnapshotContent** CRs:

```
$ oc delete volumesnapshotcontent --all
```

- d. Delete any **ReplicationDestination** resources on the namespace.

```
$ oc delete replicationdestination -n <protected_namespace> --all
```


4.10. OADP 1.3 DATA MOVER

4.10.1. About the OADP 1.3 Data Mover

OADP 1.3 includes a built-in Data Mover that you can use to move Container Storage Interface (CSI) volume snapshots to a remote object store. The built-in Data Mover allows you to restore stateful applications from the remote object store if a failure, accidental deletion, or corruption of the cluster occurs. It uses [Kopia](#) as the uploader mechanism to read the snapshot data and write to the unified repository.

OADP supports CSI snapshots on the following:

- Red Hat OpenShift Data Foundation
- Any other cloud storage provider with the Container Storage Interface (CSI) driver that supports the Kubernetes Volume Snapshot API



IMPORTANT

The OADP built-in Data Mover is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

4.10.1.1. Enabling the built-in Data Mover

To enable the built-in Data Mover, you must include the CSI plugin and enable the node agent in the **DataProtectionApplication** custom resource (CR). The node agent is a Kubernetes daemonset that hosts data movement modules. These include the Data Mover controller, uploader, and the repository.

Example DataProtectionApplication manifest

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: dpa-sample
spec:
  configuration:
    nodeAgent:
      enable: true 1
      uploaderType: kopia 2
    velero:
      defaultPlugins:
        - openshift
        - aws
        - csi 3
# ...
```

- 1** The flag to enable the node agent.
- 2** The type of uploader. The possible values are **restic** or **kopia**. The built-in Data Mover uses Kopia as the default uploader mechanism regardless of the value of the **uploaderType** field.
- 3** The CSI plugin included in the list of default plugins.

4.10.1.2. Built-in Data Mover controller and custom resource definitions (CRDs)

The built-in Data Mover feature introduces three new API objects defined as CRDs for managing backup and restore:

- **DataDownload**: Represents a data download of a volume snapshot. The CSI plugin creates one **DataDownload** object per volume to be restored. The **DataDownload** CR includes information about the target volume, the specified Data Mover, the progress of the current data download, the specified backup repository, and the result of the current data download after the process is complete.
- **DataUpload**: Represents a data upload of a volume snapshot. The CSI plugin creates one **DataUpload** object per CSI snapshot. The **DataUpload** CR includes information about the specified snapshot, the specified Data Mover, the specified backup repository, the progress of the current data upload, and the result of the current data upload after the process is complete.
- **BackupRepository**: Represents and manages the lifecycle of the backup repositories. OADP creates a backup repository per namespace when the first CSI snapshot backup or restore for a namespace is requested.

4.10.2. Backing up and restoring CSI snapshots

You can back up and restore persistent volumes by using the OADP 1.3 Data Mover.

4.10.2.1. Backing up persistent volumes with CSI snapshots

You can use the OADP Data Mover to back up Container Storage Interface (CSI) volume snapshots to a remote object store.

Prerequisites

- You have access to the cluster with the **cluster-admin** role.
- You have installed the OADP Operator.
- You have included the CSI plugin and enabled the node agent in the **DataProtectionApplication** custom resource (CR).
- You have an application with persistent volumes running in a separate namespace.
- You have added the **metadata.labels.velero.io/csi-volumesnapshot-class: "true"** key-value pair to the **VolumeSnapshotClass** CR.

Procedure

1. Create a YAML file for the **Backup** object, as in the following example:

Example Backup CR

```
kind: Backup
apiVersion: velero.io/v1
metadata:
  name: backup
  namespace: openshift-adp
spec:
  csiSnapshotTimeout: 10m0s
  defaultVolumesToFsBackup: false
  includedNamespaces:
  - mysql-persistent
  itemOperationTimeout: 4h0m0s
  snapshotMoveData: true 1
  storageLocation: default
  ttl: 720h0m0s
  volumeSnapshotLocations:
  - dpa-sample-1
# ...
```

- 1** Set to **true** to enable movement of CSI snapshots to remote object storage.

2. Apply the manifest:

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

A **DataUpload** CR is created after the snapshot creation is complete.

Verification

- Verify that the snapshot data is successfully transferred to the remote object store by

monitoring the **status.phase** field of the **DataUpload** CR. Possible values are **In Progress**, **Completed**, **Failed**, or **Canceled**. The object store is configured in the **backupLocations** stanza of the **DataProtectionApplication** CR.

- Run the following command to get a list of all **DataUpload** objects:

```
$ oc get datauploads -A
```

Example output

```

NAMESPACE   NAME                STATUS   STARTED  BYTES DONE  TOTAL
BYTES STORAGE LOCATION AGE   NODE
openshift-adp backup-test-1-sw76b Completed 9m47s  108104082 108104082
dpa-sample-1  9m47s ip-10-0-150-57.us-west-2.compute.internal
openshift-adp mongo-block-7dtpf Completed 14m    1073741824 1073741824
dpa-sample-1  14m   ip-10-0-150-57.us-west-2.compute.internal

```

- Check the value of the **status.phase** field of the specific **DataUpload** object by running the following command:

```
$ oc get datauploads <dataupload_name> -o yaml
```

Example output

```


apiVersion: velero.io/v2alpha1
kind: DataUpload
metadata:
  name: backup-test-1-sw76b
  namespace: openshift-adp
spec:
  backupStorageLocation: dpa-sample-1
  csiSnapshot:
    snapshotClass: ""
    storageClass: gp3-csi
    volumeSnapshot: velero-mysql-fq8sl
  operationTimeout: 10m0s
  snapshotType: CSI
  sourceNamespace: mysql-persistent
  sourcePVC: mysql
status:
  completionTimestamp: "2023-11-02T16:57:02Z"
  node: ip-10-0-150-57.us-west-2.compute.internal
  path: /host_pods/15116bac-cc01-4d9b-8ee7-609c3bef6bde/volumes/kubernetes.io~csi/pvc-ead8167-556b-461a-b3ec-441749e291c4/mount
  phase: Completed 1
  progress:
    bytesDone: 108104082
    totalBytes: 108104082
  snapshotID: 8da1c5febf25225f4577ada2aeb9f899
  startTimestamp: "2023-11-02T16:56:22Z"

```

- 1** Indicates that snapshot data is successfully transferred to the remote object store.

4.10.2.2. Restoring CSI volume snapshots

You can restore a volume snapshot by creating a **Restore** CR.



NOTE

You cannot restore Volsync backups from OADP 1.2 with the OADP 1.3 built-in Data Mover. It is recommended to do a file system backup of all of your workloads with Restic prior to upgrading to OADP 1.3.

Prerequisites

- You have access to the cluster with the **cluster-admin** role.
- You have an OADP **Backup** CR from which to restore the data.

Procedure

1. Create a YAML file for the **Restore** CR, as in the following example:

Example Restore CR

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: restore
  namespace: openshift-adp
spec:
  backupName: <backup>
# ...
```

2. Apply the manifest:

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

A **DataDownload** CR is created when the restore starts.

Verification

- You can monitor the status of the restore process by checking the **status.phase** field of the **DataDownload** CR. Possible values are **In Progress**, **Completed**, **Failed**, or **Canceled**.
 - To get a list of all **DataDownload** objects, run the following command:

```
$ oc get datadownloads -A
```

Example output

```
NAMESPACE   NAME                               STATUS   STARTED  BYTES DONE  TOTAL
BYTES STORAGE LOCATION AGE  NODE
openshift-adp restore-test-1-sk7lg Completed 7m11s  108104082  108104082
dpa-sample-1 7m11s ip-10-0-150-57.us-west-2.compute.internal
```

- Enter the following command to check the value of the **status.phase** field of the specific **DataDownload** object:

```
$ oc get datadownloads <datadownload_name> -o yaml
```

Example output

```
apiVersion: velero.io/v2alpha1
kind: DataDownload
metadata:
  name: restore-test-1-sk7lg
  namespace: openshift-adp
spec:
  backupStorageLocation: dpa-sample-1
  operationTimeout: 10m0s
  snapshotID: 8da1c5febf25225f4577ada2aeb9f899
  sourceNamespace: mysql-persistent
  targetVolume:
    namespace: mysql-persistent
    pv: ""
    pvc: mysql
status:
  completionTimestamp: "2023-11-02T17:01:24Z"
  node: ip-10-0-150-57.us-west-2.compute.internal
  phase: Completed 1
  progress:
    bytesDone: 108104082
    totalBytes: 108104082
  startTimestamp: "2023-11-02T17:00:52Z"
```

- 1** Indicates that the CSI snapshot data is successfully restored.

4.11. TROUBLESHOOTING

You can debug Velero custom resources (CRs) by using the [OpenShift CLI tool](#) or the [Velero CLI tool](#). The Velero CLI tool provides more detailed logs and information.

You can check [installation issues](#), [backup and restore CR issues](#), and [Restic issues](#).

You can collect logs and CR information by using the [must-gather tool](#).

You can obtain the Velero CLI tool by:

- Downloading the Velero CLI tool
- Accessing the Velero binary in the Velero deployment in the cluster

4.11.1. Downloading the Velero CLI tool

You can download and install the Velero CLI tool by following the instructions on the [Velero documentation page](#).

The page includes instructions for:

- macOS by using Homebrew
- GitHub
- Windows by using Chocolatey

Prerequisites

- You have access to a Kubernetes cluster, v1.16 or later, with DNS and container networking enabled.
- You have installed **kubectl** locally.

Procedure

1. Open a browser and navigate to ["Install the CLI" on the Velero website](#).
2. Follow the appropriate procedure for macOS, GitHub, or Windows.
3. Download the Velero version appropriate for your version of OADP and OpenShift Container Platform.

4.11.1.1. OADP-Velero-OpenShift Container Platform version relationship

OADP version	Velero version	OpenShift Container Platform version
1.1.0	1.9	4.9 and later
1.1.1	1.9	4.9 and later
1.1.2	1.9	4.9 and later
1.1.3	1.9	4.9 and later
1.1.4	1.9	4.9 and later
1.1.5	1.9	4.9 and later
1.1.6	1.9	4.11 and later
1.1.7	1.9	4.11 and later
1.2.0	1.11	4.11 and later
1.2.1	1.11	4.11 and later
1.2.2	1.11	4.11 and later
1.2.3	1.11	4.11 and later

OADP version	Velero version	OpenShift Container Platform version
1.3.0	1.12	4.12 and later

4.11.2. Accessing the Velero binary in the Velero deployment in the cluster

You can use a shell command to access the Velero binary in the Velero deployment in the cluster.

Prerequisites

- Your **DataProtectionApplication** custom resource has a status of **Reconcile complete**.

Procedure

- Enter the following command to set the needed alias:

```
$ alias velero='oc -n openshift-adp exec deployment/velero -c velero -it -- ./velero'
```

4.11.3. Debugging Velero resources with the OpenShift CLI tool

You can debug a failed backup or restore by checking Velero custom resources (CRs) and the **Velero** pod log with the OpenShift CLI tool.

Velero CRs

Use the **oc describe** command to retrieve a summary of warnings and errors associated with a **Backup** or **Restore** CR:

```
$ oc describe <velero_cr> <cr_name>
```

Velero pod logs

Use the **oc logs** command to retrieve the **Velero** pod logs:

```
$ oc logs pod/<velero>
```

Velero pod debug logs

You can specify the Velero log level in the **DataProtectionApplication** resource as shown in the following example.



NOTE

This option is available starting from OADP 1.0.3.

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: velero-sample
spec:
  configuration:
    velero:
      logLevel: warning
```

The following **logLevel** values are available:

- **trace**
- **debug**
- **info**
- **warning**
- **error**
- **fatal**
- **panic**

It is recommended to use **debug** for most logs.

4.11.4. Debugging Velero resources with the Velero CLI tool

You can debug **Backup** and **Restore** custom resources (CRs) and retrieve logs with the Velero CLI tool.

The Velero CLI tool provides more detailed information than the OpenShift CLI tool.

Syntax

Use the **oc exec** command to run a Velero CLI command:

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  <backup_restore_cr> <command> <cr_name>
```

Example

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  backup describe 0e44ae00-5dc3-11eb-9ca8-df7e5254778b-2d8ql
```

Help option

Use the **velero --help** option to list all Velero CLI commands:

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  --help
```

Describe command

Use the **velero describe** command to retrieve a summary of warnings and errors associated with a **Backup** or **Restore** CR:

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  <backup_restore_cr> describe <cr_name>
```

Example

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  backup describe 0e44ae00-5dc3-11eb-9ca8-df7e5254778b-2d8ql
```

The following types of restore errors and warnings are shown in the output of a **velero describe** request:

- **Velero:** A list of messages related to the operation of Velero itself, for example, messages related to connecting to the cloud, reading a backup file, and so on
- **Cluster:** A list of messages related to backing up or restoring cluster-scoped resources
- **Namespaces:** A list of list of messages related to backing up or restoring resources stored in namespaces

One or more errors in one of these categories results in a **Restore** operation receiving the status of **PartiallyFailed** and not **Completed**. Warnings do not lead to a change in the completion status.



IMPORTANT

- For resource-specific errors, that is, **Cluster** and **Namespaces** errors, the **restore describe --details** output includes a resource list that lists all resources that Velero succeeded in restoring. For any resource that has such an error, check to see if the resource is actually in the cluster.
- If there are **Velero** errors, but no resource-specific errors, in the output of a **describe** command, it is possible that the restore completed without any actual problems in restoring workloads, but carefully validate post-restore applications. For example, if the output contains **PodVolumeRestore** or node agent-related errors, check the status of **PodVolumeRestores** and **DataDownloads**. If none of these are failed or still running, then volume data might have been fully restored.

Logs command

Use the **velero logs** command to retrieve the logs of a **Backup** or **Restore** CR:

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  <backup_restore_cr> logs <cr_name>
```

Example

```
$ oc -n openshift-adp exec deployment/velero -c velero -- ./velero \
  restore logs ccc7c2d0-6017-11eb-afab-85d0007f5a19-x4lbf
```

4.11.5. Pods crash or restart due to lack of memory or CPU

If a Velero or Restic pod crashes due to a lack of memory or CPU, you can set specific resource requests for either of those resources.

Additional resources

- [CPU and memory requirements](#)

4.11.5.1. Setting resource requests for a Velero pod

You can use the **configuration.velero.podConfig.resourceAllocations** specification field in the **oadp_v1alpha1_dpa.yaml** file to set specific resource requests for a **Velero** pod.

Procedure

- Set the **cpu** and **memory** resource requests in the YAML file:

Example Velero file

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
configuration:
  velero:
    podConfig:
      resourceAllocations: 1
      requests:
        cpu: 200m
        memory: 256Mi
```

- 1 The **resourceAllocations** listed are for average usage.

4.11.5.2. Setting resource requests for a Restic pod

You can use the **configuration.restic.podConfig.resourceAllocations** specification field to set specific resource requests for a **Restic** pod.

Procedure

- Set the **cpu** and **memory** resource requests in the YAML file:

Example Restic file

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
configuration:
  restic:
    podConfig:
      resourceAllocations: 1
      requests:
        cpu: 1000m
        memory: 16Gi
```

- 1 The **resourceAllocations** listed are for average usage.



IMPORTANT

The values for the resource request fields must follow the same format as Kubernetes resource requirements. Also, if you do not specify **configuration.velero.podConfig.resourceAllocations** or **configuration.restic.podConfig.resourceAllocations**, the default **resources** specification for a Velero pod or a Restic pod is as follows:

```
requests:
  cpu: 500m
  memory: 128Mi
```

4.11.6. Issues with Velero and admission webhooks

Velero has limited abilities to resolve admission webhook issues during a restore. If you have workloads with admission webhooks, you might need to use an additional Velero plugin or make changes to how you restore the workload.

Typically, workloads with admission webhooks require you to create a resource of a specific kind first. This is especially true if your workload has child resources because admission webhooks typically block child resources.

For example, creating or restoring a top-level object such as **service.serving.knative.dev** typically creates child resources automatically. If you do this first, you will not need to use Velero to create and restore these resources. This avoids the problem of child resources being blocked by an admission webhook that Velero might use.

4.11.6.1. Restoring workarounds for Velero backups that use admission webhooks

This section describes the additional steps required to restore resources for several types of Velero backups that use admission webhooks.

4.11.6.1.1. Restoring Knative resources

You might encounter problems using Velero to back up Knative resources that use admission webhooks.

You can avoid such problems by restoring the top level **Service** resource first whenever you back up and restore Knative resources that use admission webhooks.

Procedure

- Restore the top level **service.serving.knative.dev Service** resource:

```
$ velero restore <restore_name> \
--from-backup=<backup_name> --include-resources \
service.serving.knative.dev
```

4.11.6.1.2. Restoring IBM AppConnect resources

If you experience issues when you use Velero to restore an IBM AppConnect resource that has an admission webhook, you can run the checks in this procedure.

Procedure

1. Check if you have any mutating admission plugins of **kind: MutatingWebhookConfiguration** in the cluster:

```
$ oc get mutatingwebhookconfigurations
```

2. Examine the YAML file of each **kind: MutatingWebhookConfiguration** to ensure that none of its rules block creation of the objects that are experiencing issues. For more information, see [the official Kubernetes documentation](#).
3. Check that any **spec.version** in **type: Configuration.appconnect.ibm.com/v1beta1** used at backup time is supported by the installed Operator.

4.11.6.2. OADP plugins known issues

The following section describes known issues in OpenShift API for Data Protection (OADP) plugins:

4.11.6.2.1. Velero plugin panics during imagestream backups due to a missing secret

When the backup and the Backup Storage Location (BSL) are managed outside the scope of the Data Protection Application (DPA), the OADP controller, meaning the DPA reconciliation does not create the relevant **oadp-<bsl_name>-<bsl_provider>-registry-secret**.

When the backup is run, the OpenShift Velero plugin panics on the imagestream backup, with the following panic error:

```
024-02-27T10:46:50.028951744Z time="2024-02-27T10:46:50Z" level=error msg="Error backing up item"
backup=openshift-adp/<backup name> error="error executing custom action
(namespace=<BSL Name>, name=postgres): rpc error: code = Aborted desc = plugin panicked:
runtime error: index out of range with length 1, stack trace: goroutine 94...
```

4.11.6.2.1.1. Workaround to avoid the panic error

To avoid the Velero plugin panic error, perform the following steps:

1. Label the custom BSL with the relevant label:

```
$ oc label BackupStorageLocation <bsl_name> app.kubernetes.io/component=bsl
```

2. After the BSL is labeled, wait until the DPA reconciles.



NOTE

You can force the reconciliation by making any minor change to the DPA itself.

3. When the DPA reconciles, confirm that the relevant **oadp-<bsl_name>-<bsl_provider>-registry-secret** has been created and that the correct registry data has been populated into it:

```
$ oc -n openshift-adp get secret/oadp-<bsl_name>-<bsl_provider>-registry-secret -o json | jq
-r '.data'
```

4.11.6.2.2. OpenShift ADP Controller segmentation fault

If you configure a DPA with both **cloudstorage** and **restic** enabled, the **openshift-adp-controller-manager** pod crashes and restarts indefinitely until the pod fails with a crash loop segmentation fault.

You can have either **velero** or **cloudstorage** defined, because they are mutually exclusive fields.

- If you have both **velero** and **cloudstorage** defined, the **openshift-adp-controller-manager** fails.
- If you have neither **velero** nor **cloudstorage** defined, the **openshift-adp-controller-manager** fails.

For more information about this issue, see [OADP-1054](#).

4.11.6.2.2.1. OpenShift ADP Controller segmentation fault workaround

You must define either **velero** or **cloudstorage** when you configure a DPA. If you define both APIs in your DPA, the **openshift-adp-controller-manager** pod fails with a crash loop segmentation fault.

4.11.6.3. Velero plugins returning "received EOF, stopping recv loop" message



NOTE

Velero plugins are started as separate processes. After the Velero operation has completed, either successfully or not, they exit. Receiving a **received EOF, stopping recv loop** message in the debug logs indicates that a plugin operation has completed. It does not mean that an error has occurred.

Additional resources

- [Admission plugins](#)
- [Webhook admission plugins](#)
- [Types of webhook admission plugins](#)

4.11.7. Installation issues

You might encounter issues caused by using invalid directories or incorrect credentials when you install the Data Protection Application.

4.11.7.1. Backup storage contains invalid directories

The **Velero** pod log displays the error message, **Backup storage contains invalid top-level directories**.

Cause

The object storage contains top-level directories that are not Velero directories.

Solution

If the object storage is not dedicated to Velero, you must specify a prefix for the bucket by setting the **spec.backupLocations.velero.objectStorage.prefix** parameter in the **DataProtectionApplication** manifest.

4.11.7.2. Incorrect AWS credentials

The `oadp-aws-registry` pod log displays the error message, **InvalidAccessKeyId: The AWS Access Key Id you provided does not exist in our records.**

The `Velero` pod log displays the error message, **NoCredentialProviders: no valid providers in chain.**

Cause

The `credentials-velero` file used to create the `Secret` object is incorrectly formatted.

Solution

Ensure that the `credentials-velero` file is correctly formatted, as in the following example:

Example `credentials-velero` file

```
[default] 1
aws_access_key_id=AKIAIOSFODNN7EXAMPLE 2
aws_secret_access_key=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

- 1 AWS default profile.
- 2 Do not enclose the values with quotation marks (" , ').

4.11.8. OADP Operator issues

The OpenShift API for Data Protection (OADP) Operator might encounter issues caused by problems it is not able to resolve.

4.11.8.1. OADP Operator fails silently

The S3 buckets of an OADP Operator might be empty, but when you run the command `oc get po -n <OADP_Operator_namespace>`, you see that the Operator has a status of **Running**. In such a case, the Operator is said to have *failed silently* because it incorrectly reports that it is running.

Cause

The problem is caused when cloud credentials provide insufficient permissions.

Solution

Retrieve a list of backup storage locations (BSLs) and check the manifest of each BSL for credential issues.

Procedure

1. Run one of the following commands to retrieve a list of BSLs:
 - a. Using the OpenShift CLI:

```
$ oc get backupstoragelocation -A
```

- b. Using the Velero CLI:


```
$ velero backup-location get -n <OADP_Operator_namespace>
```

- Using the list of BSLs, run the following command to display the manifest of each BSL, and examine each manifest for an error.

```
$ oc get backupstoragelocation -n <namespace> -o yaml
```

Example result

```
apiVersion: v1
items:
- apiVersion: velero.io/v1
  kind: BackupStorageLocation
  metadata:
    creationTimestamp: "2023-11-03T19:49:04Z"
    generation: 9703
    name: example-dpa-1
    namespace: openshift-adp-operator
    ownerReferences:
    - apiVersion: oadp.openshift.io/v1alpha1
      blockOwnerDeletion: true
      controller: true
      kind: DataProtectionApplication
      name: example-dpa
      uid: 0beeeaff-0287-4f32-bcb1-2e3c921b6e82
    resourceVersion: "24273698"
    uid: ba37cd15-cf17-4f7d-bf03-8af8655cea83
  spec:
    config:
      enableSharedConfig: "true"
      region: us-west-2
    credential:
      key: credentials
      name: cloud-credentials
    default: true
    objectStorage:
      bucket: example-oadp-operator
      prefix: example
    provider: aws
  status:
    lastValidationTime: "2023-11-10T22:06:46Z"
    message: "BackupStorageLocation \"example-dpa-1\" is unavailable: rpc
      error: code = Unknown desc = WebIdentityErr: failed to retrieve credentials\ncaused
      by: AccessDenied: Not authorized to perform sts:AssumeRoleWithWebIdentity\n\tstatus
      code: 403, request id: d3f2e099-70a0-467b-997e-ff62345e3b54"
    phase: Unavailable
  kind: List
  metadata:
    resourceVersion: ""
```

4.11.9. OADP timeouts

Extending a timeout allows complex or resource-intensive processes to complete successfully without premature termination. This configuration can reduce the likelihood of errors, retries, or failures.

Ensure that you balance timeout extensions in a logical manner so that you do not configure excessively long timeouts that might hide underlying issues in the process. Carefully consider and monitor an appropriate timeout value that meets the needs of the process and the overall system performance.

The following are various OADP timeouts, with instructions of how and when to implement these parameters:

4.11.9.1. Restic timeout

timeout defines the Restic timeout. The default value is **1h**.

Use the Restic **timeout** for the following scenarios:

- For Restic backups with total PV data usage that is greater than 500GB.
- If backups are timing out with the following error:

```
level=error msg="Error backing up item" backup=velero/monitoring error="timed out waiting
for all PodVolumeBackups to complete"
```

Procedure

- Edit the values in the **spec.configuration.restrict.timeout** block of the **DataProtectionApplication** CR manifest, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_name>
spec:
  configuration:
    restic:
      timeout: 1h
# ...
```

4.11.9.2. Velero resource timeout

resourceTimeout defines how long to wait for several Velero resources before timeout occurs, such as Velero custom resource definition (CRD) availability, **volumeSnapshot** deletion, and repository availability. The default is **10m**.

Use the **resourceTimeout** for the following scenarios:

- For backups with total PV data usage that is greater than 1TB. This parameter is used as a timeout value when Velero tries to clean up or delete the Container Storage Interface (CSI) snapshots, before marking the backup as complete.
 - A sub-task of this cleanup tries to patch VSC and this timeout can be used for that task.
- To create or ensure a backup repository is ready for filesystem based backups for Restic or Kopia.
- To check if the Velero CRD is available in the cluster before restoring the custom resource (CR) or resource from the backup.

Procedure

- Edit the values in the **spec.configuration.velero.resourceTimeout** block of the **DataProtectionApplication** CR manifest, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_name>
spec:
  configuration:
    velero:
      resourceTimeout: 10m
# ...

