



# Deployment on ROSA with FSxN

## NetApp Solutions

NetApp  
December 19, 2024

This PDF was generated from [https://docs.netapp.com/us-en/netapp-solutions/containers/rh-os-n\\_use\\_case\\_openshift\\_virtualization\\_rosa\\_overview.html](https://docs.netapp.com/us-en/netapp-solutions/containers/rh-os-n_use_case_openshift_virtualization_rosa_overview.html) on December 19, 2024. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

- Deployment on ROSA with FSxN ..... 1
- Deploy Red Hat OpenShift Virtualization with FSxN on ROSA ..... 1
- Workflows ..... 13

# Deployment on ROSA with FSxN

## Deploy Red Hat OpenShift Virtualization with FSxN on ROSA

### Overview

This section provides details for setting up FSx for NetApp ONTAP as the default Storage Class for the ROSA cluster, and then create a Virtual Machine that will leverage FSx ONTAP storage for its volumes. We will also look into connecting to the Virtual Machine using the guest credentials, and restarting the VM. And finally, we will perform a live migration of the Virtual Machine from the current node to a new node. We will examine the contents of the disk storage after a VM restart and the live migration .


### Prerequisites

- [AWS account](#)
- [A Red Hat account](#)
- IAM user [with appropriate permissions](#) to create and access ROSA cluster
- [AWS CLI](#)
- [ROSA CLI](#)
- [OpenShift command-line interface \(oc\)](#)
- [Helm 3 documentation](#)
- [A HCP ROSA cluster](#) (with at least 3 bare-metal worker nodes)
- [OpenShift Virtualization installed on ROSA Cluster](#)
- [Access to Red Hat OpenShift web console](#)

### Initial Setup

This section shows how to set up the default storage class to be trident-csi and the default VolumeSnapshotClass to be the FSx Volume Snapshot class. Then it shows how to create a VM from a template and then connect and login to it using the guest credentials.

Ensure default Storage Class is set to trident-csi



Name	Provisioner	Reclaim policy
SC gp2-csi	ebs.csi.aws.com	Delete
SC gp3-csi	ebs.csi.aws.com	Delete
SC trident-csi - Default	csi.trident.netapp.io	Retain

Ensure default VolumeSnapShotClasses is set as shown

Name	Driver	Deletion policy
VSC csi-aws-vsc	ebs.csi.aws.com	Delete
VSC fsx-snapclass - Default	csitrident.netapp.io	Delete

If the defaults are not set, you can set it up either from the console or from command line

```
$ oc patch storageclass trident-csi -p '{"metadata": {"annotations": {"storageclass.kubernetes.io/is-default-class": "true"}}}'
```

```
$ oc patch VolumeSnapshotClasses fsx-snapclass -p '{"metadata": {"annotations": {"snapshot.storage.kubernetes.io/is-default-class": "true"}}}'
```

