



OpenShift Virtualization on premises

NetApp Solutions

NetApp

January 03, 2025

This PDF was generated from https://docs.netapp.com/us-en/netapp-solutions/containers/rh-os-n_use_case_openshift_virtualization_deployment_prerequisites.html on January 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- OpenShift Virtualization on premises 1
 - Deploy Red Hat OpenShift Virtualization with NetApp ONTAP 1
 - Deploy Red Hat OpenShift Virtualization with NetApp ONTAP 1
- Workflows 5

OpenShift Virtualization on premises

Deploy Red Hat OpenShift Virtualization with NetApp ONTAP

This section details how to deploy Red Hat OpenShift Virtualization with NetApp ONTAP.

Prerequisites

- A Red Hat OpenShift cluster (later than version 4.6) installed on bare-metal infrastructure with RHCOS worker nodes
- The OpenShift cluster must be installed via installer provisioned infrastructure (IPI)
- Deploy Machine Health Checks to maintain HA for VMs
- A NetApp ONTAP cluster
- Trident installed on the OpenShift cluster
- A Trident backend configured with an SVM on ONTAP cluster
- A StorageClass configured on the OpenShift cluster with Trident as the provisioner
- Cluster-admin access to Red Hat OpenShift cluster
- Admin access to NetApp ONTAP cluster
- An admin workstation with tridentctl and oc tools installed and added to \$PATH

Because OpenShift Virtualization is managed by an operator installed on the OpenShift cluster, it imposes additional overhead on memory, CPU, and storage, which must be accounted for while planning the hardware requirements for the cluster. See the documentation [here](#) for more details.


Optionally, you can also specify a subset of the OpenShift cluster nodes to host the OpenShift Virtualization operators, controllers, and VMs by configuring node placement rules. To configure node placement rules for OpenShift Virtualization, follow the documentation [here](#).

For the storage backing OpenShift Virtualization, NetApp recommends having a dedicated StorageClass that requests storage from a particular Trident backend, which in turn is backed by a dedicated SVM. This maintains a level of multitenancy with regard to the data being served for VM-based workloads on the OpenShift cluster.


Deploy Red Hat OpenShift Virtualization with NetApp ONTAP

To install OpenShift Virtualization, complete the following steps:

1. Log into the Red Hat OpenShift bare-metal cluster with cluster-admin access.
2. Select Administrator from the Perspective drop down.
3. Navigate to Operators > OperatorHub and search for OpenShift Virtualization.



4. Select the OpenShift Virtualization tile and click Install.



OpenShift Virtualization

2.6.2 provided by Red Hat

✕

Install

Latest version

2.6.2

Capability level

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

Provider type

Red Hat

Provider

Red Hat

Requirements

Your cluster must be installed on bare metal infrastructure with Red Hat Enterprise Linux CoreOS workers.

Details

OpenShift Virtualization extends Red Hat OpenShift Container Platform, allowing you to host and manage virtualized workloads on the same platform as container-based workloads. From the OpenShift Container Platform web console, you can import a VMware virtual machine from vSphere, create new or clone existing VMs, perform live migrations between nodes, and more. You can use OpenShift Virtualization to manage both Linux and Windows VMs.

The technology behind OpenShift Virtualization is developed in the [KubeVirt](#) open source community. The KubeVirt project extends [Kubernetes](#) by adding additional virtualization resource types through [Custom Resource Definitions](#) (CRDs). Administrators can use Custom Resource Definitions to manage [VirtualMachine](#) resources alongside all other resources that Kubernetes provides.

5. On the Install Operator screen, leave all default parameters and click Install.

Update channel *

- ☐ 2.1
- ☐ 2.2
- ☐ 2.3
- ☐ 2.4
- ☒ stable

Installation mode *

- ☐ All namespaces on the cluster (default)
This mode is not supported by this Operator
- ☒ A specific namespace on the cluster
Operator will be available in a single Namespace only.

Installed Namespace *

- ☒ Operator recommended Namespace: **PR** openshift-cnv

i Namespace creation

Namespace **openshift-cnv** does not exist and will be created.


- ☐ Select a Namespace

Approval strategy *

- ☒ Automatic
- ☐ Manual

Install

Cancel

 OpenShift Virtualization
provided by Red Hat

Provided APIs

HC OpenShift
Virtualization
Deployment

Required

Represents the deployment of
OpenShift Virtualization

6. Wait for the operator installation to complete.



OpenShift Virtualization
2.6.2 provided by Red Hat



Installing Operator

The Operator is being installed. This may take a few minutes.

[View installed Operators in Namespace openshift-cnv](#)

7. After the operator has installed, click Create HyperConverged.



Installed operator – operand required

The Operator has installed successfully. Create the required custom resource to be able to use this Operator.



HyperConverged



Required

Creates and maintains an OpenShift Virtualization Deployment

Create HyperConverged

[View installed Operators in Namespace openshift-cnv](#)

- On the Create HyperConverged screen, click Create, accepting all default parameters. This step starts the installation of OpenShift Virtualization.

Name *

Labels

Infra >

infra HyperConvergedConfig influences the pod configuration (currently only placement) for all the infra components needed on the virtualization enabled cluster but not necessarily directly on each node running VMs/VMLs.

Workloads >

workloads HyperConvergedConfig influences the pod configuration (currently only placement) of components which need to be running on a node where virtualization workloads should be able to run. Changes to Workloads HyperConvergedConfig can be applied only without existing workload.

Bare Metal Platform

☒ true

BareMetalPlatform indicates whether the infrastructure is baremetal.

Feature Gates >

featureGates is a map of feature gate flags. Setting a flag to `true` will enable the feature. Setting `false` or removing the feature gate, disables the feature.

Local Storage Class Name





LocalStorageClassName the name of the local storage class.