```

4.11.9.3. Data Mover timeout

timeout is a user-supplied timeout to complete **VolumeSnapshotBackup** and **VolumeSnapshotRestore**. The default value is **10m**.

Use the Data Mover **timeout** for the following scenarios:

- If creation of **VolumeSnapshotBackups** (VSBs) and **VolumeSnapshotRestores** (VSRs), times out after 10 minutes.
- For large scale environments with total PV data usage that is greater than 500GB. Set the timeout for **1h**.
- With the **VolumeSnapshotMover** (VSM) plugin.
- Only with OADP 1.1.x.

Procedure

- Edit the values in the **spec.features.dataMover.timeout** block of the **DataProtectionApplication** CR manifest, as in the following example:

```

apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_name>
spec:
  features:
    dataMover:
      timeout: 10m
# ...

```

4.11.9.4. CSI snapshot timeout

CSISnapshotTimeout specifies the time during creation to wait until the **CSI VolumeSnapshot** status becomes **ReadyToUse**, before returning error as timeout. The default value is **10m**.

Use the **CSISnapshotTimeout** for the following scenarios:

- With the CSI plugin.

- For very large storage volumes that may take longer than 10 minutes to snapshot. Adjust this timeout if timeouts are found in the logs.



NOTE

Typically, the default value for **CSISnapshotTimeout** does not require adjustment, because the default setting can accommodate large storage volumes.

Procedure

- Edit the values in the **spec.csiSnapshotTimeout** block of the **Backup** CR manifest, as in the following example:

```
apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup_name>
spec:
  csiSnapshotTimeout: 10m
# ...
```

4.11.9.5. Velero default item operation timeout

defaultItemOperationTimeout defines how long to wait on asynchronous **BackupItemActions** and **RestoreItemActions** to complete before timing out. The default value is **1h**.

Use the **defaultItemOperationTimeout** for the following scenarios:

- Only with Data Mover 1.2.x.
- To specify the amount of time a particular backup or restore should wait for the Asynchronous actions to complete. In the context of OADP features, this value is used for the Asynchronous actions involved in the Container Storage Interface (CSI) Data Mover feature.
- When **defaultItemOperationTimeout** is defined in the Data Protection Application (DPA) using the **defaultItemOperationTimeout**, it applies to both backup and restore operations. You can use **itemOperationTimeout** to define only the backup or only the restore of those CRs, as described in the following "Item operation timeout - restore", and "Item operation timeout - backup" sections.

Procedure

- Edit the values in the **spec.configuration.velero.defaultItemOperationTimeout** block of the **DataProtectionApplication** CR manifest, as in the following example:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
metadata:
  name: <dpa_name>
spec:
  configuration:
    velero:
      defaultItemOperationTimeout: 1h
# ...
```

4.11.9.6. Item operation timeout - restore

ItemOperationTimeout specifies the time that is used to wait for **RestoreItemAction** operations. The default value is **1h**.

Use the restore **ItemOperationTimeout** for the following scenarios:

- Only with Data Mover 1.2.x.
- For Data Mover uploads and downloads to or from the **BackupStorageLocation**. If the restore action is not completed when the timeout is reached, it will be marked as failed. If Data Mover operations are failing due to timeout issues, because of large storage volume sizes, then this timeout setting may need to be increased.

Procedure

- Edit the values in the **Restore.spec.itemOperationTimeout** block of the **Restore** CR manifest, as in the following example:

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: <restore_name>
spec:
  itemOperationTimeout: 1h
# ...
```

4.11.9.7. Item operation timeout - backup

ItemOperationTimeout specifies the time used to wait for asynchronous **BackupItemAction** operations. The default value is **1h**.

Use the backup **ItemOperationTimeout** for the following scenarios:

- Only with Data Mover 1.2.x.
- For Data Mover uploads and downloads to or from the **BackupStorageLocation**. If the backup action is not completed when the timeout is reached, it will be marked as failed. If Data Mover operations are failing due to timeout issues, because of large storage volume sizes, then this timeout setting may need to be increased.

Procedure

- Edit the values in the **Backup.spec.itemOperationTimeout** block of the **Backup** CR manifest, as in the following example:

```
apiVersion: velero.io/v1
kind: Backup
metadata:
  name: <backup_name>
spec:
  itemOperationTimeout: 1h
# ...
```

4.11.10. Backup and Restore CR issues

You might encounter these common issues with **Backup** and **Restore** custom resources (CRs).

4.11.10.1. Backup CR cannot retrieve volume

The **Backup** CR displays the error message, **InvalidVolume.NotFound: The volume 'vol-xxxx' does not exist**.

Cause

The persistent volume (PV) and the snapshot locations are in different regions.

Solution

1. Edit the value of the **spec.snapshotLocations.velero.config.region** key in the **DataProtectionApplication** manifest so that the snapshot location is in the same region as the PV.
2. Create a new **Backup** CR.

4.11.10.2. Backup CR status remains in progress

The status of a **Backup** CR remains in the **InProgress** phase and does not complete.

Cause

If a backup is interrupted, it cannot be resumed.

Solution

1. Retrieve the details of the **Backup** CR:

```
$ oc -n {namespace} exec deployment/velero -c velero -- ./velero \
  backup describe <backup>
```

2. Delete the **Backup** CR:

```
$ oc delete backup <backup> -n openshift-adp
```

You do not need to clean up the backup location because a **Backup** CR in progress has not uploaded files to object storage.

3. Create a new **Backup** CR.

4.11.10.3. Backup CR status remains in PartiallyFailed

The status of a **Backup** CR without Restic in use remains in the **PartiallyFailed** phase and does not complete. A snapshot of the affiliated PVC is not created.

Cause

If the backup is created based on the CSI snapshot class, but the label is missing, CSI snapshot plugin fails to create a snapshot. As a result, the **Velero** pod logs an error similar to the following:

+

```
time="2023-02-17T16:33:13Z" level=error msg="Error backing up item" backup=openshift-adp/user1-backup-check5 error="error executing custom action (groupResource=persistentvolumeclaims, namespace=busy1, name=pvc1-user1): rpc error: code = Unknown desc = failed to get volumesnapshotclass for storageclass ocs-storagecluster-ceph-rbd: failed to get volumesnapshotclass for provisioner openshift-storage.rbd.csi.ceph.com, ensure that the desired volumesnapshot class has the velero.io/csi-volumesnapshot-class label" logSource="/remote-source/velero/app/pkg/backup/backup.go:417" name=busybox-79799557b5-vprq
```

Solution

1. Delete the **Backup** CR:

```
$ oc delete backup <backup> -n openshift-adp
```

2. If required, clean up the stored data on the **BackupStorageLocation** to free up space.
3. Apply label **velero.io/csi-volumesnapshot-class=true** to the **VolumeSnapshotClass** object:

```
$ oc label volumesnapshotclass/<snapclass_name> velero.io/csi-volumesnapshot-class=true
```

4. Create a new **Backup** CR.

4.11.11. Restic issues

You might encounter these issues when you back up applications with Restic.

4.11.11.1. Restic permission error for NFS data volumes with `root_squash` enabled

The **Restic** pod log displays the error message: **controller=pod-volume-backup error="fork/exec/usr/bin/restic: permission denied"**.

Cause

If your NFS data volumes have **root_squash** enabled, **Restic** maps to **nfsnobody** and does not have permission to create backups.

Solution

You can resolve this issue by creating a supplemental group for **Restic** and adding the group ID to the **DataProtectionApplication** manifest:

1. Create a supplemental group for **Restic** on the NFS data volume.
2. Set the **setgid** bit on the NFS directories so that group ownership is inherited.
3. Add the **spec.configuration.restic.supplementalGroups** parameter and the group ID to the **DataProtectionApplication** manifest, as in the following example:

```
spec:
  configuration:
    restic:
      enable: true
      supplementalGroups:
        - <group_id> 1
```

- 1 Specify the supplemental group ID.

4. Wait for the **Restic** pods to restart so that the changes are applied.

4.11.11.2. Restic Backup CR cannot be recreated after bucket is emptied

If you create a Restic **Backup** CR for a namespace, empty the object storage bucket, and then recreate the **Backup** CR for the same namespace, the recreated **Backup** CR fails.

The **velero** pod log displays the following error message: **stderr=Fatal: unable to open config file: Stat: The specified key does not exist.\nls there a repository at the following location?.**

Cause

Velero does not recreate or update the Restic repository from the **ResticRepository** manifest if the Restic directories are deleted from object storage. See [Velero issue 4421](#) for more information.

Solution

- Remove the related Restic repository from the namespace by running the following command:

```
$ oc delete resticrepository openshift-adp <name_of_the_restic_repository>
```

In the following error log, **mysql-persistent** is the problematic Restic repository. The name of the repository appears in italics for clarity.

```
time="2021-12-29T18:29:14Z" level=info msg="1 errors encountered backup up item" backup=velero/backup65 logSource="pkg/backup/backup.go:431" name=mysql-7d99fc949-qbkds time="2021-12-29T18:29:14Z" level=error msg="Error backing up item" backup=velero/backup65 error="pod volume backup failed: error running restic backup, stderr=Fatal: unable to open config file: Stat: The specified key does not exist.\nls there a repository at the following location?\nns3:http://minio-minio.apps.mayap-oadp-veleo-1234.qe.devcluster.openshift.com/mayapvelerooadp2/velero1/restic/mysql-persistent\n: exit status 1" error.file="/remote-source/src/github.com/vmware-tanzu/velero/pkg/restic/backupper.go:184" error.function="github.com/vmware-tanzu/velero/pkg/restic.(*backupper).BackupPodVolumes" logSource="pkg/backup/backup.go:435" name=mysql-7d99fc949-qbkds
```

4.11.12. Using the must-gather tool

You can collect logs, metrics, and information about OADP custom resources by using the **must-gather** tool.

The **must-gather** data must be attached to all customer cases.

Prerequisites

- You must be logged in to the OpenShift Container Platform cluster as a user with the **cluster-admin** role.
- You must have the OpenShift CLI (**oc**) installed.

Procedure

1. Navigate to the directory where you want to store the **must-gather** data.
2. Run the **oc adm must-gather** command for one of the following data collection options:

Additional resources

- [Gathering cluster data](#)

4.11.12.1. Using must-gather with insecure TLS connections

If a custom CA certificate is used, the **must-gather** pod fails to grab the output for **velero logs/describe**. To use the **must-gather** tool with insecure TLS connections, you can pass the **gather_without_tls** flag to the **must-gather** command.

Procedure

- Pass the **gather_without_tls** flag, with value set to **true**, to the **must-gather** tool by using the following command:

```
$ oc adm must-gather --image=registry.redhat.io/oadp/oadp-mustgather-rhel9:v1.3 --
/usr/bin/gather_without_tls <true/false>
```

By default, the flag value is set to **false**. Set the value to **true** to allow insecure TLS connections.

4.11.12.2. Combining options when using the must-gather tool

Currently, it is not possible to combine **must-gather** scripts, for example specifying a timeout threshold while permitting insecure TLS connections. In some situations, you can get around this limitation by setting up internal variables on the **must-gather** command line, such as the following example:

```
$ oc adm must-gather --image=brew.registry.redhat.io/rh-osbs/oadp-oadp-mustgather-rhel8:1.1.1-8 -
-skip_tls=true /usr/bin/gather_with_timeout <timeout_value_in_seconds>
```

In this example, set the **skip_tls** variable before running the **gather_with_timeout** script. The result is a combination of **gather_with_timeout** and **gather_without_tls**.

The only other variables that you can specify this way are the following:

- **logs_since**, with a default value of **72h**
- **request_timeout**, with a default value of **0s**

4.11.13. OADP Monitoring

The OpenShift Container Platform provides a monitoring stack that allows users and administrators to effectively monitor and manage their clusters, as well as monitor and analyze the workload performance of user applications and services running on the clusters, including receiving alerts if an event occurs.

Additional resources

- [Monitoring stack](#)

4.11.13.1. OADP monitoring setup

The OADP Operator leverages an OpenShift User Workload Monitoring provided by the OpenShift Monitoring Stack for retrieving metrics from the Velero service endpoint. The monitoring stack allows creating user-defined Alerting Rules or querying metrics by using the OpenShift Metrics query front end.

With enabled User Workload Monitoring, it is possible to configure and use any Prometheus-compatible third-party UI, such as Grafana, to visualize Velero metrics.

Monitoring metrics requires enabling monitoring for the user-defined projects and creating a **ServiceMonitor** resource to scrape those metrics from the already enabled OADP service endpoint that resides in the **openshift-adp** namespace.

Prerequisites

- You have access to an OpenShift Container Platform cluster using an account with **cluster-admin** permissions.
- You have created a cluster monitoring config map.

Procedure

1. Edit the **cluster-monitoring-config ConfigMap** object in the **openshift-monitoring** namespace:

```
$ oc edit configmap cluster-monitoring-config -n openshift-monitoring
```

2. Add or enable the **enableUserWorkload** option in the **data** section's **config.yaml** field:

```
apiVersion: v1
data:
  config.yaml: |
    enableUserWorkload: true 1
kind: ConfigMap
metadata:
# ...
```

- 1 Add this option or set to **true**

3. Wait a short period of time to verify the User Workload Monitoring Setup by checking if the following components are up and running in the **openshift-user-workload-monitoring** namespace:

```
$ oc get pods -n openshift-user-workload-monitoring
```

Example output

```
NAME                                READY STATUS RESTARTS AGE
prometheus-operator-6844b4b99c-b57j9 2/2   Running 0      43s
prometheus-user-workload-0           5/5   Running 0      32s
prometheus-user-workload-1           5/5   Running 0      32s
thanos-ruler-user-workload-0         3/3   Running 0      32s
thanos-ruler-user-workload-1         3/3   Running 0      32s
```

4. Verify the existence of the **user-workload-monitoring-config** ConfigMap in the **openshift-user-workload-monitoring**. If it exists, skip the remaining steps in this procedure.

```
$ oc get configmap user-workload-monitoring-config -n openshift-user-workload-monitoring
```

Example output

```
Error from server (NotFound): configmaps "user-workload-monitoring-config" not found
```

5. Create a **user-workload-monitoring-config ConfigMap** object for the User Workload Monitoring, and save it under the **2_configure_user_workload_monitoring.yaml** file name:

Example output

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: user-workload-monitoring-config
  namespace: openshift-user-workload-monitoring
data:
  config.yaml: |
```

6. Apply the **2_configure_user_workload_monitoring.yaml** file:

```
$ oc apply -f 2_configure_user_workload_monitoring.yaml
configmap/user-workload-monitoring-config created
```

4.11.13.2. Creating OADP service monitor

OADP provides an **openshift-adp-velero-metrics-svc** service which is created when the DPA is configured. The service monitor used by the user workload monitoring must point to the defined service.

Get details about the service by running the following commands:

Procedure

1. Ensure the **openshift-adp-velero-metrics-svc** service exists. It should contain **app.kubernetes.io/name=velero** label, which will be used as selector for the **ServiceMonitor** object.

```
$ oc get svc -n openshift-adp -l app.kubernetes.io/name=velero
```

Example output

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
openshift-adp-velero-metrics-svc	ClusterIP	172.30.38.244	<none>	8085/TCP	1h

2. Create a **ServiceMonitor** YAML file that matches the existing service label, and save the file as **3_create_oadp_service_monitor.yaml**. The service monitor is created in the **openshift-adp** namespace where the **openshift-adp-velero-metrics-svc** service resides.

Example ServiceMonitor object

```

apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
  labels:
    app: oadp-service-monitor
    name: oadp-service-monitor
    namespace: openshift-adp
spec:
  endpoints:
    - interval: 30s
      path: /metrics
      targetPort: 8085
      scheme: http
  selector:
    matchLabels:
      app.kubernetes.io/name: "velero"

```

3. Apply the **3_create_oadp_service_monitor.yaml** file:

```
$ oc apply -f 3_create_oadp_service_monitor.yaml
```

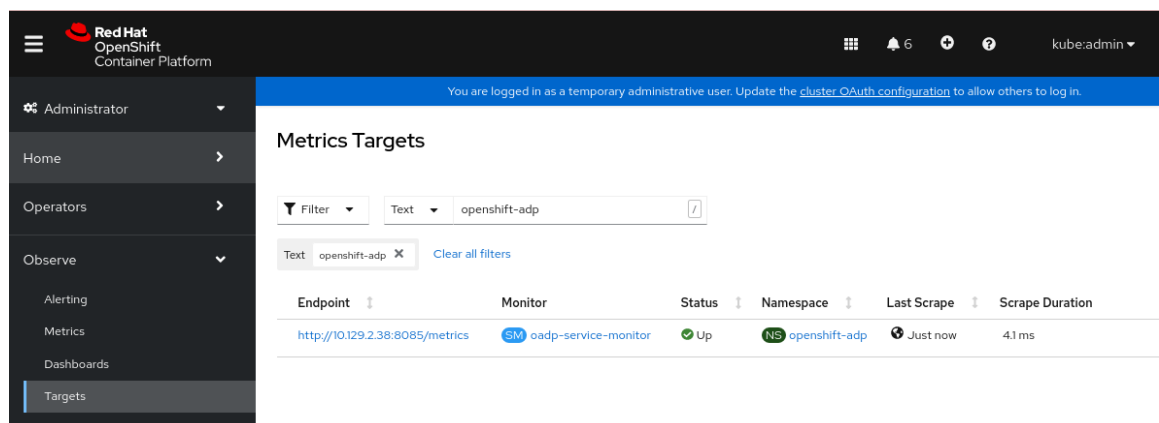
Example output

```
servicemonitor.monitoring.coreos.com/oadp-service-monitor created
```

Verification

- Confirm that the new service monitor is in an **Up** state by using the **Administrator** perspective of the OpenShift Container Platform web console:
 - a. Navigate to the **Observe** → **Targets** page.
 - b. Ensure the **Filter** is unselected or that the **User** source is selected and type **openshift-adp** in the **Text** search field.
 - c. Verify that the status for the **Status** for the service monitor is **Up**.

Figure 4.1. OADP metrics targets



4.11.13.3. Creating an alerting rule

The OpenShift Container Platform monitoring stack allows to receive Alerts configured using Alerting Rules. To create an Alerting rule for the OADP project, use one of the Metrics which are scraped with the user workload monitoring.

Procedure

1. Create a **PrometheusRule** YAML file with the sample **OADPBackupFailing** alert and save it as **4_create_oadp_alert_rule.yaml**.

Sample OADPBackupFailing alert

```
apiVersion: monitoring.coreos.com/v1
kind: PrometheusRule
metadata:
  name: sample-oadp-alert
  namespace: openshift-adp
spec:
  groups:
  - name: sample-oadp-backup-alert
    rules:
    - alert: OADPBackupFailing
      annotations:
        description: 'OADP had {{$value | humanize}} backup failures over the last 2 hours.'
        summary: OADP has issues creating backups
      expr: |
        increase(velero_backup_failure_total{job="openshift-adp-velero-metrics-svc"}[2h]) > 0
      for: 5m
      labels:
        severity: warning
```

In this sample, the Alert displays under the following conditions:

- There is an increase of new failing backups during the 2 last hours that is greater than 0 and the state persists for at least 5 minutes.
 - If the time of the first increase is less than 5 minutes, the Alert will be in a **Pending** state, after which it will turn into a **Firing** state.
2. Apply the **4_create_oadp_alert_rule.yaml** file, which creates the **PrometheusRule** object in the **openshift-adp** namespace:

```
$ oc apply -f 4_create_oadp_alert_rule.yaml
```

Example output

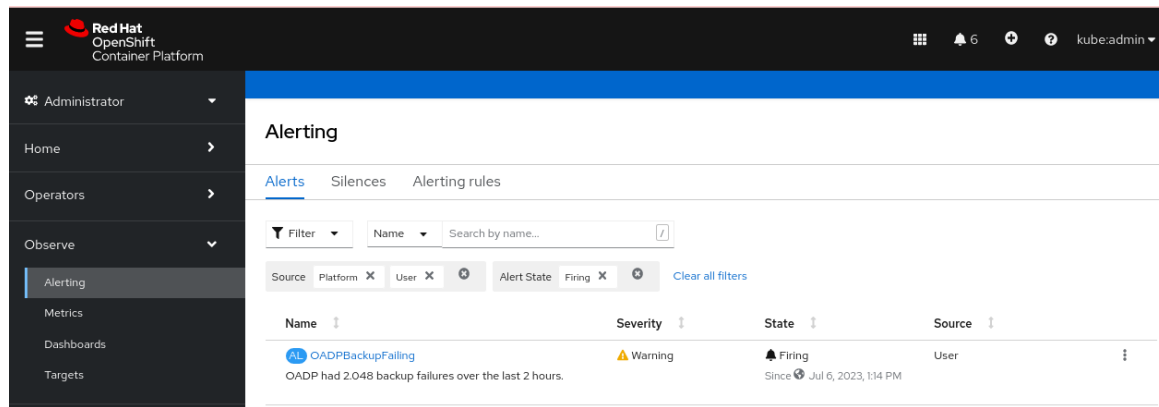
```
prometheusrule.monitoring.coreos.com/sample-oadp-alert created
```

Verification

- After the Alert is triggered, you can view it in the following ways:
 - In the **Developer** perspective, select the **Observe** menu.

- In the **Administrator** perspective under the **Observe** → **Alerting** menu, select **User** in the **Filter** box. Otherwise, by default only the **Platform Alerts** are displayed.

Figure 4.2. OADP backup failing alert



Additional resources

- [Managing alerts](#)

4.11.13.4. List of available metrics

These are the list of metrics provided by the OADP together with their [Types](#).

Metric name	Description	Type
kopia_content_cache_hit_bytes	Number of bytes retrieved from the cache	Counter
kopia_content_cache_hit_count	Number of times content was retrieved from the cache	Counter
kopia_content_cache_malformed	Number of times malformed content was read from the cache	Counter
kopia_content_cache_miss_count	Number of times content was not found in the cache and fetched	Counter
kopia_content_cache_missed_bytes	Number of bytes retrieved from the underlying storage	Counter
kopia_content_cache_miss_error_count	Number of times content could not be found in the underlying storage	Counter
kopia_content_cache_store_error_count	Number of times content could not be saved in the cache	Counter
kopia_content_get_bytes	Number of bytes retrieved using GetContent()	Counter

Metric name	Description	Type
kopia_content_get_count	Number of times GetContent() was called	Counter
kopia_content_get_error_count	Number of times GetContent() was called and the result was an error	Counter
kopia_content_get_not_found_count	Number of times GetContent() was called and the result was not found	Counter
kopia_content_write_bytes	Number of bytes passed to WriteContent()	Counter
kopia_content_write_count	Number of times WriteContent() was called	Counter
velero_backup_attempt_total	Total number of attempted backups	Counter
velero_backup_deletion_attempt_total	Total number of attempted backup deletions	Counter
velero_backup_deletion_failure_total	Total number of failed backup deletions	Counter
velero_backup_deletion_success_total	Total number of successful backup deletions	Counter
velero_backup_duration_seconds	Time taken to complete backup, in seconds	Histogram
velero_backup_failure_total	Total number of failed backups	Counter
velero_backup_items_errors	Total number of errors encountered during backup	Gauge
velero_backup_items_total	Total number of items backed up	Gauge
velero_backup_last_status	Last status of the backup. A value of 1 is success, 0.	Gauge
velero_backup_last_successful_timestamp	Last time a backup ran successfully, Unix timestamp in seconds	Gauge

Metric name	Description	Type
velero_backup_partial_failure_total	Total number of partially failed backups	Counter
velero_backup_success_total	Total number of successful backups	Counter
velero_backup_tarball_size_bytes	Size, in bytes, of a backup	Gauge
velero_backup_total	Current number of existent backups	Gauge
velero_backup_validation_failure_total	Total number of validation failed backups	Counter
velero_backup_warning_total	Total number of warned backups	Counter
velero_csi_snapshot_attempt_total	Total number of CSI attempted volume snapshots	Counter
velero_csi_snapshot_failure_total	Total number of CSI failed volume snapshots	Counter
velero_csi_snapshot_success_total	Total number of CSI successful volume snapshots	Counter
velero_restore_attempt_total	Total number of attempted restores	Counter
velero_restore_failed_total	Total number of failed restores	Counter
velero_restore_partial_failure_total	Total number of partially failed restores	Counter
velero_restore_success_total	Total number of successful restores	Counter
velero_restore_total	Current number of existent restores	Gauge
velero_restore_validation_failed_total	Total number of failed restores failing validations	Counter
velero_volume_snapshot_attempt_total	Total number of attempted volume snapshots	Counter

Metric name	Description	Type
velero_volume_snapshot_failure_total	Total number of failed volume snapshots	Counter
velero_volume_snapshot_success_total	Total number of successful volume snapshots	Counter

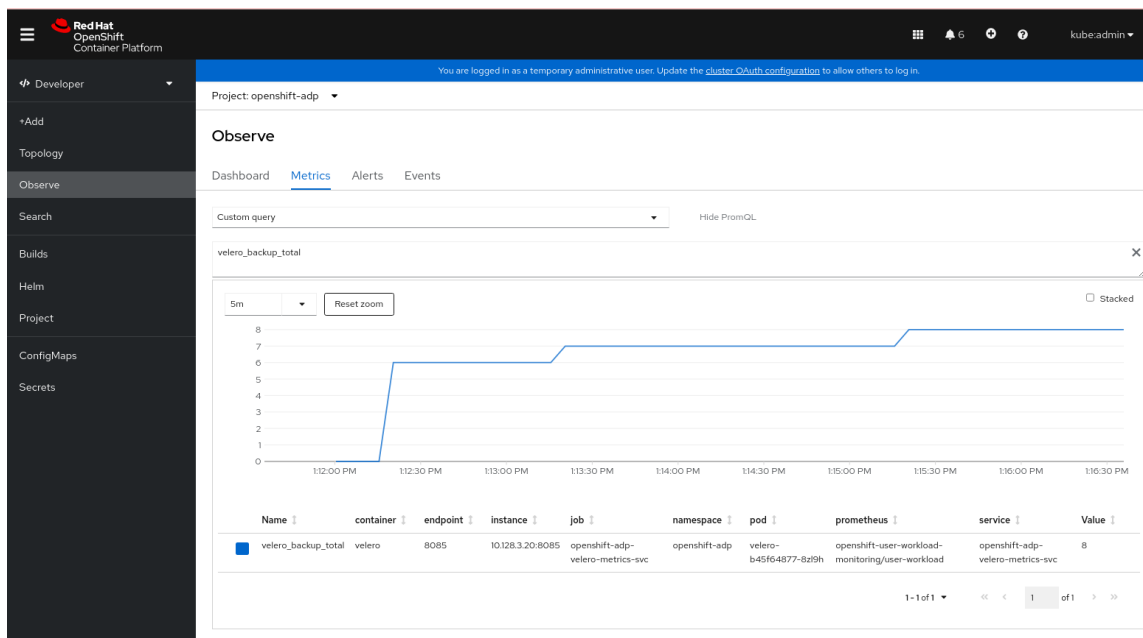
4.11.13.5. Viewing metrics using the Observe UI

You can view metrics in the OpenShift Container Platform web console from the **Administrator** or **Developer** perspective, which must have access to the **openshift-adp** project.

Procedure

- Navigate to the **Observe** → **Metrics** page:
 - If you are using the **Developer** perspective, follow these steps:
 - a. Select **Custom query**, or click on the **Show PromQL** link.
 - b. Type the query and click **Enter**.
 - If you are using the **Administrator** perspective, type the expression in the text field and select **Run Queries**.

Figure 4.3. OADP metrics query



4.12. APIS USED WITH OADP

The document provides information about the following APIs that you can use with OADP:

- Velero API
- OADP API

4.12.1. Velero API

Velero API documentation is maintained by Velero, not by Red Hat. It can be found at [Velero API types](#).

4.12.2. OADP API

The following tables provide the structure of the OADP API:

Table 4.2. DataProtectionApplicationSpec

Property	Type	Description
backupLocations	[] BackupLocation	Defines the list of configurations to use for BackupStorageLocations .
snapshotLocations	[] SnapshotLocation	Defines the list of configurations to use for VolumeSnapshotLocations .
unsupportedOverrides	map [UnsupportedImageKey] string	Can be used to override the deployed dependent images for development. Options are veleroImageFqin , awsPluginImageFqin , openshiftPluginImageFqin , azurePluginImageFqin , gcpPluginImageFqin , csiPluginImageFqin , dataMoverImageFqin , resticRestoreImageFqin , kubvirtPluginImageFqin , and operator-type .
podAnnotations	map [string] string	Used to add annotations to pods deployed by Operators.
podDnsPolicy	DNSPolicy	Defines the configuration of the DNS of a pod.
podDnsConfig	PodDNSConfig	Defines the DNS parameters of a pod in addition to those generated from DNSPolicy .
backupImages	*bool	Used to specify whether or not you want to deploy a registry for enabling backup and restore of images.

Property	Type	Description
configuration	* ApplicationConfig	Used to define the data protection application's server configuration.
features	* Features	Defines the configuration for the DPA to enable the Technology Preview features.

Complete schema definitions for the OADP API .

Table 4.3. BackupLocation

Property	Type	Description
velero	* velero.BackupStorageLocationSpec	Location to store volume snapshots, as described in Backup Storage Location .
bucket	* CloudStorageLocation	[Technology Preview] Automates creation of a bucket at some cloud storage providers for use as a backup storage location.



IMPORTANT

The **bucket** parameter is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#) .

Complete schema definitions for the type [BackupLocation](#).

Table 4.4. SnapshotLocation

Property	Type	Description
velero	* VolumeSnapshotLocationSpec	Location to store volume snapshots, as described in Volume Snapshot Location .

Complete schema definitions for the type [SnapshotLocation](#).

Table 4.5. ApplicationConfig

Property	Type	Description
velero	* VeleroConfig	Defines the configuration for the Velero server.
restic	* ResticConfig	Defines the configuration for the Restic server.

Complete schema definitions for the type [ApplicationConfig](#).

Table 4.6. VeleroConfig

Property	Type	Description
featureFlags	[] string	Defines the list of features to enable for the Velero instance.
defaultPlugins	[] string	The following types of default Velero plugins can be installed: aws , azure , csi , gcp , kubevirt , and openshift .
customPlugins	[] CustomPlugin	Used for installation of custom Velero plugins. Default and custom plugins are described in OADP plugins
restoreResourcesVersionPriority	string	Represents a config map that is created if defined for use in conjunction with the EnableAPIGroupVersions feature flag. Defining this field automatically adds EnableAPIGroupVersions to the Velero server feature flag.
noDefaultBackupLocation	bool	To install Velero without a default backup storage location, you must set the noDefaultBackupLocation flag in order to confirm installation.
podConfig	* PodConfig	Defines the configuration of the Velero pod.

Property	Type	Description
logLevel	string	Velero server's log level (use debug for the most granular logging, leave unset for Velero default). Valid options are trace , debug , info , warning , error , fatal , and panic .

Complete schema definitions for the type [VeleroConfig](#).

Table 4.7. CustomPlugin

Property	Type	Description
name	string	Name of custom plugin.
image	string	Image of custom plugin.

Complete schema definitions for the type [CustomPlugin](#).

Table 4.8. ResticConfig

Property	Type	Description
enable	*bool	If set to true , enables backup and restore using Restic. If set to false , snapshots are needed.
supplementalGroups	[]int64	Defines the Linux groups to be applied to the Restic pod.
timeout	string	A user-supplied duration string that defines the Restic timeout. Default value is 1hr (1 hour). A duration string is a possibly signed sequence of decimal numbers, each with optional fraction and a unit suffix, such as 300ms , -1.5h` or 2h45m . Valid time units are ns , us (or µs), ms , s , m , and h .
podConfig	*PodConfig	Defines the configuration of the Restic pod.

Complete schema definitions for the type [ResticConfig](#).

Table 4.9. PodConfig

Property	Type	Description
nodeSelector	map [string] string	Defines the nodeSelector to be supplied to a Velero podSpec or a Restic podSpec .
tolerations	[Toleration]	Defines the list of tolerations to be applied to a Velero deployment or a Restic daemonset .
resourceAllocations	ResourceRequirements	Set specific resource limits and requests for a Velero pod or a Restic pod as described in Setting Velero CPU and memory resource allocations .
labels	map [string] string	Labels to add to pods.

Complete schema definitions for the type [PodConfig](#).

Table 4.10. Features

Property	Type	Description
dataMover	* DataMover	Defines the configuration of the Data Mover.

Complete schema definitions for the type [Features](#).

Table 4.11. DataMover

Property	Type	Description
enable	bool	If set to true , deploys the volume snapshot mover controller and a modified CSI Data Mover plugin. If set to false , these are not deployed.
credentialName	string	User-supplied Restic Secret name for Data Mover.

Property	Type	Description
timeout	string	A user-supplied duration string for VolumeSnapshotBackup and VolumeSnapshotRestore to complete. Default is 10m (10 minutes). A duration string is a possibly signed sequence of decimal numbers, each with optional fraction and a unit suffix, such as 300ms , -1.5h or 2h45m . Valid time units are ns , us (or µs), ms , s , m , and h .

The OADP API is more fully detailed in [OADP Operator](#).

4.13. ADVANCED OADP FEATURES AND FUNCTIONALITIES

This document provides information about advanced features and functionalities of OpenShift API for Data Protection (OADP).

4.13.1. Working with different Kubernetes API versions on the same cluster

4.13.1.1. Listing the Kubernetes API group versions on a cluster

A source cluster might offer multiple versions of an API, where one of these versions is the preferred API version. For example, a source cluster with an API named **Example** might be available in the **example.com/v1** and **example.com/v1beta2** API groups.

If you use Velero to back up and restore such a source cluster, Velero backs up only the version of that resource that uses the preferred version of its Kubernetes API.

To return to the above example, if **example.com/v1** is the preferred API, then Velero only backs up the version of a resource that uses **example.com/v1**. Moreover, the target cluster needs to have **example.com/v1** registered in its set of available API resources in order for Velero to restore the resource on the target cluster.

Therefore, you need to generate a list of the Kubernetes API group versions on your target cluster to be sure the preferred API version is registered in its set of available API resources.

Procedure

- Enter the following command:

```
$ oc api-resources
```

4.13.1.2. About Enable API Group Versions

By default, Velero only backs up resources that use the preferred version of the Kubernetes API. However, Velero also includes a feature, [Enable API Group Versions](#), that overcomes this limitation. When enabled on the source cluster, this feature causes Velero to back up *all* Kubernetes API group

versions that are supported on the cluster, not only the preferred one. After the versions are stored in the backup .tar file, they are available to be restored on the destination cluster.

For example, a source cluster with an API named **Example** might be available in the **example.com/v1** and **example.com/v1beta2** API groups, with **example.com/v1** being the preferred API.

Without the Enable API Group Versions feature enabled, Velero backs up only the preferred API group version for **Example**, which is **example.com/v1**. With the feature enabled, Velero also backs up **example.com/v1beta2**.

When the Enable API Group Versions feature is enabled on the destination cluster, Velero selects the version to restore on the basis of the order of priority of API group versions.



NOTE

Enable API Group Versions is still in beta.

Velero uses the following algorithm to assign priorities to API versions, with **1** as the top priority:

1. Preferred version of the *destination* cluster
2. Preferred version of the *source_* cluster
3. Common non-preferred supported version with the highest Kubernetes version priority

Additional resources

- [Enable API Group Versions Feature](#)

4.13.1.3. Using Enable API Group Versions

You can use Velero's Enable API Group Versions feature to back up *all* Kubernetes API group versions that are supported on a cluster, not only the preferred one.



NOTE

Enable API Group Versions is still in beta.

Procedure

- Configure the **EnableAPIGroupVersions** feature flag:

```
apiVersion: oadp.openshift.io/v1alpha1
kind: DataProtectionApplication
...
spec:
  configuration:
    velero:
      featureFlags:
        - EnableAPIGroupVersions
```

Additional resources

- [Enable API Group Versions Feature](#)

4.13.2. Backing up data from one cluster and restoring it to another cluster

4.13.2.1. About backing up data from one cluster and restoring it on another cluster

OpenShift API for Data Protection (OADP) is designed to back up and restore application data in the same OpenShift Container Platform cluster. Migration Toolkit for Containers (MTC) is designed to migrate containers, including application data, from one OpenShift Container Platform cluster to another cluster.

You can use OADP to back up application data from one OpenShift Container Platform cluster and restore it on another cluster. However, doing so is more complicated than using MTC or using OADP to back up and restore on the same cluster.

To successfully use OADP to back up data from one cluster and restore it to another cluster, you must take into account the following factors, in addition to the prerequisites and procedures that apply to using OADP to back up and restore data on the same cluster:

- Operators
- Use of Velero
- UID and GID ranges

4.13.2.1.1. Operators

You must exclude Operators from the backup of an application for backup and restore to succeed.

4.13.2.1.2. Use of Velero

Velero, which OADP is built upon, does not natively support migrating persistent volume snapshots across cloud providers. To migrate volume snapshot data between cloud platforms, you must *either* enable the Velero Restic file system backup option, which backs up volume contents at the file system level, *or* use the OADP Data Mover for CSI snapshots.



NOTE

In OADP 1.1 and earlier, the Velero Restic file system backup option is called **restic**. In OADP 1.2 and later, the Velero Restic file system backup option is called **file-system-backup**.

- You must also use Velero's [File System Backup](#) to migrate data between AWS regions or between Microsoft Azure regions.
- Velero does not support restoring data to a cluster with an *earlier* Kubernetes version than the source cluster.
- It is theoretically possible to migrate workloads to a destination with a *later* Kubernetes version than the source, but you must consider the compatibility of API groups between clusters for each custom resource. If a Kubernetes version upgrade breaks the compatibility of core or native API groups, you must first update the impacted custom resources.

4.13.2.2. About determining which pod volumes to back up

Before you start a backup operation by using File System Backup (FSB), you must specify which pods contain a volume that you want to back up. Velero refers to this process as "discovering" the appropriate pod volumes.

Velero supports two approaches for determining pod volumes:

- **Opt-in approach:** The opt-in approach requires that you actively indicate that you want to include - *opt-in* - a volume in a backup. You do this by labelling each pod that contains a volume to be backed up with the name of the volume. When Velero finds a persistent volume (PV), it checks the pod that mounted the volume. If the pod is labelled with the name of the volume, Velero backs up the pod.
- **Opt-out approach:** With the opt-out approach, you must actively specify that you want to exclude a volume from a backup. You do this by labelling each pod that contains a volume you do not want to back up with the name of the volume. When Velero finds a PV, it checks the pod that mounted the volume. If the pod is labelled with the volume's name, Velero does not back up the pod.

4.13.2.2.1. Limitations

- FSB does not support backing up and restoring **hostpath** volumes. However, FSB does support backing up and restoring local volumes.
- Velero uses a static, common encryption key for all backup repositories it creates. **This static key means that anyone who can access your backup storage can also decrypt your backup data.** It is essential that you limit access to backup storage.
- For PVCs, every incremental backup chain is maintained across pod reschedules. For pod volumes that are *not* PVCs, such as **emptyDir** volumes, if a pod is deleted or recreated, for example, by a **ReplicaSet** or a deployment, the next backup of those volumes will be a full backup and not an incremental backup. It is assumed that the lifecycle of a pod volume is defined by its pod.
- Even though backup data can be kept incrementally, backing up large files, such as a database, can take a long time. This is because FSB uses deduplication to find the difference that needs to be backed up.
- FSB reads and writes data from volumes by accessing the file system of the node on which the pod is running. For this reason, FSB can only back up volumes that are mounted from a pod and not directly from a PVC. Some Velero users have overcome this limitation by running a staging pod, such as a BusyBox or Alpine container with an infinite sleep, to mount these PVC and PV pairs before performing a Velero backup..
- FSB expects volumes to be mounted under **<hostPath>/<pod UID>**, with **<hostPath>** being configurable. Some Kubernetes systems, for example, vCluster, do not mount volumes under the **<pod UID>** subdirectory, and VFSB does not work with them as expected.

4.13.2.2.2. Backing up pod volumes by using the opt-in method

You can use the opt-in method to specify which volumes need to be backed up by File System Backup (FSB). You can do this by using the **backup.velero.io/backup-volumes** command.

Procedure

- On each pod that contains one or more volumes that you want to back up, enter the following command:

```
$ oc -n <your_pod_namespace> annotate pod/<your_pod_name> \
  backup.velero.io/backup-volumes=<your_volume_name_1>, \ <your_volume_name_2>>,..., \
  <your_volume_name_n>
```

where:

<your_volume_name_x>

specifies the name of the xth volume in the pod specification.

4.13.2.2.3. Backing up pod volumes by using the opt-out method

When using the opt-out approach, all pod volumes are backed up by using File System Backup (FSB), although there are some exceptions:

- Volumes that mount the default service account token, secrets, and configuration maps.
- **hostPath** volumes

You can use the opt-out method to specify which volumes **not** to back up. You can do this by using the **backup.velero.io/backup-volumes-excludes** command.

Procedure

- On each pod that contains one or more volumes that you do not want to back up, run the following command:

```
$ oc -n <your_pod_namespace> annotate pod/<your_pod_name> \
  backup.velero.io/backup-volumes-excludes=<your_volume_name_1>, \
  <your_volume_name_2>>,...,<your_volume_name_n>
```

where:

<your_volume_name_x>

specifies the name of the xth volume in the pod specification.



NOTE

You can enable this behavior for all Velero backups by running the **velero install** command with the **--default-volumes-to-fs-backup** flag.

4.13.2.3. UID and GID ranges

If you back up data from one cluster and restore it to another cluster, problems might occur with UID (User ID) and GID (Group ID) ranges. The following section explains these potential issues and mitigations:

Summary of the issues

The namespace UID and GID ranges might change depending on the destination cluster. OADP does not back up and restore OpenShift UID range metadata. If the backed up application requires a specific UID, ensure the range is available upon restore. For more information about OpenShift's UID and GID ranges, see [A Guide to OpenShift and UIDs](#).

Detailed description of the issues

When you create a namespace in OpenShift Container Platform by using the shell command **oc**

create namespace, OpenShift Container Platform assigns the namespace a unique User ID (UID) range from its available pool of UIDs, a Supplemental Group (GID) range, and unique SELinux MCS labels. This information is stored in the **metadata.annotations** field of the cluster. This information is part of the Security Context Constraints (SCC) annotations, which comprise of the following components:

- **openshift.io/sa.scc.mcs**
- **openshift.io/sa.scc.supplemental-groups**
- **openshift.io/sa.scc.uid-range**

When you use OADP to restore the namespace, it automatically uses the information in **metadata.annotations** without resetting it for the destination cluster. As a result, the workload might not have access to the backed up data if any of the following is true:

- There is an existing namespace with other SCC annotations, for example, on another cluster. In this case, OADP uses the existing namespace during the backup instead of the namespace you want to restore.
- A label selector was used during the backup, but the namespace in which the workloads are executed does not have the label. In this case, OADP does not back up the namespace, but creates a new namespace during the restore that does not contain the annotations of the backed up namespace. This results in a new UID range being assigned to the namespace. This can be an issue for customer workloads if OpenShift Container Platform assigns a pod a **securityContext** UID to a pod based on namespace annotations that have changed since the persistent volume data was backed up.
- The UID of the container no longer matches the UID of the file owner.
- An error occurs because OpenShift Container Platform has not changed the UID range of the destination cluster to match the backup cluster data. As a result, the backup cluster has a different UID than the destination cluster, which means that the application cannot read or write data on the destination cluster.

Mitigations

You can use one or more of the following mitigations to resolve the UID and GID range issues:

- Simple mitigations:
 - If you use a label selector in the **Backup** CR to filter the objects to include in the backup, be sure to add this label selector to the namespace that contains the workspace.
 - Remove any pre-existing version of a namespace on the destination cluster before attempting to restore a namespace with the same name.
- Advanced mitigations:
 - Fix UID ranges after migration by [Resolving overlapping UID ranges in OpenShift namespaces after migration](#). Step 1 is optional.

For an in-depth discussion of UID and GID ranges in OpenShift Container Platform with an emphasis on overcoming issues in backing up data on one cluster and restoring it on another, see [A Guide to OpenShift and UIDs](#).

4.13.2.4. Backing up data from one cluster and restoring it to another cluster

In general, you back up data from one OpenShift Container Platform cluster and restore it on another OpenShift Container Platform cluster in the same way that you back up and restore data to the same cluster. However, there are some additional prerequisites and differences in the procedure when backing up data from one OpenShift Container Platform cluster and restoring it on another.

Prerequisites

- All relevant prerequisites for backing up and restoring on your platform (for example, AWS, Microsoft Azure, GCP, and so on), especially the prerequisites for the Data Protection Application (DPA), are described in the relevant sections of this guide.

Procedure

- Make the following additions to the procedures given for your platform:
 - Ensure that the backup store location (BSL) and volume snapshot location have the same names and paths to restore resources to another cluster.
 - Share the same object storage location credentials across the clusters.
 - For best results, use OADP to create the namespace on the destination cluster.
 - If you use the Velero **file-system-backup** option, enable the **--default-volumes-to-fs-backup** flag for use during backup by running the following command:

```
$ velero backup create <backup_name> --default-volumes-to-fs-backup
<any_other_options>
```



NOTE

In OADP 1.2 and later, the Velero Restic option is called **file-system-backup**.

4.13.3. Additional resources

For more information about API group versions, see [Working with different Kubernetes API versions on the same cluster](#).

For more information about OADP Data Mover, see [Using Data Mover for CSI snapshots](#).

For more information about using Restic with OADP, see [Backing up applications with Restic](#).

CHAPTER 5. CONTROL PLANE BACKUP AND RESTORE

5.1. BACKING UP ETCD

etcd is the key-value store for OpenShift Container Platform, which persists the state of all resource objects.

Back up your cluster's etcd data regularly and store in a secure location ideally outside the OpenShift Container Platform environment. Do not take an etcd backup before the first certificate rotation completes, which occurs 24 hours after installation, otherwise the backup will contain expired certificates. It is also recommended to take etcd backups during non-peak usage hours because the etcd snapshot has a high I/O cost.

Be sure to take an etcd backup after you upgrade your cluster. This is important because when you restore your cluster, you must use an etcd backup that was taken from the same z-stream release. For example, an OpenShift Container Platform 4.y.z cluster must use an etcd backup that was taken from 4.y.z.



IMPORTANT

Back up your cluster's etcd data by performing a single invocation of the backup script on a control plane host. Do not take a backup for each control plane host.

After you have an etcd backup, you can [restore to a previous cluster state](#).

5.1.1. Backing up etcd data

Follow these steps to back up etcd data by creating an etcd snapshot and backing up the resources for the static pods. This backup can be saved and used at a later time if you need to restore etcd.



IMPORTANT

Only save a backup from a single control plane host. Do not take a backup from each control plane host in the cluster.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have checked whether the cluster-wide proxy is enabled.

TIP

You can check whether the proxy is enabled by reviewing the output of **oc get proxy cluster -o yaml**. The proxy is enabled if the **httpProxy**, **httpsProxy**, and **noProxy** fields have values set.

Procedure

1. Start a debug session as root for a control plane node:

```
$ oc debug --as-root node/<node_name>
```

2. Change your root directory to **/host** in the debug shell:

```
sh-4.4# chroot /host
```

3. If the cluster-wide proxy is enabled, be sure that you have exported the **NO_PROXY**, **HTTP_PROXY**, and **HTTPS_PROXY** environment variables.
4. Run the **cluster-backup.sh** script in the debug shell and pass in the location to save the backup to.

TIP

The **cluster-backup.sh** script is maintained as a component of the etcd Cluster Operator and is a wrapper around the **etcdctl snapshot save** command.


