



Data Protection of Container Apps Using Third Party Tools

NetApp Solutions

NetApp
July 31, 2024

Table of Contents

- Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP) 1
 - Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP) 2
 - Installation of OpenShift API for Data Protection (OADP) Operator. 4
 - Creating on-demand backup for Apps in OpenShift Container Platform 14
 - Migrate an App from one cluster to another 17
 - Restore an App from a backup 22
 - Deleting backups and restores in using Velero 30

Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP)

Author: Banu Sundhar, NetApp

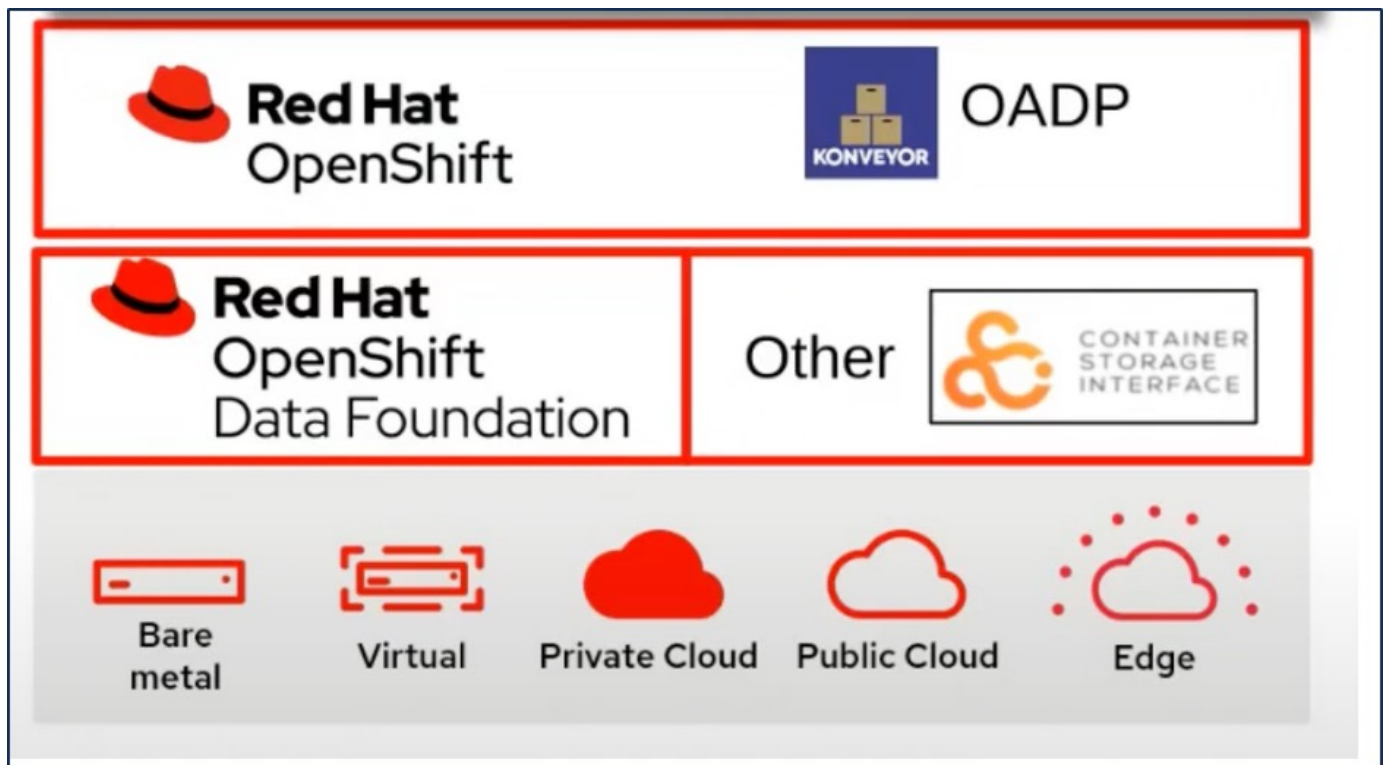
This section of the reference document provides details for creating backups of Container Apps using the OpenShift API for Data Protection (OADP) with Velero on NetApp ONTAP S3 or NetApp StorageGRID S3. The backups of namespace scoped resources including Persistent Volumes(PVs) of the app are created using CSI Astra Trident Snapshots.

The persistent storage for container apps can be backed by ONTAP storage integrated to the OpenShift Cluster using [Astra Trident CSI](#). In this section we use [OpenShift API for Data Protection \(OADP\)](#) to perform backup of apps including its data volumes to

- ONTAP Object Storage
- StorageGrid

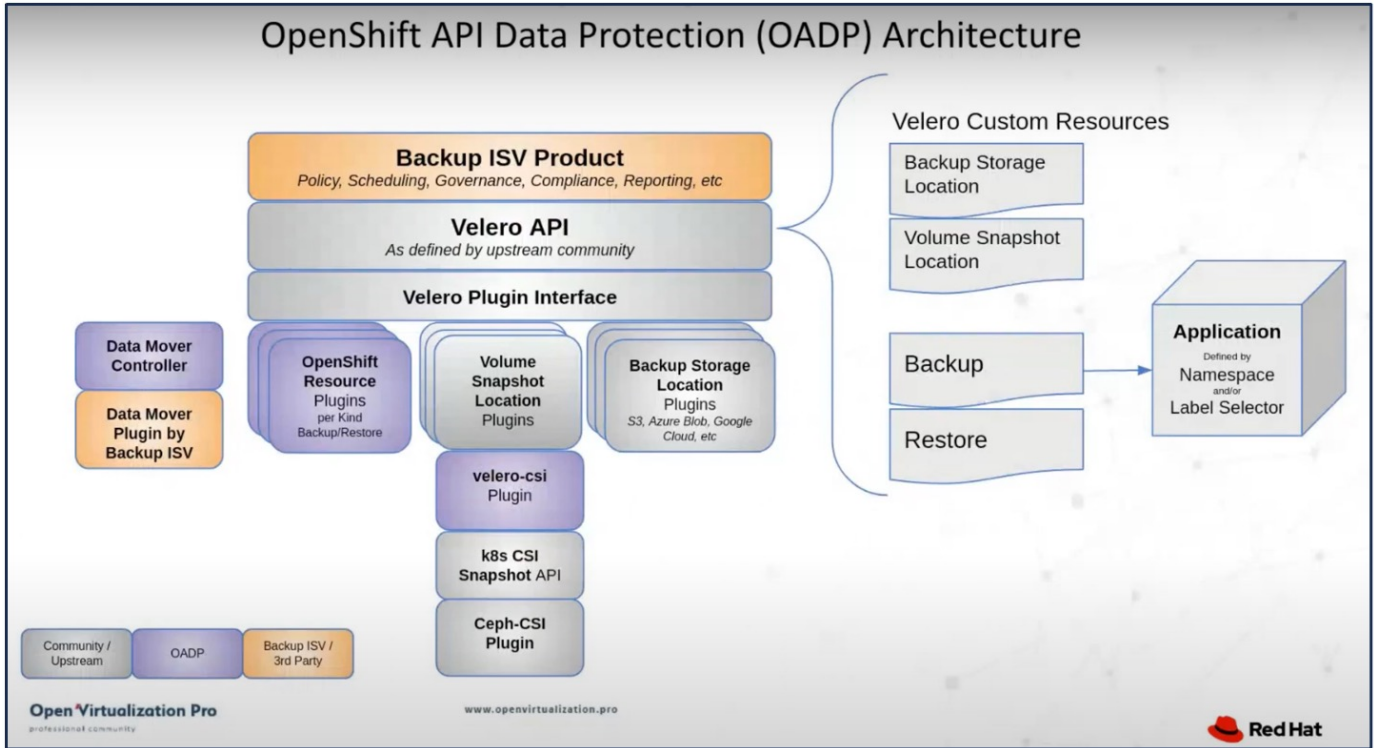
We then restore from the backup when needed. Please note that the app can be restored only to the cluster from where the backup was created.

