



基于ROSA的OpenShift虚拟化

NetApp Solutions

NetApp
December 19, 2024

目录

基于ROSA的OpenShift虚拟化	1
在ROSA上使用FSxN部署Red Hat OpenShift虚拟化	1
工作流	13

基于ROSA的OpenShift虚拟化

在ROSA上使用FSxN部署Red Hat OpenShift虚拟化

概述

本节详细介绍了如何将FSx for NetApp ONTAP设置为ROSA集群的默认存储类、然后创建一个虚拟机来将FSx ONTAP存储用于其卷。我们还将了解如何使用子系统凭据连接到虚拟机并重新启动虚拟机。最后、我们将执行虚拟机从当前节点到新节点的实时迁移。我们将在VM重新启动和实时迁移之后检查磁盘存储的内容。

前提条件

- "AWS 帐户"
- "Red Hat帐户"
- IAM用户"具有适当的权限"、用于创建和访问ROSA集群
- "AWS命令行界面"
- "罗莎命令行界面"
- "OpenShift命令行界面"(OC)
- "Helm 3文档"
- "HCP ROSA集群"(至少具有3个裸机工作节点)
- "ROSA集群上安装了OpenShift虚拟化"
- "访问Red Hat OpenShift Web控制台"

初始设置

本节介绍如何将默认存储类设置为FSX-CSI、并将默认卷快照类设置为Trident类。然后、它会显示如何使用模板创建VM、然后使用子系统凭据连接并登录到该VM。

确保默认存储类设置为
Trident



The screenshot shows the 'StorageClasses' page in the OpenShift console. A table lists three storage classes: 'gp2-csi', 'gp3-csi', and 'trident-csi - Default'. The 'trident-csi - Default' row is highlighted with a blue border, indicating it is the selected default storage class. The table columns are Name, Provisioner, and Reclaim policy.

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

确保已按所示设置默认的卷Snap肖 特克莱斯



如果未设置默认值、您可以从控制台或命令行进行设置

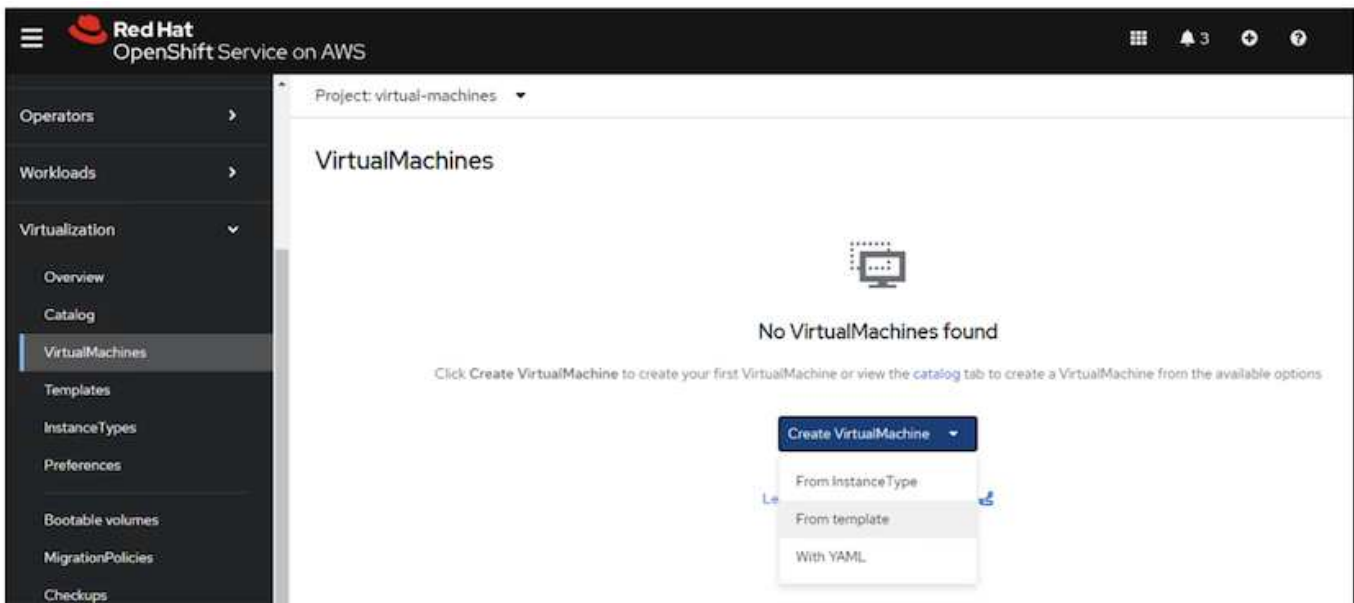
```
$ oc patch storageclass trident-csi -p '{"metadata": {"annotations": {"storageclass.kubernetes.io/is-default-class": "true"}}}'
```

