



在具有 **FSxN** 的 **ROSA** 上部署 NetApp Solutions

NetApp
December 19, 2024

目錄

在具有 FSxN 的 ROSA 上部署	1
在 ROSA 上部署 Red Hat OpenShift 虛擬化與 FSxN	1
工作流程	13

在具有 FSxN 的 ROSA 上部署

在 ROSA 上部署 Red Hat OpenShift 虛擬化與 FSxN

總覽

本節詳細說明如何將 NetApp ONTAP 的 FSX 設定為 ROSA 叢集的預設儲存類別、然後建立將 FSX ONTAP 儲存設備用於其磁碟區的虛擬機器。我們也會考慮使用來賓認證連線至虛擬機器、然後重新啟動 VM。最後、我們將執行虛擬機器從目前節點即時移轉至新節點的作業。我們將在 VM 重新啟動和實時遷移之後檢查磁盤存儲的內容。

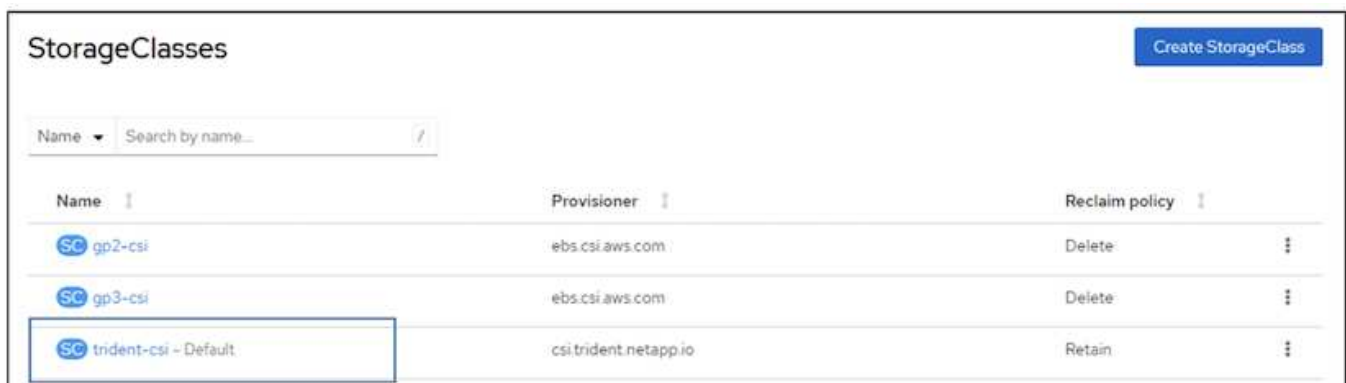
先決條件

- "AWS 帳戶"
- "Red Hat 帳戶"
- IAM 使用者"具有適當權限"可建立及存取 ROSA 叢集
- "AWS CLI"
- "ROSA CLI"
- "OpenShift 命令列介面" (OC)
- "船舵 3 文件"
- "HCP ROSA 叢集" (至少有 3 個裸機工作節點)
- "ROSA 叢集上安裝的 OpenShift 虛擬化"
- "存取 Red Hat OpenShift Web 主控台"

初始設定

本節說明如何將預設儲存類別設定為 Trident - CSI、以及將預設的 Volume SnapshotClass 設定為 FSX Volume Snapshot 類別。然後、它會示範如何從範本建立 VM、然後使用來賓認證連線並登入。

確保預設儲存類別設定為 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

確保預設的 Volume SnapShotClasses 已如圖所示設定



如果未設定預設值、您可以從主控台或命令列進行設定


```
$ 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"}}}'
```