OADP enables backup, restore, and disaster recovery of applications on an OpenShift cluster. Data that can be protected with OADP include Kubernetes resource objects, persistent volumes, and internal images.



Red Hat OpenShift has leveraged the solutions developed by the OpenSource communities for data protection. [Velero](#) is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes. To use Velero easily, OpenShift has developed the OADP operator and the Velero plugin to integrate with the CSI storage drivers. The core of the OADP APIs that

are exposed are based on the Velero APIs. After installing the OADP operator and configuring it, the backup/restore operations that can be performed are based on the operations exposed by the Velero API.



OADP 1.3 is available from the operator hub of OpenShift cluster 4.12 and later. It has a built-in Data Mover that can move CSI volume snapshots to a remote object store. This provides portability and durability by moving snapshots to an object storage location during backup. The snapshots are then available for restoration after disasters.

The following are the versions of the various components used for the examples in this section

- OpenShift Cluster 4.14
- OADP Operator 1.13 provided by Red Hat
- Velero CLI 1.13 for Linux
- Astra Trident 24.02
- ONTAP 9.12
- postgresql installed using helm.

[Astra Trident CSI](#)
[OpenShift API for Data Protection \(OADP\)](#)
[Velero](#)

Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP)

Author: Banu Sundhar, NetApp

This section of the reference document provides details for creating backups of Container Apps using the OpenShift API for Data Protection (OADP) with Velero on NetApp ONTAP


S3 or NetApp StorageGRID S3. The backups of namespace scoped resources including Persistent Volumes(PVs) of the app are created using CSI Astra Trident Snapshots.

The persistent storage for container apps can be backed by ONTAP storage integrated to the OpenShift Cluster using [Astra Trident CSI](#). In this section we use [OpenShift API for Data Protection \(OADP\)](#) to perform backup of apps including its data volumes to


- ONTAP Object Storage
- StorageGrid

We then restore from the backup when needed. Please note that the app can be restored only to the cluster from where the backup was created.

OADP enables backup, restore, and disaster recovery of applications on an OpenShift cluster. Data that can be protected with OADP include Kubernetes resource objects, persistent volumes, and internal images.



Red Hat OpenShift has leveraged the solutions developed by the OpenSource communities for data protection. [Velero](#) is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes. To use Velero easily, OpenShift has developed the OADP operator and the Velero plugin to integrate with the CSI storage drivers. The core of the OADP APIs that are exposed are based on the Velero APIs. After installing the OADP operator and configuring it, the backup/restore operations that can be performed are based on the operations exposed by the Velero API.



OADP 1.3 is available from the operator hub of OpenShift cluster 4.12 and later. It has a built-in Data Mover that can move CSI volume snapshots to a remote object store. This provides portability and durability by moving snapshots to an object storage location during backup. The snapshots are then available for restoration after disasters.

The following are the versions of the various components used for the examples in this section

- OpenShift Cluster 4.14
- OADP Operator 1.13 provided by Red Hat
- Velero CLI 1.13 for Linux
- Astra Trident 24.02
- ONTAP 9.12
- postgresql installed using helm.

[Astra Trident CSI](#)
[OpenShift API for Data Protection \(OADP\)](#)
[Velero](#)

Installation of OpenShift API for Data Protection (OADP) Operator

This section outlines the installation of OpenShift API for Data Protection (OADP) Operator.

Prerequisites


- A Red Hat OpenShift cluster (later than version 4.12) installed on bare-metal infrastructure with RHCOS

worker nodes

- A NetApp ONTAP cluster integrated with the cluster using Astra Trident
- A Trident backend configured with an SVM on ONTAP cluster
- A StorageClass configured on the OpenShift cluster with Astra Trident as the provisioner
- Trident Snapshot class created on the cluster
- Cluster-admin access to Red Hat OpenShift cluster
- Admin access to NetApp ONTAP cluster
- An application eg. postgresql deployed on the cluster
- An admin workstation with tridentctl and oc tools installed and added to \$PATH

Steps to install OADP Operator

1. Go to the Operator Hub of the cluster and select Red Hat OADP operator. In the Install page, use all the default selections and click install. On the next page, again use all the defaults and click Install. The OADP operator will be installed in the namespace openshift-adp.





OADP Operator

1.3.0 provided by Red Hat

Install

Channel

stable-1.3

OpenShift API for Data Protection (OADP) operator sets up and installs Velero on the OpenShift platform, allowing users to backup and restore applications.

Version

1.3.0

Backup and restore Kubernetes resources and internal images, at the granularity of a namespace, using a version of Velero appropriate for the installed version of OADP.

Capability level

- Basic Install
- Seamless Upgrades
- Full Lifecycle
- Deep Insights
- Auto Pilot

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 or Kopia.

- [Installing OADP for application backup and restore](#)
- [Installing OADP on a ROSA cluster and using STS, please follow the Getting Started Steps 1-3 in order to obtain the role ARN needed for using the standardized STS configuration flow via OLM](#)
- [Frequently Asked Questions](#)

Source

Red Hat

Provider

Red Hat

Infrastructure features

Disconnected

Activate Windows

Project: All Projects

Installed Operators

Installed Operators are represented by ClusterServiceVersions within this Namespace. For more information, see the [Understanding Operators documentation](#) Operator and ClusterServiceVersion using the [Operator SDK](#).

Name Search by name... /

Name	Namespace	Managed Namespaces	Status
OpenShift Virtualization 4.14.4 provided by Red Hat	NS openshift-cnv	NS openshift-cnv	✓ Succeeded Up to date
OADP Operator 1.3.0 provided by Red Hat	NS openshift-adp	NS openshift-adp	✓ Succeeded Up to date
Package Server 0.0.1-snapshot provided by	NS openshift-operator-lifecycle-manager	NS openshift-operator-lifecycle-manager	✓ Succeeded

Prerequisites for Velero configuration with Ontap S3 details

After the installation of the operator succeeds, configure the instance of Velero.

Velero can be configured to use S3 compatible Object Storage. Configure ONTAP S3 using the procedures shown in the [Object Storage Management section of ONTAP documentation](#). You will need the following information from your ONTAP S3 configuration to integrate with Velero.

- A Logical Interface (LIF) that can be used to access S3
- User credentials to access S3 that includes the access key and the secret access key
- A bucket name in S3 for backups with access permissions for the user
- For secure access to the Object storage, TLS certificate should be installed on the Object Storage server.

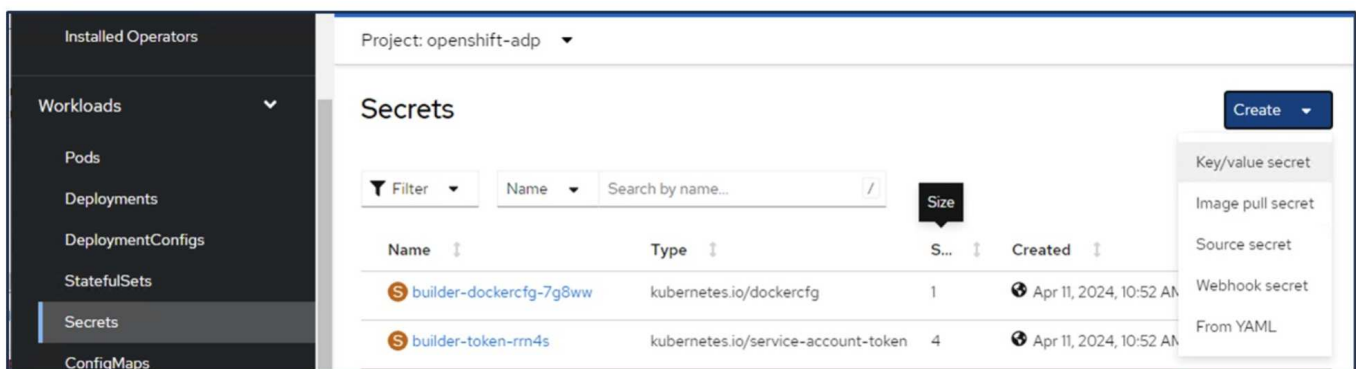
Prerequisites for Velero configuration with StorageGrid S3 details

Velero can be configured to use S3 compatible Object Storage. You can configure StorageGrid S3 using the procedures shown in the [StorageGrid documentation](#). You will need the following information from your StorageGrid S3 configuration to integrate with Velero.