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

使用模板创建虚拟机

使用Web控制台通过模板创建VM。从AWS控制台上的Red Hat OpenShift Service创建虚拟机。集群上提供了一些模板、可用于创建虚拟机。在下面的屏幕截图中、我们从该列表选择了Fedora VM。为虚拟机命名，然后单击自定义虚拟机。选择磁盘选项卡，然后单击添加磁盘。最好将磁盘名称更改为有意义的名称，确保为存储类选择 **Trident—CSI**。单击保存。单击创建虚拟机

几分钟后、虚拟机将处于running状态



Red Hat OpenShift Service on AWS

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

Administrator

Home

Operators

Workloads

Virtualization

Networking

Storage

Builds

Observe

Compute

User Management

Administration

Project: virtual-machines

Create new VirtualMachine

Select an option to create a VirtualMachine from.

InstanceTypes | **Template catalog**

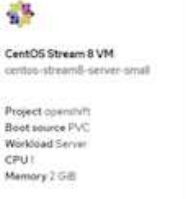
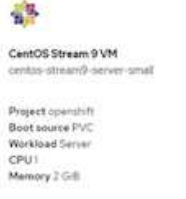






Template project: All projects

Default templates

Filter by keyword

12 items

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

 <p>CentOS Stream 8 VM centos-stream8-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>	 <p>CentOS Stream 9 VM centos-stream9-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>	 <p>CentOS 7 VM centos7-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>	 <p>Fedora VM fedora-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1 Memory 2 GiB</p>
 <p>Red Hat Enterprise Linux 7 VM rhe7-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1</p>	 <p>Red Hat Enterprise Linux 8 VM rhe8-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1</p>	 <p>Red Hat Enterprise Linux 9 VM rhe9-server-small</p> <p>Project openshift/ Boot source PVC Workload Server CPU 1</p>	 <p>Microsoft Windows 10 VM 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>



Fedora VM

fedora-server-small



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 Search by name... Mount Windows drivers disk

Name ↑	Source ↓	Size ↓	Drive ↓	Interface ↓	Storage class ↓	
cloudinitdisk	Other	-	Disk	virtio	-	⋮
rootdisk bootable	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.

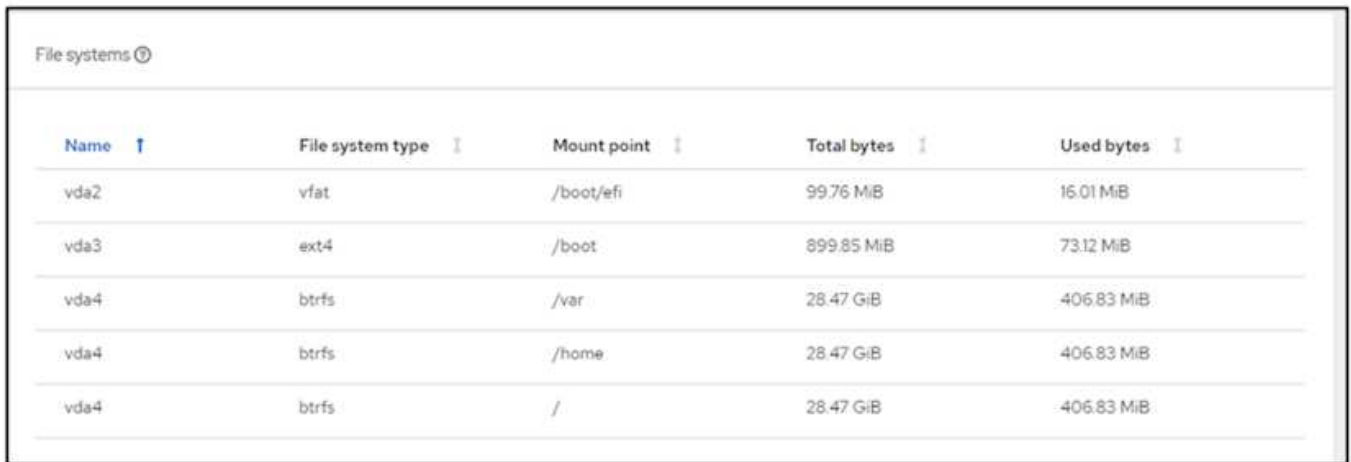
查看为虚拟机创建的所有对象

存储磁
盘。

Storage (3)