## Create a VM from the template

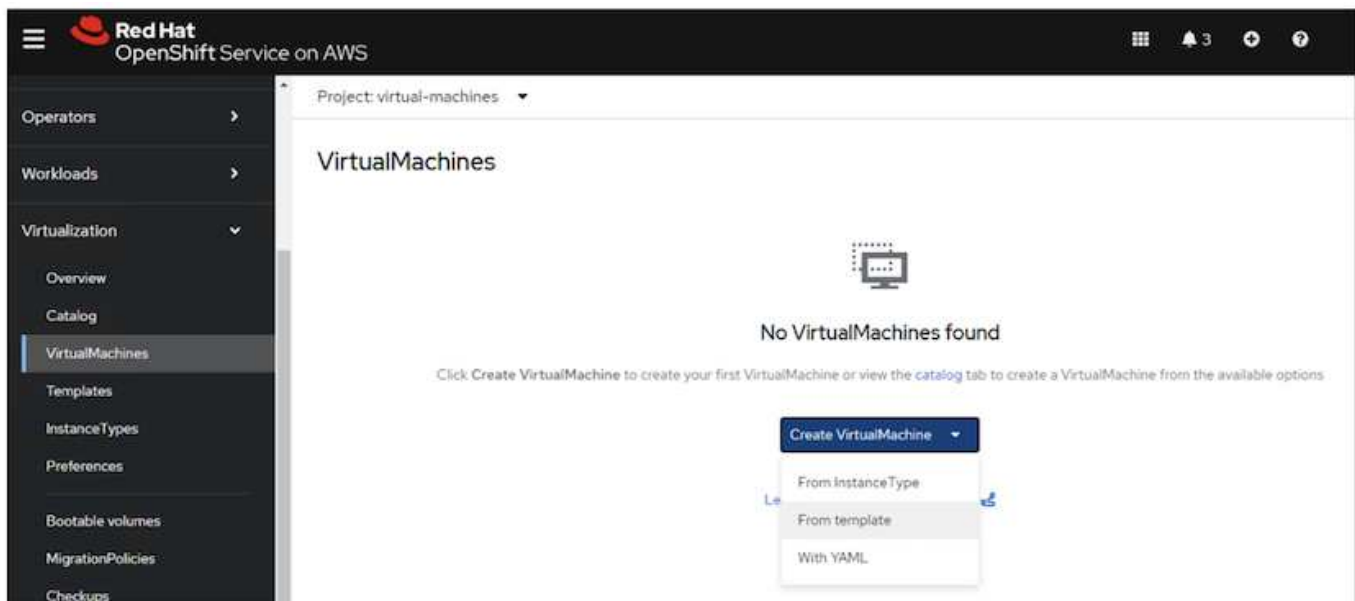
Use the web console to create a VM from a template.

From the RedHat OpenShiftService on AWS console, create a virtual machine. There are templates available on the cluster that can be used to create the VM.

In the screenshot below, we choose fedora VM from this list. Give the VM a name, and then click on **Customize Virtual Machine**. Select the **Disks** tab and click on **Add disks**.

Change the name of the disk preferably to something meaningful, ensure that **trident-csi** is selected for storage class. Click on **Save**. Click on **Create VirtualMachine**

After a few minutes, the VM is in the running state



Red Hat OpenShift Service on AWS

Exchange Password Required  
Enter your password for "samsundhar" in Internet Accounts.

Project: virtual-machines

### Create new VirtualMachine

Select an option to create a VirtualMachine from.

InstanceTypes | **Template catalog**

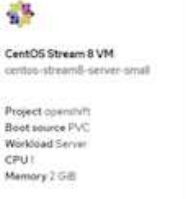
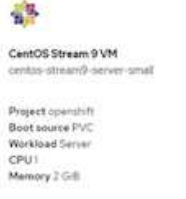






Template project: All projects

Default templates

Filter by keyword

12 items

- Boot source available
- Operating system
  - CentOS
  - Fedora
  - Other
  - RHEL
  - Windows
- Workload
  - Desktop
  - High performance
  - Server

 <p><b>CentOS Stream 8 VM</b> centos-stream8-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>	 <p><b>CentOS Stream 9 VM</b> centos-stream9-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>	 <p><b>CentOS 7 VM</b> centos7-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>	 <p><b>Fedora VM</b> fedora-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>
 <p><b>Red Hat Enterprise Linux 7 VM</b> rhe7-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1</p>	 <p><b>Red Hat Enterprise Linux 8 VM</b> rhe8-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1</p>	 <p><b>Red Hat Enterprise Linux 9 VM</b> rhe9-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1</p>	 <p><b>Microsoft Windows 10 VM</b> windows10-desktop-medium</p> <p>Activate Windows Go to Settings to activate Windows.</p> <p>Project openshift/ Boot source PVC Workload Desktop CPU 1</p>



Template info

Operating system

Fedora VM

Workload type

Server (default)

Description

Template for Fedora Linux 39 VM or newer. A PVC with the Fedora 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	-

Storage

Boot from CD

Disk source \*

Template default

Disk size \*

- 30 + GiB

Drivers

Mount Windows drivers disk

[Optional parameters](#)

Quick create VirtualMachine

VirtualMachine name \*

fedora-vm1

Project Public SSH key

default Not configured

Start this VirtualMachine after creation

Quick create VirtualMachine

Customize VirtualMachine

Activate Windows

Go to Settings to activate Windows.