```
sh-4.4# /usr/local/bin/cluster-backup.sh /home/core/assets/backup
```

Example script output

```
found latest kube-apiserver: /etc/kubernetes/static-pod-resources/kube-apiserver-pod-6
found latest kube-controller-manager: /etc/kubernetes/static-pod-resources/kube-controller-
manager-pod-7
found latest kube-scheduler: /etc/kubernetes/static-pod-resources/kube-scheduler-pod-6
found latest etcd: /etc/kubernetes/static-pod-resources/etcd-pod-3
ede95fe6b88b87ba86a03c15e669fb4aa5bf0991c180d3c6895ce72eaade54a1
etcdctl version: 3.4.14
API version: 3.4
{"level":"info","ts":1624647639.0188997,"caller":"snapshot/v3_snapshot.go:119","msg":"created
temporary db file","path":"/home/core/assets/backup/snapshot_2021-06-25_190035.db.part"}
{"level":"info","ts":"2021-06-
25T19:00:39.030Z","caller":"clientv3/maintenance.go:200","msg":"opened snapshot stream;
downloading"}
{"level":"info","ts":1624647639.0301006,"caller":"snapshot/v3_snapshot.go:127","msg":"fetching
snapshot","endpoint":"https://10.0.0.5:2379"}
{"level":"info","ts":"2021-06-
25T19:00:40.215Z","caller":"clientv3/maintenance.go:208","msg":"completed snapshot read;
closing"}
{"level":"info","ts":1624647640.6032252,"caller":"snapshot/v3_snapshot.go:142","msg":"fetched
snapshot","endpoint":"https://10.0.0.5:2379","size":"114 MB","took":1.584090459}
{"level":"info","ts":1624647640.6047094,"caller":"snapshot/v3_snapshot.go:152","msg":"saved",
"path":"/home/core/assets/backup/snapshot_2021-06-25_190035.db"}
Snapshot saved at /home/core/assets/backup/snapshot_2021-06-25_190035.db
{"hash":3866667823,"revision":31407,"totalKey":12828,"totalSize":114446336}
snapshot db and kube resources are successfully saved to /home/core/assets/backup
```

In this example, two files are created in the **/home/core/assets/backup/** directory on the control plane host:

- **snapshot_<datetimestamp>.db**: This file is the etcd snapshot. The **cluster-backup.sh** script confirms its validity.
- **static_kuberresources_<datetimestamp>.tar.gz**: This file contains the resources for the static pods. If etcd encryption is enabled, it also contains the encryption keys for the etcd snapshot.

**NOTE**

If etcd encryption is enabled, it is recommended to store this second file separately from the etcd snapshot for security reasons. However, this file is required to restore from the etcd snapshot.

Keep in mind that etcd encryption only encrypts values, not keys. This means that resource types, namespaces, and object names are unencrypted.

5.2. REPLACING AN UNHEALTHY ETCD MEMBER

This document describes the process to replace a single unhealthy etcd member.

This process depends on whether the etcd member is unhealthy because the machine is not running or the node is not ready, or whether it is unhealthy because the etcd pod is crashlooping.

**NOTE**

If you have lost the majority of your control plane hosts, follow the disaster recovery procedure to [restore to a previous cluster state](#) instead of this procedure.

If the control plane certificates are not valid on the member being replaced, then you must follow the procedure to [recover from expired control plane certificates](#) instead of this procedure.

If a control plane node is lost and a new one is created, the etcd cluster Operator handles generating the new TLS certificates and adding the node as an etcd member.

5.2.1. Prerequisites

- Take an [etcd backup](#) prior to replacing an unhealthy etcd member.

5.2.2. Identifying an unhealthy etcd member

You can identify if your cluster has an unhealthy etcd member.

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Check the status of the **EtcMembersAvailable** status condition using the following command:

```
$ oc get etcd -o=jsonpath='{range .items[0].status.conditions[?(@.type=="EtcMembersAvailable")]}{.message}{"\n"}'
```

2. Review the output:

```
2 of 3 members are available, ip-10-0-131-183.ec2.internal is unhealthy
```

This example output shows that the **ip-10-0-131-183.ec2.internal** etcd member is unhealthy.

5.2.3. Determining the state of the unhealthy etcd member

The steps to replace an unhealthy etcd member depend on which of the following states your etcd member is in:

- The machine is not running or the node is not ready
- The etcd pod is crashlooping

This procedure determines which state your etcd member is in. This enables you to know which procedure to follow to replace the unhealthy etcd member.



NOTE

If you are aware that the machine is not running or the node is not ready, but you expect it to return to a healthy state soon, then you do not need to perform a procedure to replace the etcd member. The etcd cluster Operator will automatically sync when the machine or node returns to a healthy state.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have identified an unhealthy etcd member.

Procedure

1. Determine if the **machine is not running**

```
$ oc get machines -A -ojsonpath='{range .items[*]}{@.status.nodeRef.name}{"\t"}
{@.status.providerStatus.instanceState}{"\n"}' | grep -v running
```

Example output

```
ip-10-0-131-183.ec2.internal stopped 1
```

- 1** This output lists the node and the status of the node's machine. If the status is anything other than **running**, then the **machine is not running**

If the **machine is not running** then follow the *Replacing an unhealthy etcd member whose machine is not running or whose node is not ready* procedure.

2. Determine if the **node is not ready**.

If either of the following scenarios are true, then the **node is not ready**.

- If the machine is running, then check whether the node is unreachable:

```
$ oc get nodes -o jsonpath='{range .items[*]}{"\n"}{.metadata.name}{"\t"}{range
.spec.taints[*]}{.key}{" "}' | grep unreachable
```

Example output

```
ip-10-0-131-183.ec2.internal node-role.kubernetes.io/master
node.kubernetes.io/unreachable node.kubernetes.io/unreachable 1
```

1 If the node is listed with an **unreachable** taint, then the **node is not ready**.

- If the node is still reachable, then check whether the node is listed as **NotReady**:

```
$ oc get nodes -l node-role.kubernetes.io/master | grep "NotReady"
```

Example output

```
ip-10-0-131-183.ec2.internal NotReady master 122m v1.25.0 1
```

1 If the node is listed as **NotReady**, then the **node is not ready**.

If the **node is not ready**, then follow the *Replacing an unhealthy etcd member whose machine is not running or whose node is not ready* procedure.

3. Determine if the **etcd pod is crashlooping**

If the machine is running and the node is ready, then check whether the etcd pod is crashlooping.

- Verify that all control plane nodes are listed as **Ready**:

```
$ oc get nodes -l node-role.kubernetes.io/master
```

Example output

```
NAME                                STATUS ROLES  AGE  VERSION
ip-10-0-131-183.ec2.internal Ready  master  6h13m v1.25.0
ip-10-0-164-97.ec2.internal Ready  master  6h13m v1.25.0
ip-10-0-154-204.ec2.internal Ready  master  6h13m v1.25.0
```

- Check whether the status of an etcd pod is either **Error** or **CrashloopBackoff**:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

Example output

```
etcd-ip-10-0-131-183.ec2.internal      2/3 Error 7 6h9m 1
etcd-ip-10-0-164-97.ec2.internal      3/3 Running 0 6h6m
etcd-ip-10-0-154-204.ec2.internal     3/3 Running 0 6h6m
```

1 Since this status of this pod is **Error**, then the **etcd pod is crashlooping**

If the **etcd pod is crashlooping** then follow the *Replacing an unhealthy etcd member whose etcd pod is crashlooping* procedure.

5.2.4. Replacing the unhealthy etcd member

Depending on the state of your unhealthy etcd member, use one of the following procedures:

- [Replacing an unhealthy etcd member whose machine is not running or whose node is not ready](#)
- [Replacing an unhealthy etcd member whose etcd pod is crashlooping](#)
- [Replacing an unhealthy stopped baremetal etcd member](#)

5.2.4.1. Replacing an unhealthy etcd member whose machine is not running or whose node is not ready

This procedure details the steps to replace an etcd member that is unhealthy either because the machine is not running or because the node is not ready.



NOTE

If your cluster uses a control plane machine set, see "Recovering a degraded etcd Operator" in "Troubleshooting the control plane machine set" for a more simple etcd recovery procedure.

Prerequisites

- You have identified the unhealthy etcd member.
- You have verified that either the machine is not running or the node is not ready.



IMPORTANT

You must wait if the other control plane nodes are powered off. The control plane nodes must remain powered off until the replacement of an unhealthy etcd member is complete.

- You have access to the cluster as a user with the **cluster-admin** role.
- You have taken an etcd backup.



IMPORTANT

It is important to take an etcd backup before performing this procedure so that your cluster can be restored if you encounter any issues.

Procedure

1. Remove the unhealthy member.
 - a. Choose a pod that is *not* on the affected node:
In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

Example output

```

etcd-ip-10-0-131-183.ec2.internal    3/3  Running  0    123m
etcd-ip-10-0-164-97.ec2.internal   3/3  Running  0    123m
etcd-ip-10-0-154-204.ec2.internal  3/3  Running  0    124m

```

- b. Connect to the running etcd container, passing in the name of a pod that is not on the affected node:

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc rsh -n openshift-etcd etcd-ip-10-0-154-204.ec2.internal
```

- c. View the member list:

```
sh-4.2# etcdctl member list -w table
```

Example output

```

+-----+-----+-----+-----+-----+
+-----+
|  ID   | STATUS | NAME           | PEER ADDRS   | CLIENT
ADDRS  |
+-----+-----+-----+-----+-----+
+-----+
| 6fc1e7c9db35841d | started | ip-10-0-131-183.ec2.internal | https://10.0.131.183:2380 |
https://10.0.131.183:2379 |
| 757b6793e2408b6c | started | ip-10-0-164-97.ec2.internal | https://10.0.164.97:2380 |
https://10.0.164.97:2379 |
| ca8c2990a0aa29d1 | started | ip-10-0-154-204.ec2.internal | https://10.0.154.204:2380 |
https://10.0.154.204:2379 |
+-----+-----+-----+-----+-----+
+-----+

```

Take note of the ID and the name of the unhealthy etcd member, because these values are needed later in the procedure. The **\$ etcdctl endpoint health** command will list the removed member until the procedure of replacement is finished and a new member is added.

- d. Remove the unhealthy etcd member by providing the ID to the **etcdctl member remove** command:

```
sh-4.2# etcdctl member remove 6fc1e7c9db35841d
```

Example output

```
Member 6fc1e7c9db35841d removed from cluster ead669ce1fbfb346
```

- e. View the member list again and verify that the member was removed:

```
sh-4.2# etcdctl member list -w table
```

Example output

```

+-----+-----+-----+-----+-----+
-----+
|   ID   | STATUS |   NAME   |   PEER ADDRS   |   CLIENT
ADDRS   |
+-----+-----+-----+-----+-----+
-----+
| 757b6793e2408b6c | started | ip-10-0-164-97.ec2.internal | https://10.0.164.97:2380 |
https://10.0.164.97:2379 |
| ca8c2990a0aa29d1 | started | ip-10-0-154-204.ec2.internal | https://10.0.154.204:2380 |
https://10.0.154.204:2379 |
+-----+-----+-----+-----+-----+
-----+

```

You can now exit the node shell.

2. Turn off the quorum guard by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{"spec": {"unsupportedConfigOverrides": {"useUnsupportedUnsafeNonHANonProductionUnstableEtcd": true}}}'
```

This command ensures that you can successfully re-create secrets and roll out the static pods.



IMPORTANT

After you turn off the quorum guard, the cluster might be unreachable for a short time while the remaining etcd instances reboot to reflect the configuration change.



NOTE

etcd cannot tolerate any additional member failure when running with two members. Restarting either remaining member breaks the quorum and causes downtime in your cluster. The quorum guard protects etcd from restarts due to configuration changes that could cause downtime, so it must be disabled to complete this procedure.

3. Remove the old secrets for the unhealthy etcd member that was removed.
 - a. List the secrets for the unhealthy etcd member that was removed.

```
$ oc get secrets -n openshift-etcd | grep ip-10-0-131-183.ec2.internal 1
```

- 1** Pass in the name of the unhealthy etcd member that you took note of earlier in this procedure.

There is a peer, serving, and metrics secret as shown in the following output:

Example output

```

etcd-peer-ip-10-0-131-183.ec2.internal      kubernetes.io/tls      2    47m
etcd-serving-ip-10-0-131-183.ec2.internal  kubernetes.io/tls      2    47m
etcd-serving-metrics-ip-10-0-131-183.ec2.internal  kubernetes.io/tls      2
47m

```

b. Delete the secrets for the unhealthy etcd member that was removed.

i. Delete the peer secret:

```
$ oc delete secret -n openshift-etcd etcd-peer-ip-10-0-131-183.ec2.internal
```

ii. Delete the serving secret:

```
$ oc delete secret -n openshift-etcd etcd-serving-ip-10-0-131-183.ec2.internal
```

iii. Delete the metrics secret:

```
$ oc delete secret -n openshift-etcd etcd-serving-metrics-ip-10-0-131-183.ec2.internal
```

4. Delete and recreate the control plane machine. After this machine is recreated, a new revision is forced and etcd scales up automatically.

If you are running installer-provisioned infrastructure, or you used the Machine API to create your machines, follow these steps. Otherwise, you must create the new master using the same method that was used to originally create it.

a. Obtain the machine for the unhealthy member.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output

```

NAME                                PHASE  TYPE    REGION  ZONE  AGE
NODE                                PROVIDERID  STATE
clustername-8qw5l-master-0          Running m4.xlarge us-east-1 us-east-1a
3h37m ip-10-0-131-183.ec2.internal  aws:///us-east-1a/i-0ec2782f8287dfb7e  stopped
❶
clustername-8qw5l-master-1          Running m4.xlarge us-east-1 us-east-1b
3h37m ip-10-0-154-204.ec2.internal  aws:///us-east-1b/i-096c349b700a19631  running
clustername-8qw5l-master-2          Running m4.xlarge us-east-1 us-east-1c
3h37m ip-10-0-164-97.ec2.internal  aws:///us-east-1c/i-02626f1dba9ed5bba  running
clustername-8qw5l-worker-us-east-1a-wbtgd Running m4.large us-east-1 us-east-1a
1a 3h28m ip-10-0-129-226.ec2.internal  aws:///us-east-1a/i-010ef6279b4662ced  running
clustername-8qw5l-worker-us-east-1b-lrdxb Running m4.large us-east-1 us-east-1b
3h28m ip-10-0-144-248.ec2.internal  aws:///us-east-1b/i-0cb45ac45a166173b  running
clustername-8qw5l-worker-us-east-1c-pkg26 Running m4.large us-east-1 us-east-1c
1c 3h28m ip-10-0-170-181.ec2.internal  aws:///us-east-1c/i-06861c00007751b0a  running

```

❶ This is the control plane machine for the unhealthy node, **ip-10-0-131-183.ec2.internal**.

b. Save the machine configuration to a file on your file system:

```
$ oc get machine clustername-8qw5l-master-0 \ ❶
```

```
-n openshift-machine-api \
-o yaml \
> new-master-machine.yaml
```

- 1 Specify the name of the control plane machine for the unhealthy node.

- c. Edit the **new-master-machine.yaml** file that was created in the previous step to assign a new name and remove unnecessary fields.

- i. Remove the entire **status** section:

```
status:
  addresses:
    - address: 10.0.131.183
      type: InternalIP
    - address: ip-10-0-131-183.ec2.internal
      type: InternalDNS
    - address: ip-10-0-131-183.ec2.internal
      type: Hostname
  lastUpdated: "2020-04-20T17:44:29Z"
  nodeRef:
    kind: Node
    name: ip-10-0-131-183.ec2.internal
    uid: acca4411-af0d-4387-b73e-52b2484295ad
  phase: Running
  providerStatus:
    apiVersion: awsproviderconfig.openshift.io/v1beta1
    conditions:
      - lastProbeTime: "2020-04-20T16:53:50Z"
        lastTransitionTime: "2020-04-20T16:53:50Z"
        message: machine successfully created
        reason: MachineCreationSucceeded
        status: "True"
        type: MachineCreation
    instanceId: i-0fdb85790d76d0c3f
    instanceState: stopped
    kind: AWSMachineProviderStatus
```

- ii. Change the **metadata.name** field to a new name. It is recommended to keep the same base name as the old machine and change the ending number to the next available number. In this example, **clustername-8qw5l-master-0** is changed to **clustername-8qw5l-master-3**.

For example:

```
apiVersion: machine.openshift.io/v1beta1
kind: Machine
metadata:
  ...
  name: clustername-8qw5l-master-3
  ...
```

- iii. Remove the **spec.providerID** field:

```
providerID: aws:///us-east-1a/i-0fdb85790d76d0c3f
```

- d. Delete the machine of the unhealthy member:

```
$ oc delete machine -n openshift-machine-api clustername-8qw5l-master-0 1
```

- 1** Specify the name of the control plane machine for the unhealthy node.

- e. Verify that the machine was deleted:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output

```
NAME                PHASE  TYPE    REGION  ZONE    AGE
NODE                PROVIDERID                STATE
clustername-8qw5l-master-1      Running m4.xlarge us-east-1 us-east-1b
3h37m ip-10-0-154-204.ec2.internal aws:///us-east-1b/i-096c349b700a19631 running
clustername-8qw5l-master-2      Running m4.xlarge us-east-1 us-east-1c
3h37m ip-10-0-164-97.ec2.internal aws:///us-east-1c/i-02626f1dba9ed5bba running
clustername-8qw5l-worker-us-east-1a-wbtgd Running m4.large us-east-1 us-east-
1a 3h28m ip-10-0-129-226.ec2.internal aws:///us-east-1a/i-010ef6279b4662ced
running
clustername-8qw5l-worker-us-east-1b-lrdxb Running m4.large us-east-1 us-east-1b
3h28m ip-10-0-144-248.ec2.internal aws:///us-east-1b/i-0cb45ac45a166173b running
clustername-8qw5l-worker-us-east-1c-pkg26 Running m4.large us-east-1 us-east-
1c 3h28m ip-10-0-170-181.ec2.internal aws:///us-east-1c/i-06861c00007751b0a
running
```

- f. Create the new machine using the **new-master-machine.yaml** file:

```
$ oc apply -f new-master-machine.yaml
```

- g. Verify that the new machine has been created:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output

```
NAME                PHASE  TYPE    REGION  ZONE    AGE
NODE                PROVIDERID                STATE
clustername-8qw5l-master-1      Running m4.xlarge us-east-1 us-east-1b
3h37m ip-10-0-154-204.ec2.internal aws:///us-east-1b/i-096c349b700a19631 running
clustername-8qw5l-master-2      Running m4.xlarge us-east-1 us-east-1c
3h37m ip-10-0-164-97.ec2.internal aws:///us-east-1c/i-02626f1dba9ed5bba running
clustername-8qw5l-master-3      Provisioning m4.xlarge us-east-1 us-east-1a
85s ip-10-0-133-53.ec2.internal aws:///us-east-1a/i-015b0888fe17bc2c8 running
1
clustername-8qw5l-worker-us-east-1a-wbtgd Running m4.large us-east-1 us-
east-1a 3h28m ip-10-0-129-226.ec2.internal aws:///us-east-1a/i-010ef6279b4662ced
running
clustername-8qw5l-worker-us-east-1b-lrdxb Running m4.large us-east-1 us-east-
```



```
1b 3h28m ip-10-0-144-248.ec2.internal aws:///us-east-1b/i-0cb45ac45a166173b
running
clustername-8qw5l-worker-us-east-1c-pkg26 Running m4.large us-east-1 us-
east-1c 3h28m ip-10-0-170-181.ec2.internal aws:///us-east-1c/i-06861c00007751b0a
running
```

- 1 The new machine, **clustername-8qw5l-master-3** is being created and is ready once the phase changes from **Provisioning** to **Running**.

It might take a few minutes for the new machine to be created. The etcd cluster Operator will automatically sync when the machine or node returns to a healthy state.

5. Turn the quorum guard back on by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{"spec": {"unsupportedConfigOverrides": null}}'
```

6. You can verify that the **unsupportedConfigOverrides** section is removed from the object by entering this command:

```
$ oc get etcd/cluster -oyaml
```

7. If you are using single-node OpenShift, restart the node. Otherwise, you might encounter the following error in the etcd cluster Operator:

Example output

```
EtcdCertSignerControllerDegraded: [Operation cannot be fulfilled on secrets "etcd-peer-sno-0": the object has been modified; please apply your changes to the latest version and try again, Operation cannot be fulfilled on secrets "etcd-serving-sno-0": the object has been modified; please apply your changes to the latest version and try again, Operation cannot be fulfilled on secrets "etcd-serving-metrics-sno-0": the object has been modified; please apply your changes to the latest version and try again]
```

Verification

1. Verify that all etcd pods are running properly.
In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

Example output

```
etcd-ip-10-0-133-53.ec2.internal      3/3   Running   0      7m49s
etcd-ip-10-0-164-97.ec2.internal     3/3   Running   0      123m
etcd-ip-10-0-154-204.ec2.internal    3/3   Running   0      124m
```

If the output from the previous command only lists two pods, you can manually force an etcd redeployment. In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc patch etcd cluster -p='{"spec": {"forceRedeploymentReason": "recovery-"'$( date --rfc-3339=ns )'"}}' --type=merge 1
```

- 1 The **forceRedeploymentReason** value must be unique, which is why a timestamp is appended.

2. Verify that there are exactly three etcd members.

- a. Connect to the running etcd container, passing in the name of a pod that was not on the affected node:

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc rsh -n openshift-etcd etcd-ip-10-0-154-204.ec2.internal
```

- b. View the member list:

```
sh-4.2# etcdctl member list -w table
```

Example output

```
+-----+-----+-----+-----+-----+
+-----+
| ID      | STATUS | NAME           | PEER ADDRS      | CLIENT
ADDRS    |
+-----+-----+-----+-----+-----+
+-----+
| 5eb0d6b8ca24730c | started | ip-10-0-133-53.ec2.internal | https://10.0.133.53:2380 |
https://10.0.133.53:2379 |
| 757b6793e2408b6c | started | ip-10-0-164-97.ec2.internal | https://10.0.164.97:2380 |
https://10.0.164.97:2379 |
| ca8c2990a0aa29d1 | started | ip-10-0-154-204.ec2.internal | https://10.0.154.204:2380 |
https://10.0.154.204:2379 |
+-----+-----+-----+-----+-----+
+-----+
```

If the output from the previous command lists more than three etcd members, you must carefully remove the unwanted member.



WARNING

Be sure to remove the correct etcd member; removing a good etcd member might lead to quorum loss.

Additional resources

- [Recovering a degraded etcd Operator](#)

5.2.4.2. Replacing an unhealthy etcd member whose etcd pod is crashlooping

This procedure details the steps to replace an etcd member that is unhealthy because the etcd pod is crashlooping.

Prerequisites

- You have identified the unhealthy etcd member.
- You have verified that the etcd pod is crashlooping.
- You have access to the cluster as a user with the **cluster-admin** role.
- You have taken an etcd backup.



IMPORTANT

It is important to take an etcd backup before performing this procedure so that your cluster can be restored if you encounter any issues.

Procedure

1. Stop the crashlooping etcd pod.
 - a. Debug the node that is crashlooping.
In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc debug node/ip-10-0-131-183.ec2.internal 1
```

- 1 Replace this with the name of the unhealthy node.

- b. Change your root directory to **/host**:

```
sh-4.2# chroot /host
```

- c. Move the existing etcd pod file out of the kubelet manifest directory:

```
sh-4.2# mkdir /var/lib/etcd-backup
```

```
sh-4.2# mv /etc/kubernetes/manifests/etcd-pod.yaml /var/lib/etcd-backup/
```

- d. Move the etcd data directory to a different location:

```
sh-4.2# mv /var/lib/etcd/ /tmp
```

You can now exit the node shell.

2. Remove the unhealthy member.
 - a. Choose a pod that is *not* on the affected node.
In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

-

Example output

```

etcd-ip-10-0-131-183.ec2.internal    2/3  Error    7    6h9m
etcd-ip-10-0-164-97.ec2.internal   3/3  Running  0    6h6m
etcd-ip-10-0-154-204.ec2.internal  3/3  Running  0    6h6m

```

- b. Connect to the running etcd container, passing in the name of a pod that is not on the affected node.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc rsh -n openshift-etcd etcd-ip-10-0-154-204.ec2.internal
```

- c. View the member list:

```
sh-4.2# etcdctl member list -w table
```

Example output

```

+-----+-----+-----+-----+-----+
+-----+
| ID | STATUS | NAME | PEER ADDRS | CLIENT
ADDRS |
+-----+-----+-----+-----+-----+
+-----+
| 62bcf33650a7170a | started | ip-10-0-131-183.ec2.internal | https://10.0.131.183:2380 |
https://10.0.131.183:2379 |
| b78e2856655bc2eb | started | ip-10-0-164-97.ec2.internal | https://10.0.164.97:2380 |
https://10.0.164.97:2379 |
| d022e10b498760d5 | started | ip-10-0-154-204.ec2.internal | https://10.0.154.204:2380 |
https://10.0.154.204:2379 |
+-----+-----+-----+-----+-----+
+-----+

```

Take note of the ID and the name of the unhealthy etcd member, because these values are needed later in the procedure.

- d. Remove the unhealthy etcd member by providing the ID to the **etcdctl member remove** command:

```
sh-4.2# etcdctl member remove 62bcf33650a7170a
```

Example output

```
Member 62bcf33650a7170a removed from cluster ead669ce1fbfb346
```

- e. View the member list again and verify that the member was removed:

```
sh-4.2# etcdctl member list -w table
```

Example output

```

+-----+-----+-----+-----+-----+
-----+
|   ID   | STATUS |   NAME   |   PEER ADDRS   |   CLIENT
ADDRS   |
+-----+-----+-----+-----+-----+
-----+
| b78e2856655bc2eb | started | ip-10-0-164-97.ec2.internal | https://10.0.164.97:2380 |
https://10.0.164.97:2379 |
| d022e10b498760d5 | started | ip-10-0-154-204.ec2.internal | https://10.0.154.204:2380
| https://10.0.154.204:2379 |
+-----+-----+-----+-----+-----+
-----+

```

You can now exit the node shell.

3. Turn off the quorum guard by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{"spec": {"unsupportedConfigOverrides": {"useUnsupportedUnsafeNonHANonProductionUnstableEtcd": true}}}'
```

This command ensures that you can successfully re-create secrets and roll out the static pods.

4. Remove the old secrets for the unhealthy etcd member that was removed.
 - a. List the secrets for the unhealthy etcd member that was removed.