從範本建立 VM

使用 Web 主控台從範本建立 VM。從 AWS 主控台的 RedHat OpenShiftService 建立虛擬機器。叢集上有可用於建立 VM 的範本。在下方的螢幕擷取畫面中、我們從清單中選擇 Fedora VM。命名 VM、然後按一下「自訂虛擬機器」。選擇「磁碟」標籤、然後按一下「新增磁碟」。最好將磁碟名稱變更為有意義的名稱、確定已為儲存類別選取 **Trident - CSI**。按一下「儲存」。按一下「建立虛擬機器」

幾分鐘後、虛擬機器處於執行中狀態



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**

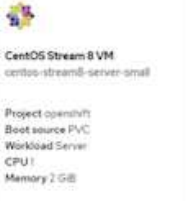
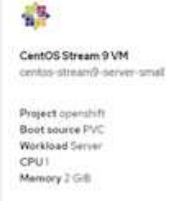
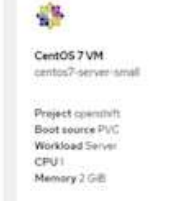




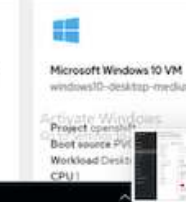
Template project: All projects

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>



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.

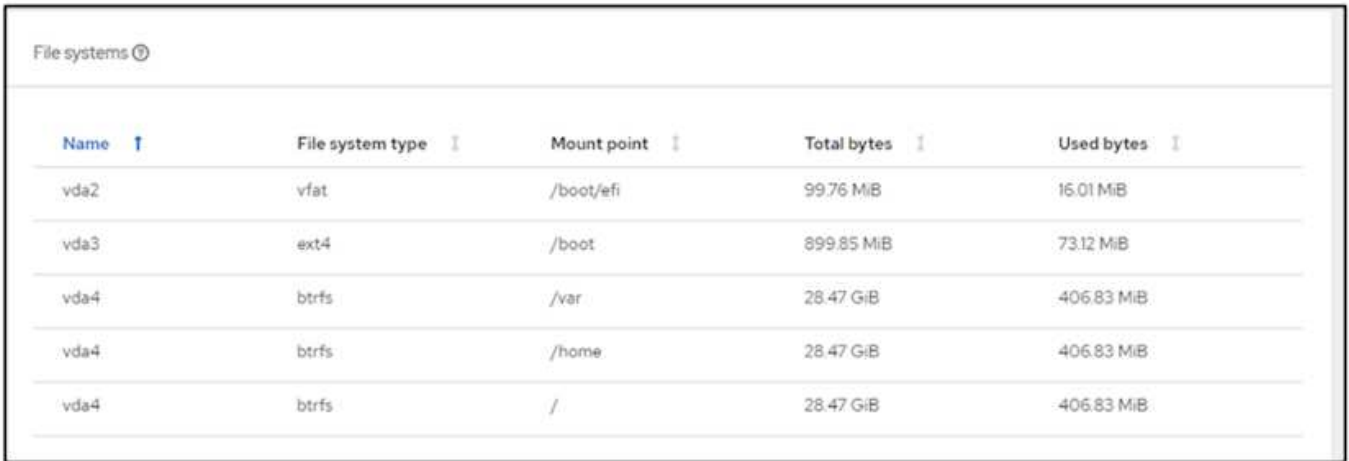
檢閱為 VM 建立的所有物件

儲存磁
碟。

Storage (3)

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

VM 的檔案系統會顯示分割區、檔案系統類型和掛載點。



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

為 VM 建立 2 個 PVC、一個從開機磁碟建立、另一個用於熱插拔磁碟。



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

開機磁碟的 PVC 顯示存取模式為 ReadWriteMany、儲存類別為 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-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-70b0a3cf-d4cc-47d5-8053-efbb6ae1035f

Activate Windows
Go to Settings to activate Windows

同樣地、熱插拔磁碟的 PVC 會顯示存取模式為 ReadWriteMany、而儲存類別則為 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