Cancel

# Customize and create VirtualMachine YAML

Template: Fedora VM

- Overview
- YAML
- Scheduling
- Environment
- Network interfaces
- Disks**
- Scripts
- Metadata


**Add disk**

Filter   Mount Windows drivers disk

Name ↑	Source ↓	Size ↓	Drive ↓	Interface ↓	Storage class ↓	
cloudinitdisk	Other	-	Disk	virtio	-	⋮
rootdisk <span>bootable</span>	Other	30 GiB	Disk	virtio	-	⋮

## Add disk



Use this disk as a boot source 

Name \*

fedora-vm1-disk1

Source \*

Empty disk (blank)

PersistentVolumeClaim size \*

-

30

+

GiB

▼

Type

Disk

Hot plug is enabled only for "Disk" type

Interface \*

VirtIO

Hot plug is enabled only for "SCSI" interface

StorageClass

 trident-csi

Save

Cancel



Project: virtual-machines

VirtualMachines > VirtualMachine details

VM fedora-vm1 Running

Overview Metrics YAML Configuration Events Console Snapshots Diagnostics

**Details**

Name: fedora-vm1

Status: Running

Created: Oct 11, 2024, 1:46 PM (4 minutes ago)

Operating system: Fedora Linux 40 (Cloud Edition)

CPU | Memory: 1 CPU | 2 GiB Memory

Time zone: UTC

Template: fedora-server-small

Hostname: fedora-vm1

Machine type: pc-q35-rhel9.4.0

VNC console

Alerts (0)

General

Namespace: virtual-machi...

Node: ip-10-10-3-191...

VirtualMachineInstance: fedora-vm1

Pod: virt-launcher-f...

Owner: No owner

Snapshots (0) [Take snapshot](#)

Activate Windows  
No snapshots found  
Go to Settings to activate Windows.

## Review all the objects created for the VM

The storage disks.

**Storage (3)**

Name	Drive	Size	Interface
rootdisk	Disk	31.75 GiB	virtio
cloudinitdisk	Disk	-	virtio
fedora-vm1-disk1	Disk	31.75 GiB	virtio

The file systems of the VM will show the Partitions, File system type and the Mount points.

File systems ⓘ

Name ↑	File system type ⓘ	Mount point ⓘ	Total bytes ⓘ	Used bytes ⓘ
vda2	vfat	/boot/efi	99.76 MiB	16.01 MiB
vda3	ext-4	/boot	899.85 MiB	73.12 MiB
vda4	btrfs	/var	28.47 GiB	406.83 MiB
vda4	btrfs	/home	28.47 GiB	406.83 MiB
vda4	btrfs	/	28.47 GiB	406.83 MiB

2 PVCs are created for the VM, one from the boot disk and one for the hot plug disk.

Project: virtual-machines ▾

### PersistentVolumeClaims

[Create PersistentVolumeClaim ▾](#)

Filter ▾ Name ▾ Search by name... /

Name ⓘ	Status ⓘ	PersistentVolumes ⓘ	Capacity ⓘ
<a href="#">PVC</a> fedora-vm1	Bound	<a href="#">PV</a> pvc-7d60a3c1-d4cc-47d5-8053-efbb6ae1135f	31.75 GiB
<a href="#">PVC</a> fedora-vm1-fedora-vm1-disk1	Bound	<a href="#">PV</a> pvc-a769e022-2ae5-43fb-b8a1-a40f4447c6c2	31.75 GiB

The PVC for the boot disk shows that the Access mode is ReadWriteMany, and the Storage Class is trident-csi.


Project: virtual-machines

PersistentVolumeClaims > PersistentVolumeClaim details

**PVC fedora-vm1** Bound

Details | YAML | Events | VolumeSnapshots

### PersistentVolumeClaim details



**Name**  
fedora-vm1

**Namespace**  
virtual-machines

**Labels** Edit

- app=containerized-data-importer
- app.kubernetes.io/part-of=hyperconverged-cluster
- instancetype.kubevirt.io/default-preference=fedora
- app.kubernetes.io/version=4.15.3
- app.kubernetes.io/component=storage
- alerts&ls.io/KubePersistentVolumeFillingUp=disabled
- app.kubernetes.io/managed-by=ncd-controller
- instancetype.kubevirt.io/default-instancetype=ul.medium
- kubevirt.io/created-by=90537934-9ba5-47b5-8caa-63c0c96e5b7f