```
$ oc get secrets -n openshift-etcd | grep ip-10-0-131-183.ec2.internal 1
```

- 1** Pass in the name of the unhealthy etcd member that you took note of earlier in this procedure.

There is a peer, serving, and metrics secret as shown in the following output:

Example output

```

etcd-peer-ip-10-0-131-183.ec2.internal      kubernetes.io/tls      2    47m
etcd-serving-ip-10-0-131-183.ec2.internal  kubernetes.io/tls      2    47m
etcd-serving-metrics-ip-10-0-131-183.ec2.internal  kubernetes.io/tls      2
47m

```

- b. Delete the secrets for the unhealthy etcd member that was removed.

- i. Delete the peer secret:

```
$ oc delete secret -n openshift-etcd etcd-peer-ip-10-0-131-183.ec2.internal
```

- ii. Delete the serving secret:

```
$ oc delete secret -n openshift-etcd etcd-serving-ip-10-0-131-183.ec2.internal
```

- iii. Delete the metrics secret:

```
$ oc delete secret -n openshift-etcd etcd-serving-metrics-ip-10-0-131-183.ec2.internal
```

- Force etcd redeployment.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc patch etcd cluster -p='{ "spec": { "forceRedeploymentReason": "single-master-recovery-$( date --rfc-3339=ns )"' --type=merge 1
```

- The **forceRedeploymentReason** value must be unique, which is why a timestamp is appended.

When the etcd cluster Operator performs a redeployment, it ensures that all control plane nodes have a functioning etcd pod.

- Turn the quorum guard back on by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{ "spec": { "unsupportedConfigOverrides": null } }
```

- You can verify that the **unsupportedConfigOverrides** section is removed from the object by entering this command:

```
$ oc get etcd/cluster -oyaml
```

- If you are using single-node OpenShift, restart the node. Otherwise, you might encounter the following error in the etcd cluster Operator:

Example output

```
EtcDCertSignerControllerDegraded: [Operation cannot be fulfilled on secrets "etcd-peer-sno-0": the object has been modified; please apply your changes to the latest version and try again, Operation cannot be fulfilled on secrets "etcd-serving-sno-0": the object has been modified; please apply your changes to the latest version and try again, Operation cannot be fulfilled on secrets "etcd-serving-metrics-sno-0": the object has been modified; please apply your changes to the latest version and try again]
```

Verification

- Verify that the new member is available and healthy.
 - Connect to the running etcd container again.
In a terminal that has access to the cluster as a cluster-admin user, run the following command:

```
$ oc rsh -n openshift-etcd etcd-ip-10-0-154-204.ec2.internal
```

- Verify that all members are healthy:

```
sh-4.2# etcdctl endpoint health
```

Example output

```
-
```

```

https://10.0.131.183:2379 is healthy: successfully committed proposal: took =
16.671434ms
https://10.0.154.204:2379 is healthy: successfully committed proposal: took =
16.698331ms
https://10.0.164.97:2379 is healthy: successfully committed proposal: took =
16.621645ms

```

5.2.4.3. Replacing an unhealthy bare metal etcd member whose machine is not running or whose node is not ready

This procedure details the steps to replace a bare metal etcd member that is unhealthy either because the machine is not running or because the node is not ready.

If you are running installer-provisioned infrastructure or you used the Machine API to create your machines, follow these steps. Otherwise you must create the new control plane node using the same method that was used to originally create it.

Prerequisites

- You have identified the unhealthy bare metal etcd member.
- You have verified that either the machine is not running or the node is not ready.
- You have access to the cluster as a user with the **cluster-admin** role.
- You have taken an etcd backup.



IMPORTANT

You must take an etcd backup before performing this procedure so that your cluster can be restored if you encounter any issues.

Procedure

1. Verify and remove the unhealthy member.
 - a. Choose a pod that is *not* on the affected node:
In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd -o wide
```

Example output

```

etcd-openshift-control-plane-0 5/5 Running 11 3h56m 192.168.10.9 openshift-
control-plane-0 <none> <none>
etcd-openshift-control-plane-1 5/5 Running 0 3h54m 192.168.10.10 openshift-
control-plane-1 <none> <none>
etcd-openshift-control-plane-2 5/5 Running 0 3h58m 192.168.10.11 openshift-
control-plane-2 <none> <none>

```

- b. Connect to the running etcd container, passing in the name of a pod that is not on the affected node:

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc rsh -n openshift-etcd etcd-openshift-control-plane-0
```

- c. View the member list:

```
sh-4.2# etcdctl member list -w table
```

Example output

```
+-----+-----+-----+-----+-----+
+-----+
| ID          | STATUS | NAME                | PEER ADDRS          | CLIENT
ADDRS        | IS LEARNER |                    |                    |
+-----+-----+-----+-----+-----+
+-----+
| 7a8197040a5126c8 | started | openshift-control-plane-2 | https://192.168.10.11:2380/ |
https://192.168.10.11:2379/ | false |
| 8d5abe9669a39192 | started | openshift-control-plane-1 | https://192.168.10.10:2380/ |
https://192.168.10.10:2379/ | false |
| cc3830a72fc357f9 | started | openshift-control-plane-0 | https://192.168.10.9:2380/ |
https://192.168.10.9:2379/ | false |
+-----+-----+-----+-----+-----+
+-----+
```

Take note of the ID and the name of the unhealthy etcd member, because these values are required later in the procedure. The **etcdctl endpoint health** command will list the removed member until the replacement procedure is completed and the new member is added.

- d. Remove the unhealthy etcd member by providing the ID to the **etcdctl member remove** command:



WARNING

Be sure to remove the correct etcd member; removing a good etcd member might lead to quorum loss.

```
sh-4.2# etcdctl member remove 7a8197040a5126c8
```

Example output

```
Member 7a8197040a5126c8 removed from cluster b23536c33f2cdd1b
```

- e. View the member list again and verify that the member was removed:

```
sh-4.2# etcdctl member list -w table
```


Example output

```

+-----+-----+-----+-----+-----+
+-----+
| ID          | STATUS | NAME                | PEER ADDRS          | CLIENT
ADDRS        | IS LEARNER |                    |                    |
+-----+-----+-----+-----+-----+
+-----+
| cc3830a72fc357f9 | started | openshift-control-plane-2 | https://192.168.10.11:2380/ |
https://192.168.10.11:2379/ | false |
| 8d5abe9669a39192 | started | openshift-control-plane-1 | https://192.168.10.10:2380/ |
https://192.168.10.10:2379/ | false |
+-----+-----+-----+-----+-----+
+-----+

```

You can now exit the node shell.

**IMPORTANT**

After you remove the member, the cluster might be unreachable for a short time while the remaining etcd instances reboot.

2. Turn off the quorum guard by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{"spec": {"unsupportedConfigOverrides": {"useUnsupportedUnsafeNonHANonProductionUnstableEtcd": true}}}'
```

This command ensures that you can successfully re-create secrets and roll out the static pods.

3. Remove the old secrets for the unhealthy etcd member that was removed by running the following commands.
 - a. List the secrets for the unhealthy etcd member that was removed.

```
$ oc get secrets -n openshift-etcd | grep openshift-control-plane-2
```

Pass in the name of the unhealthy etcd member that you took note of earlier in this procedure.

There is a peer, serving, and metrics secret as shown in the following output:

```

etcd-peer-openshift-control-plane-2      kubernetes.io/tls 2 134m
etcd-serving-metrics-openshift-control-plane-2 kubernetes.io/tls 2 134m
etcd-serving-openshift-control-plane-2    kubernetes.io/tls 2 134m

```

- b. Delete the secrets for the unhealthy etcd member that was removed.
 - i. Delete the peer secret:

```
$ oc delete secret etcd-peer-openshift-control-plane-2 -n openshift-etcd
secret "etcd-peer-openshift-control-plane-2" deleted
```

- ii. Delete the serving secret:

```
$ oc delete secret etcd-serving-metrics-openshift-control-plane-2 -n openshift-etcd
secret "etcd-serving-metrics-openshift-control-plane-2" deleted
```

- iii. Delete the metrics secret:

```
$ oc delete secret etcd-serving-openshift-control-plane-2 -n openshift-etcd
secret "etcd-serving-openshift-control-plane-2" deleted
```

4. Delete the control plane machine.

If you are running installer-provisioned infrastructure, or you used the Machine API to create your machines, follow these steps. Otherwise, you must create the new control plane node using the same method that was used to originally create it.

- a. Obtain the machine for the unhealthy member.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output

```
NAME                                PHASE  TYPE  REGION  ZONE  AGE  NODE
PROVIDERID                          STATE
examplecluster-control-plane-0      Running          3h11m openshift-control-
plane-0 baremetalhost:///openshift-machine-api/openshift-control-plane-0/da1ebe11-
3ff2-41c5-b099-0aa41222964e  externally provisioned 1
examplecluster-control-plane-1      Running          3h11m openshift-control-
plane-1 baremetalhost:///openshift-machine-api/openshift-control-plane-1/d9f9acbc-
329c-475e-8d81-03b20280a3e1  externally provisioned
examplecluster-control-plane-2      Running          3h11m openshift-control-
plane-2 baremetalhost:///openshift-machine-api/openshift-control-plane-2/3354bdac-
61d8-410f-be5b-6a395b056135  externally provisioned
examplecluster-compute-0            Running          165m openshift-compute-0
baremetalhost:///openshift-machine-api/openshift-compute-0/3d685b81-7410-4bb3-80ec-
13a31858241f  provisioned
examplecluster-compute-1            Running          165m openshift-compute-1
baremetalhost:///openshift-machine-api/openshift-compute-1/0fdae6eb-2066-4241-91dc-
e7ea72ab13b9  provisioned
```

- 1** This is the control plane machine for the unhealthy node, **examplecluster-control-plane-2**.

- b. Save the machine configuration to a file on your file system:

```
$ oc get machine examplecluster-control-plane-2 \ 1
-n openshift-machine-api \
-o yaml \
> new-master-machine.yaml
```

- 1** Specify the name of the control plane machine for the unhealthy node.

- c. Edit the **new-master-machine.yaml** file that was created in the previous step to assign a new name and remove unnecessary fields.
- i. Remove the entire **status** section:

```
status:
  addresses:
    - address: ""
      type: InternalIP
    - address: fe80::4adf:37ff:feb0:8aa1%ens1f1.373
      type: InternalDNS
    - address: fe80::4adf:37ff:feb0:8aa1%ens1f1.371
      type: Hostname
  lastUpdated: "2020-04-20T17:44:29Z"
  nodeRef:
    kind: Machine
    name: fe80::4adf:37ff:feb0:8aa1%ens1f1.372
    uid: acca4411-af0d-4387-b73e-52b2484295ad
  phase: Running
  providerStatus:
    apiVersion: machine.openshift.io/v1beta1
    conditions:
      - lastProbeTime: "2020-04-20T16:53:50Z"
        lastTransitionTime: "2020-04-20T16:53:50Z"
        message: machine successfully created
        reason: MachineCreationSucceeded
        status: "True"
        type: MachineCreation
    instanceId: i-0fdb85790d76d0c3f
    instanceState: stopped
    kind: Machine
```

5. Change the **metadata.name** field to a new name. It is recommended to keep the same base name as the old machine and change the ending number to the next available number. In this example, **examplecluster-control-plane-2** is changed to **examplecluster-control-plane-3**.

For example:

```
apiVersion: machine.openshift.io/v1beta1
kind: Machine
metadata:
  ...
  name: examplecluster-control-plane-3
  ...
```

- a. Remove the **spec.providerID** field:

```
providerID: baremetalhost:///openshift-machine-api/openshift-control-plane-2/3354bdac-61d8-410f-be5b-6a395b056135
```

- b. Remove the **metadata.annotations** and **metadata.generation** fields:

```
annotations:
  machine.openshift.io/instance-state: externally provisioned
```

```
...
generation: 2
```

- c. Remove the **spec.conditions**, **spec.lastUpdated**, **spec.nodeRef** and **spec.phase** fields:

```
lastTransitionTime: "2022-08-03T08:40:36Z"
message: 'Drain operation currently blocked by: [{Name:EtcdQuorumOperator
Owner:clusteroperator/etcd}]'
reason: HookPresent
severity: Warning
status: "False"

type: Drainable
lastTransitionTime: "2022-08-03T08:39:55Z"
status: "True"
type: InstanceExists

lastTransitionTime: "2022-08-03T08:36:37Z"
status: "True"
type: Terminable
lastUpdated: "2022-08-03T08:40:36Z"
nodeRef:
kind: Node
name: openshift-control-plane-2
uid: 788df282-6507-4ea2-9a43-24f237ccbc3c
phase: Running
```

6. Ensure that the Bare Metal Operator is available by running the following command:

```
$ oc get clusteroperator baremetal
```

Example output

```
NAME      VERSION AVAILABLE PROGRESSING DEGRADED SINCE MESSAGE
baremetal 4.12.0  True    False     False    3d15h
```

7. Remove the old **BareMetalHost** object by running the following command:

```
$ oc delete bmh openshift-control-plane-2 -n openshift-machine-api
```

Example output

```
baremetalhost.metal3.io "openshift-control-plane-2" deleted
```

8. Delete the machine of the unhealthy member by running the following command:

```
$ oc delete machine -n openshift-machine-api examplecluster-control-plane-2
```

After you remove the **BareMetalHost** and **Machine** objects, then the **Machine** controller automatically deletes the **Node** object.

If deletion of the machine is delayed for any reason or the command is obstructed and delayed, you can force deletion by removing the machine object finalizer field.



IMPORTANT

Do not interrupt machine deletion by pressing **Ctrl+c**. You must allow the command to proceed to completion. Open a new terminal window to edit and delete the finalizer fields.

- a. Edit the machine configuration by running the following command:

```
$ oc edit machine -n openshift-machine-api examplecluster-control-plane-2
```

- b. Delete the following fields in the **Machine** custom resource, and then save the updated file:

```
finalizers:
- machine.machine.openshift.io
```

Example output

```
machine.machine.openshift.io/examplecluster-control-plane-2 edited
```

9. Verify that the machine was deleted by running the following command:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output

NAME	PHASE	TYPE	REGION	ZONE	AGE	NODE
examplecluster-control-plane-0	Running				3h11m	openshift-control-plane-0
baremetalhost:///openshift-machine-api/openshift-control-plane-0/da1ebe11-3ff2-41c5-b099-0aa41222964e	externally provisioned					
examplecluster-control-plane-1	Running				3h11m	openshift-control-plane-1
baremetalhost:///openshift-machine-api/openshift-control-plane-1/d9f9acbc-329c-475e-8d81-03b20280a3e1	externally provisioned					
examplecluster-compute-0	Running				165m	openshift-compute-0
baremetalhost:///openshift-machine-api/openshift-compute-0/3d685b81-7410-4bb3-80ec-13a31858241f	provisioned					
examplecluster-compute-1	Running				165m	openshift-compute-1
baremetalhost:///openshift-machine-api/openshift-compute-1/0fdae6eb-2066-4241-91dc-e7ea72ab13b9	provisioned					

10. Verify that the node has been deleted by running the following command:

```
$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
openshift-control-plane-0	Ready	master	3h24m	v1.25.0
openshift-control-plane-1	Ready	master	3h24m	v1.25.0
openshift-compute-0	Ready	worker	176m	v1.25.0
openshift-compute-1	Ready	worker	176m	v1.25.0

11. Create the new **BareMetalHost** object and the secret to store the BMC credentials:

```
$ cat <<EOF | oc apply -f -
```

```

apiVersion: v1
kind: Secret
metadata:
  name: openshift-control-plane-2-bmc-secret
  namespace: openshift-machine-api
data:
  password: <password>
  username: <username>
type: Opaque
---
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  name: openshift-control-plane-2
  namespace: openshift-machine-api
spec:
  automatedCleaningMode: disabled
  bmc:
    address: redfish://10.46.61.18:443/redfish/v1/Systems/1
    credentialsName: openshift-control-plane-2-bmc-secret
    disableCertificateVerification: true
    bootMACAddress: 48:df:37:b0:8a:a0
    bootMode: UEFI
    externallyProvisioned: false
    online: true
    rootDeviceHints:
      deviceName: /dev/sda
    userData:
      name: master-user-data-managed
      namespace: openshift-machine-api
EOF

```



NOTE

The username and password can be found from the other bare metal host's secrets. The protocol to use in **bmc:address** can be taken from other bmh objects.



IMPORTANT

If you reuse the **BareMetalHost** object definition from an existing control plane host, do not leave the **externallyProvisioned** field set to **true**.

Existing control plane **BareMetalHost** objects may have the **externallyProvisioned** flag set to **true** if they were provisioned by the OpenShift Container Platform installation program.

After the inspection is complete, the **BareMetalHost** object is created and available to be provisioned.

- Verify the creation process using available **BareMetalHost** objects:

```
$ oc get bmh -n openshift-machine-api
```

NAME	STATE	CONSUMER	ONLINE	ERROR	AGE
------	-------	----------	--------	-------	-----

```

openshift-control-plane-0 externally provisioned examplecluster-control-plane-0 true
4h48m
openshift-control-plane-1 externally provisioned examplecluster-control-plane-1 true
4h48m
openshift-control-plane-2 available          examplecluster-control-plane-3 true    47m
openshift-compute-0    provisioned          examplecluster-compute-0    true    4h48m
openshift-compute-1    provisioned          examplecluster-compute-1    true    4h48m

```

- a. Create the new control plane machine using the **new-master-machine.yaml** file:

```
$ oc apply -f new-master-machine.yaml
```

- b. Verify that the new machine has been created:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output

```

NAME                                PHASE  TYPE  REGION  ZONE  AGE  NODE
PROVIDERID                            STATE
examplecluster-control-plane-0      Running                3h11m openshift-control-
plane-0 baremetalhost:///openshift-machine-api/openshift-control-plane-0/da1ebe11-
3ff2-41c5-b099-0aa41222964e  externally provisioned 1
examplecluster-control-plane-1      Running                3h11m openshift-control-
plane-1 baremetalhost:///openshift-machine-api/openshift-control-plane-1/d9f9acbc-
329c-475e-8d81-03b20280a3e1  externally provisioned
examplecluster-control-plane-2      Running                3h11m openshift-control-
plane-2 baremetalhost:///openshift-machine-api/openshift-control-plane-2/3354bdac-
61d8-410f-be5b-6a395b056135  externally provisioned
examplecluster-compute-0            Running                165m openshift-compute-
0 baremetalhost:///openshift-machine-api/openshift-compute-0/3d685b81-7410-
4bb3-80ec-13a31858241f    provisioned
examplecluster-compute-1            Running                165m openshift-compute-
1 baremetalhost:///openshift-machine-api/openshift-compute-1/0fdae6eb-2066-
4241-91dc-e7ea72ab13b9    provisioned

```

- 1 The new machine, **clustername-8qw5l-master-3** is being created and is ready after the phase changes from **Provisioning** to **Running**.

It should take a few minutes for the new machine to be created. The etcd cluster Operator will automatically sync when the machine or node returns to a healthy state.

- c. Verify that the bare metal host becomes provisioned and no error reported by running the following command:

```
$ oc get bmh -n openshift-machine-api
```

Example output

```

$ oc get bmh -n openshift-machine-api
NAME                                STATE          CONSUMER                                ONLINE ERROR AGE
openshift-control-plane-0 externally provisioned examplecluster-control-plane-0 true
4h48m

```

```

openshift-control-plane-1 externally provisioned examplecluster-control-plane-1 true
4h48m
openshift-control-plane-2 provisioned           examplecluster-control-plane-3 true
47m
openshift-compute-0     provisioned           examplecluster-compute-0     true
4h48m
openshift-compute-1     provisioned           examplecluster-compute-1     true
4h48m

```

- d. Verify that the new node is added and in a ready state by running this command:

```
$ oc get nodes
```

Example output

```

$ oc get nodes
NAME                                STATUS ROLES  AGE  VERSION
openshift-control-plane-0 Ready  master  4h26m v1.25.0
openshift-control-plane-1 Ready  master  4h26m v1.25.0
openshift-control-plane-2 Ready  master  12m   v1.25.0
openshift-compute-0     Ready  worker  3h58m v1.25.0
openshift-compute-1     Ready  worker  3h58m v1.25.0

```

13. Turn the quorum guard back on by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{"spec": {"unsupportedConfigOverrides": null}}'
```

14. You can verify that the **unsupportedConfigOverrides** section is removed from the object by entering this command:

```
$ oc get etcd/cluster -oyaml
```

15. If you are using single-node OpenShift, restart the node. Otherwise, you might encounter the following error in the etcd cluster Operator:

Example output

```

EtcDCertSignerControllerDegraded: [Operation cannot be fulfilled on secrets "etcd-peer-sno-0": the object has been modified; please apply your changes to the latest version and try again, Operation cannot be fulfilled on secrets "etcd-serving-sno-0": the object has been modified; please apply your changes to the latest version and try again, Operation cannot be fulfilled on secrets "etcd-serving-metrics-sno-0": the object has been modified; please apply your changes to the latest version and try again]

```

Verification

1. Verify that all etcd pods are running properly.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

Example output


```
etcd-openshift-control-plane-0 5/5 Running 0 105m
etcd-openshift-control-plane-1 5/5 Running 0 107m
etcd-openshift-control-plane-2 5/5 Running 0 103m
```

If the output from the previous command only lists two pods, you can manually force an etcd redeployment. In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc patch etcd cluster -p="{\"spec\": {\"forceRedeploymentReason\": \"recovery-\"$( date --rfc-3339=ns )\"}}\" --type=merge 1
```

- 1 The **forceRedeploymentReason** value must be unique, which is why a timestamp is appended.

To verify there are exactly three etcd members, connect to the running etcd container, passing in the name of a pod that was not on the affected node. In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc rsh -n openshift-etcd etcd-openshift-control-plane-0
```

2. View the member list:

```
sh-4.2# etcdctl member list -w table
```

Example output

```
+-----+-----+-----+-----+-----+
| ID | STATUS | NAME | PEER ADDRS | CLIENT ADDRS |
| IS LEARNER |
+-----+-----+-----+-----+-----+
| 7a8197040a5126c8 | started | openshift-control-plane-2 | https://192.168.10.11:2380 |
https://192.168.10.11:2379 | false |
| 8d5abe9669a39192 | started | openshift-control-plane-1 | https://192.168.10.10:2380 |
https://192.168.10.10:2379 | false |
| cc3830a72fc357f9 | started | openshift-control-plane-0 | https://192.168.10.9:2380 |
https://192.168.10.9:2379 | false |
+-----+-----+-----+-----+-----+
```



NOTE

If the output from the previous command lists more than three etcd members, you must carefully remove the unwanted member.

3. Verify that all etcd members are healthy by running the following command:

```
# etcdctl endpoint health --cluster
```

Example output

```
https://192.168.10.10:2379 is healthy: successfully committed proposal: took = 8.973065ms
https://192.168.10.9:2379 is healthy: successfully committed proposal: took = 11.559829ms
https://192.168.10.11:2379 is healthy: successfully committed proposal: took = 11.665203ms
```

4. Validate that all nodes are at the latest revision by running the following command:

```
$ oc get etcd -o=jsonpath='{range.items[0].status.conditions[?(@.type=="NodeInstallerProgressing")]}{.reason}\n'}{.message}\n'
```


```
AllNodesAtLatestRevision
```

5.2.5. Additional resources

- [Quorum protection with machine lifecycle hooks](#)

5.3. BACKING UP AND RESTORING ETCD ON A HOSTED CLUSTER

If you use hosted control planes on OpenShift Container Platform, the process to back up and restore etcd is different from [the usual etcd backup process](#).




IMPORTANT

Hosted control planes is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

5.3.1. Taking a snapshot of etcd on a hosted cluster

As part of the process to back up etcd for a hosted cluster, you take a snapshot of etcd. After you take the snapshot, you can restore it, for example, as part of a disaster recovery operation.



IMPORTANT

This procedure requires API downtime.

Procedure

1. Pause reconciliation of the hosted cluster by entering this command:

```
$ oc patch -n clusters hostedclusters/${CLUSTER_NAME} -p '{"spec": {"pausedUntil":"'${PAUSED_UNTIL}'"}}' --type=merge
```

2. Stop all etcd-writer deployments by entering this command:

```
$ oc scale deployment -n ${HOSTED_CLUSTER_NAMESPACE} --replicas=0 kube-apiserver openshift-apiserver openshift-oauth-apiserver
```

3. Take an etcd snapshot by using the **exec** command in each etcd container:

```
$ oc exec -it etcd-0 -n ${HOSTED_CLUSTER_NAMESPACE} -- env ETCDCCTL_API=3
/usr/bin/etcdctl --cacert /etc/etcd/tls/client/etcd-client-ca.crt --cert /etc/etcd/tls/client/etcd-
client.crt --key /etc/etcd/tls/client/etcd-client.key --endpoints=localhost:2379 snapshot save
/var/lib/data/snapshot.db
$ oc exec -it etcd-0 -n ${HOSTED_CLUSTER_NAMESPACE} -- env ETCDCCTL_API=3
/usr/bin/etcdctl -w table snapshot status /var/lib/data/snapshot.db
```

4. Copy the snapshot data to a location where you can retrieve it later, such as an S3 bucket, as shown in the following example.