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

VM的文件系统将显示分区、文件系统类型和挂载点。



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	28.47 GiB	406.83 MiB
vda4	btrfs	/home	28.47 GiB	406.83 MiB
vda4	btrfs	/	28.47 GiB	406.83 MiB

系统会为虚拟机创建2个PVC、一个来自启动磁盘、另一个用于热插拔磁盘。



Name ↓	Status ↓	PersistentVolumes ↓	Capacity ↓
PVC fedora-vm1	Bound	PV pvc-7d60a3cf-d4cc-47d5-8053-efbb6ae1135f	31.75 GiB
PVC fedora-vm1-fedora-vm1-disk1	Bound	PV pvc-a769e022-2ae5-43fb-b8a1-a40f4447c6c2	31.75 GiB

启动磁盘的PVC显示访问模式为ReadReadReadReadRead任意、存储类为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‘KubePersistentVolumeFillingUp=disabled
- app.kubernetes.io/managed-by=ncd-controller
- instancetype.kubevirt.io/default-instancetype=ul.medium
- kubevirt.io/created-by=90537934-9ba5-47b5-8caa-63c0c9e5b77

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-70b0a3cf-d4cc-47d5-8053-efbb6ae1035f

Activate Windows
Go to Settings to activate W

同样、热插拔磁盘的PVC显示、访问模式为ReadReadReadReadWeb、存储类为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=89537594-9ba5-47b8-0caa-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

在下面的屏幕截图中、我们可以看到虚拟机的POD的状态为正在运行。

Pods

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	515.5 MB	0.010 cores	Oct 11, 2024, 2:27 PM
virt-launcher-fedora-vm1-k0k99	Completed	0/1	0	VM fedora-vm1	-	-	Oct 11, 2024, 2:21 PM

在这里、我们可以看到与VM Pod关联的两个卷以及与其关联的2个PVC。

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

连接到虚拟机

单击'Open web console'(打开Web控制台)按钮，然后使用来宾凭据登录

Project: virtual-machines

VirtualMachines > VirtualMachine details

VM fedora-vm1 Running

Overview Metrics YAML Configuration Events Console Snapshots Diagnostics

Details

Name	fedora-vm1	VNC console
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	
Machine type	pc-q35-rhel9.4.0	

[Open web console](#)



发出以下命令

```
$ 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).
```

磁盘中已填充11 GB的数据。

```
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  
fedora@fedora-vm1 ~]$
```

使用vi创建示例文本文件、我们将使用该文件进行测试。

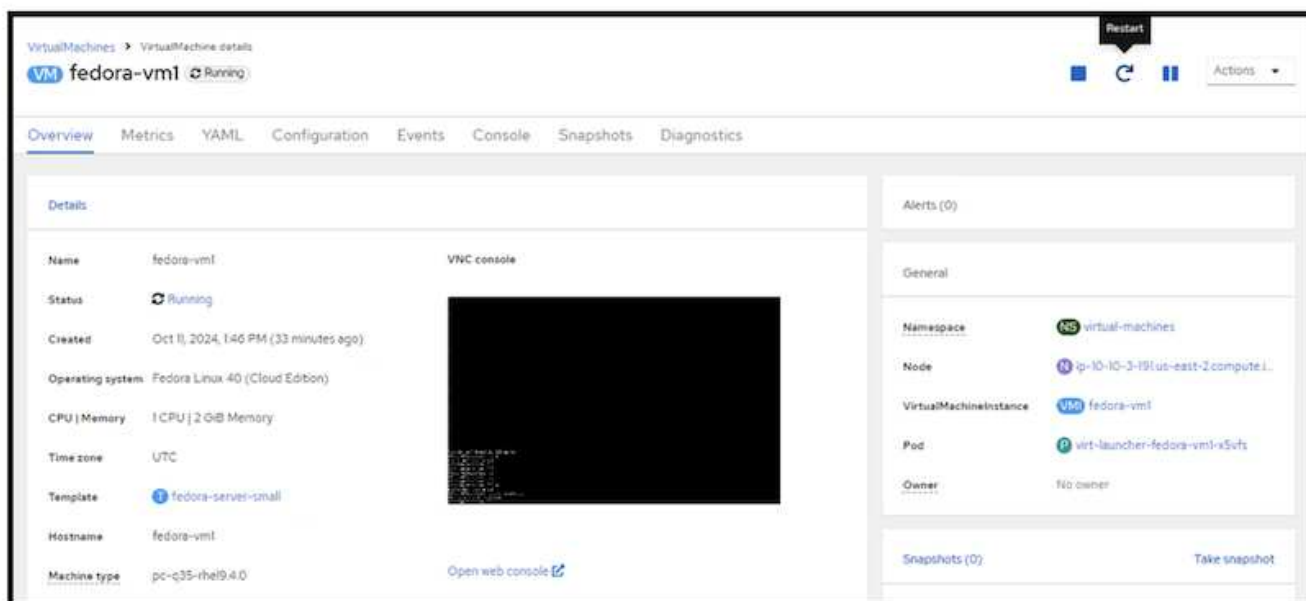
```
[fedora@fedora-vm1 ~]$ ls
random.dat  sample.txt
[fedora@fedora-vm1 ~]$ cat sample.txt
This is a sample text file.
[fedora@fedora-vm1 ~]$
```