**Annotations**  
20 annotations

**Label selector**  
No selector

**Created at**  
Oct 11, 2024, 1:46 PM

**Status**  
Bound

**Requested capacity**  
31.75 GiB

**Capacity**  
31.75 GiB

**Used**  
25.09 GiB

**Access modes**  
ReadWriteMany

**Volume mode**  
Filesystem

**StorageClasses**  
trident-csi

**PersistentVolumes**  
pvc-7db0a3cf-d4cc-47d5-8053-efbb6ae1035f

Activate Windows  
Go to Settings to activate W

Similarly, the PVC for the hot-plug disk shows that the Access mode is ReadWriteMany, and the Storage Class is trident-csi.

Project: virtual-machines

PersistentVolumeClaims > PersistentVolumeClaim details

**PVC fedora-vm1-fedora-vm1-disk1** Bound

Details | YAML | Events | VolumeSnapshots

### PersistentVolumeClaim details

**31.8 GiB Available**

**Name**  
fedora-vm1-fedora-vm1-disk1

**Namespace**  
virtual-machines

**Labels** Edit

- alerts.k8s.io/KubePersistentVolumeFillingUp=disabled
- app=containerized-data-importer
- app.kubernetes.io/component=storage
- app.kubernetes.io/managed-by=cdi-controller
- app.kubernetes.io/part-of=hyperconverged-cluster
- app.kubernetes.io/version=4.10.3
- kubevirt.io/created-by=89537594-9ba5-47bb-8caa-03c0c96e5b7f

**Annotations**  
15 annotations

**Label selector**  
No selector

**Created at**  
Oct 11, 2024, 1:46 PM

**Status**  
Bound

**Requested capacity**  
31.75 GiB

**Capacity**  
31.75 GiB

**Used**  
320 KiB

**Access modes**  
ReadWriteMany

**Volume mode**  
Filesystem

**StorageClasses**  
trident-csi

**PersistentVolumes**  
pvc-a769e022-2ae5-43fb-b8a1-a40f4447c6c2

In the screenshot below we can see that the pod for the VM has a Status of Running.

**Pods** Create Pod

Filter Name Search by name

Name	Status	Ready	Restarts	Owner	Memory	CPU	Created
virt-launcher-fedora-vm1-8fp2k	Running	1/1	0	VM fedora-vm1	535.5 MiB	0.010 cores	Oct 11, 2024, 2:27 PM
virt-launcher-fedora-vm1-ko2k9	Completed	0/1	0	VM fedora-vm1	-	-	Oct 11, 2024, 2:21 PM

Here we can see the two Volumes associated with the VM pod and the 2 PVCs associated with them.

Name	Mount path	SubPath	Type	Permissions	Utilized by
private	/var/run/kubevirt-private	No subpath		Read/Write	compute
public	/var/run/kubevirt	No subpath		Read/Write	compute
ephemeral-disks	/var/run/kubevirt-ephemeral-disks	No subpath		Read/Write	compute
container-disks	/var/run/kubevirt/container-disks	No subpath		Read/Write	compute
libvirt-runtime	/var/run/libvirt	No subpath		Read/Write	compute
sockets	/var/run/kubevirt/sockets	No subpath		Read/Write	compute
rootdisk	/var/run/kubevirt-private/vmi-disks/rootdisk	No subpath	PVC fedora-vm1	Read/Write	compute
fedora-vm1-disk1	/var/run/kubevirt-private/vmi-disks/fedora-vm1-disk1	No subpath	PVC fedora-vm1-fedora-vm1-disk1	Read/Write	compute
hotplug-disks	/var/run/kubevirt/hotplug-disks	No subpath		Read/Write	compute

## Connect to the VM

Click on 'Open web console' button, and login using Guest Credentials

The screenshot shows the OpenShift VirtualMachines console interface. At the top, it displays 'Project: virtual-machines' and 'VirtualMachines > VirtualMachine details'. The VM 'fedora-vm1' is shown as 'Running'. Below this are navigation tabs: Overview, Metrics, YAML, Configuration, Events, Console, Snapshots, and Diagnostics. The 'Details' section on the left lists various attributes: Name (fedora-vm1), Status (Running), Created (Oct 11, 2024, 1:46 PM (12 minutes ago)), Operating system (Fedora Linux 40 (Cloud Edition)), CPU | Memory (1 CPU | 2 GiB Memory), Time zone (UTC), Template (fedora-server-small), Hostname (fedora-vm1), and Machine type (pc-q35-rhel9.4.0). On the right, there is a 'VNC console' area with a play button. At the bottom right, the 'Open web console' button is highlighted with a red box.