- After all the pods move to the Running state in the openshift-cnv namespace and the OpenShift Virtualization operator is in the Succeeded state, the operator is ready to use. VMs can now be created on the OpenShift cluster.

Project: openshift-cnv ▼

Installed Operators

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

Name ▼	Managed Namespaces ↑	Status	Last updated	Provided APIs
 OpenShift Virtualization 2.6.2 provided by Red Hat	 openshift-cnv	 Succeeded Up to date	 May 18, 8:02 pm	OpenShift Virtualization Deployment HostPathProvisioner deployment

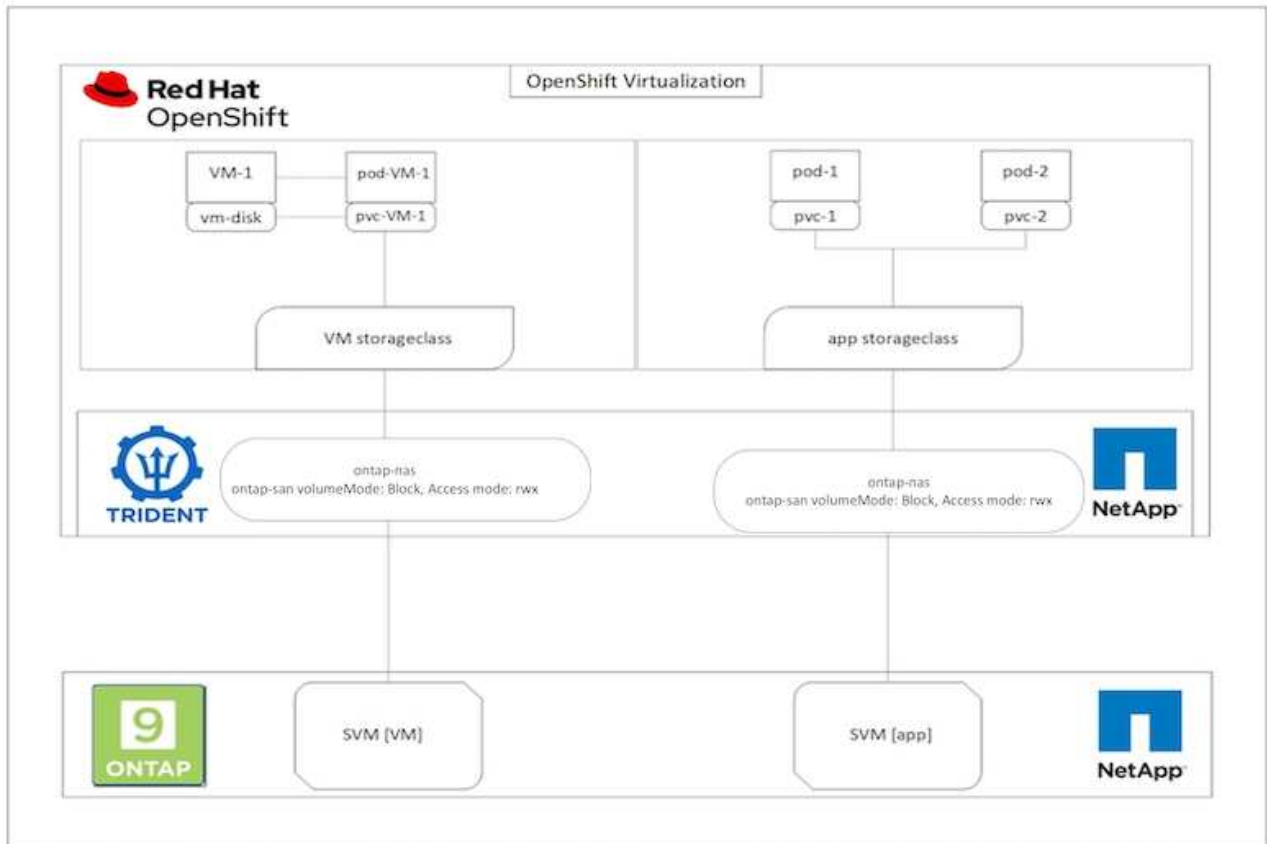
Workflows

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section covers the how to create a virtual machine with Red Hat OpenShift Virtualization.

Create VM

VMs are stateful deployments that require volumes to host the operating system and data. With CNV, because the VMs are run as pods, the VMs are backed by PVs hosted on NetApp ONTAP through Trident. These volumes are attached as disks and store the entire filesystem including the boot source of the VM.



To quickly create a virtual machine on the OpenShift cluster, complete the following steps:

1. Navigate to Virtualization > Virtual Machines and click Create.
2. Select From template.
3. Select the desired operating system for which the boot source is available.
4. Check the checkbox Start the VirtualMachine after creation.
5. Click Quick create VirtualMachine.

The virtual machine is created and started and comes to the **Running** state. It automatically creates a PVC and a corresponding PV for the boot disk using the default storage class. In order to be able to live migrate the VM in the future, you must ensure that the storage class used for the disks can support RWX volumes. This is a requirement for live migration. **ontap-nas** and **ontap-san** (volumeMode block for iSCSI and NVMe/TCP protocols) can support RWX access modes for the volumes created using the respective storage classes.

To configure **ontap-san** storage class on the cluster see the [Section for Migrating a VM from VMware to OpenShift Virtualization](#).



You can set up ontap NAS or iSCSI as the default storage class for the cluster. Clicking on Quick create VirtualMachine will use the default storage class to create the PVC and PV for the bootable root disk for the VM. If your default storage class is not ontap-nas or ontap-san, you can select the storage class for the disk, by selecting **Customize VirtualMachine > Customize VirtualMachine parameters > Disks** and then editing the disk to use the required storage class.