在下面的螢幕擷取畫面中、我們可以看到 VM 的 Pod 狀態為「執行中」。

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	515.5 MB	0.010 cores	Oct 11, 2024, 2:27 PM
virt-launcher-fedora-vm1-k2k99	Completed	0/1	0	VM fedora-vm1	-	-	Oct 11, 2024, 2:21 PM

此處我們可以看到與 VM Pod 相關聯的兩個 Volume、以及與 VM Pod 相關聯的 2 個 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

連接至 VM

按一下「開啟網路主控台」按鈕、然後使用訪客認證登入

The screenshot shows the OpenShift Virtualization 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), Operating system (Fedora Linux 40), 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, a button labeled 'Open web console' is highlighted with a blue box.



發出下列命令

```
$ 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).
```

磁碟中填滿 11 GB 的資料。

```
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 建立範例文字檔、供我們測試使用。

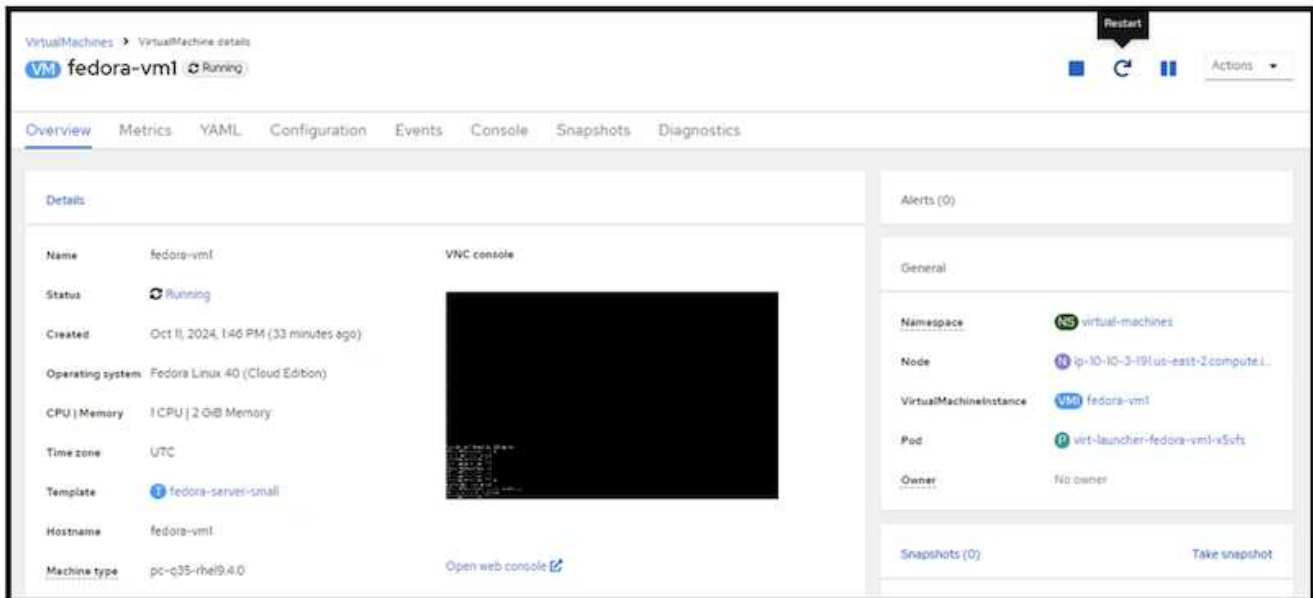
```
[fedora@fedora-vm1 ~]$ ls
random.dat  sample.txt
[fedora@fedora-vm1 ~]$ cat sample.txt
This is a sample text file.
[fedora@fedora-vm1 ~]$
```

工作流程

VM 重新啟動

在本節中、我們將執行 VM 重新啟動、然後檢查磁碟內容。

按一下重新啟動按鈕。



VM 會回到執行狀態、檔案系統中的檔案系統、PVC 和檔案完全相同

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 即時移轉

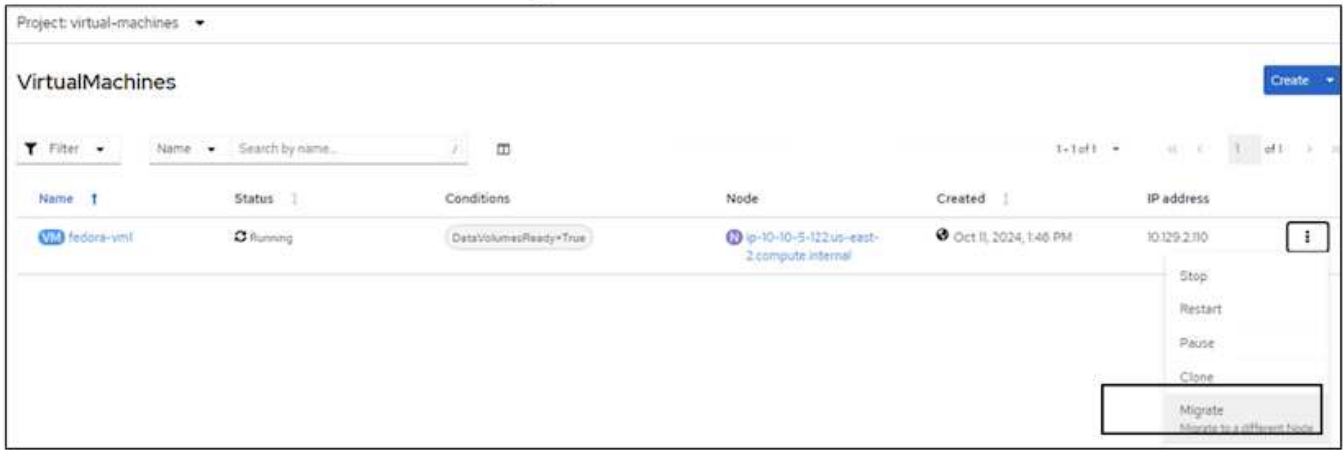
在本節中、我們將執行 VM 即時移轉、然後檢查磁碟內容。即時移轉是指將執行中的虛擬機器（VM）從一部實體主機移至另一部主機、而不會中斷正常作業或造成任何停機、或對終端使用者造成其他不良影響的程序。即時移轉被視為虛擬化的重要步驟。它可讓整個 VM 透過執行中的作業系統（OS）、記憶體、儲存設備及網路連線功能、從目前節點移至目的地。以下將說明如何從目前節點即時移轉至新節點。

記下正在執行 VM 的節點

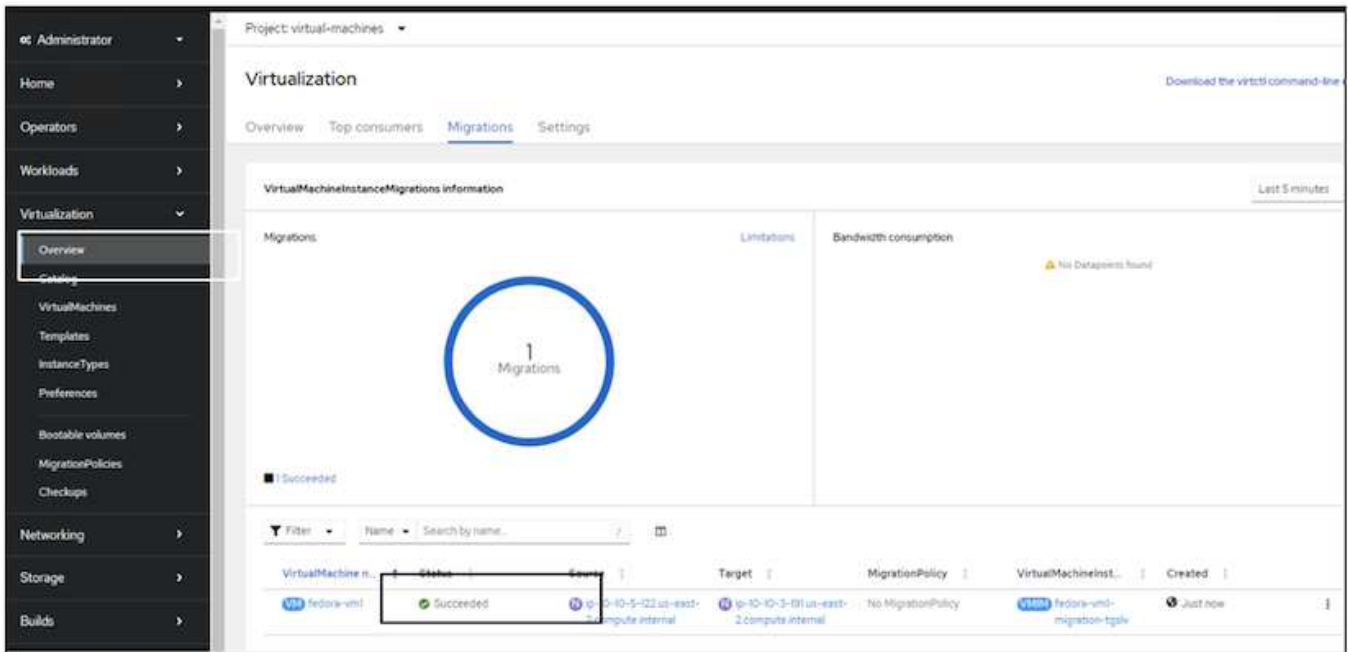