- The endpoint that can be used to access S3
- User credentials to access S3 that includes the access key and the secret access key
- A bucket name in S3 for backups with access permissions for the user
- For secure access to the Object storage, TLS certificate should be installed on the Object Storage server.

Steps to configure Velero

- First, create a secret for an ONTAP S3 user credential or StorageGrid Tenant user credentials. This will be used to configure Velero later. You can create a secret from the CLI or from the web console. To create a secret from the web console, select Secrets, then click on Key/Value Secret. Provide the values for the credential name, key and the value as shown. Be sure to use the Access Key Id and Secret Access Key of your S3 user. Name the secret appropriately. In the sample below, a secret with ONTAP S3 user credentials named `ontap-s3-credentials` is created.



The screenshot shows the OpenShift web console interface. On the left is a navigation sidebar with 'Secrets' selected. The main area displays the 'Secrets' page for the 'openshift-adp' project. A 'Create' button is visible in the top right, with a dropdown menu open showing options: 'Key/value secret', 'Image pull secret', 'Source secret', 'Webhook secret', and 'From YAML'. Below the menu is a table of existing secrets.

Name	Type	Size	Created
builder-dockercfg-7g8ww	kubernetes.io/dockercfg	1	Apr 11, 2024, 10:52 AM
builder-token-rm4s	kubernetes.io/service-account-token	4	Apr 11, 2024, 10:52 AM

Project: openshift-adp ▾

Edit key/value secret

Key/value secrets let you inject sensitive data into your application as files or environment variables.

Secret name *

 Unique name of the new secret.

Key *

Value

 Browse...

Drag and drop file with your value here or browse to upload it.

```
[default]
aws_access_key_id=
aws_secret_access_key=
```

+ Add key/value


To create a secret named sg-s3-credentials from the CLI you can use the following command.

```
# oc create secret generic sg-s3-credentials --namespace openshift-adp --from-file
cloud=cloud-credentials.txt
```

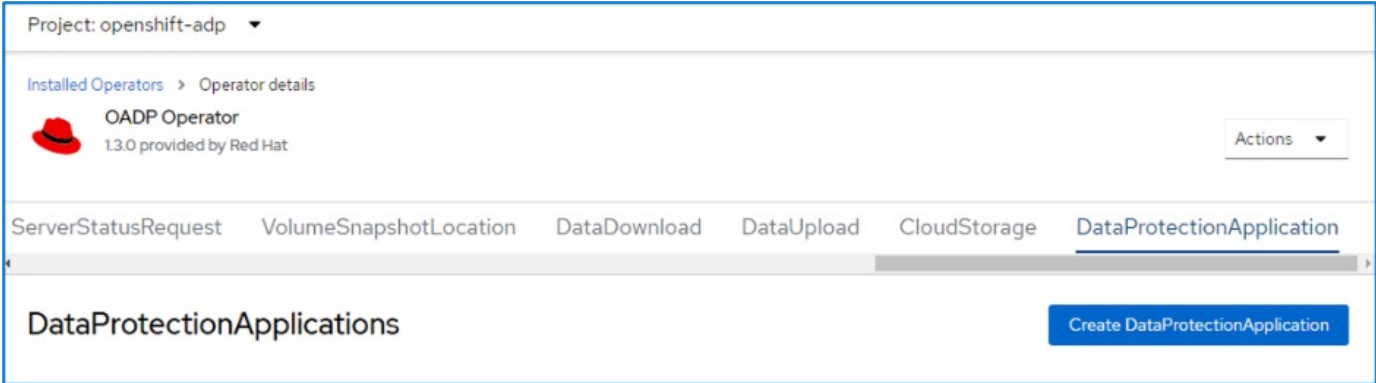
Where credentials.txt file contains the Access Key Id and the Secret Access Key of the S3 user in the following format:

```
[default]
aws_access_key_id=< Access Key ID of S3 user>
aws_secret_access_key=<Secret Access key of S3 user>
```

- Next, to configure Velero, select Installed Operators from the menu item under Operators, click on OADP operator, and then select the **DataProtectionApplication** tab.



Click on Create DataProtectionApplication. In the form view, provide a name for the DataProtection Application or use the default name.



Now go to the YAML view and replace the spec information as shown in the yaml file examples below.

Sample yaml file for configuring Velero with ONTAP S3 as the backupLocation

```

spec:
  backupLocations:
    - velero:
      config:
        insecureSkipTLSVerify: 'false' ->use this for https
communication with ONTAP S3
        profile: default
        region: us-east-1
        s3ForcePathStyle: 'true' ->This allows use of IP in s3URL
        s3Url: 'https://10.61.181.161' ->Ensure TLS certificate for S3
is configured
      credential:
        key: cloud
        name: ontap-s3-credentials -> previously created secret
        default: true
      objectStorage:
        bucket: velero -> Your bucket name previously created in S3 for
backups
        prefix: container-demo-backup ->The folder that will be created
in the bucket
        caCert: <base64 encoded CA Certificate installed on ONTAP
Cluster with the SVM Scope where the bucker exists>
        provider: aws
      configuration:
        nodeAgent:
          enable: true
          uploaderType: kopia
          #default Data Mover uses Kopia to move snapshots to Object Storage
        velero:
          defaultPlugins:
            - csi ->This plugin to use CSI snapshots
            - openshift
            - aws
            - kubevirt -> This plugin to use Velero with OIpenShift
Virtualization

```

Sample yaml file for configuring Velero with StorageGrid S3 as the backupLocation

```

spec:
  backupLocations:
    - velero:
      config:
        insecureSkipTLSVerify: 'true'
        profile: default
        region: us-east-1 ->region of your StorageGrid system
        s3ForcePathStyle: 'True'
        s3Url: 'https://172.21.254.25:10443' ->the IP used to access S3
      credential:
        key: cloud
        name: sg-s3-credentials ->secret created earlier
      default: true
      objectStorage:
        bucket: velero
        prefix: demobackup
      provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
        - aws
        - kubevirt

```

The spec section in the yaml file should be configured appropriately for the following parameters similar to the example above

backupLocations

ONTAP S3 or StorageGrid S3 (with its credentials and other information as shown in the yaml) is configured as the default BackupLocation for velero.

snapshotLocations

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. In our example, you use Astra Trident CSI and you have previously created VolumeSnapShotClass CR using the Trident CSI driver.