Issue the following commands

```
$ df (to display information about the disk space usage on a file system).
```

```
$ dd if=/dev/urandom of=random.dat bs=1M count=10240 (to create a file called random.dat in the home dir and fill it with random data).
```

The disk is filled with 11 GB of data.

```
[fedora@fedora-vm1 ~]$
[fedora@fedora-vm1 ~]$ df .
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/vda4        30327788 10939828  18943548  37% /home
[fedora@fedora-vm1 ~]$ dd if=/dev/urandom of=random.dat bs=1M count=10240
10240+0 records in
10240+0 records out
10737418240 bytes (11 GB, 10 GiB) copied, 35.8159 s, 300 MB/s
[fedora@fedora-vm1 ~]$ df
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/vda4        30327788  9699188  20190780  33% /home
[fedora@fedora-vm1 ~]$ ls
random.dat
```

Use vi to create a sample text file that we will use to test.

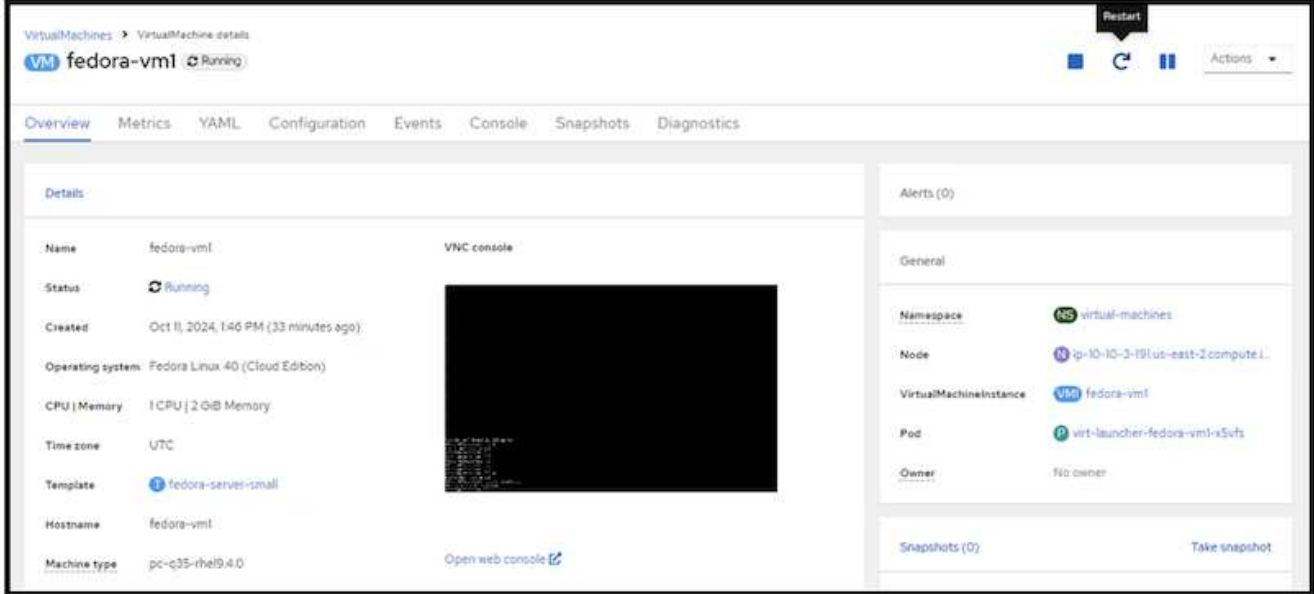
```
[fedora@fedora-vm1 ~]$ ls
random.dat sample.txt
[fedora@fedora-vm1 ~]$ cat sample.txt
This is a sample text file.
[fedora@fedora-vm1 ~]$
```

# Workflows

## VM Restart

In this sections we will perform a VM restart and then examine the contents of the disks.