Typically block access mode is preferred compared to file systems while provisioning the VM disks.

To customize the virtual machine creation after you have selected the OS template, click on **Customize VirtualMachine** instead of **Quick create**.

1. If the selected operating system has boot source configured, you can click on **Customize VirtualMachine parameters**.
2. If the selected operating system has no boot source configured, you must configure it. You can see details about the procedures shown in the [documentation](#).
3. After Configuring the boot disk, you can click on **Customize VirtualMachine parameters**.
4. You can customize the VM from the tabs on this page. For eg. click on the **Disks** tab and then click on **Add disk** to add another disk to the VM.
5. Click **Create Virtual Machine** to create the virtual machine; this spins up a corresponding pod in the background.



When a boot source is configured for a template or an operating system from an URL or from a registry, it creates a PVC in the `openshift-virtualization-os-images` project and downloads the KVM guest image to the PVC. You must make sure that template PVCs have enough provisioned space to accommodate the KVM guest image for the corresponding OS. These PVCs are then cloned and attached as rootdisk to virtual machines when they are created using the respective templates in any project.

The screenshot shows the OpenShift Virtualization console interface. On the left is a sidebar with navigation links: Administrator, Home, Operators, Workloads, Virtualization (selected), Overview, Catalog, VirtualMachines (active), Templates, and Instance Types. The main panel displays the 'VirtualMachines' page for the 'openshift-virtualization-os-images' project. It includes a table with columns: Name, Status, Conditions, Node, and IP address. There are four VMs listed, all with a status of 'Running'. A 'Create' button is visible in the top right corner, highlighted with a blue box. Below the 'Create' button, there are options: 'From template' and 'From volume With YAML'.

Name	Status	Conditions	Node	IP address
centos-stream8-hissing-artester	Running		ocp-worker3	10.130.0.143
centos-stream8-improved-kill	Running		ocp-worker3	10.130.0.145
centos-stream8-weary-toucan	Running		ocp-worker3	10.130.0.123
centos-stream8-zealous-anaconda	Running		ocp-worker3	10.130.0.117

Create new VirtualMachine

Select an option to create a VirtualMachine from.

Template catalog

InstanceTypes

Template project

All projects

All items

Default templates

User templates

☐ Boot source available

Operating system

☐ CentOS
 ☐ Fedora
 ☐ Other
 ☐ RHEL
 ☐ Windows

Workload

☐ Desktop
 ☐ High performance
 ☐ Server

Default templates

Q

Filter by keyword...

13 items

Source available

CentOS Stream 8 VM

centos-stream8-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

CentOS Stream 9 VM

centos-stream9-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

CentOS 7 VM

centos7-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

Fedora VM

fedora-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

Red Hat Enterprise Linux 7 VM

rhel7-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

Red Hat Enterprise Linux 8 VM

rhel8-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

Red Hat Enterprise Linux 9 VM

rhel9-server-small

Project openshift

Boot source PVC (auto import)

Workload Server

CPU 1

Memory 2 GiB

Source available

Microsoft Windows 10 VM

windows10-desktop-medium

Project openshift

Boot source PVC

Workload Desktop

CPU 1

Memory 4 GiB

Source available

Microsoft Windows 11 VM

windows11-desktop-medium

Project openshift

Boot source PVC

Workload Desktop

CPU 2

Memory 4 GiB

Source available

Microsoft Windows Server 2012 R2 VM

windows2k12r2-server-medium

Project openshift

Boot source PVC

Workload Server

CPU 1

Memory 4 GiB



CentOS Stream 9 VM

centos-stream9-server-small



Template info

Operating system

CentOS Stream 9 VM

Workload type

Server (default)

Description

Template for CentOS Stream 9 VM or newer. A PVC with the CentOS Stream disk image must be available.

Documentation

[Refer to documentation](#)

CPU | Memory

1 CPU | 2 GiB Memory

Network interfaces (1)

Name	Network	Type
default	Pod networking	Masquerade

Disks (2)

Name	Drive	Size
rootdisk	Disk	30 GiB
cloudinitdisk	Disk	-

Hardware devices (0)

GPU devices

Not available

Host devices

Not available

Quick create VirtualMachine

VirtualMachine name *

centos-stream9-pleased-ham...

Project

openshift-virtualization-os-images

☒ Start this VirtualMachine after creation

Quick create VirtualMachine

Customize VirtualMachine

Cancel

Activate Windows

Go to Settings to activate Windows.

Project: openshift-virtualization-os-images

Catalog > Customize template parameters > Customize VirtualMachine

Customize and create VirtualMachine

Template: CentOS Stream 9 VM

Overview YAML Scheduling Environment Network interfaces Disks Scripts Metadata

Name
centos-stream9-pleased-hamster

Namespace
openshift-virtualization-os-images

Description
Not available

Operating system
CentOS Stream 9 VM

CPU | Memory
1 CPU | 2 GiB Memory

Machine type
pc-q35-rhel9.2.0

Boot mode
BIOS

Start in pause mode
☐

Workload profile
Server

Network interfaces (1)

Name	Network	Type
default	Pod networking	Masquerade

Disks (2)

Name	Drive	Size
rootdisk	Disk	30 GiB
cloudinitdisk	Disk	-

Hardware devices

GPU devices
Not available

Host devices
Not available

Headless mode
☐

Hostname
centos-stream9-pleased-hamster

☒ Start this VirtualMachine after creation

Create VirtualMachine Cancel

VirtualMachines > VirtualMachine details

VM centos-stream9-zealous-anaconda Running VAML Actions

Overview Details Metrics VAML **Configuration** Events Console Snapshots Diagnostics