NOTE

The following example uses signature version 2. If you are in a region that supports signature version 4, such as the us-east-2 region, use signature version 4. Otherwise, if you use signature version 2 to copy the snapshot to an S3 bucket, the upload fails and signature version 2 is deprecated.

Example

```
BUCKET_NAME=somebucket
FILEPATH="/${BUCKET_NAME}/${CLUSTER_NAME}-snapshot.db"
CONTENT_TYPE="application/x-compressed-tar"
DATE_VALUE=`date -R`
SIGNATURE_STRING="PUT\n\n${CONTENT_TYPE}\n${DATE_VALUE}\n${FILEPATH}"
ACCESS_KEY=accesskey
SECRET_KEY=secret
SIGNATURE_HASH=`echo -en ${SIGNATURE_STRING} | openssl sha1 -hmac
${SECRET_KEY} -binary | base64`

oc exec -it etcd-0 -n ${HOSTED_CLUSTER_NAMESPACE} -- curl -X PUT -T
"/var/lib/data/snapshot.db" \
-H "Host: ${BUCKET_NAME}.s3.amazonaws.com" \
-H "Date: ${DATE_VALUE}" \
-H "Content-Type: ${CONTENT_TYPE}" \
-H "Authorization: AWS ${ACCESS_KEY}:${SIGNATURE_HASH}" \
https://${BUCKET_NAME}.s3.amazonaws.com/${CLUSTER_NAME}-snapshot.db
```

5. If you want to be able to restore the snapshot on a new cluster later, save the encryption secret that the hosted cluster references, as shown in this example:

Example

```
oc get hostedcluster $CLUSTER_NAME -o=jsonpath='{.spec.secretEncryption.aescbc}'
{"activeKey":{"name":"CLUSTER_NAME-etcd-encryption-key"}}

# Save this secret, or the key it contains so the etcd data can later be decrypted
oc get secret ${CLUSTER_NAME}-etcd-encryption-key -o=jsonpath='{.data.key}'
```

Next steps

Restore the etcd snapshot.

5.3.2. Restoring an etcd snapshot on a hosted cluster

If you have a snapshot of etcd from your hosted cluster, you can restore it. Currently, you can restore an etcd snapshot only during cluster creation.

To restore an etcd snapshot, you modify the output from the **create cluster --render** command and define a **restoreSnapshotURL** value in the etcd section of the **HostedCluster** specification.

Prerequisites

You took an etcd snapshot on a hosted cluster.

Procedure

1. On the **aws** command-line interface (CLI), create a pre-signed URL so that you can download your etcd snapshot from S3 without passing credentials to the etcd deployment:

```
ETCD_SNAPSHOT=${ETCD_SNAPSHOT:-"s3://${BUCKET_NAME}/${CLUSTER_NAME}-
snapshot.db"}
ETCD_SNAPSHOT_URL=$(aws s3 presign ${ETCD_SNAPSHOT})
```

2. Modify the **HostedCluster** specification to refer to the URL:

```
spec:
  etcd:
    managed:
      storage:
        persistentVolume:
          size: 4Gi
          type: PersistentVolume
        restoreSnapshotURL:
          - "${ETCD_SNAPSHOT_URL}"
        managementType: Managed
```

3. Ensure that the secret that you referenced from the **spec.secretEncryption.aescbc** value contains the same AES key that you saved in the previous steps.

5.3.3. Additional resources

- [Disaster recovery for a hosted cluster within an AWS region](#)

5.4. DISASTER RECOVERY

5.4.1. About disaster recovery

The disaster recovery documentation provides information for administrators on how to recover from several disaster situations that might occur with their OpenShift Container Platform cluster. As an administrator, you might need to follow one or more of the following procedures to return your cluster to a working state.



IMPORTANT

Disaster recovery requires you to have at least one healthy control plane host.

Restoring to a previous cluster state

This solution handles situations where you want to restore your cluster to a previous state, for example, if an administrator deletes something critical. This also includes situations where you have lost the majority of your control plane hosts, leading to etcd quorum loss and the cluster going offline. As long as you have taken an etcd backup, you can follow this procedure to restore your cluster to a previous state.

If applicable, you might also need to [recover from expired control plane certificates](#).



WARNING

Restoring to a previous cluster state is a destructive and destabilizing action to take on a running cluster. This procedure should only be used as a last resort.

Prior to performing a restore, see [About restoring cluster state](#) for more information on the impact to the cluster.



NOTE

If you have a majority of your masters still available and have an etcd quorum, then follow the procedure to [replace a single unhealthy etcd member](#).

Recovering from expired control plane certificates

This solution handles situations where your control plane certificates have expired. For example, if you shut down your cluster before the first certificate rotation, which occurs 24 hours after installation, your certificates will not be rotated and will expire. You can follow this procedure to recover from expired control plane certificates.

5.4.2. Restoring to a previous cluster state

To restore the cluster to a previous state, you must have previously [backed up etcd data](#) by creating a snapshot. You will use this snapshot to restore the cluster state.

5.4.2.1. About restoring cluster state

You can use an etcd backup to restore your cluster to a previous state. This can be used to recover from the following situations:

- The cluster has lost the majority of control plane hosts (quorum loss).
- An administrator has deleted something critical and must restore to recover the cluster.

**WARNING**

Restoring to a previous cluster state is a destructive and destabilizing action to take on a running cluster. This should only be used as a last resort.

If you are able to retrieve data using the Kubernetes API server, then etcd is available and you should not restore using an etcd backup.

Restoring etcd effectively takes a cluster back in time and all clients will experience a conflicting, parallel history. This can impact the behavior of watching components like kubelets, Kubernetes controller managers, SDN controllers, and persistent volume controllers.

It can cause Operator churn when the content in etcd does not match the actual content on disk, causing Operators for the Kubernetes API server, Kubernetes controller manager, Kubernetes scheduler, and etcd to get stuck when files on disk conflict with content in etcd. This can require manual actions to resolve the issues.

In extreme cases, the cluster can lose track of persistent volumes, delete critical workloads that no longer exist, reimagine machines, and rewrite CA bundles with expired certificates.

5.4.2.2. Restoring to a previous cluster state

You can use a saved etcd backup to restore a previous cluster state or restore a cluster that has lost the majority of control plane hosts.

**NOTE**

If your cluster uses a control plane machine set, see "Troubleshooting the control plane machine set" for a more simple etcd recovery procedure.

**IMPORTANT**

When you restore your cluster, you must use an etcd backup that was taken from the same z-stream release. For example, an OpenShift Container Platform 4.7.2 cluster must use an etcd backup that was taken from 4.7.2.

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role through a certificate-based **kubeconfig** file, like the one that was used during installation.
- A healthy control plane host to use as the recovery host.
- SSH access to control plane hosts.
- A backup directory containing both the etcd snapshot and the resources for the static pods, which were from the same backup. The file names in the directory must be in the following formats: **snapshot_<timestamp>.db** and **static_kuberresources_<timestamp>.tar.gz**.




IMPORTANT

For non-recovery control plane nodes, it is not required to establish SSH connectivity or to stop the static pods. You can delete and recreate other non-recovery, control plane machines, one by one.

Procedure

1. Select a control plane host to use as the recovery host. This is the host that you will run the restore operation on.
2. Establish SSH connectivity to each of the control plane nodes, including the recovery host. The Kubernetes API server becomes inaccessible after the restore process starts, so you cannot access the control plane nodes. For this reason, it is recommended to establish SSH connectivity to each control plane host in a separate terminal.



IMPORTANT

If you do not complete this step, you will not be able to access the control plane hosts to complete the restore procedure, and you will be unable to recover your cluster from this state.

3. Copy the etcd backup directory to the recovery control plane host. This procedure assumes that you copied the **backup** directory containing the etcd snapshot and the resources for the static pods to the **/home/core/** directory of your recovery control plane host.
4. Stop the static pods on any other control plane nodes.



NOTE

You do not need to stop the static pods on the recovery host.

- a. Access a control plane host that is not the recovery host.
- b. Move the existing etcd pod file out of the kubelet manifest directory:

```
$ sudo mv -v /etc/kubernetes/manifests/etcd-pod.yaml /tmp
```

- c. Verify that the etcd pods are stopped.

```
$ sudo crictl ps | grep etcd | egrep -v "operator|etcd-guard"
```

The output of this command should be empty. If it is not empty, wait a few minutes and check again.

- d. Move the existing Kubernetes API server pod file out of the kubelet manifest directory:

```
$ sudo mv -v /etc/kubernetes/manifests/kube-apiserver-pod.yaml /tmp
```

- e. Verify that the Kubernetes API server pods are stopped.

```
$ sudo crictl ps | grep kube-apiserver | egrep -v "operator|guard"
```

The output of this command should be empty. If it is not empty, wait a few minutes and check again.

- f. Move the etcd data directory to a different location:

```
$ sudo mv -v /var/lib/etcd/ /tmp
```

- g. If the `/etc/kubernetes/manifests/keepalived.yaml` file exists and the node is deleted, follow these steps:

- i. Move the `/etc/kubernetes/manifests/keepalived.yaml` file out of the kubelet manifest directory:

```
$ sudo mv -v /etc/kubernetes/manifests/keepalived.yaml /tmp
```

- ii. Verify that any containers managed by the **keepalived** daemon are stopped:

```
$ sudo crictl ps --name keepalived
```

The output of this command should be empty. If it is not empty, wait a few minutes and check again.

- iii. Check if the control plane has any Virtual IPs (VIPs) assigned to it:

```
$ ip -o address | egrep '<api_vip>|<ingress_vip>'
```

- iv. For each reported VIP, run the following command to remove it:

```
$ sudo ip address del <reported_vip> dev <reported_vip_device>
```

- h. Repeat this step on each of the other control plane hosts that is not the recovery host.

5. Access the recovery control plane host.

6. If the **keepalived** daemon is in use, verify that the recovery control plane node owns the VIP:

```
$ ip -o address | grep <api_vip>
```

The address of the VIP is highlighted in the output if it exists. This command returns an empty string if the VIP is not set or configured incorrectly.

7. If the cluster-wide proxy is enabled, be sure that you have exported the **NO_PROXY**, **HTTP_PROXY**, and **HTTPS_PROXY** environment variables.

TIP

You can check whether the proxy is enabled by reviewing the output of `oc get proxy cluster -o yaml`. The proxy is enabled if the `httpProxy`, `httpsProxy`, and `noProxy` fields have values set.

8. Run the restore script on the recovery control plane host and pass in the path to the etcd backup directory:


```
$ sudo -E /usr/local/bin/cluster-restore.sh /home/core/assets/backup
```


Example script output

```

...stopping kube-scheduler-pod.yaml
...stopping kube-controller-manager-pod.yaml
...stopping etcd-pod.yaml
...stopping kube-apiserver-pod.yaml
Waiting for container etcd to stop
.complete
Waiting for container etcdctl to stop
.....complete
Waiting for container etcd-metrics to stop
complete
Waiting for container kube-controller-manager to stop
complete
Waiting for container kube-apiserver to stop
.....complete
Waiting for container kube-scheduler to stop
complete
Moving etcd data-dir /var/lib/etcd/member to /var/lib/etcd-backup
starting restore-etcd static pod
starting kube-apiserver-pod.yaml
static-pod-resources/kube-apiserver-pod-7/kube-apiserver-pod.yaml
starting kube-controller-manager-pod.yaml
static-pod-resources/kube-controller-manager-pod-7/kube-controller-manager-pod.yaml
starting kube-scheduler-pod.yaml
static-pod-resources/kube-scheduler-pod-8/kube-scheduler-pod.yaml

```



NOTE

The restore process can cause nodes to enter the **NotReady** state if the node certificates were updated after the last etcd backup.

9. Check the nodes to ensure they are in the **Ready** state.
 - a. Run the following command:

```
$ oc get nodes -w
```

Sample output

```

NAME                STATUS ROLES    AGE   VERSION
host-172-25-75-28   Ready  master    3d20h v1.25.0
host-172-25-75-38   Ready  infra,worker 3d20h v1.25.0
host-172-25-75-40   Ready  master    3d20h v1.25.0
host-172-25-75-65   Ready  master    3d20h v1.25.0
host-172-25-75-74   Ready  infra,worker 3d20h v1.25.0
host-172-25-75-79   Ready  worker    3d20h v1.25.0
host-172-25-75-86   Ready  worker    3d20h v1.25.0
host-172-25-75-98   Ready  infra,worker 3d20h v1.25.0

```

It can take several minutes for all nodes to report their state.

- b. If any nodes are in the **NotReady** state, log in to the nodes and remove all of the PEM files from the `/var/lib/kubelet/pki` directory on each node. You can SSH into the nodes or use the terminal window in the web console.

```
$ ssh -i <ssh-key-path> core@<master-hostname>
```

Sample pki directory

```
sh-4.4# pwd
/var/lib/kubelet/pki
sh-4.4# ls
kubelet-client-2022-04-28-11-24-09.pem  kubelet-server-2022-04-28-11-24-15.pem
kubelet-client-current.pem            kubelet-server-current.pem
```


10. Restart the kubelet service on all control plane hosts.

- a. From the recovery host, run the following command:

```
$ sudo systemctl restart kubelet.service
```

- b. Repeat this step on all other control plane hosts.

11. Approve the pending CSRs:



NOTE

Clusters with no worker nodes, such as single-node clusters or clusters consisting of three schedulable control plane nodes, will not have any pending CSRs to approve. You can skip all the commands listed in this step.

- a. Get the list of current CSRs:

```
$ oc get csr
```

Example output

```
NAME      AGE  SIGNERNAME                                REQUESTOR
CONDITION
csr-2s94x  8m3s  kubernetes.io/kubelet-serving            system:node:<node_name>
Pending  1
csr-4bd6t  8m3s  kubernetes.io/kubelet-serving            system:node:<node_name>
Pending  2
csr-4hl85  13m   kubernetes.io/kube-apiserver-client-kubelet
system:serviceaccount:openshift-machine-config-operator:node-bootstrapper  Pending
3
csr-zhthp  3m8s  kubernetes.io/kube-apiserver-client-kubelet
system:serviceaccount:openshift-machine-config-operator:node-bootstrapper  Pending
4
...
```

1 **2** A pending kubelet service CSR (for user-provisioned installations).

3 **4** A pending **node-bootstrapper** CSR.

- b. Review the details of a CSR to verify that it is valid:

```
$ oc describe csr <csr_name> 1
```

1 <csr_name> is the name of a CSR from the list of current CSRs.

- c. Approve each valid **node-bootstrapper** CSR:

```
$ oc adm certificate approve <csr_name>
```

- d. For user-provisioned installations, approve each valid kubelet service CSR:

```
$ oc adm certificate approve <csr_name>
```

12. Verify that the single member control plane has started successfully.

- a. From the recovery host, verify that the etcd container is running.

```
$ sudo crictl ps | grep etcd | egrep -v "operator|etcd-guard"
```

Example output

```
3ad41b7908e32
36f86e2eeaaaffe662df0d21041eb22b8198e0e58abeeae8c743c3e6e977e8009
About a minute ago   Running           etcd                0
7c05f8af362f0
```

- b. From the recovery host, verify that the etcd pod is running.

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

Example output

NAME	READY	STATUS	RESTARTS	AGE
etcd-ip-10-0-143-125.ec2.internal	1/1	Running	1	2m47s

If the status is **Pending**, or the output lists more than one running etcd pod, wait a few minutes and check again.

13. If you are using the **OVNKubernetes** network plugin, delete the node objects that are associated with control plane hosts that are not the recovery control plane host.

```
$ oc delete node <non-recovery-controlplane-host-1> <non-recovery-controlplane-host-2>
```

14. Verify that the Cluster Network Operator (CNO) redeploys the OVN-Kubernetes control plane and that it no longer references the non-recovery controller IP addresses. To verify this result, regularly check the output of the following command. Wait until it returns an empty result before you proceed to restart the Open Virtual Network (OVN) Kubernetes pods on all of the hosts in the next step.

```
$ oc -n openshift-ovn-kubernetes get ds/ovnkube-master -o yaml | grep -E '<non-recovery_controller_ip_1>|<non-recovery_controller_ip_2>'
```

**NOTE**

It can take at least 5-10 minutes for the OVN-Kubernetes control plane to be redeployed and the previous command to return empty output.

15. If you are using the OVN-Kubernetes network plugin, restart the Open Virtual Network (OVN) Kubernetes pods on all of the hosts.

**NOTE**

Validating and mutating admission webhooks can reject pods. If you add any additional webhooks with the **failurePolicy** set to **Fail**, then they can reject pods and the restoration process can fail. You can avoid this by saving and deleting webhooks while restoring the cluster state. After the cluster state is restored successfully, you can enable the webhooks again.

Alternatively, you can temporarily set the **failurePolicy** to **Ignore** while restoring the cluster state. After the cluster state is restored successfully, you can set the **failurePolicy** to **Fail**.

- a. Remove the northbound database (nbdb) and southbound database (sbdb). Access the recovery host and the remaining control plane nodes by using Secure Shell (SSH) and run the following command:

```
$ sudo rm -f /var/lib/ovn/etc/*.db
```

- b. Delete all OVN-Kubernetes control plane pods by running the following command:

```
$ oc delete pods -l app=ovnkube-master -n openshift-ovn-kubernetes
```

- c. Ensure that any OVN-Kubernetes control plane pods are deployed again and are in a **Running** state by running the following command:

```
$ oc get pods -l app=ovnkube-master -n openshift-ovn-kubernetes
```

Example output

```
NAME                READY STATUS RESTARTS AGE
ovnkube-master-nb24h 4/4   Running 0      48s
```

- d. Delete all **ovnkube-node** pods by running the following command:

```
$ oc get pods -n openshift-ovn-kubernetes -o name | grep ovnkube-node | while read p ;
do oc delete $p -n openshift-ovn-kubernetes ; done
```

- e. Ensure that all the **ovnkube-node** pods are deployed again and are in a **Running** state by running the following command:

```
$ oc get pods -n openshift-ovn-kubernetes | grep ovnkube-node
```

16. Delete and re-create other non-recovery, control plane machines, one by one. After the machines are re-created, a new revision is forced and etcd automatically scales up.

- If you use a user-provisioned bare metal installation, you can re-create a control plane machine by using the same method that you used to originally create it. For more information, see "Installing a user-provisioned cluster on bare metal".



WARNING

Do not delete and re-create the machine for the recovery host.

- If you are running installer-provisioned infrastructure, or you used the Machine API to create your machines, follow these steps:



WARNING

Do not delete and re-create the machine for the recovery host.

For bare metal installations on installer-provisioned infrastructure, control plane machines are not re-created. For more information, see "Replacing a bare-metal control plane node".

- a. Obtain the machine for one of the lost control plane hosts.

In a terminal that has access to the cluster as a cluster-admin user, run the following command:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output:

```

NAME                               PHASE  TYPE      REGION  ZONE  AGE
NODE                               PROVIDERID  STATE
clustername-8qw5l-master-0        Running m4.xlarge us-east-1 us-east-1a
3h37m ip-10-0-131-183.ec2.internal  aws:///us-east-1a/i-0ec2782f8287dfb7e
stopped 1
clustername-8qw5l-master-1        Running m4.xlarge us-east-1 us-east-1b
3h37m ip-10-0-143-125.ec2.internal  aws:///us-east-1b/i-096c349b700a19631
running
clustername-8qw5l-master-2        Running m4.xlarge us-east-1 us-east-1c
3h37m ip-10-0-154-194.ec2.internal  aws:///us-east-1c/i-02626f1dba9ed5bba
running
clustername-8qw5l-worker-us-east-1a-wbtgd Running m4.large us-east-1 us-
east-1a 3h28m ip-10-0-129-226.ec2.internal  aws:///us-east-1a/i-
010ef6279b4662ced running
clustername-8qw5l-worker-us-east-1b-lrdxb Running m4.large us-east-1 us-

```

```
east-1b 3h28m ip-10-0-144-248.ec2.internal aws:///us-east-1b/i-
0cb45ac45a166173b running
clustername-8qw5l-worker-us-east-1c-pkg26 Running m4.large us-east-1 us-
east-1c 3h28m ip-10-0-170-181.ec2.internal aws:///us-east-1c/i-
06861c00007751b0a running
```

- 1 This is the control plane machine for the lost control plane host, **ip-10-0-131-183.ec2.internal**.

- b. Save the machine configuration to a file on your file system:

```
$ oc get machine clustername-8qw5l-master-0 \ 1
-n openshift-machine-api \
-o yaml \
> new-master-machine.yaml
```

- 1 Specify the name of the control plane machine for the lost control plane host.

- c. Edit the **new-master-machine.yaml** file that was created in the previous step to assign a new name and remove unnecessary fields.

- i. Remove the entire **status** section:

```
status:
  addresses:
    - address: 10.0.131.183
      type: InternalIP
    - address: ip-10-0-131-183.ec2.internal
      type: InternalDNS
    - address: ip-10-0-131-183.ec2.internal
      type: Hostname
  lastUpdated: "2020-04-20T17:44:29Z"
  nodeRef:
    kind: Node
    name: ip-10-0-131-183.ec2.internal
    uid: acca4411-af0d-4387-b73e-52b2484295ad
  phase: Running
  providerStatus:
    apiVersion: awsproviderconfig.openshift.io/v1beta1
    conditions:
      - lastProbeTime: "2020-04-20T16:53:50Z"
        lastTransitionTime: "2020-04-20T16:53:50Z"
        message: machine successfully created
        reason: MachineCreationSucceeded
        status: "True"
        type: MachineCreation
    instanceId: i-0fdb85790d76d0c3f
    instanceState: stopped
    kind: AWSMachineProviderStatus
```

- ii. Change the **metadata.name** field to a new name.

It is recommended to keep the same base name as the old machine and change the ending number to the next available number. In this example, **clustername-8qw5l-master-0** is changed to **clustername-8qw5l-master-3**:

```

apiVersion: machine.openshift.io/v1beta1
kind: Machine
metadata:
  ...
  name: clustername-8qw5l-master-3
  ...

```

- iii. Remove the **spec.providerID** field:

```

providerID: aws:///us-east-1a/i-0fdb85790d76d0c3f

```

- iv. Remove the **metadata.annotations** and **metadata.generation** fields:

```

annotations:
  machine.openshift.io/instance-state: running
  ...
generation: 2

```

- v. Remove the **metadata.resourceVersion** and **metadata.uid** fields:

```

resourceVersion: "13291"
uid: a282eb70-40a2-4e89-8009-d05dd420d31a

```

- d. Delete the machine of the lost control plane host:

```

$ oc delete machine -n openshift-machine-api clustername-8qw5l-master-0 1

```

- 1** Specify the name of the control plane machine for the lost control plane host.

- e. Verify that the machine was deleted:

```

$ oc get machines -n openshift-machine-api -o wide

```

Example output:

```

NAME                                PHASE  TYPE      REGION  ZONE  AGE
NODE                                PROVIDERID          STATE
clustername-8qw5l-master-1          Running m4.xlarge us-east-1 us-east-1b
3h37m ip-10-0-143-125.ec2.internal aws:///us-east-1b/i-096c349b700a19631
running
clustername-8qw5l-master-2          Running m4.xlarge us-east-1 us-east-1c
3h37m ip-10-0-154-194.ec2.internal aws:///us-east-1c/i-02626f1dba9ed5bba
running
clustername-8qw5l-worker-us-east-1a-wbtgd Running m4.large  us-east-1 us-
east-1a 3h28m ip-10-0-129-226.ec2.internal aws:///us-east-1a/i-
010ef6279b4662ced running
clustername-8qw5l-worker-us-east-1b-lrdxb Running m4.large  us-east-1 us-
east-1b 3h28m ip-10-0-144-248.ec2.internal aws:///us-east-1b/i-
0cb45ac45a166173b running
clustername-8qw5l-worker-us-east-1c-pkg26 Running m4.large  us-east-1 us-
east-1c 3h28m ip-10-0-170-181.ec2.internal aws:///us-east-1c/i-
06861c00007751b0a running

```

- f. Create a machine by using the **new-master-machine.yaml** file:

```
$ oc apply -f new-master-machine.yaml
```

- g. Verify that the new machine has been created:

```
$ oc get machines -n openshift-machine-api -o wide
```

Example output:

```

NAME                                PHASE    TYPE    REGION  ZONE
AGE  NODE                                PROVIDERID  STATE
clustername-8qw5l-master-1          Running   m4.xlarge us-east-1 us-east-
1b 3h37m ip-10-0-143-125.ec2.internal  aws:///us-east-1b/i-096c349b700a19631
running
clustername-8qw5l-master-2          Running   m4.xlarge us-east-1 us-east-
1c 3h37m ip-10-0-154-194.ec2.internal  aws:///us-east-1c/i-02626f1dba9ed5bba
running
clustername-8qw5l-master-3          Provisioning m4.xlarge us-east-1 us-east-
1a 85s ip-10-0-173-171.ec2.internal  aws:///us-east-1a/i-015b0888fe17bc2c8
running 1
clustername-8qw5l-worker-us-east-1a-wbtgd Running   m4.large  us-east-1
us-east-1a 3h28m ip-10-0-129-226.ec2.internal  aws:///us-east-1a/i-
010ef6279b4662ced running
clustername-8qw5l-worker-us-east-1b-lrdxb Running   m4.large  us-east-1 us-
east-1b 3h28m ip-10-0-144-248.ec2.internal  aws:///us-east-1b/i-
0cb45ac45a166173b running
clustername-8qw5l-worker-us-east-1c-pkg26 Running   m4.large  us-east-1
us-east-1c 3h28m ip-10-0-170-181.ec2.internal  aws:///us-east-1c/i-
06861c00007751b0a running

```

- 1** The new machine, **clustername-8qw5l-master-3** is being created and is ready after the phase changes from **Provisioning** to **Running**.