The screenshot shows the Google Cloud Platform console for the project 'virtual-machines'. It displays a table of Virtual Machines. The table has columns for Name, Status, Conditions, Node, Created, and IP address. The 'fedora-vm1' VM is listed with a status of 'Running' and a condition of 'DataVolumesReady=True'. The 'Node' column for this VM is highlighted with a red box, showing the node ID 'g1-10-10-5-122-us-east-2-compute-internal'.

Name	Status	Conditions	Node	Created	IP address
fedora-vm1	Running	DataVolumesReady=True	g1-10-10-5-122-us-east-2-compute-internal	Oct 11, 2024, 1:46 PM	10.129.2.180

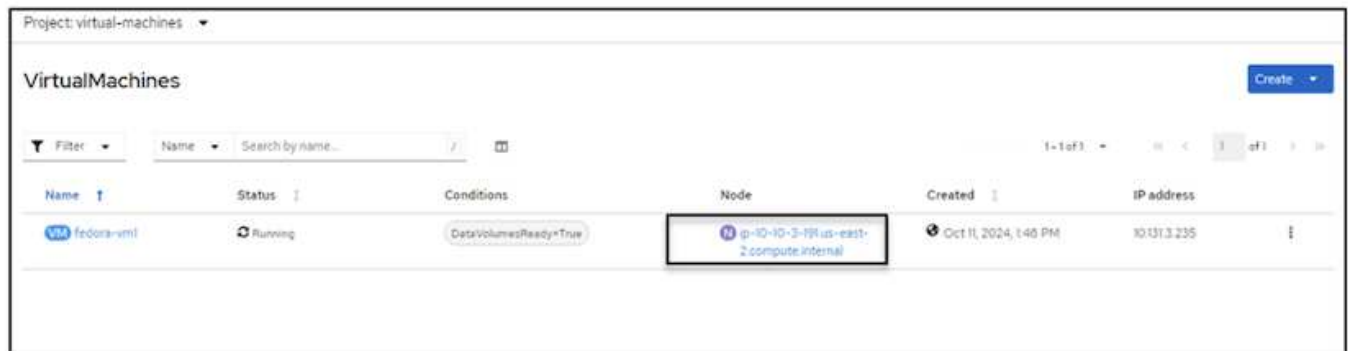
按一下 3 個點、然後選取移轉



在「總覽」頁面上、您可以看到移轉已成功、且狀態已變更為「成功」。



即時移轉完成後、虛擬機器現在位於不同的節點上。



開啟 Web 主控台並檢視磁碟內容。它仍有我們先前在線上即時移轉之前建立的 2 個檔案。