Disks +

Add disk

Network interfaces +

Scheduling +

Environment +

Scripts +

Name	Source	Size	Drive	Interface	Storage class
cloudinitdisk	Other	-	Disk	virtio	-
data-disk1 (Persistent Hotplug)	PVC centos-stream9-zealous-anaconda-data-disk1	30.00 GiB	Disk	SCSI	ontap-san-block
rootdisk (bootable)	PVC centos-stream9-zealous-anaconda	30.00 GiB	Disk	virtio	ontap-san-block

File systems +

Name	File system type	Mount point	Total bytes	Used bytes
vdal	xfs	/	29.94 GiB	1.30 GiB

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

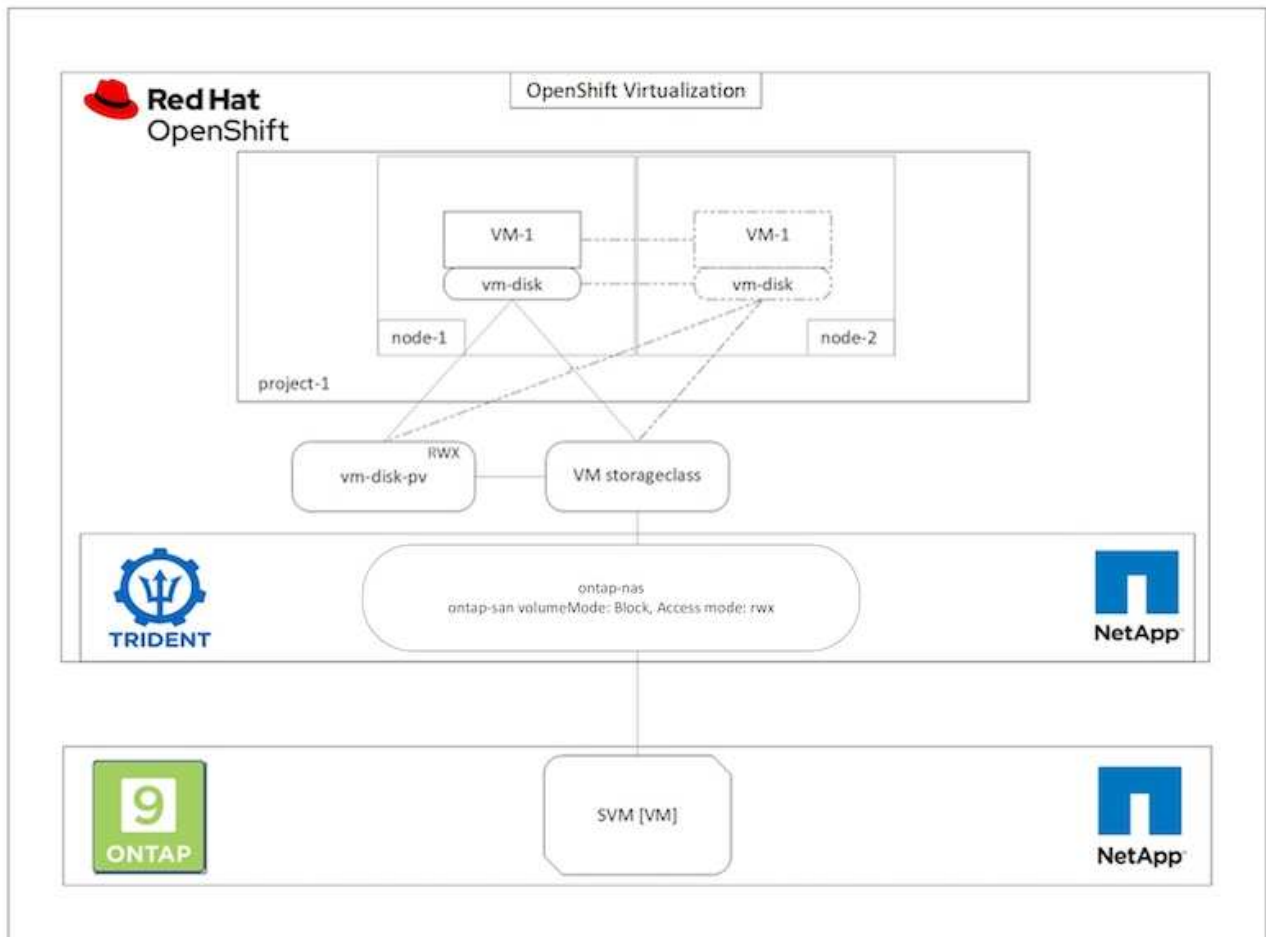
This section shows how to migrate a virtual machine in OpenShift Virtualization between nodes in the cluster .

VM Live Migration

Live Migration is a process of migrating a VM instance from one node to another in an OpenShift cluster with

no downtime. For live migration to work in an OpenShift cluster, VMs must be bound to PVCs with shared ReadWriteMany access mode. Trident backends configured using ontap-nas drivers support RWX access mode for FileSystem protocols nfs and smb. Refer to the documentation [here](#). Trident backends configured using ontap-san drivers support RWX access mode for block volumeMode for iSCSI and NVMe/TCP protocols. Refer to the documentation [here](#).

Therefore, for live migration to succeed, the VMs must be provisioned with disks (boot disks and additional hot plug disks) with PVCs using ontap-nas or ontap-san (volumeMode: Block) storage classes. When the PVCs are created, Trident creates ONTAP volumes in an SVM which is NFS-enabled or iSCSI enabled.



To perform a live migration of a VM that has been created previously and is in a Running state perform the following steps:

1. Select the VM that you want to live-migrate.
2. Click on **Configuration** tab.
3. Ensure that all the disks of the VM are created using Storage classes that can support RWX access mode.
4. Click on **Actions** on the right corner and then select **Migrate**.
5. To look at the progression of the Migration, go to Virtualization > Overview on the left hand side menu and then click on the **Migrations** tab.