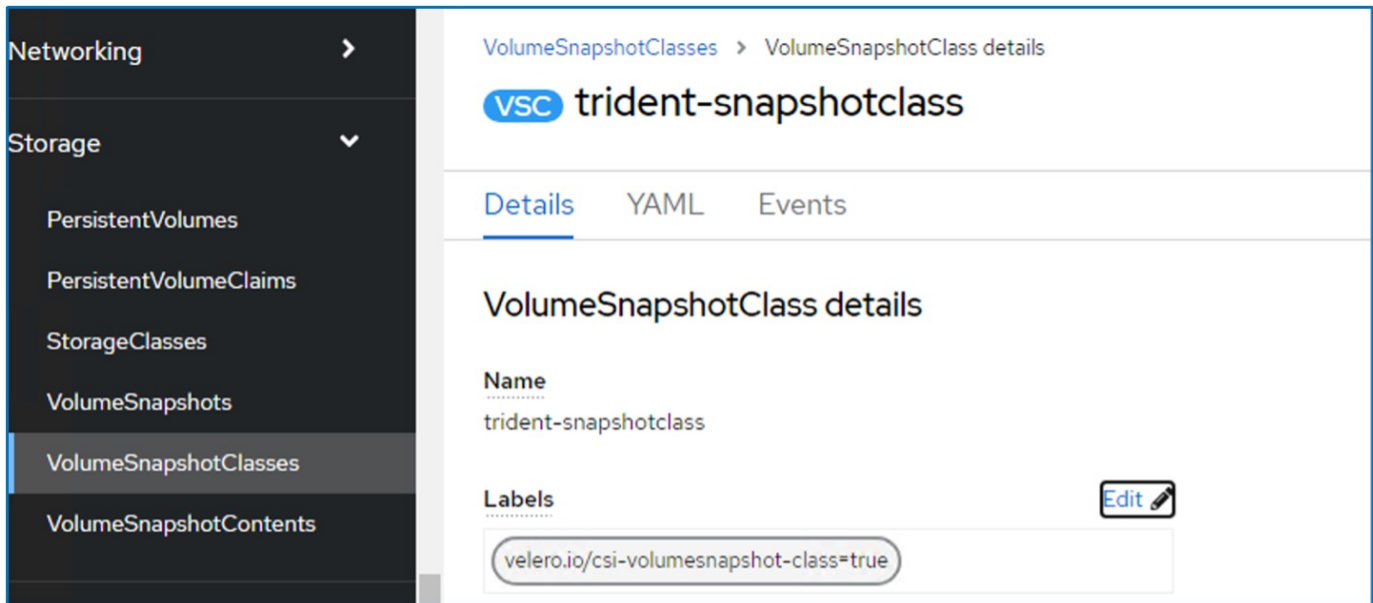
Enable CSI plugin

Add csi to the defaultPlugins for Velero to back up persistent volumes with CSI snapshots.

The Velero CSI plugins, to backup CSI backed PVCs, will choose the VolumeSnapshotClass in the cluster that has **velero.io/csi-volumesnapshot-class** label set on it. For this

- You must have the trident VolumeSnapshotClass created.
- Edit the label of the trident-snapshotclass and set it to

`velero.io/csi-volumesnapshot-class=true` as shown below.



The screenshot displays the Kubernetes dashboard interface. On the left, a dark sidebar contains a navigation menu with the following items: Networking, Storage, PersistentVolumes, PersistentVolumeClaims, StorageClasses, VolumeSnapshots, VolumeSnapshotClasses (highlighted), and VolumeSnapshotContents. The main content area shows the 'VolumeSnapshotClasses' page for the 'trident-snapshotclass'. The breadcrumb navigation is 'VolumeSnapshotClasses > VolumeSnapshotClass details'. Below the breadcrumb, there is a 'VSC' icon followed by the name 'trident-snapshotclass'. There are three tabs: 'Details' (active), 'YAML', and 'Events'. The 'Details' tab shows the 'VolumeSnapshotClass details' section. Under 'Name', the value is 'trident-snapshotclass'. Under 'Labels', there is a text input field containing 'velero.io/csi-volumesnapshot-class=true' and an 'Edit' button with a pencil icon.

Ensure that the snapshots can persist even if the VolumeSnapshot objects are deleted. This can be done by setting the **deletionPolicy** to Retain. If not, deleting a namespace will completely lose all PVCs ever backed up in it.

```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: trident-snapshotclass
driver: csi.trident.netapp.io
deletionPolicy: Retain
```

VolumeSnapshotClasses > VolumeSnapshotClass details

VSC trident-snapshotclass

Details | YAML | Events

VolumeSnapshotClass details

Name
trident-snapshotclass

Labels Edit

velero.io/csi-volumesnapshot-class=true

Annotations
1 annotation


Driver
csi.trident.netapp.io

Deletion policy
Retain

Ensure that the DataProtectionApplication is created and is in condition:Reconciled.

Project: openshift-adp

Installed Operators > Operator details


 **OADP Operator**
1.3.2 provided by Red Hat Actions

Schedule | ServerStatusRequest | VolumeSnapshotLocation | DataDownload | DataUpload | CloudStorage | DataProtectionApplication


DataProtectionApplications

Create DataProtectionApplication

Name Search by name... /

Name	Kind	Status	Labels	Last updated
 velero-container-backup-ontap	DataProtectionApplication	Condition: Reconciled	No labels	Jul 15, 2024, 2:31 PM

The OADP operator will create a corresponding BackupStorageLocation. This will be used when creating a backup.



Creating on-demand backup for Apps in OpenShift Container Platform

This section outlines how to create on-demand backup for VMs in OpenShift Virtualization.

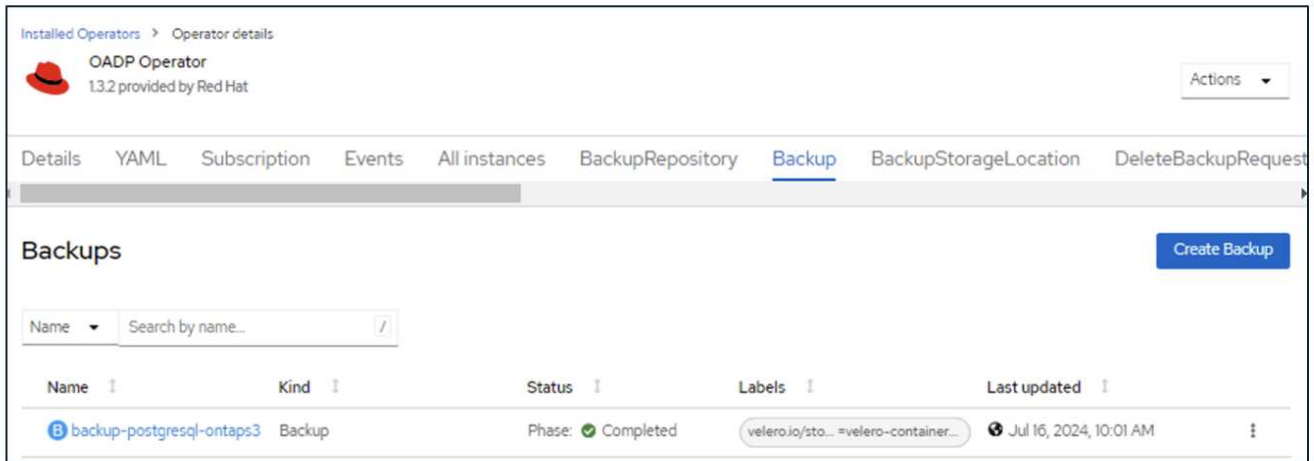
Steps to create a backup of an App

To create an on-demand backup of an app (app metadata and persistent volumes of the app), click on the **Backup** tab to create a Backup Custom Resource (CR). A sample yam1 is provided to create the Backup CR. Using this yam1, the app and its persistent storage in the specified namespace will be backed up. Additional parameters can be set as shown in the [documentation](#).