workflows

VM重新启动

在本节中、我们将执行VM重新启动、然后检查磁盘的内容。

单击重新启动按钮。



VM恢复到running状态、并在这些文件系统中使用完全相同的文件系统、PVC和文件

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 ~]$
```


虚拟机实时迁移

在本节中、我们将执行VM实时迁移、然后检查磁盘的内容。实时迁移是指将正在运行的虚拟机(Virtual Machine、VM)从一台物理主机移动到另一台主机、而不会中断正常操作、也不会导致停机或对最终用户造成其他不利影响的过程。实时迁移被视为虚拟化的一个重要步骤。它允许使用正在运行的操作系统(OS)、内存、存储和网络连接将整个VM从其当前节点移动到目标。下面我们将介绍如何将虚拟机从当前节点实时迁移到新节点。

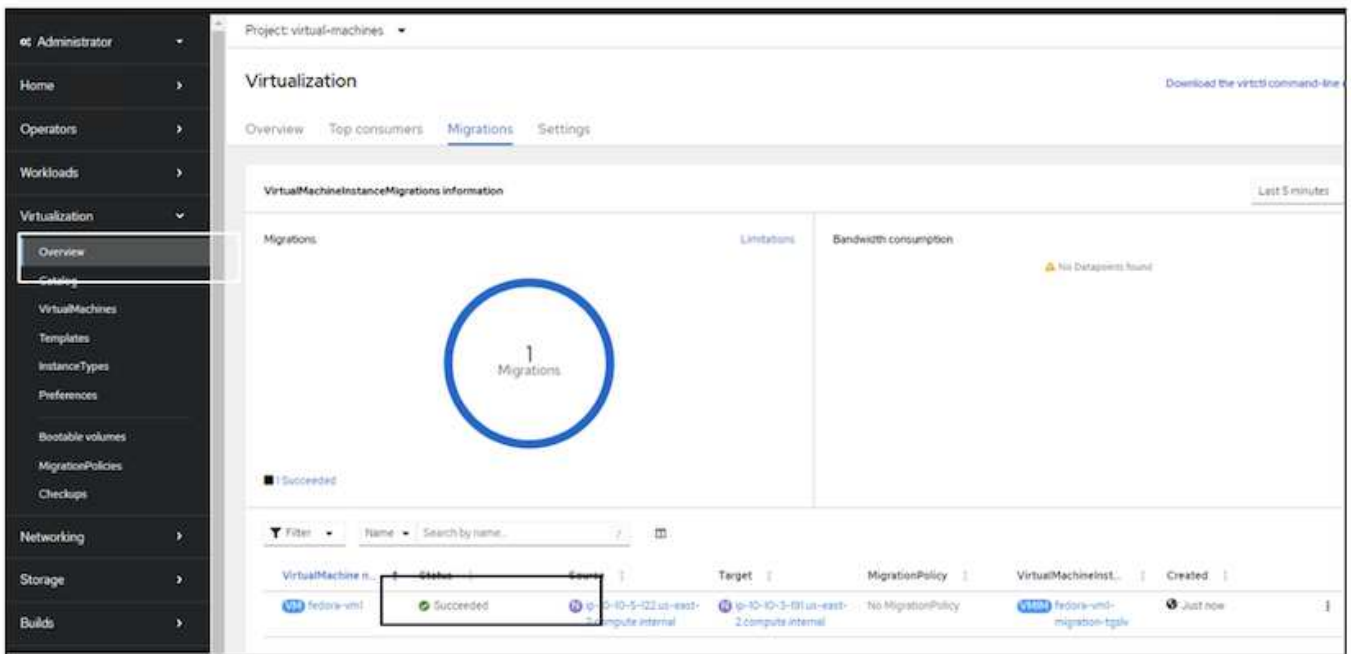
记下运行VM的节点




单击3个点、然后选择迁移



在"概述"页面上、您可以看到迁移已成功、并且"状态"已更改为"成功"



完成实时迁移后、虚拟机将位于其他节点上。



打开Web控制台并查看磁盘的内容。它仍具有与实时迁移之前创建的2个文件相同的文件。