It might take a few minutes for the new machine to be created. The etcd cluster Operator will automatically sync when the machine or node returns to a healthy state.

- h. Repeat these steps for each lost control plane host that is not the recovery host.
17. Turn off the quorum guard by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{"spec": {"unsupportedConfigOverrides": {"useUnsupportedUnsafeNonHANonProductionUnstableEtcd": true}}}'
```

This command ensures that you can successfully re-create secrets and roll out the static pods.

18. In a separate terminal window within the recovery host, export the recovery **kubeconfig** file by running the following command:

```
$ export KUBECONFIG=/etc/kubernetes/static-pod-resources/kube-apiserver-certs/secrets/node-kubeconfigs/localhost-recovery.kubeconfig
```

19. Force etcd redeployment.

In the same terminal window where you exported the recovery **kubeconfig** file, run the following command:

```
$ oc patch etcd cluster -p='{ "spec": { "forceRedeploymentReason": "recovery-"$( date --rfc-3339=ns )"' --type=merge 1
```

- 1 The **forceRedeploymentReason** value must be unique, which is why a timestamp is appended.

When the etcd cluster Operator performs a redeployment, the existing nodes are started with new pods similar to the initial bootstrap scale up.

20. Turn the quorum guard back on by entering the following command:

```
$ oc patch etcd/cluster --type=merge -p '{ "spec": { "unsupportedConfigOverrides": null } }
```

21. You can verify that the **unsupportedConfigOverrides** section is removed from the object by entering this command:

```
$ oc get etcd/cluster -oyaml
```

22. Verify all nodes are updated to the latest revision.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc get etcd -o=jsonpath='{range .items[0].status.conditions[? (@.type=="NodeInstallerProgressing")]}{.reason}{ "\n" }{.message}{ "\n" }'
```

Review the **NodeInstallerProgressing** status condition for etcd to verify that all nodes are at the latest revision. The output shows **AllNodesAtLatestRevision** upon successful update:

```
AllNodesAtLatestRevision
3 nodes are at revision 7 1
```

- 1 In this example, the latest revision number is **7**.

If the output includes multiple revision numbers, such as **2 nodes are at revision 6; 1 nodes are at revision 7**, this means that the update is still in progress. Wait a few minutes and try again.

23. After etcd is redeployed, force new rollouts for the control plane. The Kubernetes API server will reinstall itself on the other nodes because the kubelet is connected to API servers using an internal load balancer.

In a terminal that has access to the cluster as a **cluster-admin** user, run the following commands.

- a. Force a new rollout for the Kubernetes API server:

```
$ oc patch kubeapiserver cluster -p='{ "spec": { "forceRedeploymentReason": "recovery-"$( date --rfc-3339=ns )"' --type=merge
```

Verify all nodes are updated to the latest revision.

```
$ oc get kubeapiserver -o=jsonpath='{range .items[0].status.conditions[?(@.type=="NodeInstallerProgressing")]}{.reason}{"\n"}{.message}{"\n"}'
```

Review the **NodeInstallerProgressing** status condition to verify that all nodes are at the latest revision. The output shows **AllNodesAtLatestRevision** upon successful update:

```
AllNodesAtLatestRevision
3 nodes are at revision 7 1
```

1 In this example, the latest revision number is **7**.

If the output includes multiple revision numbers, such as **2 nodes are at revision 6; 1 nodes are at revision 7**, this means that the update is still in progress. Wait a few minutes and try again.

b. Force a new rollout for the Kubernetes controller manager:

```
$ oc patch kubecontrollermanager cluster -p='{ "spec": { "forceRedeploymentReason": "recovery-"$( date --rfc-3339=ns )"' }}' --type=merge
```

Verify all nodes are updated to the latest revision.

```
$ oc get kubecontrollermanager -o=jsonpath='{range .items[0].status.conditions[?(@.type=="NodeInstallerProgressing")]}{.reason}{"\n"}{.message}{"\n"}'
```

Review the **NodeInstallerProgressing** status condition to verify that all nodes are at the latest revision. The output shows **AllNodesAtLatestRevision** upon successful update:

```
AllNodesAtLatestRevision
3 nodes are at revision 7 1
```

1 In this example, the latest revision number is **7**.

If the output includes multiple revision numbers, such as **2 nodes are at revision 6; 1 nodes are at revision 7**, this means that the update is still in progress. Wait a few minutes and try again.

c. Force a new rollout for the Kubernetes scheduler:

```
$ oc patch kubescheduler cluster -p='{ "spec": { "forceRedeploymentReason": "recovery-"$( date --rfc-3339=ns )"' }}' --type=merge
```

Verify all nodes are updated to the latest revision.

```
$ oc get kubescheduler -o=jsonpath='{range .items[0].status.conditions[?(@.type=="NodeInstallerProgressing")]}{.reason}{"\n"}{.message}{"\n"}'
```

Review the **NodeInstallerProgressing** status condition to verify that all nodes are at the latest revision. The output shows **AllNodesAtLatestRevision** upon successful update:

```
AllNodesAtLatestRevision
3 nodes are at revision 7 1
```

1 In this example, the latest revision number is **7**.

If the output includes multiple revision numbers, such as **2 nodes are at revision 6; 1 nodes are at revision 7**, this means that the update is still in progress. Wait a few minutes and try again.

24. Verify that all control plane hosts have started and joined the cluster.
In a terminal that has access to the cluster as a **cluster-admin** user, run the following command:

```
$ oc -n openshift-etcd get pods -l k8s-app=etcd
```

Example output

```
etcd-ip-10-0-143-125.ec2.internal      2/2   Running   0    9h
etcd-ip-10-0-154-194.ec2.internal    2/2   Running   0    9h
etcd-ip-10-0-173-171.ec2.internal    2/2   Running   0    9h
```

To ensure that all workloads return to normal operation following a recovery procedure, restart each pod that stores Kubernetes API information. This includes OpenShift Container Platform components such as routers, Operators, and third-party components.

NOTE

On completion of the previous procedural steps, you might need to wait a few minutes for all services to return to their restored state. For example, authentication by using **oc login** might not immediately work until the OAuth server pods are restarted.

Consider using the **system:admin kubeconfig** file for immediate authentication. This method bases its authentication on SSL/TLS client certificates as against OAuth tokens. You can authenticate with this file by issuing the following command:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig
```

Issue the following command to display your authenticated user name:

```
$ oc whoami
```

5.4.2.3. Additional resources

- [Installing a user-provisioned cluster on bare metal](#)
- [Creating a bastion host to access OpenShift Container Platform instances and the control plane nodes with SSH](#)
- [Replacing a bare-metal control plane node](#)

5.4.2.4. Issues and workarounds for restoring a persistent storage state

If your OpenShift Container Platform cluster uses persistent storage of any form, a state of the cluster

is typically stored outside etcd. It might be an Elasticsearch cluster running in a pod or a database running in a **StatefulSet** object. When you restore from an etcd backup, the status of the workloads in OpenShift Container Platform is also restored. However, if the etcd snapshot is old, the status might be invalid or outdated.



IMPORTANT

The contents of persistent volumes (PVs) are never part of the etcd snapshot. When you restore an OpenShift Container Platform cluster from an etcd snapshot, non-critical workloads might gain access to critical data, or vice-versa.

The following are some example scenarios that produce an out-of-date status:

- MySQL database is running in a pod backed up by a PV object. Restoring OpenShift Container Platform from an etcd snapshot does not bring back the volume on the storage provider, and does not produce a running MySQL pod, despite the pod repeatedly attempting to start. You must manually restore this pod by restoring the volume on the storage provider, and then editing the PV to point to the new volume.
- Pod P1 is using volume A, which is attached to node X. If the etcd snapshot is taken while another pod uses the same volume on node Y, then when the etcd restore is performed, pod P1 might not be able to start correctly due to the volume still being attached to node Y. OpenShift Container Platform is not aware of the attachment, and does not automatically detach it. When this occurs, the volume must be manually detached from node Y so that the volume can attach on node X, and then pod P1 can start.
- Cloud provider or storage provider credentials were updated after the etcd snapshot was taken. This causes any CSI drivers or Operators that depend on the those credentials to not work. You might have to manually update the credentials required by those drivers or Operators.
- A device is removed or renamed from OpenShift Container Platform nodes after the etcd snapshot is taken. The Local Storage Operator creates symlinks for each PV that it manages from **/dev/disk/by-id** or **/dev** directories. This situation might cause the local PVs to refer to devices that no longer exist.

To fix this problem, an administrator must:

1. Manually remove the PVs with invalid devices.
2. Remove symlinks from respective nodes.
3. Delete **LocalVolume** or **LocalVolumeSet** objects (see *Storage → Configuring persistent storage → Persistent storage using local volumes → Deleting the Local Storage Operator Resources*).

5.4.3. Recovering from expired control plane certificates

5.4.3.1. Recovering from expired control plane certificates

The cluster can automatically recover from expired control plane certificates.

However, you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. For user-provisioned installations, you might also need to approve pending kubelet serving CSRs.

Use the following steps to approve the pending CSRs:

Procedure

1. Get the list of current CSRs:

```
$ oc get csr
```

Example output

```
NAME      AGE  SIGNERNAME                                REQUESTOR
CONDITION
csr-2s94x 8m3s kubernetes.io/kubelet-serving             system:node:<node_name>
Pending 1
csr-4bd6t 8m3s kubernetes.io/kubelet-serving             system:node:<node_name>
Pending
csr-4hl85 13m  kubernetes.io/kube-apiserver-client-kubelet
system:serviceaccount:openshift-machine-config-operator:node-bootstrapper Pending 2
csr-zhhhp 3m8s kubernetes.io/kube-apiserver-client-kubelet
system:serviceaccount:openshift-machine-config-operator:node-bootstrapper Pending
...
```

1 A pending kubelet service CSR (for user-provisioned installations).

2 A pending **node-bootstrapper** CSR.

2. Review the details of a CSR to verify that it is valid:

```
$ oc describe csr <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

3. Approve each valid **node-bootstrapper** CSR:

```
$ oc adm certificate approve <csr_name>
```

4. For user-provisioned installations, approve each valid kubelet serving CSR:

```
$ oc adm certificate approve <csr_name>
```

5.4.4. Disaster recovery for a hosted cluster within an AWS region

In a situation where you need disaster recovery (DR) for a hosted cluster, you can recover a hosted cluster to the same region within AWS. For example, you need DR when the upgrade of a management cluster fails and the hosted cluster is in a read-only state.



IMPORTANT

Hosted control planes is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

The DR process involves three main steps:

1. Backing up the hosted cluster on the source management cluster
2. Restoring the hosted cluster on a destination management cluster
3. Deleting the hosted cluster from the source management cluster

Your workloads remain running during the process. The Cluster API might be unavailable for a period, but that will not affect the services that are running on the worker nodes.



IMPORTANT

Both the source management cluster and the destination management cluster must have the **--external-dns** flags to maintain the API server URL, as shown in this example:

Example: External DNS flags

```
--external-dns-provider=aws \  
--external-dns-credentials=<AWS Credentials location> \  
--external-dns-domain-filter=<DNS Base Domain>
```

That way, the server URL ends with <https://api-sample-hosted.sample-hosted.aws.openshift.com>.

If you do not include the **--external-dns** flags to maintain the API server URL, the hosted cluster cannot be migrated.

5.4.4.1. Example environment and context

Consider an scenario where you have three clusters to restore. Two are management clusters, and one is a hosted cluster. You can restore either the control plane only or the control plane and the nodes. Before you begin, you need the following information:

- Source MGMT Namespace: The source management namespace
- Source MGMT ClusterName: The source management cluster name
- Source MGMT Kubeconfig: The source management **kubeconfig** file
- Destination MGMT Kubeconfig: The destination management **kubeconfig** file
- HC Kubeconfig: The hosted cluster **kubeconfig** file

- SSH key file: The SSH public key
- Pull secret: The pull secret file to access the release images
- AWS credentials
- AWS region
- Base domain: The DNS base domain to use as an external DNS
- S3 bucket name: The bucket in the AWS region where you plan to upload the etcd backup

This information is shown in the following example environment variables.

Example environment variables

```
SSH_KEY_FILE=${HOME}/.ssh/id_rsa.pub
BASE_PATH=${HOME}/hypershift
BASE_DOMAIN="aws.sample.com"
PULL_SECRET_FILE=${HOME}/pull_secret.json
AWS_CREDS=${HOME}/.aws/credentials"
AWS_ZONE_ID="Z02718293M33QHDEQBROL"

CONTROL_PLANE_AVAILABILITY_POLICY=SingleReplica
HYPERSHIFT_PATH=${BASE_PATH}/src/hypershift
HYPERSHIFT_CLI=${HYPERSHIFT_PATH}/bin/hypershift
HYPERSHIFT_IMAGE=${HYPERSHIFT_IMAGE:-"quay.io/${USER}/hypershift:latest"}
NODE_POOL_REPLICAS=${NODE_POOL_REPLICAS:-2}

# MGMT Context
MGMT_REGION=us-west-1
MGMT_CLUSTER_NAME="${USER}-dev"
MGMT_CLUSTER_NS=${USER}
MGMT_CLUSTER_DIR="${BASE_PATH}/hosted_clusters/${MGMT_CLUSTER_NS}-
${MGMT_CLUSTER_NAME}"
MGMT_KUBECONFIG="${MGMT_CLUSTER_DIR}/kubeconfig"

# MGMT2 Context
MGMT2_CLUSTER_NAME="${USER}-dest"
MGMT2_CLUSTER_NS=${USER}
MGMT2_CLUSTER_DIR="${BASE_PATH}/hosted_clusters/${MGMT2_CLUSTER_NS}-
${MGMT2_CLUSTER_NAME}"
MGMT2_KUBECONFIG="${MGMT2_CLUSTER_DIR}/kubeconfig"

# Hosted Cluster Context
HC_CLUSTER_NS=clusters
HC_REGION=us-west-1
HC_CLUSTER_NAME="${USER}-hosted"
HC_CLUSTER_DIR="${BASE_PATH}/hosted_clusters/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}"
HC_KUBECONFIG="${HC_CLUSTER_DIR}/kubeconfig"
BACKUP_DIR=${HC_CLUSTER_DIR}/backup

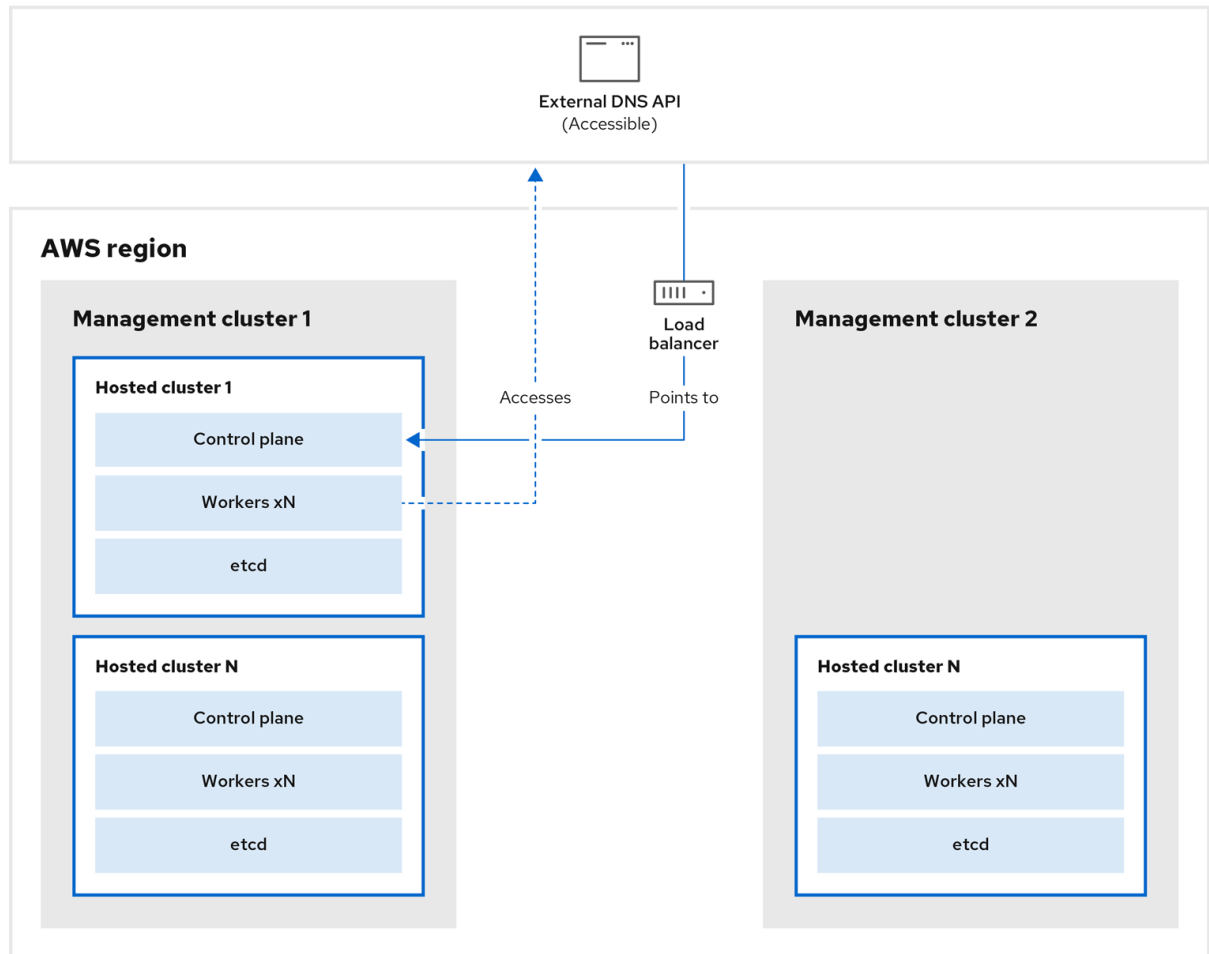
BUCKET_NAME="${USER}-hosted-${MGMT_REGION}"
```

```
# DNS
AWS_ZONE_ID="Z07342811SH9AA102K1AC"
EXTERNAL_DNS_DOMAIN="hc.jpdv.aws.kerbeross.com"
```

5.4.4.2. Overview of the backup and restore process


The backup and restore process works as follows:

1. On management cluster 1, which you can think of as the source management cluster, the control plane and workers interact by using the external DNS API. The external DNS API is accessible, and a load balancer sits between the management clusters.

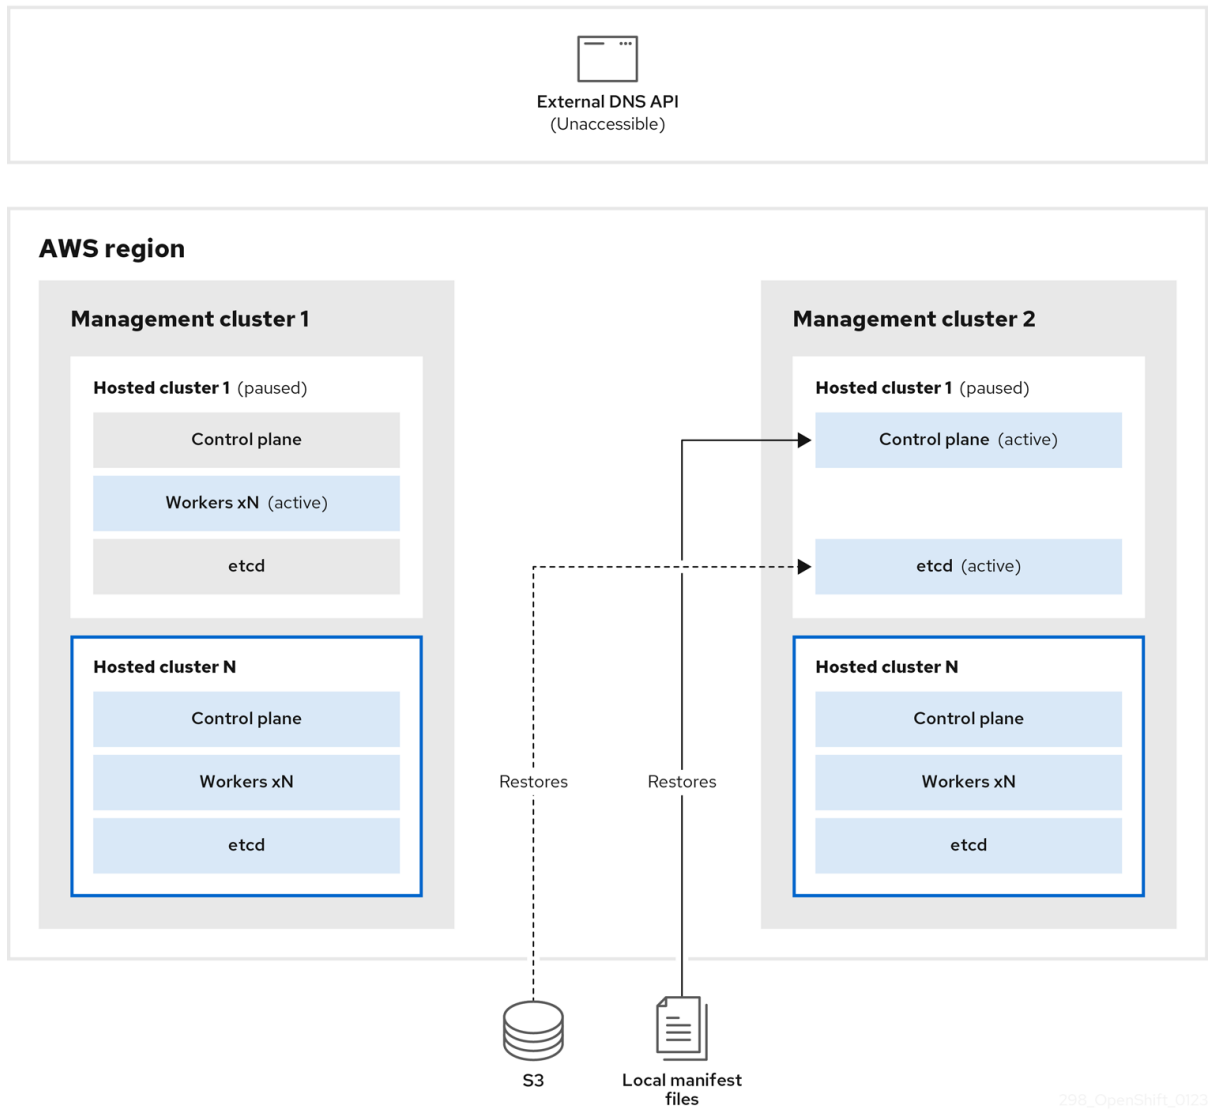


298_OpenShift_0123

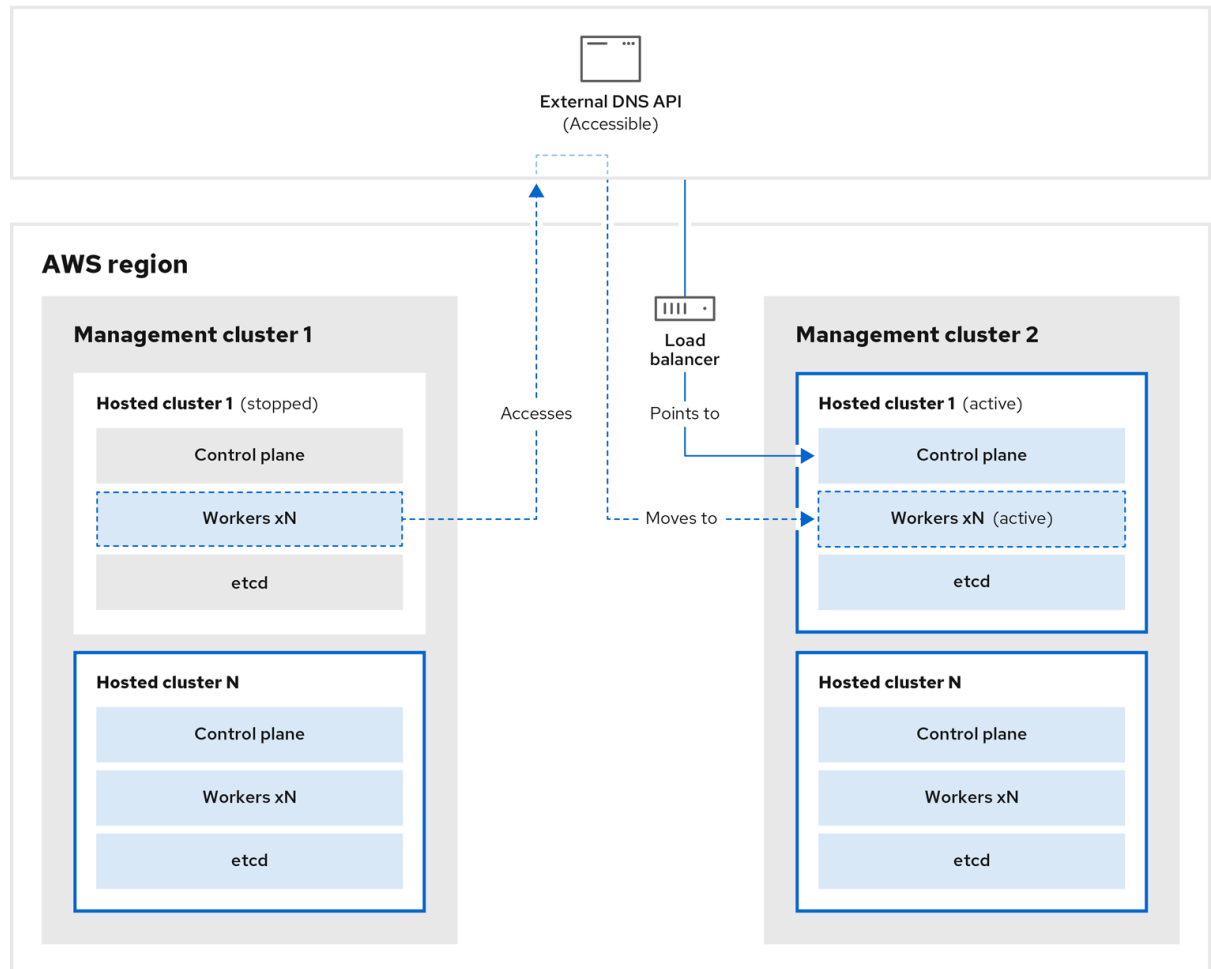
2. You take a snapshot of the hosted cluster, which includes etcd, the control plane, and the worker nodes. During this process, the worker nodes continue to try to access the external DNS API even if it is not accessible, the workloads are running, the control plane is saved in a local manifest file, and etcd is backed up to an S3 bucket. The data plane is active and the control plane is paused.