A snapshot of the persistent volumes and the app resources in the namespace specified will be created by the CSI. This snapshot will be stored in the backup location specified in the yam1. The backup will remain in the system for 30 days as specified in the ttl.

```
spec:
  csiSnapshotTimeout: 10m0s
  defaultVolumesToFsBackup: false
  includedNamespaces:
    - postgresql ->namespace of the app
  itemOperationTimeout: 4h0m0s
  snapshotMoveData: false
  storageLocation: velero-container-backup-ontap-1 -->this is the
  backupStorageLocation previously created when Velero is configured.
  ttl: 720h0m0s
```



Once the backup completes, its Phase will show as completed.



You can inspect the backup in the Object storage with the help of an S3 browser application. The path of the backup shows up in the configured bucket with the prefix name (velero/container-demo-backup). You can see the contents of the backup includes the volume snapshots, logs, and other metadata of the application.



In StorageGrid, you can also use the S3 console that is available from the Tenant Manager to view the backup objects.



Creating scheduled backups for Apps

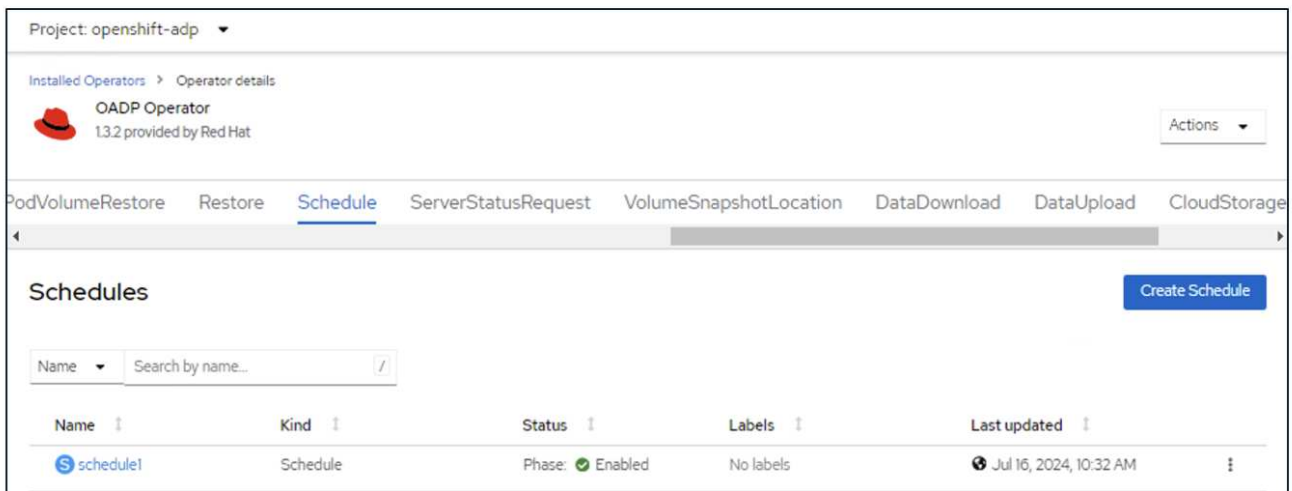
To create backups on a schedule, you need to create a Schedule CR.

The schedule is simply a Cron expression allowing you to specify the time at which you want to create the backup. A sample yaml to create a Schedule CR is shown below.

```
apiVersion: velero.io/v1
kind: Schedule
metadata:
  name: schedule1
  namespace: openshift-adp
spec:
  schedule: 0 7 * * *
  template:
    includedNamespaces:
      - postgresql
    storageLocation: velero-container-backup-ontap-1
```

The Cron expression `0 7 * * *` means a backup will be created at 7:00 every day. The namespaces to be included in the backup and the storage location for the backup are also specified. So instead of a Backup CR, Schedule CR is used to create a backup at the specified time and frequency.

Once the schedule is created, it will be Enabled.




The screenshot shows the OADP Operator interface for the 'openshift-adp' project. The 'Schedule' tab is selected, displaying a table of Schedules. A single schedule named 'schedule1' is listed with a status of 'Enabled' and a last updated time of 'Jul 16, 2024, 10:32 AM'. The interface includes a search bar and a 'Create Schedule' button.

Name	Kind	Status	Labels	Last updated
schedule1	Schedule	Phase: ✔ Enabled	No labels	Jul 16, 2024, 10:32 AM

Backups will be created according to this schedule, and can be viewed from the Backup tab.

Project: openshift-adp

Installed Operators > Operator details

 OADP Operator
1.3.2 provided by Red Hat







Actions

All instances BackupRepository **Backup** BackupStorageLocation DeleteBackupRequest DownloadRequest PodVolumeBackup

Backups

Create Backup

Name Search by name...

Name	Kind	Status	Labels	Last updated
 backup-postgresql-ontaps3	Backup	Phase:  Completed	velero.io/sto...=velero-container...	 Jul 16, 2024, 10:01 AM
 schedule1-20240717070005	Backup	Phase:  Completed	velero.io/schedule-na...=schedul... velero.io/sto...=velero-container...	 Jul 17, 2024, 3:00 AM

Migrate an App from one cluster to another

Velero's backup and restore capabilities make it a valuable tool for migrating your data between clusters. This section describes how to migrate apps(s) from one cluster to another by creating a backup of the app in Object storage from one cluster and then restoring the app from the same object storage to another cluster. .

Backup from first cluster

Prerequisites on Cluster 1

- Astra Trident must be installed on the cluster.
- A trident backend and Storage class must be created.
- OADP operator must be installed on the cluster.
- The DataProtectionApplication should be configured.

Use the following spec to configure the DataProtectionApplication object.

```
spec:
  backupLocations:
    - velero:
      config:
        insecureSkipTLSVerify: 'false'
        profile: default
        region: us-east-1
        s3ForcePathStyle: 'true'
        s3Url: 'https://10.61.181.161'
      credential:
        key: cloud
        name: ontap-s3-credentials
      default: true
      objectStorage:
        bucket: velero
        caCert: <base-64 encoded tls certificate>
        prefix: container-backup
      provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
        - aws
        - kubevirt
```

- Create an application on the cluster and take a backup of this application. As an example, install a postgres application.

```
[root@localhost ~]# oc get nodes
NAME                STATUS    ROLES    AGE     VERSION
ocp6-master1       Ready    control-plane,master  3d13h  v1.27.15+6147456
ocp6-master2       Ready    worker   3d12h  v1.27.15+6147456
ocp6-master3       Ready    control-plane,master  3d13h  v1.27.15+6147456
ocp6-worker1       Ready    worker   3d12h  v1.27.15+6147456
ocp6-worker2       Ready    worker   3d12h  v1.27.15+6147456
ocp6-worker3       Ready    control-plane,master  3d12h  v1.27.15+6147456
[root@localhost ~]# helm install postgresql bitnami/postgresql -n postgresql --create namespace^C
[root@localhost ~]# oc get pods -n postgresql
NAME                READY    STATUS    RESTARTS    AGE
postgresql-0        1/1     Running   0            4h53m
[root@localhost ~]# oc get pvc -n postgresql
NAME                STATUS    VOLUME                                     CAPACITY    ACCESS MODES    STORAGECLASS    AGE
data-postgresql-0   Bound    pvc-f7a3c772-0e61-49cb-a3d0-7c7b2ec87dc6  8Gi         RWO              ontap-nas       4h53m
[root@localhost ~]# oc get pv -n postgresql
NAME                CAPACITY    ACCESS MODES    RECLAIM POLICY    STATUS    CLAIM                                STORAGECLASS
REASON    AGE
pvc-2e9e982f-54a4-4e7b-8eae-a589e0d9d819  1Gi         RWO              Delete            Bound    trident/basic                                ontap-nas
4h55m
pvc-f7a3c772-0e61-49cb-a3d0-7c7b2ec87dc6  8Gi         RWO              Delete            Bound    postgresql/data-postgresql-0                ontap-nas
4h53m
[root@localhost ~]#
```

- Use the following spec for the backup CR:

```
spec:
  csiSnapshotTimeout: 10m0s
  defaultVolumesToFsBackup: false
  includedNamespaces:
    - postgresql
  itemOperationTimeout: 4h0m0s
  snapshotMoveData: true
  storageLocation: velero-sample-1
  ttl: 720h0m0s
```

The screenshot shows the OpenShift console interface for the OADP Operator. The breadcrumb navigation is "Installed Operators > Operator details". The operator is identified as "OADP Operator 1.4.0 provided by Red Hat". Below this, there are tabs for "Repository", "Backup", "BackupStorageLocation", "DeleteBackupRequest", "DownloadRequest", "PodVolumeBackup", and "PodVolumeRes". The "Backup" tab is active, displaying a "Backups" section with a "Create Backup" button. A search bar is present with the text "Search by name...". Below the search bar is a table with columns for "Name", "Kind", and "Status". One backup is listed with the name "backup", kind "Backup", and status "Completed". A watermark for "Activate Windows" is visible in the bottom right corner of the screenshot.