```
[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 ~]$
```

新节点上VM的存储仍会显示相同的磁盘

Storage (3)			
Name	Drive	Size	Interface
rootdisk	Disk	31.75 GiB	virtio
cloudinitdisk	Disk	-	virtio
fedora-vm1-disk1	Disk	31.75 GiB	virtio

此外、这些PVC也是相同的。

Project: virtual-machines

PersistentVolumeClaims

Filter Name Search by name...

Name	Status	PersistentVolumes	Capacity	Used	StorageClass
fedora-vm1	Bound	pvc-7d00a3cf-d4cc-47d5-8053-ef6b6ae0335f	31.75 GiB	28.12 GiB	trident-csi
fedora-vm1-fedora-vm1-disk1	Bound	pvc-a700e032-2ae5-43fb-b8a1-a40f44470bc2	31.75 GiB	320 KiB	trident-csi

与VM Pod关联的卷也与以前相同(2个PVC)

Volumes

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

演示视频

[使用Amazon FSx for NetApp ONTAP在ROSA上实时迁移OpenShift虚拟化中的虚拟机](#)

有关Red Hat OpenShift和OpenShift虚拟化解决方案的更多视频，请参见["此处"](#)。

版权信息

版权所有 © 2024 NetApp, Inc.。保留所有权利。中国印刷。未经版权所有者事先书面许可，本档中受版权保护的任何部分不得以任何形式或通过任何手段（图片、电子或机械方式，包括影印、录音、录像或存储在电子检索系统中）进行复制。

从受版权保护的 NetApp 资料派生的软件受以下许可和免责声明的约束：

本软件由 NetApp 按“原样”提供，不含任何明示或暗示担保，包括但不限于适销性以及针对特定用途的适用性的隐含担保，特此声明不承担任何责任。在任何情况下，对于因使用本软件而以任何方式造成的任何直接性、间接性、偶然性、特殊性、惩罚性或后果性损失（包括但不限于购买替代商品或服务；使用、数据或利润方面的损失；或者业务中断），无论原因如何以及基于何种责任理论，无论出于合同、严格责任或侵权行为（包括疏忽或其他行为），NetApp 均不承担责任，即使已被告知存在上述损失的可能性。

NetApp 保留在不另行通知的情况下随时对本文档所述的任何产品进行更改的权利。除非 NetApp 以书面形式明确同意，否则 NetApp 不承担因使用本文档所述产品而产生的任何责任或义务。使用或购买本产品不表示获得 NetApp 的任何专利权、商标权或任何其他知识产权许可。

本手册中描述的产品可能受一项或多项美国专利、外国专利或正在申请的专利的保护。

有限权利说明：政府使用、复制或公开本文档受 DFARS 252.227-7013（2014 年 2 月）和 FAR 52.227-19（2007 年 12 月）中“技术数据权利 — 非商用”条款第 (b)(3) 条规定的限制条件的约束。

本文档中所含数据与商业产品和/或商业服务（定义见 FAR 2.101）相关，属于 NetApp, Inc. 的专有信息。根据本协议提供的所有 NetApp 技术数据和计算机软件具有商业性质，并完全由私人出资开发。美国政府对这些数据的使用权具有非排他性、全球性、受限且不可撤销的许可，该许可既不可转让，也不可再许可，但仅限在与交付数据所依据的美国政府合同有关且受合同支持的情况下使用。除本文档规定的情形外，未经 NetApp, Inc. 事先书面批准，不得使用、披露、复制、修改、操作或显示这些数据。美国政府对国防部的授权仅限于 DFARS 的第 252.227-7015(b)（2014 年 2 月）条款中明确的权利。

商标信息

NetApp、NetApp 标识和 <http://www.netapp.com/TM> 上所列的商标是 NetApp, Inc. 的商标。其他公司和产品名称可能是其各自所有者的商标。