- On management cluster 2, which you can think of as the destination management cluster, you restore etcd from the S3 bucket and restore the control plane from the local manifest file. During this process, the external DNS API is stopped, the hosted cluster API becomes inaccessible, and any workers that use the API are unable to update their manifest files, but the workloads are still running.




4. The external DNS API is accessible again, and the worker nodes use it to move to management cluster 2. The external DNS API can access the load balancer that points to the control plane.



298_OpenShift_0123

- On management cluster 2, the control plane and worker nodes interact by using the external DNS API. The resources are deleted from management cluster 1, except for the S3 backup of etcd. If you try to set up the hosted cluster again on management cluster 1, it will not work.



298_OpenShift_0123

You can manually back up and restore your hosted cluster, or you can run a script to complete the process. For more information about the script, see "Running a script to back up and restore a hosted cluster".

5.4.4.3. Backing up a hosted cluster

To recover your hosted cluster in your target management cluster, you first need to back up all of the relevant data.

Procedure

1. Create a configmap file to declare the source management cluster by entering this command:

```
$ oc create configmap mgmt-parent-cluster -n default --from-literal=from=${MGMT_CLUSTER_NAME}
```

2. Shut down the reconciliation in the hosted cluster and in the node pools by entering these commands:

```
PAUSED_UNTIL="true"
oc patch -n ${HC_CLUSTER_NS} hostedclusters/${HC_CLUSTER_NAME} -p '{"spec": {"pausedUntil": "${PAUSED_UNTIL}"}}' --type=merge
oc scale deployment -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --replicas=0 kube-apiserver openshift-apiserver openshift-oauth-apiserver control-plane-operator
```

```

PAUSED_UNTIL="true"
oc patch -n ${HC_CLUSTER_NS} hostedclusters/${HC_CLUSTER_NAME} -p '{"spec":
{"pausedUntil":"${PAUSED_UNTIL}"}' --type=merge
oc patch -n ${HC_CLUSTER_NS} nodepools/${NODEPOOLS} -p '{"spec":
{"pausedUntil":"${PAUSED_UNTIL}"}' --type=merge
oc scale deployment -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --replicas=0 kube-
apiserver openshift-apiserver openshift-oauth-apiserver control-plane-operator

```

3. Back up etcd and upload the data to an S3 bucket by running this bash script:

TIP

Wrap this script in a function and call it from the main function.

```

# ETCD Backup
ETCD_PODS="etcd-0"
if [ "${CONTROL_PLANE_AVAILABILITY_POLICY}" = "HighlyAvailable" ]; then
  ETCD_PODS="etcd-0 etcd-1 etcd-2"
fi

for POD in ${ETCD_PODS}; do
  # Create an etcd snapshot
  oc exec -it ${POD} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -- env
ETCDCTL_API=3 /usr/bin/etcdctl --cacert /etc/etcd/tls/client/etcd-client-ca.crt --cert
/etc/etcd/tls/client/etcd-client.crt --key /etc/etcd/tls/client/etcd-client.key --
endpoints=localhost:2379 snapshot save /var/lib/data/snapshot.db
  oc exec -it ${POD} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -- env
ETCDCTL_API=3 /usr/bin/etcdctl -w table snapshot status /var/lib/data/snapshot.db

  FILEPATH="/${BUCKET_NAME}/${HC_CLUSTER_NAME}-${POD}-snapshot.db"
  CONTENT_TYPE="application/x-compressed-tar"
  DATE_VALUE=`date -R`
  SIGNATURE_STRING="PUT\n\n${CONTENT_TYPE}\n${DATE_VALUE}\n${FILEPATH}"

  set +x
  ACCESS_KEY=$(grep aws_access_key_id ${AWS_CREDS} | head -n1 | cut -d= -f2 | sed
"s/ //g")
  SECRET_KEY=$(grep aws_secret_access_key ${AWS_CREDS} | head -n1 | cut -d= -f2 |
sed "s/ //g")
  SIGNATURE_HASH=$(echo -en ${SIGNATURE_STRING} | openssl sha1 -hmac
"${SECRET_KEY}" -binary | base64)
  set -x

  # FIXME: this is pushing to the OIDC bucket
  oc exec -it etcd-0 -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -- curl -X PUT -T
"/var/lib/data/snapshot.db" \
  -H "Host: ${BUCKET_NAME}.s3.amazonaws.com" \
  -H "Date: ${DATE_VALUE}" \
  -H "Content-Type: ${CONTENT_TYPE}" \
  -H "Authorization: AWS ${ACCESS_KEY}:${SIGNATURE_HASH}" \
  https://${BUCKET_NAME}.s3.amazonaws.com/${HC_CLUSTER_NAME}-${POD}-
snapshot.db
done

```

For more information about backing up etcd, see "Backing up and restoring etcd on a hosted cluster".

4. Back up Kubernetes and OpenShift Container Platform objects by entering the following commands. You need to back up the following objects:

- **HostedCluster** and **NodePool** objects from the HostedCluster namespace
- **HostedCluster** secrets from the HostedCluster namespace
- **HostedControlPlane** from the Hosted Control Plane namespace
- **Cluster** from the Hosted Control Plane namespace
- **AWSCluster**, **AWSMachineTemplate**, and **AWSMachine** from the Hosted Control Plane namespace
- **MachineDeployments**, **MachineSets**, and **Machines** from the Hosted Control Plane namespace
- **ControlPlane** secrets from the Hosted Control Plane namespace

```
mkdir -p ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}
chmod 700 ${BACKUP_DIR}/namespaces/

# HostedCluster
echo "Backing Up HostedCluster Objects:"
oc get hc ${HC_CLUSTER_NAME} -n ${HC_CLUSTER_NS} -o yaml >
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/hc-${HC_CLUSTER_NAME}.yaml
echo "--> HostedCluster"
sed -i " -e '/^status:$/, $d' ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/hc-
${HC_CLUSTER_NAME}.yaml

# NodePool
oc get np ${NODEPOOLS} -n ${HC_CLUSTER_NS} -o yaml >
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/np-${NODEPOOLS}.yaml
echo "--> NodePool"
sed -i " -e '/^status:$/, $d' ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/np-
${NODEPOOLS}.yaml

# Secrets in the HC Namespace
echo "--> HostedCluster Secrets:"
for s in $(oc get secret -n ${HC_CLUSTER_NS} | grep "^${HC_CLUSTER_NAME}" |
awk '{print $1}'); do
    oc get secret -n ${HC_CLUSTER_NS} $s -o yaml >
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/secret-${s}.yaml
done

# Secrets in the HC Control Plane Namespace
echo "--> HostedCluster ControlPlane Secrets:"
for s in $(oc get secret -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} | egrep -v
"docker|service-account-token|oauth-openshift|NAME|token-${HC_CLUSTER_NAME}" |
awk '{print $1}'); do
    oc get secret -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} $s -o yaml >
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/secret-
${s}.yaml
```

```

done

# Hosted Control Plane
echo "--> HostedControlPlane:"
oc get hcp ${HC_CLUSTER_NAME} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}
-o yaml > ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/hcp-${HC_CLUSTER_NAME}.yaml

# Cluster
echo "--> Cluster:"
CL_NAME=$(oc get hcp ${HC_CLUSTER_NAME} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o jsonpath={.metadata.labels.*} | grep
${HC_CLUSTER_NAME})
oc get cluster ${CL_NAME} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o yaml
> ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/cl-
${HC_CLUSTER_NAME}.yaml

# AWS Cluster
echo "--> AWS Cluster:"
oc get awscluster ${HC_CLUSTER_NAME} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o yaml >
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/awscl-
${HC_CLUSTER_NAME}.yaml

# AWS MachineTemplate
echo "--> AWS Machine Template:"
oc get awsmachinetemplate ${NODEPOOLS} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o yaml >
${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/awsmt-
${HC_CLUSTER_NAME}.yaml

# AWS Machines
echo "--> AWS Machine:"
CL_NAME=$(oc get hcp ${HC_CLUSTER_NAME} -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o jsonpath={.metadata.labels.*} | grep
${HC_CLUSTER_NAME})
for s in $(oc get awsmachines -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --no-headers | grep ${CL_NAME} | cut -f1 -d\ ); do
    oc get -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} awsmachines $s -o yaml >
    ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/awsm-
    ${s}.yaml
done

# MachineDeployments
echo "--> HostedCluster MachineDeployments:"
for s in $(oc get machinedeployment -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o name); do
    mdp_name=$(echo $s | cut -f 2 -d /)
    oc get -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} $s -o yaml >
    ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}/machinedeployment-${mdp_name}.yaml
done

# MachineSets
echo "--> HostedCluster MachineSets:"
for s in $(oc get machineset -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o

```

```

name); do
    ms_name=$(echo ${s} | cut -f 2 -d /)
    oc get -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} $s -o yaml >
    ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
    ${HC_CLUSTER_NAME}/machineset-${ms_name}.yaml
done

# Machines
echo "--> HostedCluster Machine:"
for s in $(oc get machine -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o name);
do
    m_name=$(echo ${s} | cut -f 2 -d /)
    oc get -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} $s -o yaml >
    ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
    ${HC_CLUSTER_NAME}/machine-${m_name}.yaml
done

```

- Clean up the **ControlPlane** routes by entering this command:

```
$ oc delete routes -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --all
```

By entering that command, you enable the ExternalDNS Operator to delete the Route53 entries.

- Verify that the Route53 entries are clean by running this script:

```

function clean_routes() {
    if [[ -z "${1}" ]];then
        echo "Give me the NS where to clean the routes"
        exit 1
    fi

    # Constants
    if [[ -z "${2}" ]];then
        echo "Give me the Route53 zone ID"
        exit 1
    fi

    ZONE_ID=${2}
    ROUTES=10
    timeout=40
    count=0

    # This allows us to remove the ownership in the AWS for the API route
    oc delete route -n ${1} --all

    while [ ${ROUTES} -gt 2 ]
    do
        echo "Waiting for ExternalDNS Operator to clean the DNS Records in AWS Route53
where the zone id is: ${ZONE_ID}..."
        echo "Try: (${count}/${timeout})"
        sleep 10
        if [[ $count -eq timeout ]];then
            echo "Timeout waiting for cleaning the Route53 DNS records"
            exit 1

```



```

    fi
    count=$((count+1))
    ROUTES=$(aws route53 list-resource-record-sets --hosted-zone-id ${ZONE_ID} --max-
items 10000 --output json | grep -c ${EXTERNAL_DNS_DOMAIN})
    done
}

# SAMPLE: clean_routes "<HC ControlPlane Namespace>" "<AWS_ZONE_ID>"
clean_routes "${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}" "${AWS_ZONE_ID}"

```

Verification

Check all of the OpenShift Container Platform objects and the S3 bucket to verify that everything looks as expected.

Next steps

Restore your hosted cluster.

5.4.4.4. Restoring a hosted cluster

Gather all of the objects that you backed up and restore them in your destination management cluster.

Prerequisites

You backed up the data from your source management cluster.

TIP

Ensure that the **kubeconfig** file of the destination management cluster is placed as it is set in the **KUBECONFIG** variable or, if you use the script, in the **MGMT2_KUBECONFIG** variable. Use **export KUBECONFIG=<Kubeconfig FilePath>** or, if you use the script, use **export KUBECONFIG=\${MGMT2_KUBECONFIG}**.

Procedure

1. Verify that the new management cluster does not contain any namespaces from the cluster that you are restoring by entering these commands:

```

# Just in case
export KUBECONFIG=${MGMT2_KUBECONFIG}
BACKUP_DIR=${HC_CLUSTER_DIR}/backup

# Namespace deletion in the destination Management cluster
$ oc delete ns ${HC_CLUSTER_NS} || true
$ oc delete ns ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} || true

```

2. Re-create the deleted namespaces by entering these commands:

```

# Namespace creation
$ oc new-project ${HC_CLUSTER_NS}
$ oc new-project ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}

```

3. Restore the secrets in the HC namespace by entering this command:

```
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/secret-*
```

- Restore the objects in the **HostedCluster** control plane namespace by entering these commands:

```
# Secrets
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/secret-*

# Cluster
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/hcp-*
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/cl-*
```

- If you are recovering the nodes and the node pool to reuse AWS instances, restore the objects in the HC control plane namespace by entering these commands:

```
# AWS
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/awscl-*
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/awsmt-*
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/awsm-*

# Machines
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/machinedeployment-*
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/machineset-*
$ oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}/machine-*
```

- Restore the etcd data and the hosted cluster by running this bash script:

```
ETCD_PODS="etcd-0"
if [ "${CONTROL_PLANE_AVAILABILITY_POLICY}" = "HighlyAvailable" ]; then
  ETCD_PODS="etcd-0 etcd-1 etcd-2"
fi

HC_RESTORE_FILE=${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/hc-
${HC_CLUSTER_NAME}-restore.yaml
HC_BACKUP_FILE=${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/hc-
${HC_CLUSTER_NAME}.yaml
HC_NEW_FILE=${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/hc-
${HC_CLUSTER_NAME}-new.yaml
cat ${HC_BACKUP_FILE} > ${HC_NEW_FILE}
cat > ${HC_RESTORE_FILE} <<EOF
  restoreSnapshotURL:
EOF

for POD in ${ETCD_PODS}; do
  # Create a pre-signed URL for the etcd snapshot
  ETCD_SNAPSHOT="s3://${BUCKET_NAME}/${HC_CLUSTER_NAME}-${POD}-
```

```

snapshot.db"
ETCD_SNAPSHOT_URL=$(AWS_DEFAULT_REGION=${MGMT2_REGION} aws s3
presign ${ETCD_SNAPSHOT})

# FIXME no CLI support for restoreSnapshotURL yet
cat >> ${HC_RESTORE_FILE} <<EOF
- "${ETCD_SNAPSHOT_URL}"
EOF
done

cat ${HC_RESTORE_FILE}

if ! grep ${HC_CLUSTER_NAME}-snapshot.db ${HC_NEW_FILE}; then
  sed -i " -e '/type: PersistentVolume/r ${HC_RESTORE_FILE}'" ${HC_NEW_FILE}
  sed -i " -e '/pausedUntil:/d'" ${HC_NEW_FILE}
fi

HC=$(oc get hc -n ${HC_CLUSTER_NS} ${HC_CLUSTER_NAME} -o name || true)
if [[ ${HC} == "" ]];then
  echo "Deploying HC Cluster: ${HC_CLUSTER_NAME} in ${HC_CLUSTER_NS}
namespace"
  oc apply -f ${HC_NEW_FILE}
else
  echo "HC Cluster ${HC_CLUSTER_NAME} already exists, avoiding step"
fi

```

7. If you are recovering the nodes and the node pool to reuse AWS instances, restore the node pool by entering this command:

```
oc apply -f ${BACKUP_DIR}/namespaces/${HC_CLUSTER_NS}/np-*
```

Verification

- To verify that the nodes are fully restored, use this function:

```

timeout=40
count=0
NODE_STATUS=$(oc get nodes --kubeconfig=${HC_KUBECONFIG} | grep -v NotReady |
grep -c "worker") || NODE_STATUS=0

while [ ${NODE_POOL_REPLICAS} != ${NODE_STATUS} ]
do
  echo "Waiting for Nodes to be Ready in the destination MGMT Cluster:
${MGMT2_CLUSTER_NAME}"
  echo "Try: (${count}/${timeout})"
  sleep 30
  if [[ $count -eq timeout ]];then
    echo "Timeout waiting for Nodes in the destination MGMT Cluster"
    exit 1
  fi
  count=$((count+1))
  NODE_STATUS=$(oc get nodes --kubeconfig=${HC_KUBECONFIG} | grep -v NotReady |
grep -c "worker") || NODE_STATUS=0
done

```

Next steps

Shut down and delete your cluster.

5.4.4.5. Deleting a hosted cluster from your source management cluster

After you back up your hosted cluster and restore it to your destination management cluster, you shut down and delete the hosted cluster on your source management cluster.

Prerequisites

You backed up your data and restored it to your source management cluster.

TIP

Ensure that the **kubeconfig** file of the destination management cluster is placed as it is set in the **KUBECONFIG** variable or, if you use the script, in the **MGMT_KUBECONFIG** variable. Use **export KUBECONFIG=<Kubeconfig FilePath>** or, if you use the script, use **export KUBECONFIG=\${MGMT_KUBECONFIG}**.

Procedure

1. Scale the **deployment** and **statefulset** objects by entering these commands:

```
# Just in case
export KUBECONFIG=${MGMT_KUBECONFIG}

# Scale down deployments
oc scale deployment -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --replicas=0 --all
oc scale statefulset.apps -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --replicas=0 --all
sleep 15
```

2. Delete the **NodePool** objects by entering these commands:

```
NODEPOOLS=$(oc get nodepools -n ${HC_CLUSTER_NS} -o=jsonpath='{.items[?(@.spec.clusterName=="${HC_CLUSTER_NAME}").metadata.name]}')
if [[ ! -z "${NODEPOOLS}" ]];then
  oc patch -n "${HC_CLUSTER_NS}" nodepool ${NODEPOOLS} --type=json --patch='[ {
"op":"remove", "path": "/metadata/finalizers" }]'
  oc delete np -n ${HC_CLUSTER_NS} ${NODEPOOLS}
fi
```

3. Delete the **machine** and **machineset** objects by entering these commands:

```
# Machines
for m in $(oc get machines -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o name); do
  oc patch -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} ${m} --type=json --patch='[ {
"op":"remove", "path": "/metadata/finalizers" }]' || true
  oc delete -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} ${m} || true
done

oc delete machineset -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --all || true
```

4. Delete the cluster object by entering these commands:

```
# Cluster
C_NAME=$(oc get cluster -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} -o name)
oc patch -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} ${C_NAME} --type=json --
patch='[ { "op": "remove", "path": "/metadata/finalizers" } ]'
oc delete cluster.cluster.x-k8s.io -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} --all
```

5. Delete the AWS machines (Kubernetes objects) by entering these commands. Do not worry about deleting the real AWS machines. The cloud instances will not be affected.

```
# AWS Machines
for m in $(oc get awsmachine.infrastructure.cluster.x-k8s.io -n ${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME} -o name)
do
  oc patch -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} ${m} --type=json --patch='[ {
"op": "remove", "path": "/metadata/finalizers" } ]' || true
  oc delete -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} ${m} || true
done
```

6. Delete the **HostedControlPlane** and **ControlPlane** HC namespace objects by entering these commands:

```
# Delete HCP and ControlPlane HC NS
oc patch -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}
hostedcontrolplane.hypershift.openshift.io ${HC_CLUSTER_NAME} --type=json --patch='[ {
"op": "remove", "path": "/metadata/finalizers" } ]'
oc delete hostedcontrolplane.hypershift.openshift.io -n ${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME} --all
oc delete ns ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME} || true
```

7. Delete the **HostedCluster** and HC namespace objects by entering these commands:

```
# Delete HC and HC Namespace
oc -n ${HC_CLUSTER_NS} patch hostedclusters ${HC_CLUSTER_NAME} -p '{"metadata":
{"finalizers": null}}' --type merge || true
oc delete hc -n ${HC_CLUSTER_NS} ${HC_CLUSTER_NAME} || true
oc delete ns ${HC_CLUSTER_NS} || true
```

Verification

- To verify that everything works, enter these commands:

```
# Validations
export KUBECONFIG=${MGMT2_KUBECONFIG}

oc get hc -n ${HC_CLUSTER_NS}
oc get np -n ${HC_CLUSTER_NS}
oc get pod -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}
oc get machines -n ${HC_CLUSTER_NS}-${HC_CLUSTER_NAME}

# Inside the HostedCluster
export KUBECONFIG=${HC_KUBECONFIG}
oc get clusterversion
oc get nodes
```

Next steps

Delete the OVN pods in the hosted cluster so that you can connect to the new OVN control plane that runs in the new management cluster:

1. Load the **KUBECONFIG** environment variable with the hosted cluster's kubeconfig path.
2. Enter this command:

```
$ oc delete pod -n openshift-ovn-kubernetes --all
```

5.4.4.6. Running a script to back up and restore a hosted cluster

To expedite the process to back up a hosted cluster and restore it within the same region on AWS, you can modify and run a script.

Procedure

1. Replace the variables in the following script with your information:

```
# Fill the Common variables to fit your environment, this is just a sample
SSH_KEY_FILE=${HOME}/.ssh/id_rsa.pub
BASE_PATH=${HOME}/hypershift
BASE_DOMAIN="aws.sample.com"
PULL_SECRET_FILE="${HOME}/pull_secret.json"
AWS_CREDS="${HOME}/.aws/credentials"
CONTROL_PLANE_AVAILABILITY_POLICY=SingleReplica
HYPERSHIFT_PATH=${BASE_PATH}/src/hypershift
HYPERSHIFT_CLI=${HYPERSHIFT_PATH}/bin/hypershift
HYPERSHIFT_IMAGE=${HYPERSHIFT_IMAGE:-"quay.io/${USER}/hypershift:latest"}
NODE_POOL_REPLICAS=${NODE_POOL_REPLICAS:-2}

# MGMT Context
MGMT_REGION=us-west-1
MGMT_CLUSTER_NAME="${USER}-dev"
MGMT_CLUSTER_NS=${USER}
MGMT_CLUSTER_DIR="${BASE_PATH}/hosted_clusters/${MGMT_CLUSTER_NS}-
${MGMT_CLUSTER_NAME}"
MGMT_KUBECONFIG="${MGMT_CLUSTER_DIR}/kubeconfig"

# MGMT2 Context
MGMT2_CLUSTER_NAME="${USER}-dest"
MGMT2_CLUSTER_NS=${USER}
MGMT2_CLUSTER_DIR="${BASE_PATH}/hosted_clusters/${MGMT2_CLUSTER_NS}-
${MGMT2_CLUSTER_NAME}"
MGMT2_KUBECONFIG="${MGMT2_CLUSTER_DIR}/kubeconfig"

# Hosted Cluster Context
HC_CLUSTER_NS=clusters
HC_REGION=us-west-1
HC_CLUSTER_NAME="${USER}-hosted"
HC_CLUSTER_DIR="${BASE_PATH}/hosted_clusters/${HC_CLUSTER_NS}-
${HC_CLUSTER_NAME}"
HC_KUBECONFIG="${HC_CLUSTER_DIR}/kubeconfig"
BACKUP_DIR=${HC_CLUSTER_DIR}/backup
```

```
BUCKET_NAME="${USER}-hosted-${MGMT_REGION}"  
  
# DNS  
AWS_ZONE_ID="Z026552815SS3YPH9H6MG"  
EXTERNAL_DNS_DOMAIN="guest.jpdv.aws.kerbeross.com"
```

2. Save the script to your local file system.
3. Run the script by entering the following command:

```
source <env_file>
```

where: **env_file** is the name of the file where you saved the script.

The migration script is maintained at the following repository:

<https://github.com/openshift/hypershift/blob/main/contrib/migration/migrate-hcp.sh>.