The Migration of the VM will transition from **Pending** to **Scheduling** to **Succeeded**



A VM instance in an OpenShift cluster automatically migrates to another node when the original node is placed into maintenance mode if the evictionStrategy is set to LiveMigrate.

VirtualMachines > VirtualMachine details

VM centos-stream9-zealous-anaconda Running

Overview Details Metrics YAML Configuration Events Console Snapshots Diagnostics

Disks

Network interfaces

Scheduling

Environment

Scripts

File systems

Name	Source	Size	Drive	Interface	Storage class
cloudinitdisk	Other	-	Disk	virtio	-
data-disk (Persistent Hotplug)	PVC centos-stream9-zealous-anaconda-data-disk	30.00 GiB	Disk	SCSI	ontap-san-bk
rootdisk (bootable)	PVC centos-stream9-zealous-anaconda	30.00 GiB	Disk	virtio	ontap-san-block

File systems

Name	File system type	Mount point	Total bytes	Used bytes
vdal	xfs	/	29.94 GiB	130 GiB

Home Operators Workloads Virtualization Overview Catalog VirtualMachines Templates InstanceTypes Preferences Bootable volumes MigrationPolicies Networking Storage Builds Observe Compute

Virtualization Overview Top consumer Migrations Settings

VirtualMachineInstanceMigrations information

Migrations

Limitations

Bandwidth consumption

1 Migrations

1 Succeeded


VirtualMachine name	Status	Source	Target	MigrationPolicy	VirtualMachineInstanceMigration...
centos-stream9-zealous-anaconda	Succeeded	ocp-worker2	ocp-worker3	No MigrationPolicy	centos-stream9-zealous-anaconda-migration-cp0tc

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section covers the how to clone a virtual machine with Red Hat OpenShift Virtualization.

VM cloning

Cloning an existing VM in OpenShift is achieved with the support of Trident's Volume CSI cloning feature. CSI volume cloning allows for creation of a new PVC using an existing PVC as the data source by duplicating its PV. After the new PVC is created, it functions as a separate entity and without any link to or dependency on the source PVC.



There are certain restrictions with CSI volume cloning to consider:

1. Source PVC and destination PVC must be in the same project.
2. Cloning is supported within the same storage class.
3. Cloning can be performed only when source and destination volumes use the same VolumeMode setting; for example, a block volume can only be cloned to another block volume.

VMs in an OpenShift cluster can be cloned in two ways:

1. By shutting down the source VM
2. By keeping the source VM live

By Shutting down the source VM

Cloning an existing VM by shutting down the VM is a native OpenShift feature that is implemented with support from Trident. Complete the following steps to clone a VM.

1. Navigate to Workloads > Virtualization > Virtual Machines and click the ellipsis next to the virtual machine you wish to clone.
2. Click Clone Virtual Machine and provide the details for the new VM.

Clone Virtual Machine

Name *

rhel8-short-frog-clone

Description

Namespace *

default



Start virtual machine on clone

Configuration

Operating System

Red Hat Enterprise Linux 8.0 or higher

Flavor

Small: 1 CPU | 2 GiB Memory

Workload Profile

server

NICs

default - virtio

Disks

cloudinitdisk - cloud-init disk

rootdisk - 20Gi - basic



The VM rhel8-short-frog is still running. It will be powered off while cloning.

Cancel

Clone Virtual Machine

3. Click Clone Virtual Machine; this shuts down the source VM and initiates the creation of the clone VM.
4. After this step is completed, you can access and verify the content of the cloned VM.

By keeping the source VM live

An existing VM can also be cloned by cloning the existing PVC of the source VM and then creating a new VM using the cloned PVC. This method does not require you to shut down the source VM. Complete the following steps to clone a VM without shutting it down.

1. Navigate to Storage > PersistentVolumeClaims and click the ellipsis next to the PVC that is attached to the source VM.
2. Click Clone PVC and furnish the details for the new PVC.

Clone

Name *

rhel8-short-frog-rootdisk-28dvv-clone

Access Mode *

☐ Single User (RWO) ☒ Shared Access (RWX) ☐ Read Only (ROX)

Size *

20

GiB ▼

PVC details

Namespace

 default

Requested capacity

20 GiB

Access mode

Shared Access (RWX)

Storage Class

 basic

Used capacity

2.2 GiB

Volume mode

Filesystem

Cancel

Clone

3. Then click Clone. This creates a PVC for the new VM.
4. Navigate to Workloads > Virtualization > Virtual Machines and click Create > With YAML.
5. In the spec > template > spec > volumes section, attach the cloned PVC instead of the container disk. Provide all other details for the new VM according to your requirements.

```
- name: rootdisk
  persistentVolumeClaim:
    claimName: rhel8-short-frog-rootdisk-28dvvb-clone
```

6. Click Create to create the new VM.
7. After the VM is created successfully, access and verify that the new VM is a clone of the source VM.