You can click on the **All instances** tab to see the different objects being created and moving through different phases to finally come to the backup **completed** phase.

A backup of the resources in the namespace postgresql will be stored in the Object Storage location (ONTAP S3) specified in the backupLocation in the OADP spec.

Restore to a second cluster

Prerequisites on Cluster 2

- Astra Trident must be installed on cluster 2.
- The postgresql app must NOT be already installed in the postgresql namespace.
- OADP operator must be installed on cluster 2, and the BackupStorage Location must be pointing to the same object storage location where the backup was stored from the first cluster.
- The Backup CR must be visible from the second cluster.


```
[root@localhost ~]# oc get pods -n trident
NAME                                READY   STATUS    RESTARTS   AGE
trident-controller-6799cfb77f-8rzvk 6/6     Running   6           2d7h
trident-node-linux-7wvjz            2/2     Running   2           2d7h
trident-node-linux-8vvm2            2/2     Running   0           2d7h
trident-node-linux-bgs6f            2/2     Running   2           2d7h
trident-node-linux-njwb8            2/2     Running   0           2d7h
trident-node-linux-scqjl            2/2     Running   0           2d7h
trident-node-linux-swr69            2/2     Running   2           2d7h
trident-operator-b88b86fc8-7fk68    1/1     Running   1           2d7h
[root@localhost ~]#
```

```
[root@localhost ~]# oc get nodes
NAME                STATUS   ROLES                    AGE   VERSION
ocp7-master1       Ready   control-plane,master    3d    v1.27.15+6147456
ocp7-master2       Ready   control-plane,master    3d    v1.27.15+6147456
ocp7-master3       Ready   control-plane,master    3d    v1.27.15+6147456
ocp7-worker1       Ready   worker                   3d    v1.27.15+6147456
ocp7-worker2       Ready   worker                   3d    v1.27.15+6147456
ocp7-worker3       Ready   worker                   3d    v1.27.15+6147456
[root@localhost ~]# oc get pods -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]# oc get pvc -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]# oc get pv -n postgresql
NAME                CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM                STORAGECLASS   REASON   AGE
pvc-c6660630-0cfe-484b-aaa3-5ada54c8b9a7 1Gi        RWO            Delete           Bound   trident/basic        Available 11m
pvc-edcc6551-81b0-40b4-8547-e9df70c1740d 10Gi       RWO            Delete           Bound   default/test-pvc     vsphere-sc   2d7h
[root@localhost ~]#
```

The screenshot shows the OpenShift console interface. At the top, the project is set to 'openshift-adp'. Below this, the 'Installed Operators' section is expanded to show 'Operator details' for the 'OADP Operator', version 1.4.0, provided by Red Hat. A navigation bar contains several tabs: 'Backup', 'BackupStorageLocation' (which is selected), 'DeleteBackupRequest', 'DownloadRequest', 'PodVolumeBackup', 'PodVolumeRestore', and 'Res'. Below the navigation bar, the 'BackupStorageLocations' section is displayed, featuring a 'Create BackupStorageLocation' button. A search bar is present with the text 'Search by name...'. A table lists the BackupStorageLocations:

Name	Kind	Status
BSL velero-container-demo-1	BackupStorageLocation	Phase: Available

Installed Operators > Operator details

 **OADP Operator**
1.4.0 provided by Red Hat




Actions

Details | YAML | Subscription | Events | All instances | BackupRepository | Backup | BackupStorageLocation | DeleteBackupRequest | DownloadRequest

Backups

Create Backup

Name Search by name... /

Name	Kind	Status	Labels	Last updated
 backup	Backup	Phase:  Completed	velero.io/storage-locati...=velero-sampl...	 Jul 25, 2024, 8:39 PM

Restore the app on this cluster from the backup. Use the following yaml to create the Restore CR.

```


apiVersion: velero.io/v1
kind: Restore
apiVersion: velero.io/v1
metadata:
  name: restore
  namespace: openshift-adp
spec:
  backupName: backup
  restorePVs: true

```

When the restore is completed, you will see that the postgresql app is running on this cluster and is associated with the pvc and a corresponding pv. The state of the app is the same as when the backup was taken.

Project: openshift-adp

Installed Operators > Operator details

 **OADP Operator**
1.4.0 provided by Red Hat



Actions

eLocation | DeleteBackupRequest | DownloadRequest | PodVolumeBackup | PodVolumeRestore | Restore | Schedule | Server


Restores

Create Restore

Name Search by name... /

Name	Kind	Status
 restore	Restore	Phase:  Completed

Activate Windows
Go to Settings to activate Windows.



Restore an App from a backup

This section describes how to restore app(s) from a backup.

Prerequisites

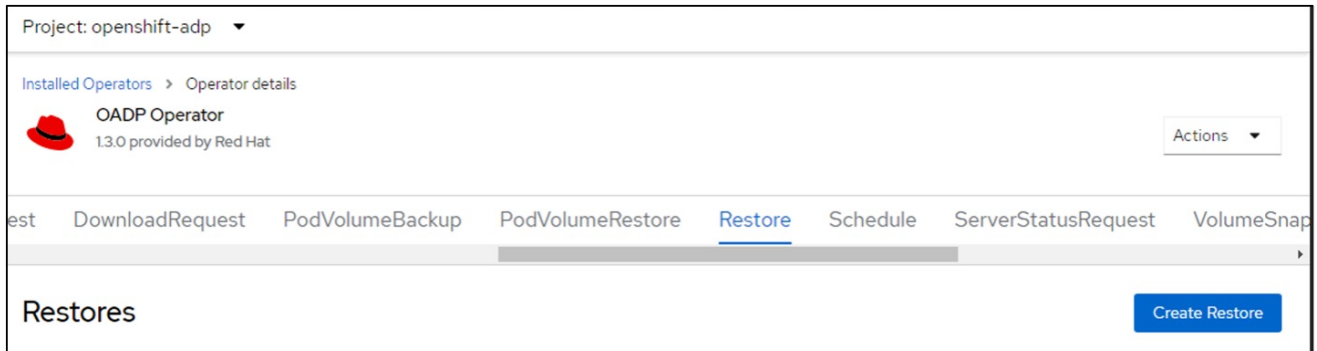
To restore from a backup, let us assume that the namespace where the app existed got accidentally deleted.

```
[root@localhost ~]# oc get pods -n postgresql
NAME          READY   STATUS    RESTARTS   AGE
postgresql-0  1/1    Running   0          102s
[root@localhost ~]# oc delete ns postgresql
namespace "postgresql" deleted

[root@localhost ~]#
[root@localhost ~]#
[root@localhost ~]# oc get pods -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]#
```