Click on the restart button.



The VM comes back to the running state with the exact same filesystems, PVCs and files in the filesystems

File systems ⓘ

Name ↑	File system type ↓	Mount point ↓	Total bytes ↓	Used bytes ↓
vda2	vfat	/boot/efi	99.76 MiB	16.01 MiB
vda3	ext4	/boot	899.85 MiB	73.12 MiB
vda4	btrfs	/var	28.50 GiB	10.43 GiB
vda4	btrfs	/home	28.50 GiB	10.43 GiB
vda4	btrfs	/	28.50 GiB	10.43 GiB

```
[fedora@fedora-vm1 ~]$ ls
random.dat  sample.txt
[fedora@fedora-vm1 ~]$ df .
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/vda4      30327788 10948176  18935632  37% /home
[fedora@fedora-vm1 ~]$ _
```

```
[fedora@fedora-vm1 ~]$ ls
random.dat  sample.txt
[fedora@fedora-vm1 ~]$ cat sample.txt
This is a sample text file.
[fedora@fedora-vm1 ~]$
```

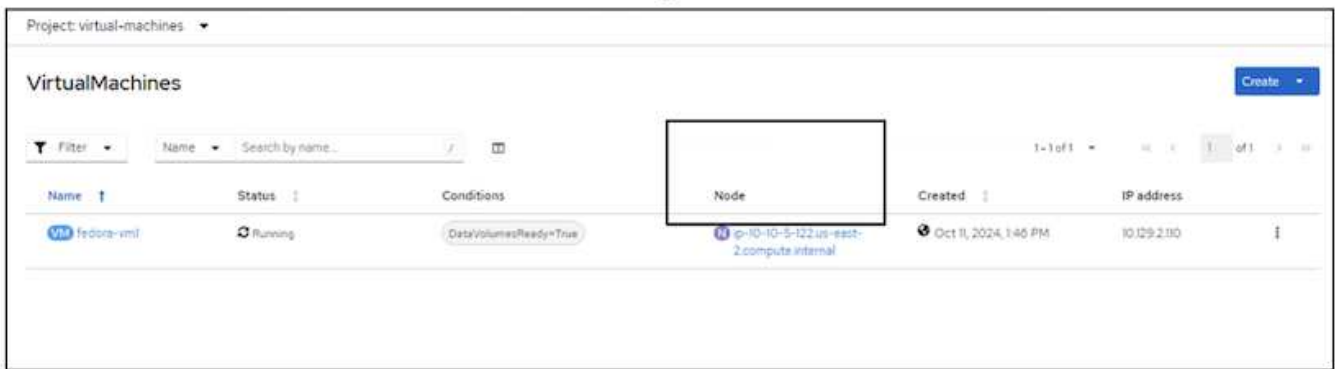
## VM live migration

In this sections we will perform a VM live migration and then examine the contents of the disks.

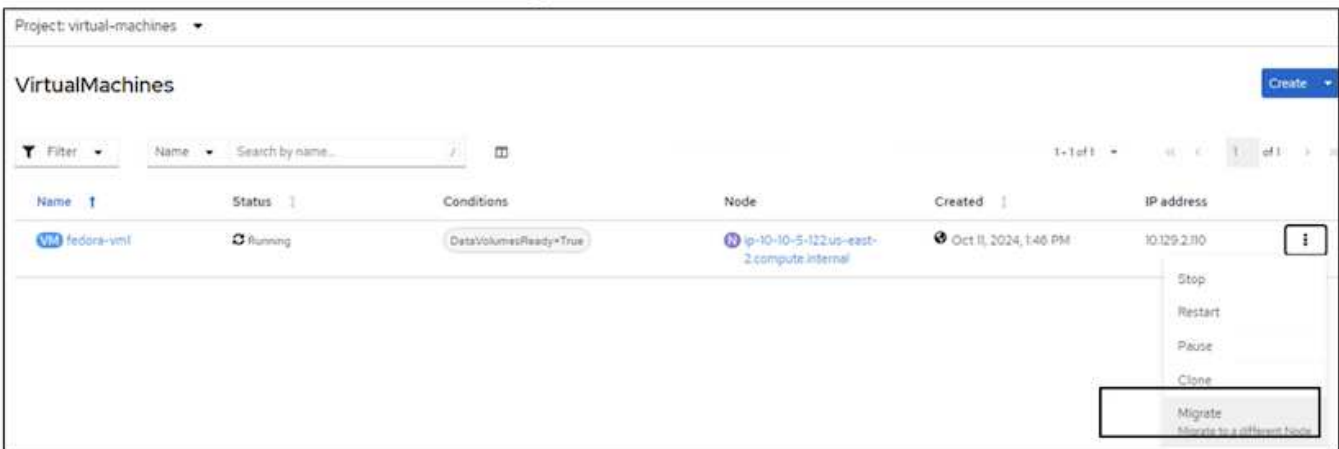
Live migration refers to the process of moving a running Virtual Machine (VM) from one physical host to another host without disrupting normal operations or causing any downtime, or other adverse effects for the end user.