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

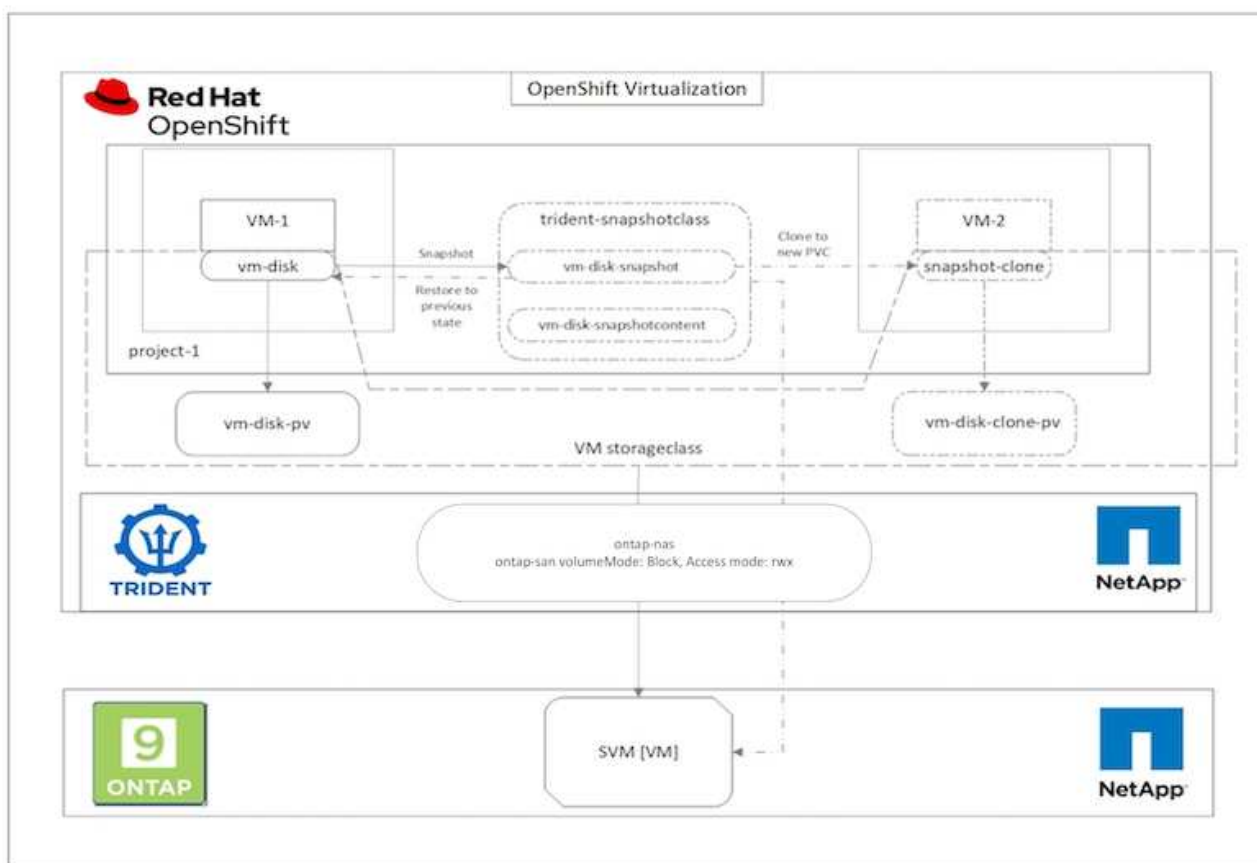
This section shows how to create a virtual machine from a Snapshot with Red Hat OpenShift Virtualization.

Create VM from a Snapshot

With Trident and Red Hat OpenShift, users can take a snapshot of a persistent volume on Storage Classes provisioned by it. With this feature, users can take a point-in-time copy of a volume and use it to create a new volume or restore the same volume back to a previous state. This enables or supports a variety of use-cases, from rollback to clones to data restore.

For Snapshot operations in OpenShift, the resources `VolumeSnapshotClass`, `VolumeSnapshot`, and `VolumeSnapshotContent` must be defined.

- A `VolumeSnapshotContent` is the actual snapshot taken from a volume in the cluster. It is cluster-wide resource analogous to `PersistentVolume` for storage.
- A `VolumeSnapshot` is a request for creating the snapshot of a volume. It is analogous to a `PersistentVolumeClaim`.
- `VolumeSnapshotClass` lets the administrator specify different attributes for a `VolumeSnapshot`. It allows you to have different attributes for different snapshots taken from the same volume.



To create Snapshot of a VM, complete the following steps:

1. Create a VolumeSnapshotClass that can then be used to create a VolumeSnapshot. Navigate to Storage > VolumeSnapshotClasses and click Create VolumeSnapshotClass.
2. Enter the name of the Snapshot Class, enter `csi.trident.netapp.io` for the driver, and click Create.

```
1 apiVersion: snapshot.storage.k8s.io/v1
2 kind: VolumeSnapshotClass
3 metadata:
4   name: trident-snapshot-class
5 driver: csi.trident.netapp.io
6 deletionPolicy: Delete
7
```

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

3. Identify the PVC that is attached to the source VM and then create a Snapshot of that PVC. Navigate to Storage > VolumeSnapshots and click Create VolumeSnapshots.
4. Select the PVC that you want to create the Snapshot for, enter the name of the Snapshot or accept the default, and select the appropriate VolumeSnapshotClass. Then click Create.

Create VolumeSnapshot

[Edit YAML](#)

PersistentVolumeClaim *

PVC rhel8-short-frog-rootdisk-28dvb ▼

Name *

rhel8-short-frog-rootdisk-28dvb-snapshot

Snapshot Class *

VSC trident-snapshot-class ▼

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

5. This creates the snapshot of the PVC at that point in time.

Create a new VM from the snapshot

1. First, restore the Snapshot into a new PVC. Navigate to Storage > VolumeSnapshots, click the ellipsis next to the Snapshot that you wish to restore, and click Restore as new PVC.
2. Enter the details of the new PVC and click Restore. This creates a new PVC.

Restore as new PVC

When restore action for snapshot **rhel8-short-frog-rootdisk-28dvb-snapshot** is finished a new crash-consistent PVC copy will be created.