```

[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 ~]#

```

新節點上的 VM 儲存設備仍顯示相同的磁碟

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、PVC 也是一樣的。

Project: virtual-machines

PersistentVolumeClaims

Filter Name Search by name...

Name	Status	PersistentVolumes	Capacity	Used	StorageClass
fedora-vm1	Bound	pvc-7d00a3cf-d4cc-47d5-8053-ef6b6ae033f	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

與 VM Pod 相關聯的磁碟區也與之前相同（2 個 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

示範影片

透過適用於 NetApp ONTAP 的 Amazon FSX、在 ROSA 上以 OpenShift 虛擬化技術即時移轉虛擬機器

有關 Red Hat OpenShift 和 OpenShift 虛擬化解決方案的更多影片 ["請按這裡"](#)、請參閱。

版權資訊

Copyright © 2024 NetApp, Inc. 版權所有。台灣印製。非經版權所有人事先書面同意，不得將本受版權保護文件的任何部分以任何形式或任何方法（圖形、電子或機械）重製，包括影印、錄影、錄音或儲存至電子檢索系統中。

由 NetApp 版權資料衍伸之軟體必須遵守下列授權和免責聲明：

此軟體以 NETAPP「原樣」提供，不含任何明示或暗示的擔保，包括但不限於有關適售性或特定目的適用性之擔保，特此聲明。於任何情況下，就任何已造成或基於任何理論上責任之直接性、間接性、附隨性、特殊性、懲罰性或衍生性損害（包括但不限於替代商品或服務之採購；使用、資料或利潤上的損失；或企業營運中斷），無論是在使用此軟體時以任何方式所產生的契約、嚴格責任或侵權行為（包括疏忽或其他）等方面，NetApp 概不負責，即使已被告知有前述損害存在之可能性亦然。

NetApp 保留隨時變更本文所述之任何產品的權利，恕不另行通知。NetApp 不承擔因使用本文所述之產品而產生的責任或義務，除非明確經過 NetApp 書面同意。使用或購買此產品並不會在依據任何專利權、商標權或任何其他 NetApp 智慧財產權的情況下轉讓授權。

本手冊所述之產品受到一項（含）以上的美國專利、國外專利或申請中專利所保障。

有限權利說明：政府機關的使用、複製或公開揭露須受 DFARS 252.227-7013（2014 年 2 月）和 FAR 52.227-19（2007 年 12 月）中的「技術資料權利 - 非商業項目」條款 (b)(3) 小段所述之限制。

此處所含屬於商業產品和 / 或商業服務（如 FAR 2.101 所定義）的資料均為 NetApp, Inc. 所有。根據本協議提供的所有 NetApp 技術資料和電腦軟體皆屬於商業性質，並且完全由私人出資開發。美國政府對於該資料具有非專屬、非轉讓、非轉授權、全球性、有限且不可撤銷的使用權限，僅限於美國政府為傳輸此資料所訂合約所允許之範圍，並基於履行該合約之目的方可使用。除非本文另有規定，否則未經 NetApp Inc. 事前書面許可，不得逕行使用、揭露、重製、修改、履行或展示該資料。美國政府授予國防部之許可權利，僅適用於 DFARS 條款 252.227-7015(b)（2014 年 2 月）所述權利。

商標資訊

NETAPP、NETAPP 標誌及 <http://www.netapp.com/TM> 所列之標章均為 NetApp, Inc. 的商標。文中所涉及的所有其他公司或產品名稱，均為其各自所有者的商標，不得侵犯。