Restore to the same namespace

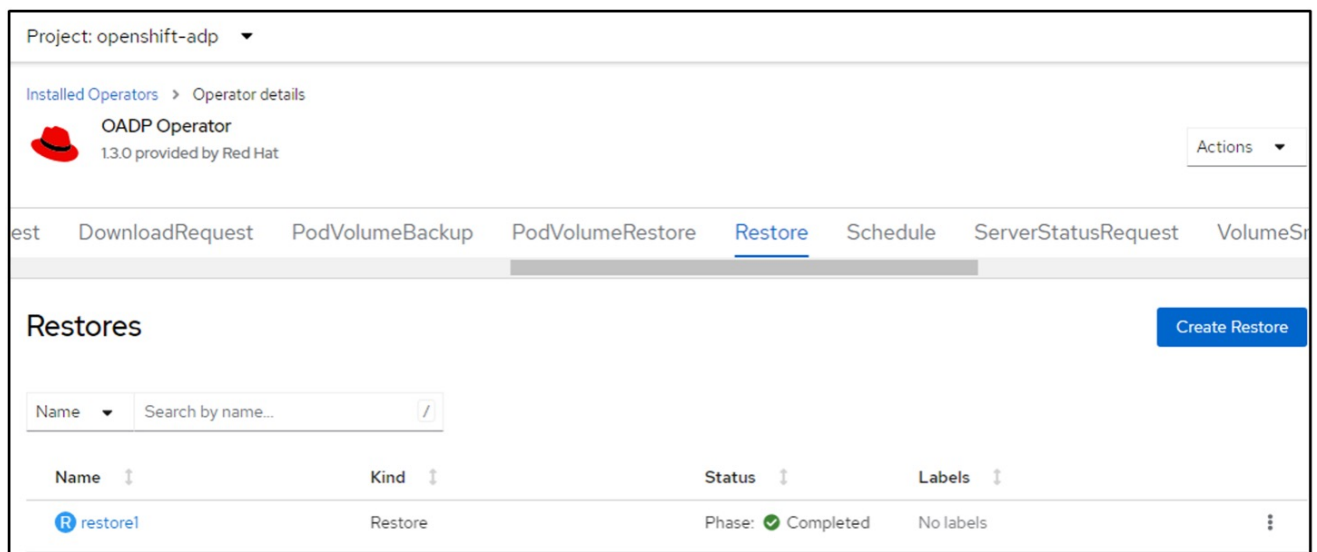
To restore from the backup that we just created, we need to create a Restore Custom Resource (CR). We need to provide it a name, provide the name of the backup that we want to restore from and set the restorePVs to true. Additional parameters can be set as shown in the [documentation](#). Click on Create button.



The screenshot shows the OADP Operator interface for the 'openshift-adp' project. The 'Restore' tab is selected, and a 'Create Restore' button is visible in the top right corner.

```
apiVersion: velero.io/v1
kind: Restore
apiVersion: velero.io/v1
metadata:
  name: restore
  namespace: openshift-adp
spec:
  backupName: backup-postgresql-ontaps3
  restorePVs: true
```

When the phase shows completed, you can see that the app has been restored to the state when the snapshot was taken. The app is restored to the same namespace.



The screenshot shows the OADP Operator interface with a table of restores. The 'restore1' entry is shown with a status of 'Completed'.

Name	Kind	Status	Labels
restore1	Restore	Phase: ✔ Completed	No labels

```
[root@localhost ~]#  
[root@localhost ~]# oc get pods -n postgresql  
No resources found in postgresql namespace.  
[root@localhost ~]# oc get pods -n postgresql  
NAME          READY   STATUS             RESTARTS   AGE  
postgresql-0  0/1    ContainerCreating  0          16s  
[root@localhost ~]# oc get pods -n postgresql  
NAME          READY   STATUS    RESTARTS   AGE  
postgresql-0  0/1    Running   0          22s  
[root@localhost ~]# oc get pods -n postgresql  
NAME          READY   STATUS    RESTARTS   AGE  
postgresql-0  0/1    Running   0          29s  
[root@localhost ~]# oc get pods -n postgresql  
NAME          READY   STATUS    RESTARTS   AGE  
postgresql-0  1/1    Running   0          37s  
[root@localhost ~]#
```

Restore to a different namespace

To restore the App to a different namespace, you can provide a namespaceMapping in the yaml definition of the Restore CR.

The following sample yaml file creates a Restore CR to restore an App and its persistent storage from the postgresql namespace, to the new namespace postgresql-restored.

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: restore-to-different-ns
  namespace: openshift-adp
spec:
  backupName: backup-postgresql-ontaps3
  restorePVs: true
  includedNamespaces:
  - postgresql
  namespaceMapping:
    postgresql: postgresql-restored
```

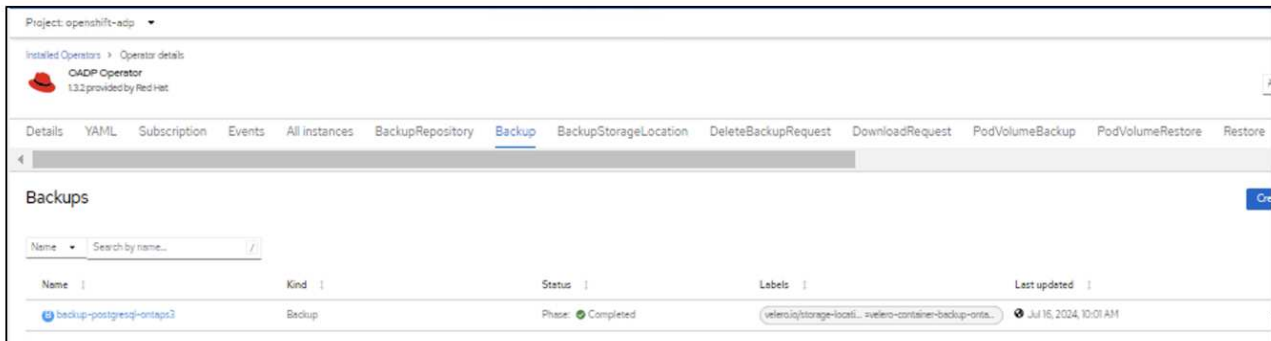

When the phase shows completed, you can see that the app has been restored to the state when the snapshot was taken. The App is restored to a different namespace as specified in the yaml.

```
[root@localhost ~]# oc get pods -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]# oc get pods -n postgresql-restored
NAME          READY   STATUS    RESTARTS   AGE
postgresql-0  0/1    Running   0          19s
[root@localhost ~]# oc get pods -n postgresql-restored
NAME          READY   STATUS    RESTARTS   AGE
postgresql-0  0/1    Running   0          22s
[root@localhost ~]# oc get pods -n postgresql-restored
NAME          READY   STATUS    RESTARTS   AGE
postgresql-0  1/1    Running   0          36s
[root@localhost ~]#
```

Restore to a different storage class

Velero provides a generic ability to modify the resources during restore by specifying json patches. The json patches are applied to the resources before they are restored. The json patches are specified in a configmap and the configmap is referenced in the restore command. This feature enables you to restore using different storage class.

In the example below, the app, during deployment uses ontap-nas as the storage class for its persistent volumes. A backup of the app named backup-postgresql-ontaps3 is created.



Simulate a loss of the app by uninstalling the app.

To restore the VM using a different storage class, for example, ontap-nas-eco storage class, you need to do the following two steps:

Step 1

Create a config map (console) in the openshift-adp namespace as follows:

Fill in the details as shown in the screenshot:

Select namespace : openshift-adp

Name: change-ontap-sc (can be any name)

Key: change-ontap-sc-config.yaml:

Value:

```
version: v1
resourceModifierRules:
- conditions:
  groupResource: persistentvolumeclaims
  resourceNameRegex: "data-postgresql*"
  namespaces:
  - postgresql
patches:
- operation: replace
  path: "/spec/storageClassName"
  value: "ontap-nas-eco"
```

Project: openshift-adp ▾

Edit ConfigMap

Config maps hold key-value pairs that can be used in pods to read application configuration.

Configure via: Form view YAML view

Name *

change-ontap-sc

A unique name for the ConfigMap within the project

Immutable
Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated

Data

Data contains the configuration data that is in UTF-8 range

[Remove key/value](#)

Key *

change-ontap-sc.yaml

Value

[Browse...](#)

Drag and drop file with your value here or browse to upload it.