Name *

rhel8-short-frog-rootdisk-28dvb-snapshot-restore

Storage Class *

SC basic

Access Mode *

☐ Single User (RWO) ☒ Shared Access (RWX) ☐ Read Only (ROX)

Size *

20

GiB

VolumeSnapshot details

Created at

 May 21, 12:46 am

Namespace

NS default

Status

 Ready

API version

snapshot.storage.k8s.io/v1

Size

20 GiB

3. Next, create a new VM from this PVC. Navigate to Virtualization > Virtual Machines and click Create > With YAML.
4. In the spec > template > spec > volumes section, specify the new PVC created from Snapshot instead of

from the container disk. Provide all other details for the new VM according to your requirements.

```
- name: rootdisk
  persistentVolumeClaim:
    claimName: rhel8-short-frog-rootdisk-28dvh-snapshot-restore
```

5. Click Create to create the new VM.
6. After the VM is created successfully, access and verify that the new VM has the same state as that of the VM whose PVC was used to create the snapshot at the time when the snapshot was created.

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section covers the how to migrate a virtual machine between from VMware to an OpenShift Cluster using Red Hat OpenShift Virtualization migration toolkit.

Migration of VM from VMware to OpenShift Virtualization using Migration Toolkit for Virtualization

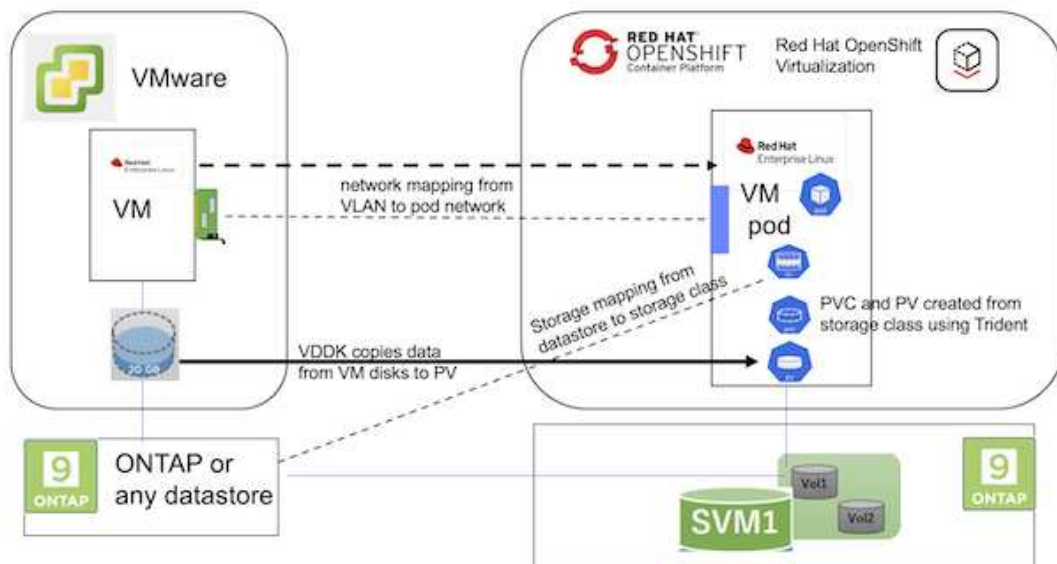
In this section, we will see how to use the Migration Toolkit for Virtualization (MTV) to migrate virtual machines from VMware to OpenShift Virtualization running on OpenShift Container platform and integrated with NetApp ONTAP storage using Trident.

The following video shows a demonstration of the migration of a RHEL VM from VMware to OpenShift Virtualization using ontap-san storage class for persistent storage.

[Using Red Hat MTV to migrate VMs to OpenShift Virtualization with NetApp ONTAP Storage](#)

The following diagram shows a high level view of the migration of a VM from VMware to Red Hat OpenShift Virtualization.

Migration of VM from VMware to OpenShift Virtualization



Prerequisites for the sample migration

On VMware

- A RHEL 9 VM using rhel 9.3 with the following configurations were installed:
 - CPU: 2, Memory: 20 GB, Hard disk: 20 GB
 - user credentials: root user and an admin user credentials
- After the VM was ready, postgresql server was installed.
 - postgresql server was started and enabled to start on boot

```
systemctl start postgresql.service`  
systemctl enable postgresql.service  
The above command ensures that the server can start in the VM in  
OpenShift Virtualization after migration
```

- Added 2 databases, 1 table and 1 row in the table were added. Refer [here](#) for the instructions for installing postgresql server on RHEL and creating database and table entries.



Ensure that you start the postgresql server and enable the service to start at boot.

On OpenShift Cluster

The following installations were completed before installing MTV:

- OpenShift Cluster 4.13.34
- [Trident 23.10](#)
- Multipath on the cluster nodes enabled for iSCSI (for ontap-san storage class). See the provided yaml to create a daemon set that enables iSCSI on each node in the cluster.
- Trident backend and Storage class for ontap SAN using iSCSI. See the provided yaml files for trident backend and storage class.
- [OpenShift Virtualization](#)

To install iscsi and multipath on the OpenShift Cluster nodes use the yaml file given below

Preparing the cluster nodes for iSCSI

```
apiVersion: apps/v1  
kind: DaemonSet  
metadata:  
  namespace: trident  
  name: trident-iscsi-init  
  labels:  
    name: trident-iscsi-init  
spec:  
  selector:  
    matchLabels:  
      name: trident-iscsi-init
```

```

template:
  metadata:
    labels:
      name: trident-iscsi-init
  spec:
    hostNetwork: true
    serviceAccount: trident-node-linux
    initContainers:
      - name: init-node
        command:
          - nsenter
          - --mount=/proc/1/ns/mnt
          - --
          - sh
          - -c
        args: ["$(STARTUP_SCRIPT)"]
        image: alpine:3.7
        env:
          - name: STARTUP_SCRIPT
            value: |
              #! /bin/bash
              sudo yum install -y lsscsi iscsi-initiator-utils sg3_utils
              device-mapper-multipath
              rpm -q iscsi-initiator-utils
              sudo sed -i 's/^\(node.session.scan\).*$/\1 = manual/'
              /etc/iscsi/iscsid.conf
              cat /etc/iscsi/initiatorname.iscsi
              sudo mpathconf --enable --with_multipathd y --find_multipaths
n
              sudo systemctl enable --now iscsid multipathd
              sudo systemctl enable --now iscsi
        securityContext:
          privileged: true
    hostPID: true
    containers:
      - name: wait
        image: k8s.gcr.io/pause:3.1
    hostPID: true
    hostNetwork: true
    tolerations:
      - effect: NoSchedule
        key: node-role.kubernetes.io/master
    updateStrategy:
      type: RollingUpdate