Live migration is considered a major step in Virtualization. It allows an entire VM to be moved with a running operating system (OS), Memory, storage, and network connectivity from their current node to the destination. Below we will see how to perform a Live Migration of the VM from the current node to a new node.

Note the node on which the VM is running




Click on the 3 dots and select Migrate




On the Overview page, you can see that the migration has succeeded, and the Status has changed to Succeeded.





After the completion of the Live Migration, the VM is now on a different node.



Open the web console and view the contents of the disks. It still has the same 2 files that we previously created before the Live Migration.

```
[fedora@fedora-vm1 ~]# df .
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/vda1       30327788 10956768  18927040  37% /home
[fedora@fedora-vm1 ~]#
[fedora@fedora-vm1 ~]#
[fedora@fedora-vm1 ~]# ls
random.dat  sample.txt
[fedora@fedora-vm1 ~]#
```

```
[fedora@fedora-vm1 ~]$ ls
random.dat  sample.txt
[fedora@fedora-vm1 ~]$ cat sample.txt
This is a sample text file.
[fedora@fedora-vm1 ~]$
```

The storage for the VM on the new node still shows the same disks

Storage (3)

Name	Drive	Size	Interface
rootdisk	Disk	31.75 GiB	virtio
cloudinitdisk	Disk	-	virtio
fedora-vm1-disk1	Disk	31.75 GiB	virtio

Also, the PVCs are the same.

Project: virtual-machines

PersistentVolumeClaims Create PersistentVolumeClaim

Filter Name Search by name

Name	Status	PersistentVolumes	Capacity	Used	StorageClass
fedora-vm1	Bound	pvc-7d00a3cf-d4cc-47d5-8053-efbb0ae1135f	31.75 GiB	28.12 GiB	trident-csi
fedora-vm1-fedora-vm1-disk1	Bound	pvc-a709e022-2ae5-43fb-b8af-a40f4447c6c2	31.75 GiB	320 KiB	trident-csi

Volumes associated with the VM pod are also the same (2 PVCs) as before.

Name	Mount path	SubPath	Type	Permissions	Utilized by
private	/var/run/kubevirt-private	No subpath		Read/Write	compute
public	/var/run/kubevirt	No subpath		Read/Write	compute
ephemeral-disks	/var/run/kubevirt-ephemeral-disks	No subpath		Read/Write	compute
container-disks	/var/run/kubevirt/container-disks	No subpath		Read/Write	compute
libvirt-runtime	/var/run/libvirt	No subpath		Read/Write	compute
sockets	/var/run/kubevirt/sockets	No subpath		Read/Write	compute
rootdisk	/var/run/kubevirt-private/vmi-disks/rootdisk	No subpath	PV/D fedora-vmi	Read/Write	compute
fedora-vmi-disk1	/var/run/kubevirt-private/vmi-disks/fedora-vmi-disk1	No subpath	PV/D fedora-vmi-fedora-vmi-disk1	Read/Write	compute
hotplug-disks	/var/run/kubevirt/hotplug-disks	No subpath		Read/Write	compute

## Demo video

[Live migration of virtual machines in OpenShift Virtualization on ROSA with Amazon FSx for NetApp ONTAP](#)

More videos on Red Hat OpenShift and OpenShift Virtualization solutions can be found [here](#).

## 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.