```

version: v1
resourceModifierRules:
- conditions:
  groupResource: persistentvolumeclaims
  resourceNameRegex: "data-postgresql*"
  namespaces:
  - postgresql
patches:
- operation: replace
  path: "/spec/storageClassName"
  value: "ontap-nas-eco"

```

The resulting config map object should look like this (CLI):

```

[root@localhost ~]# kubectl describe cm/change-ontap-sc -n openshift-adp
Name:          change-ontap-sc
Namespace:     openshift-adp
Labels:        <none>
Annotations:   <none>

Data
====
change-ontap-sc.yaml:
----
version: v1
resourceModifierRules:
- conditions:
  groupResource: persistentvolumeclaims
  resourceNameRegex: "data-postgresql*"
  namespaces:
  - postgresql
  patches:
  - operation: replace
    path: "/spec/storageClassName"
    value: "ontap-nas-eco"

BinaryData
====

Events: <none>
[root@localhost ~]# █

```

This config map will apply the resource modifier rule when the restore is created. A patch will be applied to replace the storage class name to ontap-nas-eco for all persistent volume claims starting with rhel.

Step 2

To restore the VM use the following command from the Velero CLI:

```
#velero restore create restore1 --from-backup backup1 --resource
-modifier-configmap change-storage-class-config -n openshift-adp
```

The app is restored in the same namespace with the persistent volume claims created using the storage class ontap-nas-eco.

```

[root@localhost ~]# oc get pods -n postgresql
NAME          READY  STATUS   RESTARTS  AGE
postgresql-0  1/1    Running  0          11m
[root@localhost ~]# oc get pvc -n postgresql
NAME          STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
data-postgresql-0  Bound  pvc-33526ea4-37c2-4180-a9f6-fb47aea3b4e2  8Gi       RWO           ontap-nas-eco  11m
[root@localhost ~]# █

```

Deleting backups and restores in using Velero

This section outlines how to delete backups and restores of Apps in OpenShift container platform using Velero.

List all backups

You can list all Backup CRs by using the OC CLI tool or the Velero CLI tool. Download the Velero CLI as given in the instructions in the [Velero documentation](#).

```
[root@localhost ~]# oc get backups -n openshift-adp
NAME                AGE
backup-postgresql-ontaps3 23h
backup2              26s
schedule1-20240717070005 6h42m
[root@localhost ~]# velero get backups -n openshift-adp
NAME                STATUS    ERRORS    WARNINGS    CREATED                EXPIRES    STORAGE LOCATION    SELECTOR
backup-postgresql-ontaps3  Completed  0         0           2024-07-16 10:01:08 -0400 EDT  29d        velero-container-backup-ontap-1  <none>
backup2                Completed  0         0           2024-07-17 09:42:32 -0400 EDT  29d        velero-container-backup-ontap-1  <none>
schedule1-20240717070005  Completed  0         0           2024-07-17 03:00:05 -0400 EDT  29d        velero-container-backup-ontap-1  <none>
[root@localhost ~]#
```

Deleting a backup

You can delete a Backup CR without deleting the Object Storage data by using the OC CLI tool. The backup will be removed from the CLI/Console output. However, since the corresponding backup is not removed from the object storage, it will re-appear in the CLI/console output.

```
[root@localhost ~]# oc delete backup backup2 -n openshift-adp
backup.velero.io "backup2" deleted
[root@localhost ~]# oc get backups -n openshift-adp
NAME                AGE
backup-postgresql-ontaps3 23h
schedule1-20240717070005 6h49m
[root@localhost ~]# oc get backups -n openshift-adp
NAME                AGE
backup-postgresql-ontaps3 23h
backup2              24s
schedule1-20240717070005 6h50m
[root@localhost ~]#
```

If you want to delete the Backup CR AND the associated object storage data, you can do so by using the Velero CLI tool.


```
[root@localhost ~]# velero get backups -n openshift-adp
```

NAME	STATUS	ERRORS	WARNINGS	CREATED	EXPIRES	STORAGE LOCATION	SELECTOR
backup-postgresql-ontaps3	Completed	0	0	2024-07-16 10:01:08 -0400 EDT	29d	velero-container-backup-ontap-1	<none>
backup2	Completed	0	0	2024-07-17 09:42:32 -0400 EDT	29d	velero-container-backup-ontap-1	<none>
schedule1-20240717070005	Completed	0	0	2024-07-17 03:00:05 -0400 EDT	29d	velero-container-backup-ontap-1	<none>

```
[root@localhost ~]# velero delete backup backup2 -n openshift-adp
Are you sure you want to continue (Y/N)? Y
Request to delete backup "backup2" submitted successfully.
The backup will be fully deleted after all associated data (disk snapshots, backup files, restores) are removed.
[root@localhost ~]# velero get backups -n openshift-adp
```

NAME	STATUS	ERRORS	WARNINGS	CREATED	EXPIRES	STORAGE LOCATION	SELECTOR
backup-postgresql-ontaps3	Completed	0	0	2024-07-16 10:01:08 -0400 EDT	29d	velero-container-backup-ontap-1	<none>
schedule1-20240717070005	Completed	0	0	2024-07-17 03:00:05 -0400 EDT	29d	velero-container-backup-ontap-1	<none>

```
[root@localhost ~]#
```

Deleting the Restore

You can delete the Restore CR Object by using either the OC CLI or the Velero CLI

```
[root@localhost ~]# velero get restore -n openshift-adp
```

NAME	BACKUP	STATUS	STARTED	COMPLETED	ERRORS	WARNINGS	CREATED	SELECTOR
restore	backup-postgresql-ontaps3	Completed	2024-07-16 14:59:22 -0400 EDT	2024-07-16 14:59:45 -0400 EDT	0	10	2024-07-16 14:59:22 -0400 EDT	<none>
restore1	backup-postgresql-ontaps3	Completed	2024-07-16 16:36:37 -0400 EDT	2024-07-16 16:36:59 -0400 EDT	0	9	2024-07-16 16:36:37 -0400 EDT	<none>

```
[root@localhost ~]# velero restore delete restore1 -n openshift-adp
Are you sure you want to continue (Y/N)? Y
Request to delete restore "restore1" submitted successfully.
The restore will be fully deleted after all associated data (restore files in object storage) are removed.
[root@localhost ~]# velero get restore -n openshift-adp
```

NAME	BACKUP	STATUS	STARTED	COMPLETED	ERRORS	WARNINGS	CREATED	SELECTOR
restore	backup-postgresql-ontaps3	Completed	2024-07-16 14:59:22 -0400 EDT	2024-07-16 14:59:45 -0400 EDT	0	10	2024-07-16 14:59:22 -0400 EDT	<none>

```
[root@localhost ~]#
[root@localhost ~]# oc delete restore restore -n openshift-adp
restore.velero.io "restore" deleted
[root@localhost ~]# oc get restore -n openshift-adp
No resources found in openshift-adp namespace.
[root@localhost ~]# velero get restore -n openshift-adp
[root@localhost ~]#
```

Activate Windows

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.