```

Use the following yaml file to create trident backend configuration for using ontap san storage

Trident backend for iSCSI

```
apiVersion: v1
kind: Secret
metadata:
  name: backend-tbc-ontap-san-secret
type: Opaque
stringData:
  username: <username>
  password: <password>
---
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: ontap-san
spec:
  version: 1
  storageDriverName: ontap-san
  managementLIF: <management LIF>
  backendName: ontap-san
  svm: <SVM name>
  credentials:
    name: backend-tbc-ontap-san-secret
```

Use the following yaml file to create trident storage class configuration for using ontap san storage

Trident storage class for iSCSI

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ontap-san
provisioner: csi.trident.netapp.io
parameters:
  backendType: "ontap-san"
  media: "ssd"
  provisioningType: "thin"
  snapshots: "true"
allowVolumeExpansion: true
```

Install MTV

Now you can install the Migration Toolkit for virtualization (MTV). Refer to the instructions provided [here](#) for help with the installation.

The Migration Toolkit for Virtualization (MTV) user interface is integrated into the OpenShift web console. You can refer [here](#) to start using the user interface for various tasks.

Create Source Provider

In order to migrate the RHEL VM from VMware to OpenShift Virtualization, you need to first create the source provider for VMware. Refer to the instructions [here](#) to create the source provider.

You need the following to create your VMware source provider:

- VCenter url
- VCenter Credentials
- VCenter server thumbprint
- VDDK image in a repository

Sample source provider creation:

Select provider type *

vm vSphere

Provider resource name *

vmware-source

Unique Kubernetes resource name identifier

URL *

URL of the vCenter SDK endpoint. Ensure the URL includes the "/sdk" path. For example: https://vCenter-host-example.com/sdk

VDDK init image

docker.repo.eng.netapp.com/banum/vddk:801

VDDK container image of the provider, when left empty some functionality will not be available

Username *

administrator@vsphere.local

vSphere REST API user name.

Password *

vSphere REST API password credentials.

SSHA-1 fingerprint *

The provider currently requires the SHA-1 fingerprint of the vCenter Server's TLS certificate in all circumstances. vSphere calls this the server's thumbprint.

Skip certificate validation

☒



The Migration Toolkit for Virtualization (MTV) uses the VMware Virtual Disk Development Kit (VDDK) SDK to accelerate transferring virtual disks from VMware vSphere. Therefore, creating a VDDK image, although optional, is highly recommended. To make use of this feature, you download the VMware Virtual Disk Development Kit (VDDK), build a VDDK image, and push the VDDK image to your image registry.

Follow the instructions provided [here](#) to create and push the VDDK image to a registry accessible from the OpenShift Cluster.

Create Destination provider

The host cluster is automatically added as the OpenShift virtualization provider is the source provider.

Create Migration Plan

Follow the instructions provided [here](#) to create a migration plan.

While creating a plan, you need to create the following if not already created:

- A network mapping to map the source network to the target network.
- A storage mapping to map the source datastore to the target storage class. For this you can choose ontap-san storage class.

Once the migration plan is created, the status of the plan should show **Ready** and you should now be able to **Start** the plan.

The screenshot shows the Red Hat OpenShift MTV interface. The left sidebar contains the navigation menu with 'Migration' expanded. The main panel displays a table of migration plans. The first plan, 'mtv-migration-demo', is in 'Ready' status and has a 'Start' button. The other three plans are in 'Succeeded' status.

Name	Source	Target	VMs	Status	Description
mtv-migration-demo	vmware	host	1	Ready	Plan for migrating VM to OpenShift Virt...
vmware-ovs-migration	vmware2	host	1	Succeeded	Migrating RHEL 9 vm to OpenShift Virtu...
vmware-ovs-migration-plan1	vmware2	host	1	Succeeded	
vmware-ovs-migration-plan2	vmware2	host	1	Succeeded	migrating RHEL 9 vm using ONTAP NFS...

Clicking on **Start** will run through a sequence of steps to complete the migration of the VM.

Red Hat OpenShift

You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.

Migration plans > mig-migration-bene

Migration details by VM

Filter by name: [Search] [Cancel]

1-1 of 1

Name	Start time	End time	Data copied	Status
oig-source-rhe9...	06 Mar 2024, 09:42...	06 Mar 2024, 09:52...	20.00 / 20.00 GB	Complete

Get logs

Step	Elapsed time	State
Initialize migration	00:00:25	Completed
Allocate disks	00:00:00	Completed
Convert image to kubevirt	00:02:45	Completed
Copy disks	00:04:58	Completed
Create VM	00:00:00	Completed

1-1 of 1

Activate Windows
Go to Settings to activate Windows.

When all steps are completed, you can see the migrated VMs by clicking on the **virtual machines** under **Virtualization** in the left-side navigation menu.

Instructions to access the virtual machines are provided [here](#).

You can log into the virtual machine and verify the contents of the postgresql databases. The databases, tables and the entries in the table should be the same as what was created on the source VM.

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.